

Data-Engineer-Associate Reliable Braindumps Questions & Latest Data-Engineer-Associate Test Vce

Appian ACD100 Appian Certified Associate Developer 2

and efficiency when preparing for ACD100 Exams, thus inspiring them obtain the targeted ACD100 certificate successfully. There are many advantages of our ACD100 question torrent that we are happy to introduce you and you can pass the exam for sure.

Appian Certified Associate Developer Sample Questions (Q16-Q21):

NEW QUESTION # 16
You want to retrieve data from the database to show on your form. Which option should you use?

- A. `alqueryColumn()`
- B. `alquerySelection()`
- C. `alquery()`
- D. `alqueryEntity()`

Answer: D

NEW QUESTION # 17
When designing a new interface, you have to create a table populated with record data. The table needs to link to the relevant record. With component should you use?

- A. Editable Grid
- B. Rich Text Display
- C. Text Layout
- D. Read-Only Grid

Answer: D

NEW QUESTION # 18
A form has 5 rule inputs, as follows:
1 CDT variable
3 Text variables
1 Integer variable
According to best practices, how many process variables are required in the process model's user input task?

- A. 0
- B. 1
- C. 2
- D. 3

Answer: A

NEW QUESTION # 19

New ACD100 Reliable Braindumps Ppt Free PDF Valid ACD100 Sample Questions Answers: Appian Certified Associate Developer

For candidates who want to obtain the certification for Data-Engineer-Associate exam, passing the exam is necessary. We will help you pass the exam just one time. Data-Engineer-Associate training materials are high-quality, since we have experienced experts who are quite familiar with exam center to compile and verify the exam dumps. In addition, we offer you free update for 365 days after payment, and the latest version for Data-Engineer-Associate Training Materials will be sent to your email automatically. We have online and offline chat service and if you have any questions for Data-Engineer-Associate exam materials, you can have a chat with us.

About Amazon Data-Engineer-Associate Exam, each candidate is very confused. Everyone has their own different ideas. But the same idea is that this is a very difficult exam. We are all aware of Amazon Data-Engineer-Associate exam is a difficult exam. But as long as we believe EduDump, this will not be a problem. EduDump's Amazon Data-Engineer-Associate exam training materials is an essential product for each candidate. It is tailor-made for the candidates who will participate in the exam. You will absolutely pass the exam. If you do not believe, then take a look into the website of EduDump. You will be surprised, because its daily purchase rate is the highest. Do not miss it, and add to your shoppingcart quickly.

>> Data-Engineer-Associate Reliable Braindumps Questions <<

Latest Amazon Data-Engineer-Associate Test Vce - Data-Engineer-Associate Test Dumps Demo

If you want to pass the exam quickly, Data-Engineer-Associate prep guide is your best choice. We know that many users do not have a large amount of time to learn. In response to this, we have scientifically set the content of the data. You can use your piecemeal time to learn, and every minute will have a good effect. In order for you to really absorb the content of Data-Engineer-Associate Exam Questions, we will tailor a learning plan for you. This study plan may also have a great impact on your work and life. As long as you carefully study the Data-Engineer-Associate study guide for twenty to thirty hours, you can go to the Data-Engineer-Associate exam.

Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q102-Q107):

NEW QUESTION # 102

A company uses Amazon RDS to store transactional data. The company runs an RDS DB instance in a private subnet. A developer wrote an AWS Lambda function with default settings to insert, update, or delete data in the DB instance.

The developer needs to give the Lambda function the ability to connect to the DB instance privately without using the public internet. Which combination of steps will meet this requirement with the LEAST operational overhead? (Choose two.)

- A. Update the security group of the DB instance to allow only Lambda function invocations on the database port.
- **B. Configure the Lambda function to run in the same subnet that the DB instance uses.**
- **C. Attach the same security group to the Lambda function and the DB instance. Include a self-referencing rule that allows access through the database port.**
- D. Update the network ACL of the private subnet to include a self-referencing rule that allows access through the database port.
- E. Turn on the public access setting for the DB instance.

Answer: B,C

Explanation:

To enable the Lambda function to connect to the RDS DB instance privately without using the public internet, the best combination of steps is to configure the Lambda function to run in the same subnet that the DB instance uses, and attach the same security group to the Lambda function and the DB instance. This way, the Lambda function and the DB instance can communicate within the same private network, and the security group can allow traffic between them on the database port. This solution has the least operational overhead, as it does not require any changes to the public access setting, the network ACL, or the security group of the DB instance.

The other options are not optimal for the following reasons:

* A. Turn on the public access setting for the DB instance. This option is not recommended, as it would expose the DB instance to the public internet, which can compromise the security and privacy of the data. Moreover, this option would not enable the Lambda function to connect to the DB instance privately, as it would still require the Lambda function to use the public internet to access the DB instance.

* B. Update the security group of the DB instance to allow only Lambda function invocations on the database port. This option is not sufficient, as it would only modify the inbound rules of the security group of the DB instance, but not the outbound rules of the security group of the Lambda function.

Moreover, this option would not enable the Lambda function to connect to the DB instance privately, as it would still require the Lambda function to use the public internet to access the DB instance.

* E. Update the network ACL of the private subnet to include a self-referencing rule that allows access through the database port. This option is not necessary, as the network ACL of the private subnet already allows all traffic within the subnet by default.

Moreover, this option would not enable the Lambda function to connect to the DB instance privately, as it would still require the Lambda function to use the public internet to access the DB instance.

1: Connecting to an Amazon RDS DB instance

2: Configuring a Lambda function to access resources in a VPC

3: Working with security groups

4: Network ACLs

NEW QUESTION # 103

The company stores a large volume of customer records in Amazon S3. To comply with regulations, the company must be able to access new customer records immediately for the first 30 days after the records are created. The company accesses records that are older than 30 days infrequently.

The company needs to cost-optimize its Amazon S3 storage.

Which solution will meet these requirements MOST cost-effectively?

- A. Use S3 Standard-Infrequent Access (S3 Standard-IA) storage for all customer records.
- B. Use S3 Intelligent-Tiering storage.
- C. Transition records to S3 Glacier Deep Archive storage after 30 days.
- **D. Apply a lifecycle policy to transition records to S3 Standard Infrequent-Access (S3 Standard-IA) storage after 30 days.**

Answer: D

Explanation:

The most cost-effective solution in this case is to apply a lifecycle policy to transition records to Amazon S3 Standard-IA storage after 30 days. Here's why:

* **Amazon S3 Lifecycle Policies:** Amazon S3 offers lifecycle policies that allow you to automatically transition objects between different storage classes to optimize costs. For data that is frequently accessed in the first 30 days and infrequently accessed after that, transitioning from the S3 Standard storage class to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days makes the most sense. S3 Standard-IA is designed for data that is accessed less frequently but still needs to be retained, offering lower storage costs than S3 Standard with a retrieval cost for access.

* **Cost Optimization:** S3 Standard-IA offers a lower price per GB than S3 Standard. Since the data will be accessed infrequently after 30 days, using S3 Standard-IA will lower storage costs while still allowing for immediate retrieval when necessary.

* **Compliance with Regulations:** Since the records need to be immediately accessible for the first 30 days, the use of S3 Standard for that period ensures compliance with regulatory requirements. After 30 days, transitioning to S3 Standard-IA continues to meet access requirements for infrequent access while reducing storage costs.

* **Alternatives Considered:**

* **Option B (S3 Intelligent-Tiering):** While S3 Intelligent-Tiering automatically moves data between access tiers based on access patterns, it incurs a small monthly monitoring and automation charge per object. It could be a viable option, but transitioning data to S3 Standard-IA directly would be more cost-effective since the pattern of access is well-known (frequent for 30 days, infrequent thereafter).

* **Option C (S3 Glacier Deep Archive):** Glacier Deep Archive is the lowest-cost storage class, but it is not suitable in this case because the data needs to be accessed immediately within 30 days and on an infrequent basis thereafter. Glacier Deep Archive requires hours for data retrieval, which is not acceptable for infrequent access needs.

* **Option D (S3 Standard-IA for all records):** Using S3 Standard-IA for all records would result in higher costs for the first 30 days, as the data is frequently accessed. S3 Standard-IA incurs retrieval charges, making it less suitable for frequently accessed data.

References:

- * Amazon S3 Lifecycle Policies
- * S3 Storage Classes
- * Cost Management and Data Optimization Using Lifecycle Policies
- * AWS Data Engineering Documentation

NEW QUESTION # 104

A company has a production AWS account that runs company workloads. The company's security team created a security AWS account to store and analyze security logs from the production AWS account. The security logs in the production AWS account are stored in Amazon CloudWatch Logs.

The company needs to use Amazon Kinesis Data Streams to deliver the security logs to the security AWS account.

Which solution will meet these requirements?

- A. Create a destination data stream in the production AWS account. In the security AWS account, create an IAM role that has cross-account permissions to Kinesis Data Streams in the production AWS account.
- B. Create a destination data stream in the production AWS account. In the production AWS account, create an IAM role that has cross-account permissions to Kinesis Data Streams in the security AWS account.
- C. Create a destination data stream in the security AWS account. Create an IAM role and a trust policy to grant CloudWatch Logs the permission to put data into the stream. Create a subscription filter in the security AWS account.
- **D. Create a destination data stream in the security AWS account. Create an IAM role and a trust policy to grant CloudWatch Logs the permission to put data into the stream. Create a subscription filter in the production AWS account.**

Answer: D

Explanation:

Amazon Kinesis Data Streams is a service that enables you to collect, process, and analyze real-time streaming data. You can use Kinesis Data Streams to ingest data from various sources, such as Amazon CloudWatch Logs, and deliver it to different destinations, such as Amazon S3 or Amazon Redshift. To use Kinesis Data Streams to deliver the security logs from the production AWS account to the security AWS account, you need to create a destination data stream in the security AWS account. This data stream will receive the log data from the CloudWatch Logs service in the production AWS account. To enable this cross-account

data delivery, you need to create an IAM role and a trust policy in the security AWS account. The IAM role defines the permissions that the CloudWatch Logs service needs to put data into the destination data stream. The trust policy allows the production AWS account to assume the IAM role. Finally, you need to create a subscription filter in the production AWS account. A subscription filter defines the pattern to match log events and the destination to send the matching events. In this case, the destination is the destination data stream in the security AWS account. This solution meets the requirements of using Kinesis Data Streams to deliver the security logs to the security AWS account. The other options are either not possible or not optimal.

You cannot create a destination data stream in the production AWS account, as this would not deliver the data to the security AWS account. You cannot create a subscription filter in the security AWS account, as this would not capture the log events from the production AWS account. References:

* Using Amazon Kinesis Data Streams with Amazon CloudWatch Logs

* AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide, Chapter 3: Data Ingestion and Transformation, Section 3.3: Amazon Kinesis Data Streams

NEW QUESTION # 105

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 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.
- 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 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.

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 # 106

A company needs to generate a one-time performance report by joining data that is stored in Amazon DynamoDB, Amazon RDS, Amazon Redshift, and Amazon S3. The company wants to avoid unnecessary data movement and to minimize query execution time.

Which solution will meet these requirements?

- **A. Use Amazon Athena Federated Query to perform one-time joins and analysis across DynamoDB, Amazon RDS, Amazon Redshift, and Amazon S3.**
- B. Set up an AWS Glue ETL pipeline to extract, transform, and centralize data in Amazon S3. Use Amazon Athena to run analytical queries.
- C. Capture data from DynamoDB by using DynamoDB Streams. Migrate data from Amazon RDS by using AWS DMS. Export Amazon Redshift data. Store all data in Amazon S3. Use Redshift Spectrum to run queries.
- D. Deploy an Amazon EMR cluster powered by Apache Spark to ingest, process, and merge datasets from multiple sources. Run analytical workloads on the merged data.

Answer: A

Explanation:

Amazon Athena Federated Query enables SQL queries across multiple data sources without moving data.

Athena can directly query Amazon DynamoDB, Amazon RDS, Amazon Redshift, and Amazon S3 using data source connectors.

This capability is ideal for one-time or ad hoc analytical queries, where setting up ingestion pipelines would introduce unnecessary cost and delay.

Because the requirement is for a one-time performance report, the solution should minimize operational overhead and data duplication. Athena Federated Query executes queries in place, eliminating ETL processes, data exports, and staging layers. This significantly reduces query preparation time and avoids storage costs associated with copying data into a centralized location. The other options require extensive data movement and setup. DynamoDB Streams, AWS DMS, Glue ETL pipelines, and Amazon EMR clusters are appropriate for ongoing analytics pipelines, not for a single analytical report. These approaches increase operational complexity and execution time.

Athena Federated Query is purpose-built for this exact scenario and provides the fastest path to insight with minimal overhead.

NEW QUESTION # 107

.....

now our Data-Engineer-Associate training materials have become the most popular Data-Engineer-Associate practice materials in the international market. There are so many advantages of our study materials, and will show you some of them for your reference. First and foremost, our company has prepared Data-Engineer-Associate free demo in this website for our customers. Second, it is convenient for you to read and make notes with our PDF version. So let our Data-Engineer-Associate practice materials to be your learning partner in the course of preparing for the Data-Engineer-Associate exam, especially the PDF version is really a wise choice for you.

Latest Data-Engineer-Associate Test Vce: <https://www.edudump.com/exams/Amazon/Data-Engineer-Associate/>

With our Data-Engineer-Associate study materials, you can make full use of those time originally spent in waiting for the delivery of exam files so that you can get preparations as early as possible, Amazon Data-Engineer-Associate Reliable Braindumps Questions Get Complete and Professional premium VCE file, Each version's using method and functions are different but the questions and answers of our Data-Engineer-Associate study quiz is the same, In order to provide the top service on our Data-Engineer-Associate training prep, our customer agents will work 24/7.

By John Ray, Sean Johnson, What percentage of the defects do you find Data-Engineer-Associate in compile, and what percent do you find in the testing, what percent do you find in inspection or a code review, and that sort of thing.

Data-Engineer-Associate Original Questions & Data-Engineer-Associate Training Online & Data-Engineer-Associate Dumps Torrent

With our Data-Engineer-Associate study materials, you can make full use of those time originally spent in waiting for the delivery of exam files so that you can get preparations as early as possible.

Get Complete and Professional premium VCE file, Each version's using method and functions are different but the questions and answers of our Data-Engineer-Associate study quiz is the same.

In order to provide the top service on our Data-Engineer-Associate training prep, our customer agents will work 24/7, Data-Engineer-Associate sure pass torrent is the latest and edited and checked by Data-Engineer-Associate Preparation Store our professional experts, which always can cover all the topics in the actual test.

- Amazon Data-Engineer-Associate Questions Obtain Exam Results Simply 2026 Easily obtain Data-Engineer-Associate for free download through 「 www.prepawayexam.com 」 Latest Data-Engineer-Associate Test Fee
- 100% Pass 2026 Data-Engineer-Associate: AWS Certified Data Engineer - Associate (DEA-C01) Authoritative Reliable Braindumps Questions Easily obtain Data-Engineer-Associate for free download through 「 www.pdfvce.com 」 Data-Engineer-Associate Exam Prep
- Latest Data-Engineer-Associate Exam Questions form the Most Valid Preparation Brain Dumps - www.examcollectionpass.com Search for Data-Engineer-Associate on www.examcollectionpass.com immediately to obtain a free download Test Data-Engineer-Associate Dumps Demo
- Data-Engineer-Associate Braindumps Torrent Test Data-Engineer-Associate Passing Score Data-Engineer-Associate Exam Demo Search for 【 Data-Engineer-Associate 】 and download exam materials for free through www.pdfvce.com Data-Engineer-Associate Latest Exam Online
- Perfect Data-Engineer-Associate – 100% Free Reliable Braindumps Questions | Latest Data-Engineer-Associate Test Vce The page for free download of Data-Engineer-Associate on www.examcollectionpass.com will open immediately Test Data-Engineer-Associate Dumps Demo
- Amazon Data-Engineer-Associate Exam Dumps - Key To Getting Success Easily obtain Data-Engineer-Associate for free download through “ www.pdfvce.com ” Valid Data-Engineer-Associate Test Book
- Test Data-Engineer-Associate Dump Reliable Data-Engineer-Associate Exam Preparation Latest Data-Engineer-Associate Test Guide Search for Data-Engineer-Associate and easily obtain a free download on

