Free PDF Valid Data-Engineer-Associate - AWS Certified Data Engineer - Associate (DEA-C01) Reliable Real Exam

P.S. Free & New Data-Engineer-Associate dumps are available on Google Drive shared by BraindumpsVCE: https://drive.google.com/open?id=1ztL3dwJgxbuiIQBZqTeETlQdZOjqX6Z0

IT exam become more important than ever in today's highly competitive world, these things mean a different future. Amazon Data-Engineer-Associate exam will be a milestone in your career, and may dig into new opportunities, but how do you pass Amazon Data-Engineer-Associate Exam? Do not worry, help is at hand, with BraindumpsVCE you no longer need to be afraid. BraindumpsVCE Amazon Data-Engineer-Associate exam questions and answers is the pioneer in exam preparation.

If you are still troubled for the Amazon Data-Engineer-Associate Certification Exam, then select the BraindumpsVCE's training materials please. BraindumpsVCE's Amazon Data-Engineer-Associate exam training materials is the best training materials, this is not doubt. Select it will be your best choice. It can guarantee you 100% pass the exam. Come on, you will be the next best IT experts.

>> Data-Engineer-Associate Reliable Real Exam <<

Mock Data-Engineer-Associate Exam & New Data-Engineer-Associate Exam Discount

Our Amazon Data-Engineer-Associate PDF dumps format has actual Data-Engineer-Associate questions which are printable and portable. Hence, you can go through these Amazon Data-Engineer-Associate questions via your smart devices like smartphones, laptops, and tablets. The AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) dumps PDF file can be used from any location and at any time. Furthermore, you can take print of Amazon Questions PDF to do an off-screen study.

Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q86-Q91):

NEW QUESTION # 86

A security company stores IoT data that is in JSON format in an Amazon S3 bucket. The data structure can change when the

company upgrades the IoT devices. The company wants to create a data catalog that includes the IoT data. The company's analytics department will use the data catalog to index the data.
Which solution will meet these requirements MOST cost-effectively?

- A. Create an Amazon Athena workgroup. Explore the data that is in Amazon S3 by using Apache Spark through Athena. Provide the Athena workgroup schema and tables to the analytics department.
- B. Create an Amazon Redshift provisioned cluster. Create an Amazon Redshift Spectrum database for the analytics department to explore the data that is in Amazon S3. Create Redshift stored procedures to load the data into Amazon Redshift.
- C. Create an AWS Glue Data Catalog. Configure an AWS Glue Schema Registry. Create AWS Lambda user defined functions (UDFs) by using the Amazon Redshift Data API. Create an AWS Step Functions job to orchestrate the ingestion of the data that the analytics department will use into Amazon Redshift Serverless.
- D. Create an AWS Glue Data Catalog. Configure an AWS Glue Schema Registry. Create a new AWS Glue workload to orchestrate the ingestion of the data that the analytics department will use into Amazon Redshift Serverless.

Answer: A

Explanation:
The best solution to meet the requirements of creating a data catalog that includes the IoT data, and allowing the analytics department to index the data, most cost-effectively, is to create an Amazon Athena workgroup, explore the data that is in Amazon S3 by using Apache Spark through Athena, and provide the Athena workgroup schema and tables to the analytics department.
Amazon Athena is a serverless, interactive query service that makes it easy to analyze data directly in Amazon S3 using standard SQL or Python1. Amazon Athena also supports Apache Spark, an open-source distributed processing framework that can run large-scale data analytics applications across clusters of servers2. You can use Athena to run Spark code on data in Amazon S3 without having to set up, manage, or scale any infrastructure. You can also use Athena to create and manage external tables that point to your data in Amazon S3, and store them in an external data catalog, such as AWS Glue Data Catalog, Amazon Athena Data Catalog, or your own Apache Hive metastore3. You can create Athena workgroups to separate query execution and resource allocation based on different criteria, such as users, teams, or applications4. You can share the schemas and tables in your Athena workgroup with other users or applications, such as Amazon QuickSight, for data visualization and analysis5.
Using Athena and Spark to create a data catalog and explore the IoT data in Amazon S3 is the most cost-effective solution, as you pay only for the queries you run or the compute you use, and you pay nothing when the service is idle1. You also save on the operational overhead and complexity of managing data warehouse infrastructure, as Athena and Spark are serverless and scalable. You can also benefit from the flexibility and performance of Athena and Spark, as they support various data formats, including JSON, and can handle schema changes and complex queries efficiently.
Option A is not the best solution, as creating an AWS Glue Data Catalog, configuring an AWS Glue Schema Registry, creating a new AWS Glue workload to orchestrate theingestion of the data that the analytics department will use into Amazon Redshift Serverless, would incur more costs and complexity than using Athena and Spark. AWS Glue Data Catalog is a persistent metadata store that contains table definitions, job definitions, and other control information to help you manage your AWS Glue components6. AWS Glue Schema Registry is a service that allows you to centrally store and manage the schemas of your streaming data in AWS Glue Data Catalog7. AWS Glue is a serverless data integration service that makes it easy to prepare, clean, enrich, and move data between data stores8. Amazon Redshift Serverless is a feature of Amazon Redshift, a fully managed data warehouse service, that allows you to run and scale analytics without having to manage data warehouse infrastructure9. While these services are powerful and useful for many data engineering scenarios, they are not necessary or cost-effective for creating a data catalog and indexing the IoT data in Amazon S3. AWS Glue Data Catalog and Schema Registry charge you based on the number of objects stored and the number of requests made67. AWS Glue charges you based on the compute time and the data processed by your ETL jobs8. Amazon Redshift Serverless charges you based on the amount of data scanned by your queries and the compute time used by your workloads9. These costs can add up quickly, especially if you have large volumes of IoT data and frequent schema changes. Moreover, using AWS Glue and Amazon Redshift Serverless would introduce additional latency and complexity, as you would have to ingest the data from Amazon S3 to Amazon Redshift Serverless, and then query it from there, instead of querying it directly from Amazon S3 using Athena and Spark.
Option B is not the best solution, as creating an Amazon Redshift provisioned cluster, creating an Amazon Redshift Spectrum database for the analytics department to explore the data that is in Amazon S3, and creating Redshift stored procedures to load the data into Amazon Redshift, would incur more costs and complexity than using Athena and Spark. Amazon Redshift provisioned clusters are clusters that you create and manage by specifying the number and type of nodes, and the amount of storage and compute capacity10. Amazon Redshift Spectrum is a feature of Amazon Redshift that allows you to query and join data across your data warehouse and your data lake using standard SQL11. Redshift stored procedures are SQL statements that you can define and store in Amazon Redshift, and then call them by using the CALL command12. While these features are powerful and useful for many data warehousing scenarios, they are not necessary or cost-effective for creating a data catalog and indexing the IoT data in Amazon S3. Amazon Redshift provisioned clusters charge you based on the node type, the number of nodes, and the duration of the cluster10. Amazon Redshift Spectrum charges you based on the amount of data scanned by your queries11. These costs can add up quickly, especially if you have large volumes of IoT data and frequent schema changes. Moreover, using Amazon Redshift

provisioned clusters and Spectrum would introduce additional latency and complexity, as you would have to provision andmanage the cluster, create an external schema and database for the data in Amazon S3, and load the data into the cluster using stored procedures, instead of querying it directly from Amazon S3 using Athena and Spark.
Option D is not the best solution, as creating an AWS Glue Data Catalog, configuring an AWS Glue Schema Registry, creating AWS Lambda user defined functions (UDFs) by using the Amazon Redshift Data API, and creating an AWS Step Functions job to orchestrate the ingestion of the data that the analytics department will use into Amazon Redshift Serverless, would incur more costs and complexity than using Athena and Spark. AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers13. AWS Lambda UDFs are Lambda functions that you can invoke from within an Amazon Redshift query. Amazon Redshift Data API is a service that allows you to run SQL statements on Amazon Redshift clusters using HTTP requests, without needing a persistent connection. AWS Step Functions is a service that lets you coordinate multiple AWS services into serverless workflows. While these services are powerful and useful for many data engineering scenarios, they are not necessary or cost-effective for creating a data catalog and indexing the IoT data in Amazon S3. AWS Glue Data Catalog and Schema Registry charge you based on the number of objects stored and the number of requests made67. AWS Lambda charges you based on the number of requests and the duration of your functions13. Amazon Redshift Serverless charges you based on the amount of data scanned by your queries and the compute time used by your workloads9. AWS Step Functions charges you based on the number of state transitions in your workflows. These costs can add up quickly, especially if you have large volumes of IoT data and frequent schema changes. Moreover, using AWS Glue, AWS Lambda, Amazon Redshift Data API, and AWS Step Functions would introduce additional latency and complexity, as you would have to create and invoke Lambda functions to ingest the data from Amazon S3 to Amazon Redshift Serverless using the Data API, and coordinate the ingestion process using Step Functions, instead of querying it directly from Amazon S3 using Athena and Spark. References:
What is Amazon Athena?
Apache Spark on Amazon Athena
Creating tables, updating the schema, and adding new partitions in the Data Catalog from AWS Glue ETL jobs Managing Athena workgroups Using Amazon QuickSight to visualize data in Amazon Athena AWS Glue Data Catalog AWS Glue Schema Registry
What is AWS Glue?
Amazon Redshift Serverless
Amazon Redshift provisioned clusters
Querying external data using Amazon Redshift Spectrum
Using stored procedures in Amazon Redshift
What is AWS Lambda?
[Creating and using AWS Lambda UDFs]
[Using the Amazon Redshift Data API]
[What is AWS Step Functions?]
AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide

NEW QUESTION # 87
A company is planning to upgrade its Amazon Elastic Block Store (Amazon EBS) General Purpose SSD storage from gp2 to gp3. The company wants to prevent any interruptions in its Amazon EC2 instances that will cause data loss during the migration to the upgraded storage.
Which solution will meet these requirements with the LEAST operational overhead?

- A. Change the volume type of the existing gp2 volumes to gp3. Enter new values for volume size, IOPS, and throughput.
- B. Create snapshots of the gp2 volumes. Create new gp3 volumes from the snapshots. Attach the new gp3 volumes to the EC2 instances.
- C. Create new gp3 volumes. Gradually transfer the data to the new gp3 volumes. When the transfer is complete, mount the new gp3 volumes to the EC2 instances to replace the gp2 volumes.
- D. Use AWS DataSync to create new gp3 volumes. Transfer the data from the original gp2 volumes to the new gp3 volumes.

Answer: A

Explanation:
Changing the volume type of the existing gp2 volumes to gp3 is the easiest and fastest way to migrate to the new storage type without any downtime or data loss. You can use the AWS Management Console, the AWS CLI, or the Amazon EC2 API to modify the volume type, size, IOPS, and throughput of your gp2 volumes.
The modification takes effect immediately, and you can monitor the progress of the modification using CloudWatch. The other options are either more complex or require additional steps, such as creating snapshots, transferring data, or attaching new volumes, which can increase the operational overhead and the risk of errors. References:
* Migrating Amazon EBS volumes from gp2 to gp3 and save up to 20% on costs (Section: How to migrate from gp2 to gp3)
* Switching from gp2 Volumes to gp3 Volumes to Lower AWS EBS Costs (Section: How to Switch from GP2 Volumes to GP3 Volumes)

* Modifying the volume type, IOPS, or size of an EBS volume - Amazon Elastic Compute Cloud (Section: Modifying the volume type)

NEW QUESTION # 88

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

Answer: C,D

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.
References:
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 # 89

A company uses Amazon S3 to store data and Amazon QuickSight to create visualizations.
The company has an S3 bucket in an AWS account named Hub-Account. The S3 bucket is encrypted by an AWS Key Management Service (AWS KMS) key. The company's QuickSight instance is in a separate account named BI-Account The company updates the S3 bucket policy to grant access to the QuickSight service role. The company wants to enable cross-account access to allow QuickSight to interact with the S3 bucket.
Which combination of steps will meet this requirement? (Select TWO.)

- A. Add an IAM policy to the QuickSight service role to give QuickSight access to the KMS key that encrypts the S3 bucket.
- B. Use the existing AWS KMS key to encrypt connections from QuickSight to the S3 bucket.
- C. Add the 53 bucket as a resource that the QuickSight service role can access.
- D. Use AWS Resource Access Manager (AWS RAM) to share the S3 bucket with the BI-Account account.

- E. Add the KMS key as a resource that the QuickSight service role can access.

Answer: A,E

Explanation:
Problem Analysis:
The company needs cross-account access to allow QuickSight in BI-Account to interact with an S3 bucket in Hub-Account.
The bucket is encrypted with an AWS KMS key.
Appropriate permissions must be set for both S3 access and KMS decryption.
Key Considerations:
QuickSight requires IAM permissions to access S3 data and decrypt files using the KMS key.
Both S3 and KMS permissions need to be properly configured across accounts.
Solution Analysis:
Option A: Use Existing KMS Key for Encryption
While the existing KMS key is used for encryption, it must also grant decryption permissions to QuickSight.
Option B: Add S3 Bucket to QuickSight Role
Granting S3 bucket access to the QuickSight service role is necessary for cross-account access.
Option C: AWS RAM for Bucket Sharing
AWS RAM is not required; bucket policies and IAM roles suffice for granting cross-account access.
Option D: IAM Policy for KMS Access
QuickSight's service role in BI-Account needs explicit permissions to use the KMS key for decryption.
Option E: Add KMS Key as Resource for Role
The KMS key must explicitly list the QuickSight role as an entity that can access it.
Implementation Steps:
S3 Bucket Policy in Hub-Account:
Add a policy to the S3 bucket granting the QuickSight service role access:
json

```
{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Allow",
"Principal": { "AWS": "arn:aws:iam::<BI-Account-ID>:role/service-role/QuickSightRole" },
"Action": "s3:GetObject",
"Resource": "arn:aws:s3:::<Bucket-Name>/*"
}
]
}
```

KMS Key Policy in Hub-Account:
Add permissions for the QuickSight role:

```
{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Allow",
"Principal": { "AWS": "arn:aws:iam::<BI-Account-ID>:role/service-role/QuickSightRole" },
"Action": [
"kms:Decrypt",
"kms:DescribeKey"
],
"Resource": "*"
}
]
}
```

IAM Policy for QuickSight Role in BI-Account:
Attach the following policy to the QuickSight service role:

```
{
"Version": "2012-10-17",
"Statement": [
{
"Effect": "Allow",
```

```
"Action": [
"s3:GetObject",
"kms:Decrypt"
],
"Resource": [
"arn:aws:s3:::<Bucket-Name>/*",
"arn:aws:kms:<region>:<Hub-Account-ID>:key/<KMS-Key-ID>"
]
}
]
}
```

Reference:
Setting Up Cross-Account S3 Access
AWS KMS Key Policy Examples
Amazon QuickSight Cross-Account Access

NEW QUESTION # 90
A company needs to store semi-structured transactional data for an application in a database. The database must be serverless. The application writes the data infrequently, but it reads the data frequently. The application must retrieve the data within milliseconds.
Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: C

Explanation:
Option D is correct because Amazon DynamoDB is a serverless NoSQL database that is well suited for semi- structured data, and DynamoDB Accelerator (DAX) is designed to improve read performance for read-heavy applications. AWS documentation states that DAX can improve DynamoDB performance from milliseconds to microseconds, and the DynamoDB guide says DAX is most effective for applications with a high read-to- write ratio, which exactly matches this workload of infrequent writes and frequent reads.
Option A and B are incorrect because Amazon S3 and Iceberg are object storage and table-format solutions, not low-latency transactional databases for millisecond retrieval. Option C is not serverless and requires managing a database cluster, even if optimized reads improve performance. By contrast, DynamoDB is fully managed and serverless, and DAX adds the lowest-latency read path for hot data. The uploaded study guide also identifies DynamoDB as the right fit for dynamic or semi-structured data that requires low-latency access at scale.

NEW QUESTION # 91
......

Relationships and Service Social media and ubiquitous mobile technologies will become even more pervasive, Data-Engineer-Associate Reliable Real Exam changing the way accounting and tax professionals conduct business and attract new clients.

Data-Engineer-Associate – 100% Free Reliable Real Exam | Reliable Mock AWS Certified Data Engineer - Associate (DEA-C01) Exam

The exam comes with the features, which enables you to Pass Data-Engineer-Associate Exam with good grades, In cases where payment is made other than credit card it may take up to 24 Hours to activate order.

You do not need to spend money; because our Data-Engineer-Associate test questions provide you with the demo for free, Don't worry about how difficult the exam will be, our Data-Engineer-Associate preparation labs will help you clear exam easily.

They are perfect Data-Engineer-Associate pass-sure torrent for you without defects.

P.S. Free 2026 Amazon Data-Engineer-Associate dumps are available on Google Drive shared by BraindumpsVCE: https://drive.google.com/open?id=1ztL3dwJgxbuiIQBZqTeETlQdZOjqX6Z0