

Free PDF Quiz Amazon - Data-Engineer-Associate-Valid Reliable Dump

exam and clear it with a brilliant score. You can easily get these AWS Certified SysOps Administrator - Associate (SOA-C02) (SOA-C02) exam dumps from VCETorrent that are helping candidates achieve their goals. As a working person, the Amazon SOA-C02 Practice Exam will be a great help because you are left with little time to prepare for the SOA-C02 certification exam which you cannot waste to make time for the SOA-C02 exam questions.

The Amazon SOA-C02 exam consists of 65 multiple-choice and multiple-response questions, and candidates have 130 minutes to complete it. The passing score for the exam is 720 out of 1000. SOA-C02 exam fee is \$150 USD, and it can be taken at a testing center or online through Pearson VUE.

Amazon AWS Certified SysOps Administrator - Associate (SOA-C02) Sample Questions (Q325-Q330):

NEW QUESTION # 325

A company hosts its website on Amazon EC2 instances behind an Application Load Balancer. The company manages its DNS with Amazon Route 53, and wants to point its domain's zone apex to the website.

Which type of record should be used to meet these requirements?

- A. An AAAA record for the domain's zone apex
- B. A CNAME record for the domain's zone apex
- C. An A record for the domain's zone apex
- D. An alias record for the domain's zone apex

Answer: D

Explanation:

Reference:

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resource-record-sets-choosing-alias-non-alias.html>

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-elb-load-balancer.html>

NEW QUESTION # 326

A company has set up an IPsec tunnel between its AWS environment and its on-premises data center. The tunnel is reporting as UP, but the Amazon EC2 instances are not able to ping any on-premises resources.

What should a SysOps administrator do to resolve this issue?

- A. Create a new inbound rule on the EC2 instances' security groups to allow ICMP traffic from the on-premises CIDR.
- B. Modify the VPC's DHCP options set. Add the IPsec tunnel to the VPN section.
- C. Create a peering connection between the IPsec tunnel and the subnet of the EC2 instances.
- D. Enable route propagation for the virtual private gateway in the route table that is assigned to the subnet of the EC2 instances.

Answer: A

Explanation:

https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html

DOWNLOAD the newest Exam4Labs Data-Engineer-Associate PDF dumps from Cloud Storage for free:
<https://drive.google.com/open?id=1FCJURvYsiMnsSMA5cdut-e93pudli-S3>

We all know that Data-Engineer-Associate learning guide can help us solve learning problems. But if it is too complex, not only can't we get good results, but also the burden of students' learning process will increase largely. Unlike those complex and esoteric materials, our Data-Engineer-Associate Preparation prep is not only of high quality, but also easy to learn. For our professional experts simplified the content of the Data-Engineer-Associate exam questions for all our customers to be understood.

As we all know, respect and power is gained through knowledge or skill. The society will never welcome lazy people. Do not satisfy what you have owned. Challenge some fresh and meaningful things, and when you complete Data-Engineer-Associate exam, you will find you have reached a broader place where you have never reach. There must be one that suits you best. Your life will become more meaningful because of your new change, and our Data-Engineer-Associate question torrents will be your first step.

>> **Data-Engineer-Associate Reliable Dump** <<

Amazon Data-Engineer-Associate Test Certification Cost - Data-Engineer-Associate Test Simulator

Our website experts simplify complex concepts of the Data-Engineer-Associate exam questions and add examples, simulations, and

diagrams to explain anything that might be difficult to understand. Therefore, even ordinary examiners can master all the Data-Engineer-Associate learning materials without difficulty. And the price of our Data-Engineer-Associate Study Guide is reasonable for even the students can afford it. At the same time, we give some discounts from time to time, you can buy our Data-Engineer-Associate practice engine at a favorable price.

Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q14-Q19):

NEW QUESTION # 14

A company receives call logs as Amazon S3 objects that contain sensitive customer information. The company must protect the S3 objects by using encryption. The company must also use encryption keys that only specific employees can access. Which solution will meet these requirements with the LEAST effort?

- A. Use server-side encryption with customer-provided keys (SSE-C) to encrypt the objects that contain customer information. Restrict access to the keys that encrypt the objects.
- B. Use an AWS CloudHSM cluster to store the encryption keys. Configure the process that writes to Amazon S3 to make calls to CloudHSM to encrypt and decrypt the objects. Deploy an IAM policy that restricts access to the CloudHSM cluster.
- C. Use server-side encryption with Amazon S3 managed keys (SSE-S3) to encrypt the objects that contain customer information. Configure an IAM policy that restricts access to the Amazon S3 managed keys that encrypt the objects.
- **D. Use server-side encryption with AWS KMS keys (SSE-KMS) to encrypt the objects that contain customer information. Configure an IAM policy that restricts access to the KMS keys that encrypt the objects.**

Answer: D

Explanation:

Option C is the best solution to meet the requirements with the least effort because server-side encryption with AWS KMS keys (SSE-KMS) is a feature that allows you to encrypt data at rest in Amazon S3 using keys managed by AWS Key Management Service (AWS KMS). AWS KMS is a fully managed service that enables you to create and manage encryption keys for your AWS services and applications. AWS KMS also allows you to define granular access policies for your keys, such as who can use them to encrypt and decrypt data, and under what conditions. By using SSE-KMS, you can protect your S3 objects by using encryption keys that only specific employees can access, without having to manage the encryption and decryption process yourself.

Option A is not a good solution because it involves using AWS CloudHSM, which is a service that provides hardware security modules (HSMs) in the AWS Cloud. AWS CloudHSM allows you to generate and use your own encryption keys on dedicated hardware that is compliant with various standards and regulations. However, AWS CloudHSM is not a fully managed service and requires more effort to set up and maintain than AWS KMS. Moreover, AWS CloudHSM does not integrate with Amazon S3, so you have to configure the process that writes to S3 to make calls to CloudHSM to encrypt and decrypt the objects, which adds complexity and latency to the data protection process.

Option B is not a good solution because it involves using server-side encryption with customer-provided keys (SSE-C), which is a feature that allows you to encrypt data at rest in Amazon S3 using keys that you provide and manage yourself. SSE-C requires you to send your encryption key along with each request to upload or retrieve an object. However, SSE-C does not provide any mechanism to restrict access to the keys that encrypt the objects, so you have to implement your own key management and access control system, which adds more effort and risk to the data protection process.

Option D is not a good solution because it involves using server-side encryption with Amazon S3 managed keys (SSE-S3), which is a feature that allows you to encrypt data at rest in Amazon S3 using keys that are managed by Amazon S3. SSE-S3 automatically encrypts and decrypts your objects as they are uploaded and downloaded from S3. However, SSE-S3 does not allow you to control who can access the encryption keys or under what conditions. SSE-S3 uses a single encryption key for each S3 bucket, which is shared by all users who have access to the bucket. This means that you cannot restrict access to the keys that encrypt the objects by specific employees, which does not meet the requirements.

Reference:

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

Protecting Data Using Server-Side Encryption with AWS KMS-Managed Encryption Keys (SSE-KMS) - Amazon Simple Storage Service
What is AWS Key Management Service? - AWS Key Management Service
What is AWS CloudHSM? - AWS CloudHSM
Protecting Data Using Server-Side Encryption with Customer-Provided Encryption Keys (SSE-C) - Amazon Simple Storage Service
Protecting Data Using Server-Side Encryption with Amazon S3-Managed Encryption Keys (SSE-S3) - Amazon Simple Storage Service

NEW QUESTION # 15

An ecommerce company collects daily customer transaction logs in CSV format and stores the logs in Amazon S3. The company uses Amazon Athena to scan a subset of attributes from the logs on the same day the company receives each log.

Query times are increasing because of increasing transaction volume. The company wants to improve query performance. Which solution will meet these requirements with the SHORTEST query times?

- A. Convert the CSV logs to JSON. Partition by date in Amazon S3. Use Athena with dynamic filtering to reduce data scans.
- **B. Convert the CSV logs to a single Apache Parquet file for each day. Partition the data by date in Amazon S3. Use Athena with predicate pushdown filters.**
- C. Convert the CSV logs to Avro. Partition by date in Amazon S3. Use Athena with projection-based partitioning.
- D. Convert the CSV logs into multiple ORC files for better parallelism in Athena. Partition by date in Amazon S3. Use columnar pushdown filters.

Answer: B

Explanation:

Amazon Athena achieves the fastest query performance when data is stored in columnar formats such as Apache Parquet and when queries can take advantage of partition pruning and predicate pushdown.

Converting CSV files to Parquet significantly reduces the amount of data scanned because Parquet stores data in a column-oriented layout. Since Athena queries only a subset of attributes, it reads only the required columns instead of scanning entire rows, which dramatically improves performance. Predicate pushdown further reduces query time by filtering data at the storage layer.

Partitioning the data by date ensures that Athena scans only the relevant partitions for same-day queries, minimizing unnecessary data reads. Storing one Parquet file per day is efficient and avoids the overhead of managing excessive small files.

ORC is also a columnar format, but Parquet is more commonly optimized and recommended for Athena workloads in AWS exam guidance. JSON and Avro are row-based or semi-row-based formats and result in larger scan sizes and slower query execution. Therefore, Option D provides the shortest query times and aligns with Athena performance best practices.

NEW QUESTION # 16

A data engineer must orchestrate a series of Amazon Athena queries that will run every day. Each query can run for more than 15 minutes.

Which combination of steps will meet these requirements MOST cost-effectively? (Choose two.)

- A. Use an AWS Glue Python shell script to run a sleep timer that checks every 5 minutes to determine whether the current Athena query has finished running successfully. Configure the Python shell script to invoke the next query when the current query has finished running.
- **B. Create an AWS Step Functions workflow and add two states. Add the first state before the Lambda function. Configure the second state as a Wait state to periodically check whether the Athena query has finished using the Athena Boto3 `get_query_execution` API call. Configure the workflow to invoke the next query when the current query has finished running.**
- C. Use an AWS Glue Python shell job and the Athena Boto3 client `start_query_execution` API call to invoke the Athena queries programmatically.
- **D. Use an AWS Lambda function and the Athena Boto3 client `start_query_execution` API call to invoke the Athena queries programmatically.**
- E. Use Amazon Managed Workflows for Apache Airflow (Amazon MWAA) to orchestrate the Athena queries in AWS Batch.

Answer: B,D

Explanation:

Option A and B are the correct answers because they meet the requirements most cost-effectively. Using an AWS Lambda function and the Athena Boto3 client `start_query_execution` API call to invoke the Athena queries programmatically is a simple and scalable way to orchestrate the queries. Creating an AWS Step Functions workflow and adding two states to check the query status and invoke the next query is a reliable and efficient way to handle the long-running queries.

Option C is incorrect because using an AWS Glue Python shell job to invoke the Athena queries programmatically is more expensive than using a Lambda function, as it requires provisioning and running a Glue job for each query.

Option D is incorrect because using an AWS Glue Python shell script to run a sleep timer that checks every 5 minutes to determine whether the current Athena query has finished running successfully is not a cost-effective or reliable way to orchestrate the queries, as it wastes resources and time.

Option E is incorrect because using Amazon Managed Workflows for Apache Airflow (Amazon MWAA) to orchestrate the Athena queries in AWS Batch is an overkill solution that introduces unnecessary complexity and cost, as it requires setting up and managing an Airflow environment and an AWS Batch compute environment.

References:

* AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide, Chapter 5: Data Orchestration, Section 5.2: AWS Lambda, Section 5.3: AWS Step Functions, Pages 125-135

- * Building Batch Data Analytics Solutions on AWS, Module 5: Data Orchestration, Lesson 5.1: AWS Lambda, Lesson 5.2: AWS Step Functions, Pages 1-15
- * AWS Documentation Overview, AWS Lambda Developer Guide, Working with AWS Lambda Functions, Configuring Function Triggers, Using AWS Lambda with Amazon Athena, Pages 1-4
- * AWS Documentation Overview, AWS Step Functions Developer Guide, Getting Started, Tutorial: Create a Hello World Workflow, Pages 1-8

NEW QUESTION # 17

A company runs concurrent analytical queries on Amazon Redshift tables multiple times each day. The queries require consistent data views three times each day. The company runs extract, transform, and load (ETL) operations that update dimension tables while the queries run. The company has noticed that the queries cause table-level locks during the ETL operations. The company's current solution experiences query timeouts and deadlocks during peak processing hours, which affects analytical reporting and on-demand analysis.

Which solution will fix this issue?

- A. Use Amazon Redshift Spectrum to query data in Amazon S3 for analytical workloads. Maintain ETL operations on Amazon Redshift tables with transaction isolation.
- **B. Deploy separate Amazon Redshift clusters for ETL and analytics workloads. Use cross-database queries and data sharing to maintain data consistency.**
- C. Use Amazon Redshift materialized views for analytical queries. Schedule ETL operations during off-peak hours to minimize lock contention.
- D. Configure Amazon Redshift federated queries to access source data directly. Use read replicas to isolate analytical workloads from ETL operations.

Answer: B

Explanation:

Option D is correct because the problem is lock contention between ETL writes and analytical reads on the same Redshift tables. AWS documents that table locks can block both reads and writes, and that Redshift data sharing lets you share live data across Redshift clusters and workgroups without manually moving or copying the data. Because the shared data is live, consumer users see the most up-to-date information as soon as it is updated. Separating ETL and analytics onto different Redshift clusters while using data sharing is therefore the cleanest way to isolate workloads and still provide consistent data views.

Option A reduces contention only by changing schedules and does not fully solve concurrent peak-hour conflicts. Option B is not correct because Redshift federated queries are read-only against external sources and do not provide a Redshift "read replica" design for this case. Option C does not address the locking problem on the dimension tables that ETL is updating. The official Redshift guidance and service capabilities point to workload isolation plus data sharing as the best fix. This also matches the exam guide's emphasis on choosing architectures that maintain performance and consistency without unnecessary copying.

NEW QUESTION # 18

An ecommerce company stores sales data in an AWS Glue table named sales_data. The company stores the sales_data table in an Amazon S3 Standard bucket. The table contains columns named order_id, customer_id, product_id, order_date, shipping_date, and order_amount.

The company wants to improve query performance by partitioning the sales_data table by order_date. The company needs to add the partition to the existing sales_data table in AWS Glue.

Which solution will meet these requirements?

- A. Manually modify the S3 bucket to use the new partition.
- B. Update the AWS Glue table's schema to include the new partition.
- C. Edit the AWS Glue table's metadata file directly in Amazon S3.
- **D. Use the AWS Glue Data Catalog API to add the new partition to the table.**

Answer: D

Explanation:

In AWS Glue, table partitions are managed as metadata objects within the AWS Glue Data Catalog. To add a new partition to an existing table, the correct and supported approach is to use the AWS Glue Data Catalog API, such as the CreatePartition operation, or equivalent console or SDK actions.

Updating the table schema alone does not create partitions or inform Glue about new partition values. Editing metadata files directly in Amazon S3 is unsupported and can corrupt the Data Catalog. Manually modifying the S3 bucket structure without registering

partitions in Glue will result in Athena and other query engines being unable to recognize the partitions. By adding partitions through the Glue Data Catalog API, query engines such as Amazon Athena and Amazon Redshift Spectrum can perform partition pruning, which significantly improves query performance by scanning only relevant data. This method aligns with AWS best practices, ensures metadata consistency, and avoids unnecessary operational risk. Therefore, Option C is the correct solution.

NEW QUESTION # 19

.....

In order to make all customers feel comfortable, our company will promise that we will offer the perfect and considerate service for all customers. If you buy the Data-Engineer-Associate study materials from our company, you will have the right to enjoy the perfect service. We have employed a lot of online workers to help all customers solve their problem. If you have any questions about the Data-Engineer-Associate Study Materials, do not hesitate and ask us in your anytime, we are glad to answer your questions and help you use our Data-Engineer-Associate study materials well.

Data-Engineer-Associate Test Certification Cost: <https://www.exam4labs.com/Data-Engineer-Associate-practice-torrent.html>

Amazon Data-Engineer-Associate Reliable Dump Everything you need to prepare and quickly pass the tough certification exams the first time, Amazon Data-Engineer-Associate Reliable Dump However, you can't just take it for granted, Our company has been engaged in compiling the Data-Engineer-Associate latest dumps: AWS Certified Data Engineer - Associate (DEA-C01) for workers more than 10 years, and our products has become the rage at the market, Amazon Data-Engineer-Associate Reliable Dump The content is the best way to help you get to know the knowledge in depth.

Minimize all but the active window, She can also look for communities Data-Engineer-Associate Reliable Dump based on keywords, Everything you need to prepare and quickly pass the tough certification exams the first time.

However, you can't just take it for granted, Our company has been engaged in compiling the Data-Engineer-Associate Latest Dumps: AWS Certified Data Engineer - Associate (DEA-C01) for workers more than 10 years, and our products has become the rage at the market.

Explore the Amazon Data-Engineer-Associate Online Practice Test Engine

The content is the best way to help you get to know the knowledge Data-Engineer-Associate in depth, The APP on-line test engine are available in all operate system and can be used on any electronic products.

- Free Updates For Amazon Data-Engineer-Associate PDF Questions Immediately open www.dumpsmaterials.com and search for “Data-Engineer-Associate” to obtain a free download Data-Engineer-Associate New Question
- 100% Pass 2026 Amazon The Best Data-Engineer-Associate: AWS Certified Data Engineer - Associate (DEA-C01) Reliable Dump The page for free download of (Data-Engineer-Associate) on www.pdfvce.com will open immediately Data-Engineer-Associate Latest Practice Questions
- Amazon Data-Engineer-Associate - First-grade AWS Certified Data Engineer - Associate (DEA-C01) Reliable Dump **【 www.examcollectionpass.com 】** is best website to obtain { Data-Engineer-Associate } for free download Data-Engineer-Associate Latest Practice Questions
- Data-Engineer-Associate Latest Practice Questions Study Data-Engineer-Associate Reference Test Data-Engineer-Associate Preparation Enter ➡ www.pdfvce.com and search for ⇒ Data-Engineer-Associate ⇐ to download for free Study Data-Engineer-Associate Reference
- 100% Pass 2026 Amazon The Best Data-Engineer-Associate: AWS Certified Data Engineer - Associate (DEA-C01) Reliable Dump Open ➡ www.examcollectionpass.com and search for **【 Data-Engineer-Associate 】** to download exam materials for free Study Data-Engineer-Associate Reference
- Data-Engineer-Associate Latest Version Study Data-Engineer-Associate Reference Best Data-Engineer-Associate Practice Search for ➡ Data-Engineer-Associate on ✓ www.pdfvce.com ✓ immediately to obtain a free download Data-Engineer-Associate Valid Test Guide
- Quiz Data-Engineer-Associate - Marvelous AWS Certified Data Engineer - Associate (DEA-C01) Reliable Dump Search for ⇒ Data-Engineer-Associate ⇐ and download it for free immediately on 《 www.dumpsmaterials.com 》 Data-Engineer-Associate Latest Practice Questions
- Data-Engineer-Associate Exam Topics Pdf Data-Engineer-Associate New Question Best Data-Engineer-Associate Practice Open website (www.pdfvce.com) and search for Data-Engineer-Associate for free download Data-Engineer-Associate Exam Vce Format
- Free PDF High Hit-Rate Amazon - Data-Engineer-Associate Reliable Dump Search for Data-Engineer-Associate on (www.troytecdumps.com) immediately to obtain a free download Exam Data-Engineer-Associate Outline

