

# Free PDF Quiz Amazon - Fantastic Data-Engineer-Associate Interactive Testing Engine



P.S. Free 2026 Amazon Data-Engineer-Associate dumps are available on Google Drive shared by DumpsKing:  
<https://drive.google.com/open?id=1rjLVpIUrmidr8HwSY4pwS3xs7Mj31mvb>

The language in our Amazon Data-Engineer-Associate test guide is easy to understand that will make any learner without any learning disabilities, whether you are a student or a in-service staff, whether you are a novice or an experienced staff who has abundant experience for many years. It should be a great wonderful idea to choose our Data-Engineer-Associate Guide Torrent for sailing through the difficult test.

As we all know, if candidates fail to pass the exam, time and energy you spend on the practicing will be returned nothing. If you choose us, we will let your efforts be payed off. Data-Engineer-Associate learning materials are edited and reviewed by professional experts who possess the professional knowledge for the exam, and therefore you can use them at ease. Besides, we are pass guarantee and money back guarantee for Data-Engineer-Associate Exam Materials. If you fail to pass the exam, we will give you full refund. We offer you free update for 365 days for Data-Engineer-Associate exam materials, and the update version will be sent to you automatically.

>> Data-Engineer-Associate Interactive Testing Engine <<

## Actual Amazon Data-Engineer-Associate Exam Dumps – Pass Exam With Good Scores

More and more people hope to enhance their professional competitiveness by obtaining Data-Engineer-Associate certification. However, under the premise that the pass rate is strictly controlled, fierce competition makes it more and more difficult to pass the Data-Engineer-Associate examination. Whether you are the first or the second or even more taking Data-Engineer-Associate examination, our Data-Engineer-Associate exam prep not only can help you to save much time and energy but also can help you pass the exam. In the other words, passing the exam once will no longer be a dream.

## Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q42-Q47):

## NEW QUESTION # 42

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

### Answer: A

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 # 43

A hotel management company receives daily data files from each of its hotels. The company wants to upload its data to AWS. The company plans to use Amazon Athena to access the files. The company needs to protect the files from accidental deletion. The company will develop an application on its on-premises servers to automatically forward the files to a fully managed AWS ingestion service.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use the Amazon Kinesis Agent on the on-premises servers to send data to Amazon Data Firehose. Store the data in an Amazon S3 bucket that has versioning enabled.
- B. Use a self-managed Apache Kafka agent on the on-premises servers to stream data to Amazon Managed Streaming for Apache Kafka (Amazon MSK). Store the data in an Amazon S3 bucket with versioning enabled.
- C. Use AWS Glue jobs to ingest data from the on-premises servers into Amazon RDS. Enable automated backups for data protection.
- D. Use AWS DataSync to replicate data from the on-premises servers to Amazon Elastic File System (Amazon EFS). Configure automatic backups in AWS Backup.

### Answer: A

Explanation:

Amazon Kinesis Data Firehose is a fully managed data ingestion service that enables reliable and scalable delivery of streaming and batch data into Amazon S3 with minimal operational overhead. This directly satisfies the requirement for a fully managed AWS ingestion service while avoiding the need to provision, scale, or manage infrastructure.

By using the Amazon Kinesis Agent on the on-premises servers, the company can automatically forward daily data files to Kinesis Data Firehose. Firehose handles buffering, retry logic, scaling, and delivery without requiring administrative effort. Delivering the data to Amazon S3 allows seamless integration with Amazon Athena, which natively queries data stored in S3 without requiring data movement or transformation.

Enabling Amazon S3 versioning protects files from accidental deletion by preserving previous versions of objects. This aligns with AWS best practices for data durability and governance, especially for analytics workloads and compliance requirements.

Other options introduce unnecessary operational complexity. AWS DataSync with Amazon EFS is not optimized for Athena-based analytics. AWS Glue jobs and Amazon RDS are unsuitable for file-based analytical access. A self-managed Apache Kafka solution with Amazon MSK significantly increases operational overhead.

Therefore, option B is the most efficient, scalable, and operationally optimal solution according to AWS Certified Data Engineer - Associate best practices.

#### NEW QUESTION # 44

A company stores data from an application in an Amazon DynamoDB table that operates in provisioned capacity mode. The workloads of the application have predictable throughput load on a regular schedule.

Every Monday, there is an immediate increase in activity early in the morning. The application has very low usage during weekends. The company must ensure that the application performs consistently during peak usage times.

Which solution will meet these requirements in the MOST cost-effective way?

- A. Divide the table into two tables. Provision each table with half of the provisioned capacity of the original table. Spread queries evenly across both tables.
- B. Increase the provisioned capacity to the maximum capacity that is currently present during peak load times.
- C. Use AWS Application Auto Scaling to schedule higher provisioned capacity for peak usage times.  
Schedule lower capacity during off-peak times.
- D. Change the capacity mode from provisioned to on-demand. Configure the table to scale up and scale down based on the load on the table.

**Answer: C**

Explanation:

Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. DynamoDB offers two capacity modes for throughput capacity:

provisioned and on-demand. In provisioned capacity mode, you specify the number of read and write capacity units per second that you expect your application to require. DynamoDB reserves the resources to meet your throughput needs with consistent performance. In on-demand capacity mode, you pay per request and DynamoDB scales the resources up and down automatically based on the actual workload. On-demand capacity mode is suitable for unpredictable workloads that can vary significantly over time1.

The solution that meets the requirements in the most cost-effective way is to use AWS Application Auto Scaling to schedule higher provisioned capacity for peak usage times and lower capacity during off-peak times. This solution has the following advantages:

It allows you to optimize the cost and performance of your DynamoDB table by adjusting the provisioned capacity according to your predictable workload patterns. You can use scheduled scaling to specify the date and time for the scaling actions, and the new minimum and maximum capacity limits. For example, you can schedule higher capacity for every Monday morning and lower capacity for weekends2.

It enables you to take advantage of the lower cost per unit of provisioned capacity mode compared to on-demand capacity mode. Provisioned capacity mode charges a flat hourly rate for the capacity you reserve, regardless of how much you use. On-demand capacity mode charges for each read and write request you consume, with no minimum capacity required. For predictable workloads, provisioned capacity mode can be more cost-effective than on-demand capacity mode1.

It ensures that your application performs consistently during peak usage times by having enough capacity to handle the increased load. You can also use auto scaling to automatically adjust the provisioned capacity based on the actual utilization of your table, and set a target utilization percentage for your table or global secondary index. This way, you can avoid under-provisioning or over-provisioning your table2.

Option A is incorrect because it suggests increasing the provisioned capacity to the maximum capacity that is currently present during peak load times. This solution has the following disadvantages:

It wastes money by paying for unused capacity during off-peak times. If you provision the same high capacity for all times, regardless of the actual workload, you are over-provisioning your table and paying for resources that you don't need1.

It does not account for possible changes in the workload patterns over time. If your peak load times increase or decrease in the future, you may need to manually adjust the provisioned capacity to match the new demand. This adds operational overhead and complexity to your application2.

Option B is incorrect because it suggests dividing the table into two tables and provisioning each table with half of the provisioned capacity of the original table. This solution has the following disadvantages:

It complicates the data model and the application logic by splitting the data into two separate tables. You need to ensure that the queries are evenly distributed across both tables, and that the data is consistent and synchronized between them. This adds extra development and maintenance effort to your application3.

It does not solve the problem of adjusting the provisioned capacity according to the workload patterns.

You still need to manually or automatically scale the capacity of each table based on the actual utilization and demand. This may result in under-provisioning or over-provisioning your tables2.

Option D is incorrect because it suggests changing the capacity mode from provisioned to on-demand. This solution has the following disadvantages:

It may incur higher costs than provisioned capacity mode for predictable workloads. On-demand capacity mode charges for each read and write request you consume, with no minimum capacity required. For predictable workloads, provisioned capacity mode can be more cost-effective than on-demand capacity mode, as you can reserve the capacity you need at a lower rate1.

It may not provide consistent performance during peak usage times, as on-demand capacity mode may take some time to scale up

the resources to meet the sudden increase in demand. On-demand capacity mode uses adaptive capacity to handle bursts of traffic, but it may not be able to handle very large spikes or sustained high throughput. In such cases, you may experience throttling or increased latency.

References:

- 1: Choosing the right DynamoDB capacity mode - [Amazon DynamoDB](#)
- 2: Managing throughput capacity automatically with DynamoDB auto scaling - [Amazon DynamoDB](#)
- 3: Best practices for designing and using partition keys effectively - [Amazon DynamoDB](#)
- [4]: On-demand mode guidelines - [Amazon DynamoDB](#)
- [5]: How to optimize Amazon DynamoDB costs - [AWS Database Blog](#)
- [6]: DynamoDB adaptive capacity: How it works and how it helps - [AWS Database Blog](#)
- [7]: Amazon DynamoDB pricing - [Amazon Web Services \(AWS\)](#)

## NEW QUESTION # 45

A company stores its processed data in an S3 bucket. The company has a strict data access policy. The company uses IAM roles to grant teams within the company different levels of access to the S3 bucket.

The company wants to receive notifications when a user violates the data access policy. Each notification must include the username of the user who violated the policy.

Which solution will meet these requirements?

- A. Use Amazon S3 server access logs to monitor access to the bucket. Forward the access logs to an Amazon CloudWatch log group. Use metric filters on the log group to set up CloudWatch alarms.
- B. Use Amazon CloudWatch metrics to gather object-level metrics. Set up CloudWatch alarms.
- C. Use AWS CloudTrail to track object-level events for the S3 bucket. Forward events to Amazon CloudWatch to set up CloudWatch alarms.
- D. Use AWS Config rules to detect violations of the data access policy. Set up compliance alarms.

**Answer: C**

Explanation:

The requirement is to detect violations of data access policies and receive notifications with the username of the violator. AWS CloudTrail can provide object-level tracking for S3 to capture detailed API actions on specific S3 objects, including the user who performed the action.

\* AWS CloudTrail:

- \* CloudTrail can monitor API calls made to an S3 bucket, including object-level API actions such as GetObject, PutObject, and DeleteObject. This will help detect access violations based on the API calls made by different users.
- \* CloudTrail logs include details such as the user identity, which is essential for meeting the requirement of including the username in notifications.
- \* The CloudTrail logs can be forwarded to Amazon CloudWatch to trigger alarms based on certain access patterns (e.g., violations of specific policies).

## NEW QUESTION # 46

A data engineer needs to use Amazon Neptune to develop graph applications.

Which programming languages should the engineer use to develop the graph applications? (Select TWO.)

- A. ANSI SQL
- B. Spark SQL
- C. SPARQL
- D. Gremlin
- E. SQL

**Answer: C,D**

Explanation:

Amazon Neptune supports graph applications using Gremlin and SPARQL as query languages. Neptune is a fully managed graph database service that supports both property graph and RDF graph models.

Option A: Gremlin

Gremlin is a query language for property graph databases, which is supported by Amazon Neptune. It allows the traversal and manipulation of graph data in the property graph model.

Option D: SPARQL

SPARQL is a query language for querying RDF graph data in Neptune. It is used to query, manipulate, and retrieve information stored in RDF format.

Other options:

SQL (Option B) and ANSI SQL (Option C) are traditional relational database query languages and are not used for graph databases.

Spark SQL (Option E) is related to Apache Spark for big data processing, not for querying graph databases.

Reference:

Amazon Neptune Documentation

Gremlin Documentation

SPARQL Documentation

## NEW QUESTION # 47

.....

The AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) certification is one of the hottest career advancement credentials in the modern Amazon world. The Amazon Data-Engineer-Associate certification can help you to demonstrate your expertise and knowledge level. With only one badge of Data-Engineer-Associate Certification, successful candidates can advance their careers and increase their earning potential.

**Data-Engineer-Associate Reliable Exam Sims:** <https://www.dumpsking.com/Data-Engineer-Associate-testking-dumps.html>

Our company does not only give consideration to improve the quality of our product but also take into account the fact that many people always feel nervous in the exam and cannot perform well in the real Data-Engineer-Associate Reliable Exam Sims - AWS Certified Data Engineer - Associate (DEA-C01) exam. As the Data-Engineer-Associate exam continues to update, our software will be always updating with it. We are committed to let all of the candidates pass Amazon exam and gain the IT certification successfully, but if you unfortunately failed the exam even with the help of Data-Engineer-Associate exam questions: AWS Certified Data Engineer - Associate (DEA-C01), we will promise a full refund for you, but you need to show your report card to us, and as soon as we confirm it we will give you a full refund, so just do not worry about your money of buying the Data-Engineer-Associate study materials.

Again, we'll program that particular event momentarily, To apply Data-Engineer-Associate fixes to a specific area of the image, use the Quick Selection tool to select an active area, Our company does not only give consideration to improve the quality of our product but also take Data-Engineer-Associate Reliable Exam Sims into account the fact that many people always feel nervous in the exam and cannot perform well in the real AWS Certified Data Engineer - Associate (DEA-C01) exam.

## 100% Pass Perfect Data-Engineer-Associate - AWS Certified Data Engineer - Associate (DEA-C01) Intereactive Testing Engine

As the Data-Engineer-Associate Exam continues to update, our software will be always updating with it. We are committed to let all of the candidates pass Amazon exam and gain the IT certification successfully, but if you unfortunately failed the exam even with the help of Data-Engineer-Associate exam questions: AWS Certified Data Engineer - Associate (DEA-C01), we will promise a full refund for you, but you need to show your report card to us, and as soon as we confirm it we will give you a full refund, so just do not worry about your money of buying the Data-Engineer-Associate study materials.

You can receive your downloading link and password within ten minutes after payment, so that you can start your learning by using Data-Engineer-Associate exam dumps, It gives you confidence and boosts your morale.

- Get 1 year Of Updated Amazon Data-Engineer-Associate Exam Question Dumps □ Search for □ Data-Engineer-Associate □ on ► [www.troytecdumps.com](http://www.troytecdumps.com) ▲ immediately to obtain a free download □ Test Data-Engineer-Associate Answers
- Three Top Amazon Data-Engineer-Associate Dumps Formats □ Open ► [www.pdfvce.com](http://www.pdfvce.com) □ and search for 《 Data-Engineer-Associate 》 to download exam materials for free □ New Exam Data-Engineer-Associate Braindumps
- New Data-Engineer-Associate Test Vce □ Exam Data-Engineer-Associate Cost □ Valid Data-Engineer-Associate Exam Voucher □ Go to website ► [www.verifieddumps.com](http://www.verifieddumps.com) □ □ □ open and search for ► Data-Engineer-Associate □ to download for free □ Premium Data-Engineer-Associate Exam
- Three Top Amazon Data-Engineer-Associate Dumps Formats □ Download ► Data-Engineer-Associate □ for free by simply searching on “ [www.pdfvce.com](http://www.pdfvce.com) ” □ Data-Engineer-Associate Latest Material
- Data-Engineer-Associate Exam Bootcamp: AWS Certified Data Engineer - Associate (DEA-C01) - Data-Engineer-Associate Original Questions - Data-Engineer-Associate Exam Prep □ Download “ Data-Engineer-Associate ” for free by simply searching on □ [www.testkingpass.com](http://www.testkingpass.com) □ □ Reliable Data-Engineer-Associate Real Exam

2026 Latest DumpsKing Data-Engineer-Associate PDF Dumps and Data-Engineer-Associate Exam Engine Free Share:

<https://drive.google.com/open?id=1rjLVplUrmidr8HwSY4pwS3xs7Mj31mvb>