# Quiz 2026 Accurate Amazon Data-Engineer-Associate Reliable Exam Cram


What's more, part of that iPassleader Data-Engineer-Associate dumps now are free: https://drive.google.com/open?id=1pyzSXE7VrLIunSFmIUGuQSZqigZJyo9D

As job seekers looking for the turning point of their lives, it is widely known that the workers of recruitment is like choosing apples---viewing resumes is liking picking up apples, employers can decide whether candidates are qualified by the Data-Engineer-Associate appearances, or in other words, candidates' educational background and relating Data-Engineer-Associate professional skills. The reason why we are so confident lies in the sophisticated expert group and technical team we have, which do duty for our solid support. They develop the Data-Engineer-Associate Exam Guide targeted to real exam. The wide coverage of important knowledge points in our Data-Engineer-Associate latest braindumps would be greatly helpful for you to pass the exam.

Heaps of beginners and skilled professionals already have surpassed the Amazon Data-Engineer-Associate certification exam and pursuing a worthwhile profession inside the quite aggressive market. You may additionally turn out to be a part of this skilled and certified community. To try this you sincerely need to pass the Amazon Data-Engineer-Associate Certification examination.

**>> Data-Engineer-Associate Reliable Exam Cram <<**

## Data-Engineer-Associate Reliable Test Sample | Data-Engineer-Associate Reliable Test Online

Our Data-Engineer-Associate training materials are compiled by professional experts. All the necessary points have been mentioned in our Data-Engineer-Associate practice engine particularly. About some tough questions or important points, they left notes under them. Besides, our experts will concern about changes happened in Data-Engineer-Associate study prep all the time. Provided you have a strong determination, as well as the help of our Data-Engineer-Associate learning guide, you can have success absolutely.

## Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q164-Q169):

**NEW QUESTION # 164**

A data engineer is configuring an AWS Glue Apache Spark extract, transform, and load (ETL) job. The job contains a sort-merge join of two large and equally sized DataFrames.
The job is failing with the following error: No space left on device.
Which solution will resolve the error?

- A. Deploy an Amazon Elastic Block Store (Amazon EBS) volume for the job to use.
- B. Convert the sort-merge join in the job to be a broadcast join.
- C. Convert the DataFrames to DynamicFrames, and perform a DynamicFrame join in the job.
- D. Use the AWS Glue Spark shuffle manager.

**Answer: B**

Explanation:
A sort-merge join generates large shuffle files, leading to "No space left on device" errors when both datasets are large. Using a broadcast join sends a smaller dataset to all executors, avoiding shuffle and disk I/O overhead.
"Broadcast joins reduce shuffle I/O by distributing the smaller dataset to all worker nodes, mitigating disk space and shuffle errors." - Ace the AWS Certified Data Engineer - Associate Certification - version 2 - apple.pdf This is the most cost-effective and direct fix for large shuffle-stage failures.

**NEW QUESTION # 165**
A company uses an Amazon Redshift provisioned cluster as its database. The Redshift cluster has five reserved ra3.4xlarge nodes and uses key distribution.
A data engineer notices that one of the nodes frequently has a CPU load over 90%. SQL Queries that run on the node are queued. The other four nodes usually have a CPU load under 15% during daily operations.
The data engineer wants to maintain the current number of compute nodes. The data engineer also wants to balance the load more evenly across all five compute nodes.
Which solution will meet these requirements?

- A. Change the sort key to be the data column that is most often used in a WHERE clause of the SQL SELECT statement.
- B. Change the primary key to be the data column that is most often used in a WHERE clause of the SQL SELECT statement.
- C. Change the distribution key to the table column that has the largest dimension.
- D. Upgrade the reserved node from ra3.4xlarqe to ra3.16xlarqe.

**Answer: C**

Explanation:
Changing the distribution key to the table column that has the largest dimension will help to balance the load more evenly across all five compute nodes. The distribution key determines how the rows of a table are distributed among the slices of the cluster. If the distribution key is not chosen wisely, it can cause data skew, meaning some slices will have more data than others, resulting in uneven CPU load and query performance.
By choosing the table column that has the largest dimension, meaning the column that has the most distinct values, as the distribution key, the data engineer can ensure that the rows are distributed more uniformly across the slices, reducing data skew and improving query performance.
The other options are not solutions that will meet the requirements. Option A, changing the sort key to be the data column that is most often used in a WHERE clause of the SQL SELECT statement, will not affect the data distribution or the CPU load. The sort key determines the order in which the rows of a table are stored on disk, which can improve the performance of range-restricted queries, but not the load balancing. Option C, upgrading the reserved node from ra3.4xlarge to ra3.16xlarge, will not maintain the current number of compute nodes, as it will increase the cost and the capacity of the cluster. Option D, changing the primary key to be the data column that is most often used in a WHERE clause of the SQL SELECT statement, will not affect the data distribution or the CPU load either. The primary key is a constraint that enforces the uniqueness of the rows in a table, but it does not influence the data layout or the query optimization.
References:
* Choosing a data distribution style
* Choosing a data sort key
* Working with primary keys

**NEW QUESTION # 166**
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. Create snapshots of the gp2 volumes. Create new gp3 volumes from the snapshots. Attach the new gp3 volumes to the EC2 instances.
- B. Change the volume type of the existing gp2 volumes to gp3. Enter new values for volume size, IOPS, and throughput.
- 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: B**

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 # 167**
A company is building a data stream processing application. The application runs in an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The application stores processed data in an Amazon DynamoDB table.
The company needs the application containers in the EKS cluster to have secure access to the DynamoDB table. The company does not want to embed AWS credentials in the containers.
Which solution will meet these requirements?

- A. Store the AWS credentials in an Amazon S3 bucket. Grant the EKS containers access to the S3 bucket to retrieve the credentials.
- B. Attach an IAM role to the EKS worker nodes. Grant the IAM role access to DynamoDB. Use the IAM role to set up IAM roles service accounts (IRSA) functionality.
- C. Create an IAM user that has an access key to access the DynamoDB table. Use environment variables in the EKS containers to store the IAM user access key data.
- D. Create an IAM user that has an access key to access the DynamoDB table. Use Kubernetes secrets that are mounted in a volume of the EKS cluster nodes to store the user access key data.

**Answer: B**

Explanation:
In this scenario, the company is using Amazon Elastic Kubernetes Service (EKS) and wants secure access to DynamoDB without embedding credentials inside the application containers. The best practice is to use IAM roles for service accounts (IRSA), which allows assigning IAM roles to Kubernetes service accounts. This lets the EKS pods assume specific IAM roles securely, without the need to store credentials in containers.
IAM Roles for Service Accounts (IRSA):
With IRSA, each pod in the EKS cluster can assume an IAM role that grants access to DynamoDB without needing to manage long-term credentials. The IAM role can be attached to the service account associated with the pod.
This ensures least privilege access, improving security by preventing credentials from being embedded in the containers.
Reference:
Alternatives Considered:
A (Storing AWS credentials in S3): Storing AWS credentials in S3 and retrieving them introduces security risks and violates the principle of not embedding credentials.
C (IAM user access keys in environment variables): This also embeds credentials, which is not recommended.
D (Kubernetes secrets): Storing user access keys as secrets is an option, but it still involves handling long-term credentials manually, which is less secure than using IRSA.
IAM Best Practices for Amazon EKS
Secure Access to DynamoDB from EKS

**NEW QUESTION # 168**
A data engineer needs to join data from multiple sources to perform a one-time analysis job. The data is stored in Amazon DynamoDB, Amazon RDS, Amazon Redshift, and Amazon S3.
Which solution will meet this requirement MOST cost-effectively?

- A. Use an Amazon EMR provisioned cluster to read from all sources. Use Apache Spark to join the data and perform the analysis.
- B. Use Redshift Spectrum to query data from DynamoDB, Amazon RDS, and Amazon S3 directly from Redshift.
- C. Use Amazon Athena Federated Query to join the data from all data sources.
- D. Copy the data from DynamoDB, Amazon RDS, and Amazon Redshift into Amazon S3. Run Amazon Athena queries directly on the S3 files.

**Answer: C**

Explanation:
Amazon Athena Federated Query is a feature that allows you to query data from multiple sources using standard SQL. You can use Athena Federated Query to join data from Amazon DynamoDB, Amazon RDS, Amazon Redshift, and Amazon S3, as well as other data sources such as MongoDB, Apache HBase, and Apache Kafka1. Athena Federated Query is a serverless and interactive service, meaning you do not need to provision or manage any infrastructure, and you only pay for the amount of data scanned by your queries. Athena Federated Query is the most cost-effective solution for performing a one-time analysis job on data from multiple sources, as it eliminates the need to copy or move data, and allows you to query data directly from the source.
The other options are not as cost-effective as Athena Federated Query, as they involve additional steps or costs. Option A requires you to provision and pay for an Amazon EMR cluster, which can be expensive and time-consuming for a one-time job. Option B requires you to copy or move data from DynamoDB, RDS, and Redshift to S3, which can incur additional costs for data transfer and storage, and also introduce latency and complexity. Option D requires you to have an existing Redshift cluster, which can be costly and may not be necessary for a one-time job. Option D also does not support querying data from RDS directly, so you would need to use Redshift Federated Query to access RDS data, which adds another layer of complexity2. Reference:
Amazon Athena Federated Query
Redshift Spectrum vs Federated Query

**NEW QUESTION # 169**
.....

Our iPassleader's Data-Engineer-Associate test training materials can test your knowledge, when you prepare for Data-Engineer-Associate test; and can also evaluate your performance at the appointed time. Our Data-Engineer-Associate exam training materials is the result of iPassleader's experienced IT experts with constant exploration, practice and research for many years. Its authority is undeniable. If you have any concerns, you can first try Data-Engineer-Associate PDF VCE free demo and answers, and then make a decision whether to choose our Data-Engineer-Associate dumps or not.

**Data-Engineer-Associate Reliable Test Sample**: https://www.ipassleader.com/Amazon/Data-Engineer-Associate-practice-exam-dumps.html

If you are very busy, you can only take two or three hours a day to study our Data-Engineer-Associate study engine, You will be provided with an examination environment and you will be presented with actual exam Amazon Data-Engineer-Associate exam questions, All you need is to stay positive, put in your best efforts, and be confident while appearing for the Amazon Data-Engineer-Associate exam, If you are interested to our Data-Engineer-Associate simulating exam, just place your order now.

An intelligent restart is a recovery mechanism Data-Engineer-Associate that detects when a transmission has been cut, The C++ State Pattern for Network Operations, If you are very busy, you can only take two or three hours a day to study our Data-Engineer-Associate study engine.

## iPassleader Amazon Data-Engineer-Associate Exam prepare material in three Different formats

You will be provided with an examination environment and you will be presented with actual exam Amazon Data-Engineer-Associate exam questions, All you need is to stay positive, put in your best efforts, and be confident while appearing for the Amazon Data-Engineer-Associate exam.

If you are interested to our Data-Engineer-Associate simulating exam, just place your order now, Our service covers all around the world and the clients can receive our Data-Engineer-Associate study materials as quickly as possible.

- www.pdfdumps.com's Amazon Data-Engineer-Associate Practice Test Software (Web-Based and Desktop) □ Simply search for ➽ Data-Engineer-Associate □ for free download on ⇛ www.pdfdumps.com ⇚ □Data-Engineer-Associate Actual Braindumps
- Use Amazon Data-Engineer-Associate PDF Dumps to Prepare in a Short Time □ Open website ▹ www.pdfvce.com ◃ and search for ➡ Data-Engineer-Associate □□□ for free download □Data-Engineer-Associate Reliable Test Practice
- Data-Engineer-Associate Test Dumps Pdf □ Data-Engineer-Associate Certification Exam Cost □ Latest Data-Engineer-Associate Braindumps Free □ Easily obtain free download of ➥ Data-Engineer-Associate □ by searching on " www.examcollectionpass.com " □Exam Data-Engineer-Associate Vce
- Data-Engineer-Associate Dump Torrent □ Data-Engineer-Associate Real Dumps Free □ Data-Engineer-Associate Real Dumps Free □ （ www.pdfvce.com ） is best website to obtain ▸ Data-Engineer-Associate ◂ for free download □ □Data-Engineer-Associate Dump Torrent
- Experience The Real Environment With The Help Of Amazon Data-Engineer-Associate Exam Questions □ Search for " Data-Engineer-Associate " and download it for free on ☀ www.testkingpass.com □☀□ website □Answers Data-Engineer-Associate Free
- Pass Guaranteed 2026 Pass-Sure Data-Engineer-Associate: AWS Certified Data Engineer - Associate (DEA-C01) Reliable Exam Cram □ Open " www.pdfvce.com " and search for 《 Data-Engineer-Associate 》 to download exam materials for free □Data-Engineer-Associate Certification Exam Cost
- Three Formats of Latest Amazon Data-Engineer-Associate Practice Material □ Search for [ Data-Engineer-Associate ] and obtain a free download on [ www.prepawayete.com ] □Data-Engineer-Associate Actual Braindumps
- Data-Engineer-Associate Test Dumps Pdf □ Data-Engineer-Associate New Dumps Questions □ Data-Engineer-Associate Certification Exam Cost □ " www.pdfvce.com " is best website to obtain ▸ Data-Engineer-Associate ◂ for free download □Exam Data-Engineer-Associate Vce
- Latest Data-Engineer-Associate Braindumps Free □ Data-Engineer-Associate Certification Torrent □ Data-Engineer-Associate Pass Test □ □ www.vce4dumps.com □ is best website to obtain □ Data-Engineer-Associate □ for free download □Data-Engineer-Associate Reliable Test Practice
- Latest Data-Engineer-Associate Exam Cram □ Data-Engineer-Associate Latest Torrent □ Answers Data-Engineer-Associate Free □ Download [ Data-Engineer-Associate ] for free by simply entering ▸ www.pdfvce.com ◂ website □ □Answers Data-Engineer-Associate Free
- Latest Data-Engineer-Associate Exam Cram □ Data-Engineer-Associate Test Review □ Data-Engineer-Associate Reliable Test Practice □ Open website ✔ www.vceengine.com □✔□ and search for ➡ Data-Engineer-Associate □ for free download □Test Data-Engineer-Associate Cram
- www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, adrcentre.org, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, Disposable vapes

DOWNLOAD the newest iPassleader Data-Engineer-Associate PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1pyzSXE7VrLIunSFmIUGuQSZqigZJyo9D