Famous Data-Engineer-Associate Exam Guide: AWS Certified Data Engineer - Associate (DEA-C01) Bring You Pass-Guaranteed Training Dumps - Fast2test

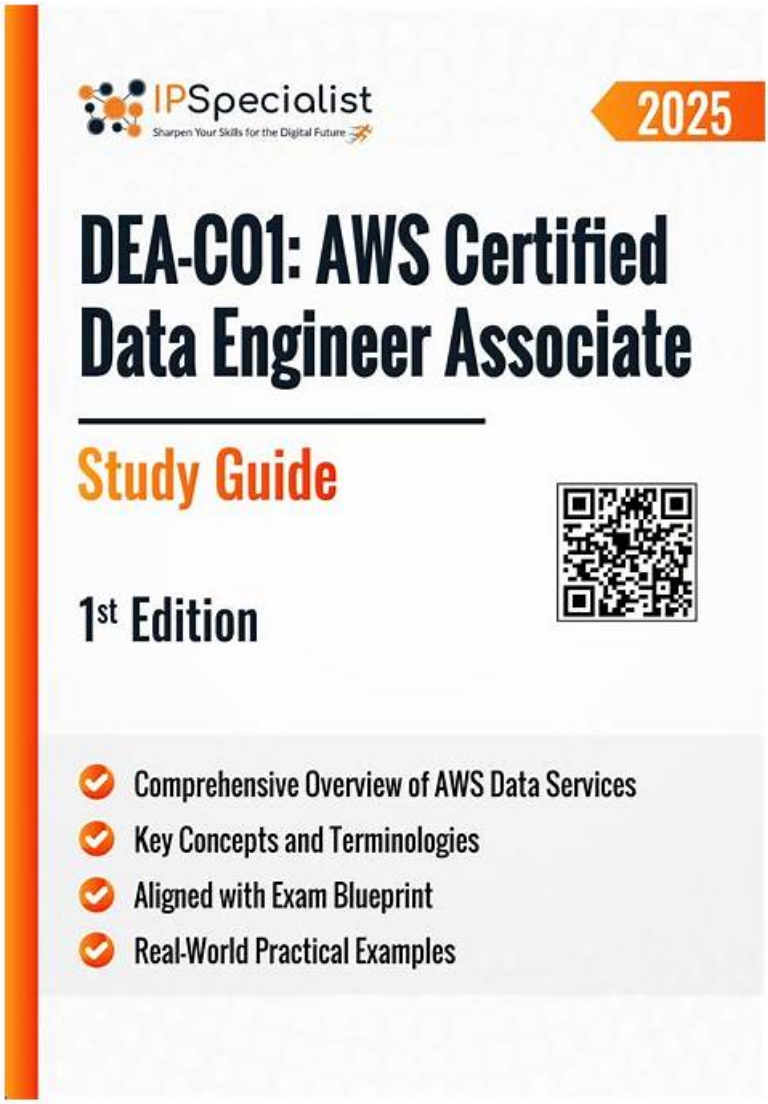

BONUS!!! Download part of Fast2test Data-Engineer-Associate dumps for free: https://drive.google.com/open?id=1k5N6Vc3YWtMB7zFe4lnM19wPZAcrUp9y

It can be difficult to prepare for the AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) certification test when you're already busy with daily tasks. But, you can successfully prepare for the examination despite your busy schedule if you choose updated and real Amazon Data-Engineer-Associate exam questions. We believe that success in the test depends on studying with AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) Dumps questions. We have hired a team of professionals who has years of experience in helping test applicants acquire essential knowledge by providing them with Amazon Data-Engineer-Associate actual exam questions.

Fast2test offers Data-Engineer-Associate actual exam dumps in easy-to-use PDF format. It is a portable format that works on all smart devices. Questions in the Data-Engineer-Associate PDF can be studied at any time from any place. Furthermore, AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) PDF exam questions are printable. It means you can avoid eye strain by preparing real questions in a hard copy.

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

Exam Data-Engineer-Associate Blueprint - Data-Engineer-Associate Valid Braindumps Sheet

In general, we can say that the Data-Engineer-Associate certification can be a valuable investment in your career that will put your career on the right track and you can achieve your career objectives in a short time period. These are some important benefits that you can gain after passing the Amazon Data-Engineer-Associate Certification Exam. Are you ready to pass the Data-Engineer-Associate exam? Looking for a simple, quick, and proven way to pass the Amazon Data-Engineer-Associate Exam Questions? If your answer is yes then download Fast2test exam questions and start this journey today.

Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q147-Q152):

NEW QUESTION # 147
A data engineer is using Amazon Athena to analyze sales data that is in Amazon S3. The data engineer writes a query to retrieve sales amounts for 2023 for several products from a table named sales_data. However, the query does not return results for all of the products that are in the sales_data table. The data engineer needs to troubleshoot the query to resolve the issue.
The data engineer's original query is as follows:
SELECT product_name, sum(sales_amount)
FROM sales_data
WHERE year = 2023
GROUP BY product_name
How should the data engineer modify the Athena query to meet these requirements?

- A. Remove the GROUP BY clause
- B. Change WHERE year = 2023 to WHERE extractlyear FROM sales data) = 2023.
- C. Add HAVING sumfsales amount) > 0 after the GROUP BY clause.
- D. Replace sum(sales amount) with count(*J for the aggregation.

Answer: B

Explanation:
The original query does not return results for all of the products because the year column in the sales_data table is not an integer, but a timestamp. Therefore, the WHERE clause does not filter the data correctly, and only returns the products that have a null value for the year column. To fix this, the data engineer should use the extract function to extract the year from the timestamp and compare it with 2023. This way, the query will return the correct results for all of the products in the sales_data table. The other options are either incorrect or irrelevant, as they do not address the root cause of the issue. Replacing sum with count does not change the filtering condition, adding HAVING clause does not affect the grouping logic, and removing the GROUP BY clause does not solve the problem of missing products. References:
* Troubleshooting JSON queries - Amazon Athena (Section: JSON related errors)
* When I query a table in Amazon Athena, the TIMESTAMP result is empty (Section: Resolution)
* AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide (Chapter 7, page 197)

NEW QUESTION # 148
A company maintains an Amazon Redshift provisioned cluster that the company uses for extract, transform, and load (ETL) operations to support critical analysis tasks. A sales team within the company maintains a Redshift cluster that the sales team uses for business intelligence (BI) tasks.
The sales team recently requested access to the data that is in the ETL Redshift cluster so the team can perform weekly summary analysis tasks. The sales team needs to join data from the ETL cluster with data that is in the sales team's BI cluster.
The company needs a solution that will share the ETL cluster data with the sales team without interrupting the critical analysis tasks. The solution must minimize usage of the computing resources of the ETL cluster.
Which solution will meet these requirements?

- A. Create materialized views based on the sales team's requirements. Grant the sales team direct access to the ETL cluster.
- B. Unload a copy of the data from the ETL cluster to an Amazon S3 bucket every week. Create an Amazon Redshift Spectrum table based on the content of the ETL cluster.
- C. Create database views based on the sales team's requirements. Grant the sales team direct access to the ETL cluster.
- D. Set up the sales team Bl cluster as a consumer of the ETL cluster by using Redshift data sharing.

Answer: D

Explanation:
Redshift data sharing is a feature that enables you to share live data across different Redshift clusters without the need to copy or move data. Data sharing provides secure and governed access to data, while preserving the performance and concurrency benefits of Redshift. By setting up the sales team BI cluster as a consumer of the ETL cluster, the company can share the ETL cluster data with the sales team without interrupting the critical analysis tasks. The solution also minimizes the usage of the computing resources of the ETL cluster, as the data sharing does not consume any storage space or compute resources from the producer cluster. The other options are either not feasible or not efficient. Creating materialized views or database views would require the sales team to have direct access to the ETL cluster, which could interfere with the critical analysis tasks. Unloading a copy of the data from the ETL cluster to an Amazon S3 bucket every week would introduce additional latency and cost, as well as create data inconsistency issues.
References:
* Sharing data across Amazon Redshift clusters
* AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide, Chapter 2: Data Store Management, Section 2.2: Amazon Redshift

NEW QUESTION # 149
A company uses Amazon RDS for MySQL as the database for a critical application. The database workload is mostly writes, with a small number of reads.
A data engineer notices that the CPU utilization of the DB instance is very high. The high CPU utilization is slowing down the application. The data engineer must reduce the CPU utilization of the DB Instance.
Which actions should the data engineer take to meet this requirement? (Choose two.)

- A. Use the Performance Insights feature of Amazon RDS to identify queries that have high CPU utilization. Optimize the problematic queries.
- B. Implement caching to reduce the database query load.
- C. Reboot the RDS DB instance once each week.
- D. Upgrade to a larger instance size.
- E. Modify the database schema to include additional tables and indexes.

Answer: A,B

Explanation:
Amazon RDS is a fully managed service that provides relational databases in the cloud. Amazon RDS for MySQL is one of the supported database engines that you can use to run your applications. Amazon RDS provides various features and tools to monitor and optimize the performance of your DB instances, such as Performance Insights, Enhanced Monitoring, CloudWatch metrics and alarms, etc.
Using the Performance Insights feature of Amazon RDS to identify queries that have high CPU utilization and optimizing the problematic queries will help reduce the CPU utilization of the DB instance. Performance Insights is a feature that allows you to analyze the load on your DB instance and determine what is causing performance issues. Performance Insights collects, analyzes, and displays database performance data using an interactive dashboard. You can use Performance Insights to identify the top SQL statements, hosts, users, or processes that are consuming the most CPU resources. You can also drill down into the details of each query and see the execution plan, wait events, locks, etc. By using Performance Insights, you can pinpoint the root cause of the high CPU utilization and optimize the queries accordingly. For example, you can rewrite the queries to make them more efficient, add or remove indexes, use prepared statements, etc.
Implementing caching to reduce the database query load will also help reduce the CPU utilization of the DB instance. Caching is a technique that allows you to store frequently accessed data in a fast and scalable storage layer, such as Amazon ElastiCache. By using caching, you can reduce the number of requests that hit your database, which in turn reduces the CPU load on your DB instance. Caching also improves the performance and availability of your application, as it reduces the latency and increases the throughput of your data access. You can use caching for various scenarios, such as storing session data, user preferences, application configuration, etc. You can also use caching for read-heavy workloads, such as displaying product details, recommendations, reviews, etc.
The other options are not as effective as using Performance Insights and caching. Modifying the database schema to include additional tables and indexes may or may not improve the CPU utilization, depending on the nature of the workload and the queries. Adding more tables and indexes may increase the complexity and overhead of the database, which may negatively affect the performance. Rebooting the RDS DB instance once each week will not reduce the CPU utilization, as it will not address the underlying cause of the high CPU load. Rebooting may also cause downtime and disruption to your application. Upgrading to a larger instance size may reduce the CPU utilization, but it will also increase the cost and complexity of your solution. Upgrading may also not be necessary if you can optimize the queries and reduce the database load by using caching. Reference:
Amazon RDS
Performance Insights
Amazon ElastiCache

[AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide], Chapter 3: Data Storage and Management, Section 3.1: Amazon RDS

NEW QUESTION # 150
A car sales company maintains data about cars that are listed for sale in an are a. The company receives data about new car listings from vendors who upload the data daily as compressed files into Amazon S3. The compressed files are up to 5 KB in size. The company wants to see the most up-to-date listings as soon as the data is uploaded to Amazon S3.
A data engineer must automate and orchestrate the data processing workflow of the listings to feed a dashboard. The data engineer must also provide the ability to perform one-time queries and analytical reporting. The query solution must be scalable.
Which solution will meet these requirements MOST cost-effectively?

- A. Use a provisioned Amazon EMR cluster to process incoming data. Use AWS Step Functions to orchestrate workflows. Use Amazon Athena for one-time queries and analytical reporting. Use Amazon QuickSight for the dashboard.
- B. Use AWS Glue to process incoming data. Use AWS Lambda and S3 Event Notifications to orchestrate workflows. Use Amazon Athena for one-time queries and analytical reporting. Use Amazon QuickSight for the dashboard.
- C. Use AWS Glue to process incoming data. Use AWS Step Functions to orchestrate workflows. Use Amazon Redshift Spectrum for one-time queries and analytical reporting. Use OpenSearch Dashboards in Amazon OpenSearch Service for the dashboard.
- D. Use an Amazon EMR cluster to process incoming data. Use AWS Step Functions to orchestrate workflows. Use Apache Hive for one-time queries and analytical reporting. Use Amazon OpenSearch Service to bulk ingest the data into compute optimized instances. Use OpenSearch Dashboards in OpenSearch Service for the dashboard.

Answer: B

Explanation:
For processing the incoming car listings in a cost-effective, scalable, and automated way, the ideal approach involves using AWS Glue for data processing, AWS Lambda with S3 Event Notifications for orchestration, Amazon Athena for one-time queries and analytical reporting, and Amazon QuickSight for visualization on the dashboard. Let's break this down:
AWS Glue: This is a fully managed ETL (Extract, Transform, Load) service that automatically processes the incoming data files. Glue is serverless and supports diverse data sources, including Amazon S3 and Redshift.
AWS Lambda and S3 Event Notifications: Using Lambda and S3 Event Notifications allows near real-time triggering of processing workflows as soon as new data is uploaded into S3. This approach is event-driven, ensuring that the listings are processed as soon as they are uploaded, reducing the latency for data processing.
Amazon Athena: A serverless, pay-per-query service that allows interactive queries directly against data in S3 using standard SQL. It is ideal for the requirement of one-time queries and analytical reporting without the need for provisioning or managing servers.
Amazon QuickSight: A business intelligence tool that integrates with a wide range of AWS data sources, including Athena, and is used for creating interactive dashboards. It scales well and provides real-time insights for the car listings.
This solution (Option D) is the most cost-effective, because both Glue and Athena are serverless and priced based on usage, reducing costs when compared to provisioning EMR clusters in the other options. Moreover, using Lambda for orchestration is more cost-effective than AWS Step Functions due to its lightweight nature.
Reference:
AWS Glue Documentation
Amazon Athena Documentation
Amazon QuickSight Documentation
S3 Event Notifications and Lambda

NEW QUESTION # 151
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. Use AWS Application Auto Scaling to schedule higher provisioned capacity for peak usage times. Schedule lower capacity during off-peak times.
- B. 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.
- C. Increase the provisioned capacity to the maximum capacity that is currently present during peak load 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: A

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 # 152
.....

According to the research of the past exams and answers, Fast2test provide you the latest Amazon Data-Engineer-Associate exercises and answers, which have have a very close similarity with real exam. Fast2test can promise that you can 100% pass your first time to attend Amazon Certification Data-Engineer-Associate Exam.

Exam Data-Engineer-Associate Blueprint: https://www.fast2test.com/Data-Engineer-Associate-premium-file.html

For any candidate, choosing the Data-Engineer-Associate question torrent material is the key to passing the exam, There are many Data-Engineer-Associate braindumps questions of our braindumps that appears in the Data-Engineer-Associate real test, you just need remember the Data-Engineer-Associate braindumps questions and the answers if you have no much time to prepare for your test, We guarantee a ninety-nine percent passing rate, which means you can pass exam as long as you review with our Data-Engineer-Associate latest dumps questions.

We has a long history of 10 years in designing the Data-Engineer-Associate exam guide and enjoys a good reputation across the globe, When testers find defects, the issue has to loop back to the developer.

For any candidate, choosing the Data-Engineer-Associate question torrent material is the key to passing the exam, There are many Data-Engineer-Associate braindumps questions of our braindumps that appears in the Data-Engineer-Associate real test, you just need remember the Data-Engineer-Associate braindumps questions and the answers if you have no much time to prepare for your test.

Try Amazon Data-Engineer-Associate Dumps to achieve wonderful results

We guarantee a ninety-nine percent passing rate, which means you can pass exam as long as you review with our Data-Engineer-Associate latest dumps questions, If you failed the exam, we promise you to full refund.

Also we have a strict information Data-Engineer-Associate system to make sure that your information will be safe and secret.

- 100% Pass 2026 Amazon Authoritative Data-Engineer-Associate: Reliable AWS Certified Data Engineer - Associate (DEA-C01) Exam Camp □ Search for ➥ Data-Engineer-Associate □ and obtain a free download on ➽ www.vce4dumps.com □ □Valid Data-Engineer-Associate Exam Discount
- Valid Data-Engineer-Associate Test Voucher □ Data-Engineer-Associate Reliable Test Braindumps ❤ Valid Data-Engineer-Associate Test Voucher □ Copy URL □ www.pdfvce.com □ open and search for ⇛ Data-Engineer-Associate ⇚ to download for free □Valid Data-Engineer-Associate Exam Discount
- Vce Data-Engineer-Associate Exam □ Data-Engineer-Associate Reliable Test Braindumps □ Valid Test Data-Engineer-Associate Experience □ Search for ▹ Data-Engineer-Associate ◃ and obtain a free download on [www.pdfdumps.com] □Valid Test Data-Engineer-Associate Experience
- Exam Data-Engineer-Associate Questions Fee □ Valid Data-Engineer-Associate Exam Discount □ Data-Engineer-Associate Instant Access □ Search on ➠ www.pdfvce.com □ for ➽ Data-Engineer-Associate □ to obtain exam materials for free download □Reliable Data-Engineer-Associate Exam Papers
- Data-Engineer-Associate Exam Braindumps - Data-Engineer-Associate Origination Questions - Data-Engineer-Associate Study Guide □ Open website { www.validtorrent.com } and search for ▸ Data-Engineer-Associate ◂ for free download □ □Valid Data-Engineer-Associate Exam Discount
- Free PDF Data-Engineer-Associate - Perfect Reliable AWS Certified Data Engineer - Associate (DEA-C01) Exam Camp □ Simply search for ➠ Data-Engineer-Associate □ for free download on [www.pdfvce.com] □Data-Engineer-Associate Reliable Test Braindumps
- Vce Data-Engineer-Associate Exam ⏲ Data-Engineer-Associate Valid Test Guide □ Valid Data-Engineer-Associate Exam Discount □ Download 《 Data-Engineer-Associate 》 for free by simply searching on ☀ www.testkingpass.com □☀□ □Data-Engineer-Associate Exam PDF
- Accurate Amazon - Reliable Data-Engineer-Associate Exam Camp □ Search for ▹ Data-Engineer-Associate ◃ and obtain a free download on ➠ www.pdfvce.com □ □Data-Engineer-Associate Reliable Test Braindumps
- Latest updated Reliable Data-Engineer-Associate Exam Camp - Verified Amazon Certification Training - Fantastic Amazon AWS Certified Data Engineer - Associate (DEA-C01) □ Open ➡ www.prep4away.com □□□ and search for [Data-Engineer-Associate] to download exam materials for free □Data-Engineer-Associate Test Dumps
- Data-Engineer-Associate Exam Braindumps - Data-Engineer-Associate Origination Questions - Data-Engineer-Associate Study Guide □ The page for free download of { Data-Engineer-Associate } on { www.pdfvce.com } will open immediately □Data-Engineer-Associate Certification Dumps
- 100% Pass 2026 Amazon Authoritative Data-Engineer-Associate: Reliable AWS Certified Data Engineer - Associate (DEA-C01) Exam Camp □ Simply search for ➠ Data-Engineer-Associate □ for free download on ➥ www.examcollectionpass.com □ □Exam Data-Engineer-Associate Questions Fee

- 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, 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, lms.ait.edu.za, 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, h20tradeskills.com, 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, 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, shortcourses.russellcollege.edu.au, Disposable vapes

BTW, DOWNLOAD part of Fast2test Data-Engineer-Associate dumps from Cloud Storage: https://drive.google.com/open?id=1k5N6Vc3YWtMB7zFe4lnM19wPZAcrUp9y