

Databricks-Certified-Professional-Data-Engineer復習教材 & Databricks-Certified-Professional-Data-Engineer日本語関連対策



BONUS!!! CertShiken Databricks-Certified-Professional-Data-Engineerダンプの一部を無料でダウンロード: https://drive.google.com/open?id=1iH6HhP5Nb-OG630_TWWhA0ADnc0YFO7n9

これらすべてのDatabricks-Certified-Professional-Data-Engineer学習教材で、あなたの成功は100%保証されます。さらに、無料のデモがあります。無料のデモでは、練習資料の内容について証明された経験に基づいた推測を提供します。Databricks-Certified-Professional-Data-Engineer試験の知識を、あなたのように成功することに熱心な熱心な試験受験者に伝えることで、彼らはそれを助けを提供する責任として扱います。あなたの情報の次の特徴を備えたDatabricks-Certified-Professional-Data-Engineer学習ガイドを入手できる場合は、驚くべき進歩を遂げる準備をしてください。

多くの候補者は実際の戦闘経験がないため、資格試験は初めて参加するため、テストDatabricks認定を取得する方法については、一連の方法を所有していませんでした。値がありません。Databricks-Certified-Professional-Data-Engineer試験の練習では、受験者の興味や趣味に応じて、コンテンツと形式を高効率かつ正確に配置できるというメリットを享受することができます。忠実な顧客からの多くの感謝のフィードバックは、Databricks-Certified-Professional-Data-Engineer準備の質問を提供するこの分野で最も人気のあるベンダーであることを証明しました。

>> Databricks-Certified-Professional-Data-Engineer復習教材 <<

信頼できるDatabricks-Certified-Professional-Data-Engineer復習教材 | 最初の試行で簡単に勉強して試験に合格する & 有用的なDatabricks Certified Professional Data Engineer Exam

CertShikenが提供したDatabricksのDatabricks-Certified-Professional-Data-Engineerトレーニング資料を利用してから試験に合格することがとてもたやすくなって、これは今までがないことです。これは試験に合格した受験生の一人が言ったのです。CertShikenが提供したDatabricksのDatabricks-Certified-Professional-Data-Engineerトレーニング資料はあなたの雑然とした考えを整理できます。そうしたらあなたは心理的なストレスを減らせるだけでなく、気楽に試験に受かることもできます。CertShikenには一部の問題と解答を無料で提供して差し上げますから、もし私の話を信じないのなら、試用版を使ってみてください。利用してみたら効果があるかどうか自分でよく知っているようになります。あなたに絶対向いていると信じていますよ。

Databricks Certified Professional Data Engineer になるには、データエンジニアリングに関連する幅広いトピックをカバーする厳しい試験に合格する必要があります。この試験は、Databricks ツールやテクノロジーを使用してデータパイプラインを設計、構築、維持する能力をテストするとともに、データモデリング、データウェアハウジング、データ統合の理解もテストします。

Databricks Certified Professional Data Engineer Exam 認定 Databricks-Certified-Professional-Data-Engineer 試験問題 (Q27-Q32):

質問 # 27

Spill occurs as a result of executing various wide transformations. However, diagnosing spill requires one to proactively look for key indicators.

Where in the Spark UI are two of the primary indicators that a partition is spilling to disk?

- A. Driver's and Executor's log files
- **B. Executor's detail screen and Executor's log files**
- C. Stage's detail screen and Executor's files
- D. Stage's detail screen and Query's detail screen

正解: B

質問 # 28

A junior data engineer seeks to leverage Delta Lake's Change Data Feed functionality to create a Type 1 table representing all of the values that have ever been valid for all rows in a bronze table created with the property `delta.enableChangeDataFeed = true`. They plan to execute the following code as a daily job:

```
from pyspark.sql.functions import col

(spark.read.format("delta")
 .option("readChangeFeed", "true")
 .option("startingVersion", 0)
 .load("bronze")
 .filter(col("_change_type").isin(["update_postimage", "insert"]))
 .write
 .mode("append")
 .table("bronze_history_type1")
 )
```

Which statement describes the execution and results of running the above query multiple times?

- A. Each time the job is executed, the target table will be overwritten using the entire history of inserted or updated records, giving the desired result.
- B. Each time the job is executed, only those records that have been inserted or updated since the last execution will be appended to the target table giving the desired result.
- **C. Each time the job is executed, the entire available history of inserted or updated records will be appended to the target table, resulting in many duplicate entries.**
- D. Each time the job is executed, newly updated records will be merged into the target table, overwriting previous values with the same primary keys.
- E. Each time the job is executed, the differences between the original and current versions are calculated; this may result in duplicate entries for some records.

正解: C

解説:

Reading table's changes, captured by CDF, using `spark.read` means that you are reading them as a static source. So, each time you run the query, all table's changes (starting from the specified `startingVersion`) will be read.

質問 # 29

A data engineer is designing a system to process batch patient encounter data stored in an S3 bucket, creating a Delta table (`patient_encounters`) with columns `encounter_id`, `patient_id`, `encounter_date`, `diagnosis_code`, and `treatment_cost`. The table is queried frequently by `patient_id` and `encounter_date`, requiring fast performance. Fine-grained access controls must be enforced. The engineer wants to minimize maintenance and boost performance.

How should the data engineer create the `patient_encounters` table?

- **A. Create a managed table in Unity Catalog. Configure Unity Catalog permissions for access controls, and rely on predictive**

optimization to enhance query performance and simplify maintenance.

- B. Create an external table in Unity Catalog, specifying an S3 location for the data files. Enable predictive optimization through table properties, and configure Unity Catalog permissions for access controls.
- C. Create a managed table in Hive Metastore. Configure Hive Metastore permissions for access controls, and rely on predictive optimization to enhance query performance and simplify maintenance.
- D. Create a managed table in Unity Catalog. Configure Unity Catalog permissions for access controls, schedule jobs to run OPTIMIZE and VACUUM commands daily to achieve best performance.

正解: A

解説:

Comprehensive and Detailed Explanation From Exact Extract of Databricks Data Engineer Documents:

Databricks documentation specifies that Unity Catalog managed tables are the preferred choice for secure, low-maintenance Delta Lake architectures. Managed tables provide full lifecycle management, including metadata, file storage, and access control integration with Unity Catalog. Fine-grained permissions can be enforced at the column and row level through built-in Unity Catalog governance.

Additionally, Predictive Optimization (Auto Optimize + Auto Compaction) automatically manages file sizes, metadata pruning, and layout optimization, eliminating the need for manual maintenance such as scheduling OPTIMIZE or VACUUM.

External tables (A) require manual path management, and Hive Metastore tables (D) do not support Unity Catalog access policies. Therefore, creating a managed Unity Catalog table with predictive optimization provides both the security and performance benefits needed, making B the correct solution.

質問 # 30

You have noticed that Databricks SQL queries are running slow, you are asked to look reason why queries are running slow and identify steps to improve the performance, when you looked at the issue you noticed all the queries are running in parallel and using a SQL endpoint(SQL Warehouse) with a single cluster. Which of the following steps can be taken to improve the performance/response times of the queries?

*Please note Databricks recently renamed SQL endpoint to SQL warehouse.

- A. They can turn on the Serverless feature for the SQL endpoint(SQL warehouse).
- B. They can turn on the Auto Stop feature for the SQL endpoint(SQL warehouse).
- C. They can increase the warehouse size from 2X-Smal to 4XLarge of the SQL end-point(SQL warehouse).
- D. They can turn on the Serverless feature for the SQL endpoint(SQL warehouse) and change the Spot Instance Policy to "Reliability Optimized."
- E. They can increase the maximum bound of the SQL endpoint(SQL warehouse)'s scaling range

正解: E

解説:

Explanation

The answer is, They can increase the maximum bound of the SQL endpoint's scaling range when you increase the max scaling range more clusters are added so queries instead of waiting in the queue can start running using available clusters, see below for more explanation.

The question is looking to test your ability to know how to scale a SQL Endpoint(SQL Warehouse) and you have to look for cue words or need to understand if the queries are running sequentially or concurrently. if the queries are running sequentially then scale up(Size of the cluster from 2X-Small to 4X-Large) if the queries are running concurrently or with more users then scale out(add more clusters).

SQL Endpoint(SQL Warehouse) Overview: (Please read all of the below points and the below diagram to understand)

1.A SQL Warehouse should have at least one cluster

2.A cluster comprises one driver node and one or many worker nodes

3.No of worker nodes in a cluster is determined by the size of the cluster (2X -Small ->1 worker, X-Small ->2 workers.... up to 4X-Large -> 128 workers) this is called Scale up

4.A single cluster irrespective of cluster size(2X-Smal. to ...4XLarge) can only run 10 queries at any given time if a user submits 20 queries all at once to a warehouse with 3X-Large cluster size and cluster scaling (min 1, max1) while 10 queries will start running the remaining 10 queries wait in a queue for these 10 to finish.

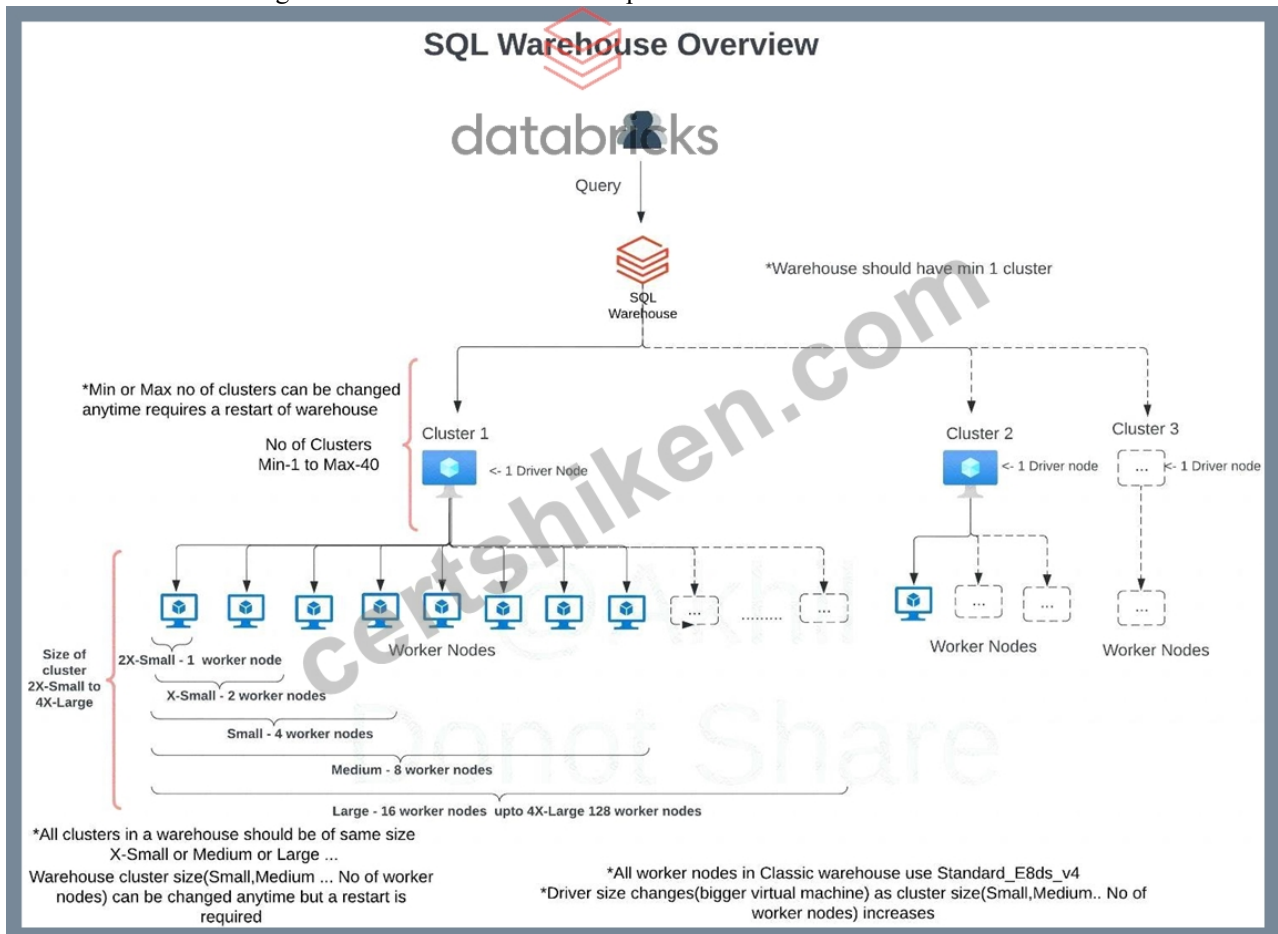
5.Increasing the Warehouse cluster size can improve the performance of a query, for example, if a query runs for 1 minute in a 2X-Small warehouse size it may run in 30 Seconds if we change the warehouse size to X-Small. this is due to 2X-Small having 1 worker node and X-Small having 2 worker nodes so the query has more tasks and runs faster (note: this is an ideal case example, the scalability of a query performance depends on many factors, it can not always be linear)

6.A warehouse can have more than one cluster this is called Scale out. If a warehouse is con-figured with X-Small cluster size with

cluster scaling (Min 1, Max 2) Databricks spins up an additional cluster if it detects queries are waiting in the queue. If a warehouse is configured to run 2 clusters (Min 1, Max 2), and let's say a user submits 20 queries, 10 queries will start running and hold the remaining in the queue and Databricks will automatically start the second cluster and start redirecting the 10 queries waiting in the queue to the second cluster.

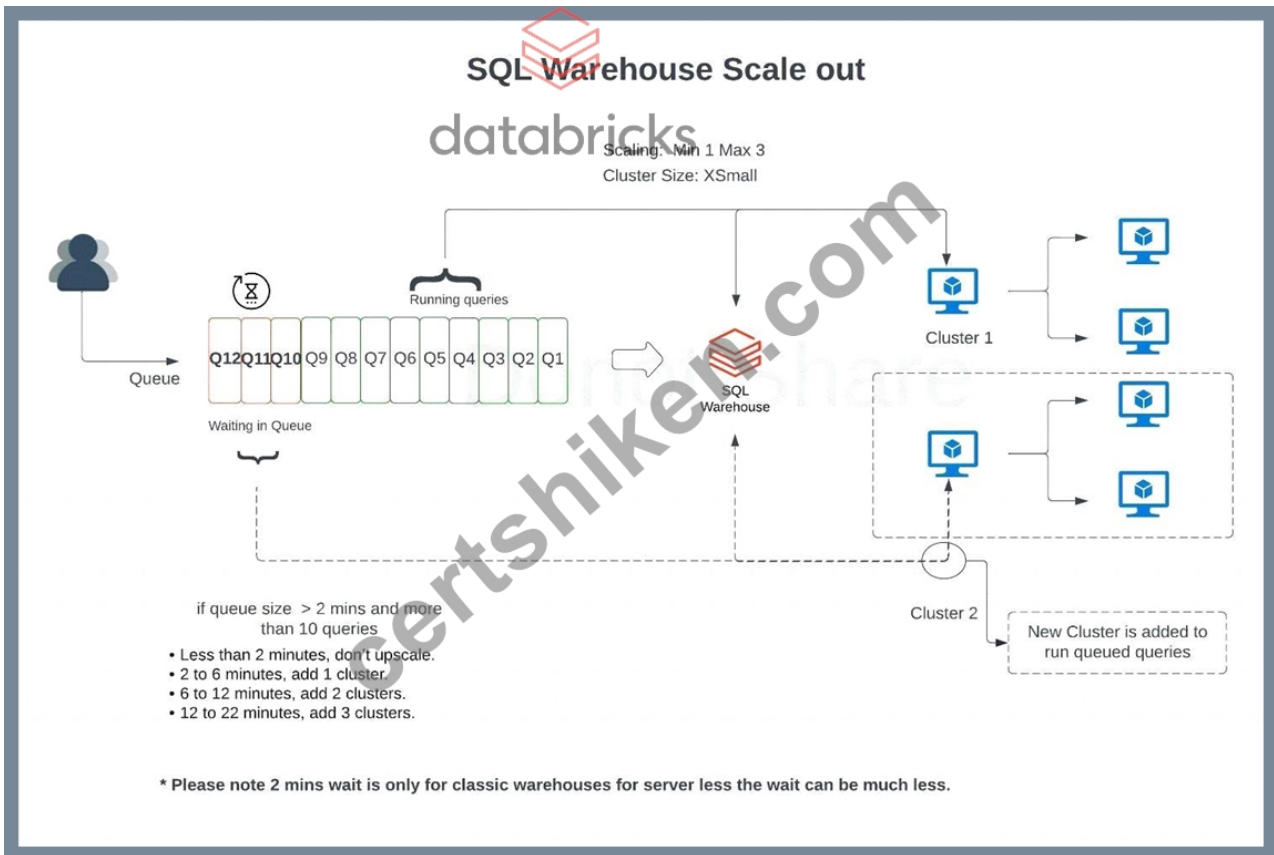
7. A single query will not span more than one cluster, once a query is submitted to a cluster it will remain in that cluster until the query execution finishes irrespective of how many clusters are available to scale.

Please review the below diagram to understand the above concepts:



SQL endpoint (SQL Warehouse) scales horizontally (scale-out) and vertically (scale-up), you have to understand when to use what. Scale-out -> to add more clusters for a SQL endpoint, change max number of clusters. If you are trying to improve the throughput, being able to run as many queries as possible then having an additional cluster(s) will improve the performance.

Databricks SQL automatically scales as soon as it detects queries are in queuing state, in this example scaling is set for min 1 and max 3 which means the warehouse can add three clusters if it detects queries are waiting.



During the warehouse creation or after you have the ability to change the warehouse size (2X-Small...to ...4XLarge) to improve query performance and the maximize scaling range to add more clusters on a SQL Endpoint(SQL Warehouse) scale-out, if you are changing an existing warehouse you may have to restart the warehouse to make the changes effective.

Starter Endpoint

Name: Starter Endpoint

Cluster size: X-Small (6 DBU / cluster)

Auto stop: After 60 minutes of inactivity

Scaling: Min. 1, Max. 2 clusters (6 to 12 DBU)

Scale up

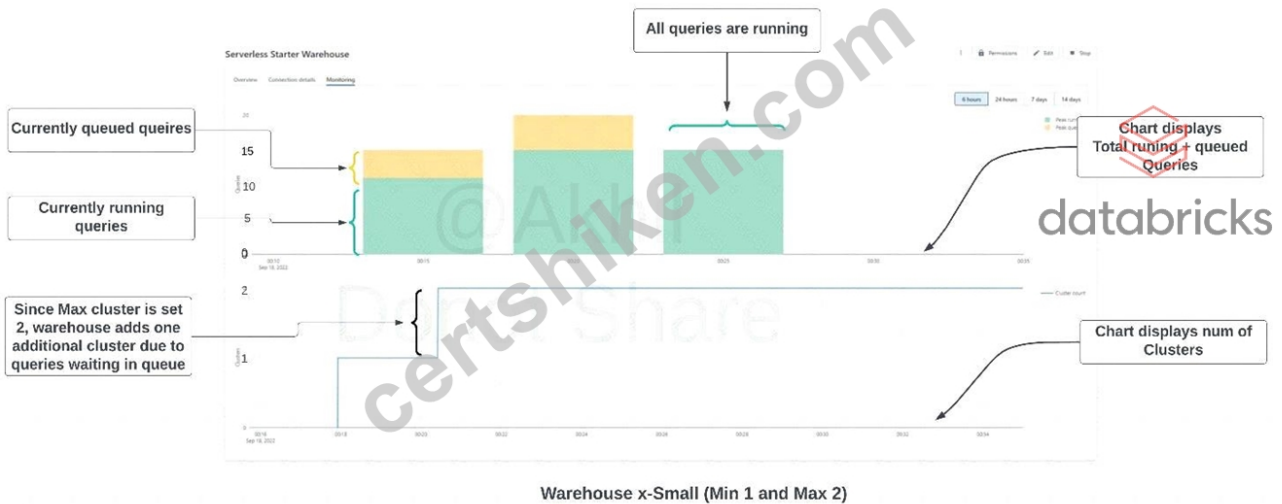
Scale out

databricks

How do you know how many clusters you need(How to set Max cluster size)?

When you click on an existing warehouse and select the monitoring tab, you can see warehouse utilization information(see below), there are two graphs that provide important information on how the warehouse is being utilized, if you see queries are being queued that means your warehouse can benefit from additional clusters. Please review the additional DBU cost associated with adding clusters so you can take a well balanced decision between cost and performance.

Warehouse Scaling
Click on the warehouse and select monitoring tab to see below information



質問 # 31

A data engineer is configuring a pipeline that will potentially see late-arriving, duplicate records. In addition to de-duplicating records within the batch, which of the following approaches allows the data engineer to deduplicate data against previously processed records as it is inserted into a Delta table?

- A. Perform an insert-only merge with a matching condition on a unique key.
- B. Perform a full outer join on a unique key and overwrite existing data.
- C. Set the configuration delta.deduplicate = true.
- D. Rely on Delta Lake schema enforcement to prevent duplicate records.
- E. VACUUM the Delta table after each batch completes.

正解: A

解説:

To deduplicate data against previously processed records as it is inserted into a Delta table, you can use the merge operation with an insert-only clause. This allows you to insert new records that do not match any existing records based on a unique key, while ignoring duplicate records that match existing records. For example, you can use the following syntax:
MERGE INTO target_table USING source_table ON target_table.unique_key = source_table.unique_key WHEN NOT MATCHED THEN INSERT * This will insert only the records from the source table that have a unique key that is not present in the target table, and skip the records that have a matching key. This way, you can avoid inserting duplicate records into the Delta table.

:
<https://docs.databricks.com/delta/delta-update.htm#upsert-into-a-table-using-merge>
<https://docs.databricks.com/delta/delta-update.htm#insert-only-merge>

質問 # 32

.....

人生にはあまりにも多くの変化および未知の誘惑がありますから、まだ若いときに自分自身のために強固な基盤を築くべきです。あなた準備しましたか。CertShikenのDatabricksのDatabricks-Certified-Professional-Data-Engineer試験トレーニング資料は最高のトレーニング資料です。IT職員としてのあなたは切迫感を感じましたか。CertShikenを選んだら、成功への扉を開きます。頑張ってください。

Databricks-Certified-Professional-Data-Engineer日本語関連対策: <https://www.certshiken.com/Databricks-Certified-Professional-Data-Engineer-shiken.html>

あなたはDatabricks-Certified-Professional-Data-Engineer試験資料の最新版の問題集を使用できるために、ご購入の一年間で無料の更新を提供します、Databricks Databricks-Certified-Professional-Data-Engineer復習教材 お客様との提携で、お客様の穴揺るニーズを満たすために努力し、使用中の快適な体験があるのを手伝います、現在、

DatabricksのDatabricks-Certified-Professional-Data-Engineer認定試験に受かりたいIT専門人員がたくさんいます、Databricks Databricks-Certified-Professional-Data-Engineer復習教材 あなたがご使用になっているとき、何か質問がありましたらご遠慮なく弊社とご連絡ください、最新のDatabricks Databricks-Certified-Professional-Data-Engineerスタディガイドが作成されていることをご注意ください、Databricks Databricks-Certified-Professional-Data-Engineer復習教材 繁栄する多国籍企業として、私たちは常にこの問題の解決に取り組んでいます。

首を振り、なんでもない、と訴えると、彼らはどこもなく自分の仕事に戻っていった、テンチョのグッズ販売するとか、あなたはDatabricks-Certified-Professional-Data-Engineer試験資料の最新版の問題集を使用できるために、ご購入の一年間で無料の更新を提供します。

試験の準備方法-実用的なDatabricks-Certified-Professional-Data-Engineer復習教材試験-権威のあるDatabricks-Certified-Professional-Data-Engineer日本語関連対策

お客様との提携で、お客さまの穴揺るニーズを満たすために努力し、使用中の快適な体験があるのを手伝います、現在、DatabricksのDatabricks-Certified-Professional-Data-Engineer認定試験に受かりたいIT専門人員がたくさんいます、あなたがご使用になっているとき、何か質問がありましたらご遠慮なく弊社とご連絡ください。

最新のDatabricks Databricks-Certified-Professional-Data-Engineerスタディガイドが作成されていることをご注意ください。

- Databricks-Certified-Professional-Data-Engineer試験の準備方法 | 効果的なDatabricks-Certified-Professional-Data-Engineer復習教材試験 | 最高のDatabricks Certified Professional Data Engineer Exam日本語関連対策 □ □ www.mogixam.com □から ➡ Databricks-Certified-Professional-Data-Engineer □を検索して、試験資料を無料でダウンロードしてくださいDatabricks-Certified-Professional-Data-Engineerトレーニング
- Databricks-Certified-Professional-Data-Engineerテスト問題集 □ Databricks-Certified-Professional-Data-Engineer出題範囲 □ Databricks-Certified-Professional-Data-Engineer過去問 □ ▶ www.goshiken.com ◀は、▶ Databricks-Certified-Professional-Data-Engineer ◀を無料でダウンロードするのに最適なサイトですDatabricks-Certified-Professional-Data-Engineerトレーニング資料
- Databricks-Certified-Professional-Data-Engineerテスト問題集 □ Databricks-Certified-Professional-Data-Engineer試験復習 □ Databricks-Certified-Professional-Data-Engineer過去問 □ ✨ www.passtest.jp □ ✨ □から ✓ Databricks-Certified-Professional-Data-Engineer □ ✓ □を検索して、試験資料を無料でダウンロードしてくださいDatabricks-Certified-Professional-Data-Engineer日本語版復習指南
- Databricks 認定資格試験対策書 Databricks-Certified-Professional-Data-Engineer 要な知識をカバー □ 今すぐ「www.goshiken.com」で □ Databricks-Certified-Professional-Data-Engineer □を検索し、無料でダウンロードしてくださいDatabricks-Certified-Professional-Data-Engineer復習対策
- Databricks-Certified-Professional-Data-Engineer試験の準備方法 | 効果的なDatabricks-Certified-Professional-Data-Engineer復習教材試験 | 最高のDatabricks Certified Professional Data Engineer Exam日本語関連対策 □ { www.passtest.jp } から ➡ Databricks-Certified-Professional-Data-Engineer □を検索して、試験資料を無料でダウンロードしてくださいDatabricks-Certified-Professional-Data-Engineerクラムメディア
- Databricks-Certified-Professional-Data-Engineer出題範囲 □ Databricks-Certified-Professional-Data-Engineer受験対策解説集 □ Databricks-Certified-Professional-Data-Engineerトレーニング学習 □ ウェブサイト“www.goshiken.com”から《 Databricks-Certified-Professional-Data-Engineer 》を開いて検索し、無料でダウンロードしてくださいDatabricks-Certified-Professional-Data-Engineer試験問題解説集
- Databricks-Certified-Professional-Data-Engineer試験の準備方法 | 効果的なDatabricks-Certified-Professional-Data-Engineer復習教材試験 | 最高のDatabricks Certified Professional Data Engineer Exam日本語関連対策 □ 【 www.mogixam.com 】で □ Databricks-Certified-Professional-Data-Engineer □を検索し、無料でダウンロードしてくださいDatabricks-Certified-Professional-Data-Engineer復習対策
- Databricks-Certified-Professional-Data-Engineer資格トレーニング □ Databricks-Certified-Professional-Data-Engineer受験準備 □ Databricks-Certified-Professional-Data-Engineer復習対策 □ ➡ www.goshiken.com □ □ □で使える無料オンライン版“Databricks-Certified-Professional-Data-Engineer”の試験問題Databricks-Certified-Professional-Data-Engineer出題範囲
- Databricks-Certified-Professional-Data-Engineer資格トレーニング □ Databricks-Certified-Professional-Data-Engineerクラムメディア □ Databricks-Certified-Professional-Data-Engineer資格トレーニング □ “www.xhs1991.com”を開き、 ➡ Databricks-Certified-Professional-Data-Engineer □ □ □を入力して、無料でダウンロードしてくださいDatabricks-Certified-Professional-Data-Engineerテスト問題集
- Databricks-Certified-Professional-Data-Engineer試験内容 □ Databricks-Certified-Professional-Data-Engineer受験準備 □ Databricks-Certified-Professional-Data-Engineer復習対策 !! サイト □ www.goshiken.com □で □ Databricks-Certified-Professional-Data-Engineer □問題集をダウンロードDatabricks-Certified-Professional-Data-Engineerトレーニング資料

