GoogleのProfessional-Data-Engineer認定試験の一番新しい問題集の登場

さらに、Fast2test Professional-Data-Engineerダンプの一部が現在無料で提供されています：https://drive.google.com/open?id=1cIC1E7akzgrHSRItnvyUvBQ7MpVo8LJD

Professional-Data-Engineerの学習教材は、テストの迅速な合格に役立ちます。認証を利用できます。多くの人が、Professional-Data-Engineer試験問題の助けを借りて、日々の仕事でより効率的に行動する能力を向上させています。弊社のProfessional-Data-Engineer学習教材を選択すると、あなたの夢がより明確に提示されます。次に、私の紹介を通じて、Professional-Data-Engineer学習クイズをより深く理解していただければ幸いです。 Professional-Data-Engineerの学習教材が試験に合格するための手助けになることを本当に願っています。

「あきらめたら そこで試合終了ですよ」という『スラムダンク』の中の安西監督が言った名言があります。この文は人々に知られています。試合と同じ、試験もそのどおりですよ。試験に準備する時間が十分ではないから、Professional-Data-Engineer認定試験を諦めた人がたくさんいます。しかし、優秀な資料を利用すれば、短時間の準備をしても、高得点で試験に合格することができます。信じないでしょうか。Fast2testの試験問題集はそのような資料ですよ。はやく試してください。

>> Professional-Data-Engineerソフトウエア <<

Professional-Data-Engineer問題集、Professional-Data-Engineer赤本勉強

私たちのProfessional-Data-Engineer試験参考書の品質は一番良いと言えます。そして、Professional-Data-Engineer試験参考書はすごく人気があります。まず、Professional-Data-Engineer試験参考書は専門家が作られました。また、専門家はProfessional-Data-Engineer試験参考書の更新に対して、定期的に検査を行います。だから、あなたはProfessional-Data-Engineer試験参考書の更新版を定期的に入手できます。

Google Certified Professional Data Engineer Exam 認定 Professional-Data-Engineer 試験問題 (Q173-Q178):

質問 # 173
You want to process payment transactions in a point-of-sale application that will run on Google Cloud Platform. Your user base could grow exponentially, but you do not want to manage infrastructure scaling.
Which Google database service should you use?

- A. Cloud Datastore
- B. Cloud SQL
- C. Cloud Bigtable
- D. BigQuery

正解：A

解説：
https://cloud.google.com/datastore/docs/concepts/overview

質問 # 174
Your company's customer_order table in BigOuery stores the order history for 10 million customers, with a table size of 10 PB. You need to create a dashboard for the support team to view the order history. The dashboard has two filters, countryname and username. Both are string data types in the BigQuery table. When a filter is applied, the dashboard fetches the order history from the table and displays the query results. However, the dashboard is slow to show the results when applying the filters to the following query:

```
SELECT date, order, status FROM customer_order
WHERE country = '<country_name>' AND username = '<username>'
```

How should you redesign the BigQuery table to support faster access?

- A. Partition the table by country and username fields.
- B. Partition the table by _PARTITIONTIME.
- C. Cluster the table by country and username fields
- D. Cluster the table by country field, and partition by username field.

正解：C

解説：
To improve the performance of querying a large BigQuery table with filters on countryname and username, clustering the table by these fields is the most effective approach. Here's why option C is the best choice:
Clustering in BigQuery:
Clustering organizes data based on the values in specified columns. This can significantly improve query performance by reducing the amount of data scanned during query execution.
Clustering by countryname and username means that data is physically sorted and stored together based on these fields, allowing BigQuery to quickly locate and read only the relevant data for queries using these filters.
Filter Efficiency:
With the table clustered by countryname and username, queries that filter on these columns can benefit from efficient data retrieval, reducing the amount of data processed and speeding up query execution.
This directly addresses the performance issue of the dashboard queries that apply filters on these fields.
Steps to Implement:
Redesign the Table:
Create a new table with clustering on countryname and username:
CREATE TABLE project.dataset.new_table
CLUSTER BY countryname, username AS
SELECT * FROM project.dataset.customer_order;
Migrate Data:
Transfer the existing data from the original table to the new clustered table.
Update Queries:
Modify the dashboard queries to reference the new clustered table.
Reference:
BigQuery Clustering Documentation
Optimizing Query Performance

質問 # 175
Your company's customer and order databases are often under heavy load. This makes performing analytics against them difficult without harming operations. The databases are in a MySQL cluster, with nightly backups taken using mysqldump. You want to perform analytics with minimal impact on operations. What should you do?

- A. Use an ETL tool to load the data from MySQL into Google BigQuery.
- B. Connect an on-premises Apache Hadoop cluster to MySQL and perform ETL.
- C. Mount the backups to Google Cloud SQL, and then process the data using Google Cloud Dataproc.
- D. Add a node to the MySQL cluster and build an OLAP cube there.

正解：B

解説：
Topic 2, Flowlogistic Case Study
Company Overview
Flowlogistic is a leading logistics and supply chain provider. They help businesses throughout the world manage their resources and transport them to their final destination. The company has grown rapidly, expanding their offerings to include rail, truck, aircraft, and oceanic shipping.
Company Background
The company started as a regional trucking company, and then expanded into other logistics market. Because they have not updated their infrastructure, managing and tracking orders and shipments has become a bottleneck. To improve operations, Flowlogistic developed proprietary technology for tracking shipments in real time at the parcel level. However, they are unable to deploy it because their technology stack, based on Apache Kafka, cannot support the processing volume. In addition, Flowlogistic wants to further analyze their orders and shipments to determine how best to deploy their resources.
Solution Concept
Flowlogistic wants to implement two concepts using the cloud:
* Use their proprietary technology in a real-time inventory-tracking system that indicates the location of their loads
* Perform analytics on all their orders and shipment logs, which contain both structured and unstructured data, to determine how best to deploy resources, which markets to expand info. They also want to use predictive analytics to learn earlier when a shipment will be delayed.
Existing Technical Environment
Flowlogistic architecture resides in a single data center:
* Databases
* 8 physical servers in 2 clusters
* SQL Server - user data, inventory, static data
* 3 physical servers
* Cassandra - metadata, tracking messages
10 Kafka servers - tracking message aggregation and batch insert
* Application servers - customer front end, middleware for order/customs
* 60 virtual machines across 20 physical servers
* Tomcat - Java services
* Nginx - static content
* Batch servers
Storage appliances
* iSCSI for virtual machine (VM) hosts
* Fibre Channel storage area network (FC SAN) - SQL server storage
* Network-attached storage (NAS) image storage, logs, backups
* Apache Hadoop /Spark servers
* Core Data Lake
* Data analysis workloads
* 20 miscellaneous servers
* Jenkins, monitoring, bastion hosts,
Business Requirements
* Build a reliable and reproducible environment with scaled panty of production.
* Aggregate data in a centralized Data Lake for analysis
* Use historical data to perform predictive analytics on future shipments
* Accurately track every shipment worldwide using proprietary technology
* Improve business agility and speed of innovation through rapid provisioning of new resources
* Analyze and optimize architecture for performance in the cloud
* Migrate fully to the cloud if all other requirements are met
Technical Requirements
* Handle both streaming and batch data
* Migrate existing Hadoop workloads
* Ensure architecture is scalable and elastic to meet the changing demands of the company.
* Use managed services whenever possible
* Encrypt data flight and at rest
* Connect a VPN between the production data center and cloud environment SEO Statement We have grown so quickly that our inability to upgrade our infrastructure is really hampering further growth and efficiency. We are efficient at moving shipments around the world, but we are inefficient at moving data around.
We need to organize our information so we can more easily understand where our customers are and what they are shipping.
CTO Statement

IT has never been a priority for us, so as our data has grown, we have not invested enough in our technology. I have a good staff to manage IT, but they are so busy managing our infrastructure that I cannot get them to do the things that really matter, such as organizing our data, building the analytics, and figuring out how to implement the CFO' s tracking technology.
CFO Statement
Part of our competitive advantage is that we penalize ourselves for late shipments and deliveries. Knowing where out shipments are at all times has a direct correlation to our bottom line and profitability. Additionally, I don't want to commit capital to building out a server environment.

質問 # 176
Does Dataflow process batch data pipelines or streaming data pipelines?

- A. None of the above
- B. Only Batch Data Pipelines
- C. Both Batch and Streaming Data Pipelines
- D. Only Streaming Data Pipelines

正解：C

解説：
Dataflow is a unified processing model, and can execute both streaming and batch data pipelines Reference:
https://cloud.google.com/dataflow/

質問 # 177
Case Study: 2 - MJTelco
Company Overview
MJTelco is a startup that plans to build networks in rapidly growing, underserved markets around the world. The company has patents for innovative optical communications hardware. Based on these patents, they can create many reliable, high-speed backbone links with inexpensive hardware.
Company Background
Founded by experienced telecom executives, MJTelco uses technologies originally developed to overcome communications challenges in space. Fundamental to their operation, they need to create a distributed data infrastructure that drives real-time analysis and incorporates machine learning to continuously optimize their topologies. Because their hardware is inexpensive, they plan to overdeploy the network allowing them to account for the impact of dynamic regional politics on location availability and cost. Their management and operations teams are situated all around the globe creating many-to- many relationship between data consumers and provides in their system. After careful consideration, they decided public cloud is the perfect environment to support their needs.
Solution Concept
MJTelco is running a successful proof-of-concept (PoC) project in its labs. They have two primary needs:
Scale and harden their PoC to support significantly more data flows generated when they ramp to more than 50,000 installations.
Refine their machine-learning cycles to verify and improve the dynamic models they use to control topology definition.
MJTelco will also use three separate operating environments ?development/test, staging, and production ?
to meet the needs of running experiments, deploying new features, and serving production customers.
Business Requirements
Scale up their production environment with minimal cost, instantiating resources when and where needed in an unpredictable, distributed telecom user community. Ensure security of their proprietary data to protect their leading-edge machine learning and analysis.
Provide reliable and timely access to data for analysis from distributed research workers Maintain isolated environments that support rapid iteration of their machine-learning models without affecting their customers.
Technical Requirements
Ensure secure and efficient transport and storage of telemetry data Rapidly scale instances to support between 10,000 and 100,000 data providers with multiple flows each.
Allow analysis and presentation against data tables tracking up to 2 years of data storing approximately
100m records/day
Support rapid iteration of monitoring infrastructure focused on awareness of data pipeline problems both in telemetry flows and in production learning cycles.
CEO Statement
Our business model relies on our patents, analytics and dynamic machine learning. Our inexpensive hardware is organized to be highly reliable, which gives us cost advantages. We need to quickly stabilize our large distributed data pipelines to meet our reliability and capacity commitments.

CTO Statement
Our public cloud services must operate as advertised. We need resources that scale and keep our data secure. We also need environments in which our data scientists can carefully study and quickly adapt our models. Because we rely on automation to process our data, we also need our development and test environments to work as we iterate.
CFO Statement
The project is too large for us to maintain the hardware and software required for the data and analysis.
Also, we cannot afford to staff an operations team to monitor so many data feeds, so we will rely on automation and infrastructure. Google Cloud's machine learning will allow our quantitative researchers to work on our high-value problems instead of problems with our data pipelines.
MJTelco needs you to create a schema in Google Bigtable that will allow for the historical analysis of the last 2 years of records. Each record that comes in is sent every 15 minutes, and contains a unique identifier of the device and a data record. The most common query is for all the data for a given device for a given day. Which schema should you use?

- A. Rowkey: dateColumn data: device_id, data_point
- B. Rowkey: date#device_idColumn data: data_point
- C. Rowkey: device_idColumn data: date, data_point
- D. Rowkey: date#data_pointColumn data: device_id
- E. Rowkey: data_pointColumn data: device_id, date

正解：E

質問 # 178
.....

IT認証試験に合格したい受験生の皆さんはきっと試験の準備をするために大変悩んでいるでしょう。しかし準備しなければならないのですから、落ち着かない心理になりました。しかし、Fast2testのGoogleのProfessional-Data-Engineerトレーニング資料を利用してから、その落ち着かない心はなくなった人がたくさんいます。Fast2testのGoogleのProfessional-Data-Engineerトレーニング資料を持っていたら、自信を持つようになります。試験に合格しない心配する必要がないですから、気楽に試験を受けることができます。これは心のヘルプだけではなく、試験に合格することで、明るい明日を持つこともできるようになります。

Professional-Data-Engineer問題集: https://jp.fast2test.com/Professional-Data-Engineer-premium-file.html

顧客の信頼を確立し、間違った試験問題を選択することによる損失を避けるために、購入前にダウンロードできるProfessional-Data-Engineer試験問題の関連する無料デモを提供しています、また、当社のProfessional-Data-Engineerテストガイドに連絡する機会もあります、Fast2testのProfessional-Data-Engineer問題集を購入する前に、問題集の無料なサンプルをダウンロードして試用してもいいです、Professional-Data-Engineer試験問題を購入するだけで、Professional-Data-Engineer試験に簡単に合格できます、ただし、PayPalはProfessional-Data-Engineerの最新の試験問題集に追加の税金を支払って、売り手と買い手のアカウントが安全であることを保証できます、しかし、Professional-Data-Engineer試験のために、どんな資料がいいですか、でもFast2test Professional-Data-Engineer問題集の最新問題集がこの問題を解決できますよ。

凄味まで見えそうな完璧な笑顔で言葉もなく坂崎が運転手Professional-Data-Engineerを睨みつけていると、車はいつの間にか琥珀へと到着していた、監督としてしっかりするべきだったんだ、顧客の信頼を確立し、間違った試験問題を選択することによる損失を避けるために、購入前にダウンロードできるProfessional-Data-Engineer試験問題の関連する無料デモを提供しています。

信頼的なProfessional-Data-Engineerソフトウエア & 合格スムーズProfessional-Data-Engineer問題集 | ユニークなProfessional-Data-Engineer赤本勉強

また、当社のProfessional-Data-Engineerテストガイドに連絡する機会もあります、Fast2testのProfessional-Data-Engineer問題集を購入する前に、問題集の無料なサンプルをダウンロードして試用してもいいです、Professional-Data-Engineer試験問題を購入するだけで、Professional-Data-Engineer試験に簡単に合格できます。

ただし、PayPalはProfessional-Data-Engineerの最新の試験問題集に追加の税金を支払って、売り手と買い手のアカウントが安全であることを保証できます。

- Professional-Data-Engineer全真模擬試験 □ Professional-Data-Engineer無料問題 □ Professional-Data-Engineerブロンズ教材 □ （ jp.fast2test.com ）を開いて➥ Professional-Data-Engineer □を検索し、試験資料を無料でダウ

ンロードしてくださいProfessional-Data-Engineer関連問題資料

- 真実的なProfessional-Data-Engineerソフトウエア試験-試験の準備方法-効率的なProfessional-Data-Engineer問題集 □ □ Professional-Data-Engineer □を無料でダウンロード□ www.goshiken.com □ウェブサイトを入力するだけProfessional-Data-Engineer受験料過去問
- Professional-Data-Engineer模擬資料 □ Professional-Data-Engineer練習問題 □ Professional-Data-Engineer試験勉強書 〜 "jp.fast2test.com"サイトにて➽ Professional-Data-Engineer □問題集を無料で使おうProfessional-Data-Engineer対応内容
- 試験の準備方法-一番優秀なProfessional-Data-Engineerソフトウエア試験-効果的なProfessional-Data-Engineer問題集 □ ▹ www.goshiken.com ◃サイトにて[Professional-Data-Engineer]問題集を無料で使おうProfessional-Data-Engineer基礎訓練
- 真実的なProfessional-Data-Engineerソフトウエア試験-試験の準備方法-効率的なProfessional-Data-Engineer問題集 □ ➤ www.japancert.com □に移動し、（ Professional-Data-Engineer ）を検索して無料でダウンロードしてくださいProfessional-Data-Engineer試験解説
- Professional-Data-Engineer関連問題資料 □ Professional-Data-Engineerキャリアパス □ Professional-Data-Engineer日本語版と英語版 □ サイト「 www.goshiken.com 」で➽ Professional-Data-Engineer □問題集をダウンロードProfessional-Data-Engineer復習対策
- 信頼的なProfessional-Data-Engineerソフトウエア一回合格-一番優秀なProfessional-Data-Engineer問題集 ↗ ➠ Professional-Data-Engineer □を無料でダウンロード➥ www.mogiexam.com □で検索するだけProfessional-Data-Engineer認証pdf資料
- 信頼的なProfessional-Data-Engineerソフトウエア一回合格-一番優秀なProfessional-Data-Engineer問題集 □ （ www.goshiken.com ）は、《 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Professional-Data-Engineer日本語版と英語版
- Professional-Data-Engineer受験料過去問 □ Professional-Data-Engineerテスト模擬問題集 □ Professional-Data-Engineer試験解説 ➡ Open Webサイト□ www.passtest.jp □検索➤ Professional-Data-Engineer □無料ダウンロードProfessional-Data-Engineer対応内容
- ユニークなProfessional-Data-Engineerソフトウエア試験-試験の準備方法-最高のProfessional-Data-Engineer問題集 □ 今すぐ（ www.goshiken.com ）を開き、▹ Professional-Data-Engineer ◃を検索して無料でダウンロードしてくださいProfessional-Data-Engineer受験資格
- Professional-Data-Engineerテスト模擬問題集 □ Professional-Data-Engineerテスト模擬問題集 ➡□ Professional-Data-Engineer日本語対策 □ ➥ www.mogiexam.com □に移動し、✔ Professional-Data-Engineer □✔□を検索して、無料でダウンロード可能な試験資料を探しますProfessional-Data-Engineer受験資格
- bbs.longmenshentu.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, www.stes.tyc.edu.tw, online.citinstitute.org, 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, 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, www.stes.tyc.edu.tw, Disposable vapes

ちなみに、Fast2test Professional-Data-Engineerの一部をクラウドストレージからダウンロードできます：https://drive.google.com/open?id=1cIC1E7akzgrHSRItnvyUvBQ7MpVo8LJD