

Professional-Data-Engineer Exam Torrents: Google Certified Professional Data Engineer Exam Prepare Torrents & Professional-Data-Engineer Test Braindumps

Google Data Engineer Certification Exam Blueprint	
Domain	Subject Areas
1. Design Data Processing Systems	<ul style="list-style-type: none">• Design flexible data representations• Design data processing infrastructure• Design data pipelines
2. Build and Maintain Data Structures and Databases	<ul style="list-style-type: none">• Build and maintain flexible data representations• Build and maintain processing infrastructure• Build and maintain pipelines
3. Analyze Data and Enable Machine Learning	<ul style="list-style-type: none">• Analyze data• Machine learning• Machine learning model deployment
4. Modeling Business Processes for Optimization and Analysis	<ul style="list-style-type: none">• Map business requirements to data presentations• Optimize data representations
5. Ensure Reliability	<ul style="list-style-type: none">• Perform quality control• Assess, improve, and troubleshoot data representations• Assess, improve, and troubleshoot data processing infrastructure• Recovery of data
6. Visualize Data and Advocate Policy	<ul style="list-style-type: none">• Build or select data visualization and reporting tools• Advocate policies and publish data and reports
7. Design for Security and Compliance	<ul style="list-style-type: none">• Design secure data infrastructure and processes• Design for legal compliance

BTW, DOWNLOAD part of Pass4Leader Professional-Data-Engineer dumps from Cloud Storage: https://drive.google.com/open?id=1TDIsNY_1et0a4fg2EejzC8OS7cbGJYwJ

Obtaining the certification may be not an easy thing for some candidates. If you choose us, we can help you pass the exam and obtain corresponding certification easily. Professional-Data-Engineer learning materials are edited by professional experts, and you can use them at ease. Furthermore, Professional-Data-Engineer exam braindumps have the most of the knowledge points for the exam, and you can learn a lot in the process of learning. We offer you free update for 365 days after payment for Professional-Data-Engineer Exam Dumps, and our system will send you the latest version automatically. We have online and offline service, if you have any questions, you can consult us.

Among all substantial practice materials with similar themes, our Professional-Data-Engineer practice materials win a majority of credibility for promising customers who are willing to make progress in this line. With excellent quality at attractive price, our Professional-Data-Engineer practice materials get high demand of orders in this fierce market with passing rate up to 98 to 100 percent all these years. We shall highly appreciate your acceptance of our Professional-Data-Engineer practice materials and your decision will lead you to bright future with highly useful certificates.

>> Professional-Data-Engineer Valid Vce <<

TOP Professional-Data-Engineer Valid Vce - Google Google Certified Professional Data Engineer Exam - Latest Professional-Data-Engineer Dumps Collection

Revised and updated according to the syllabus changes and all the latest developments in theory and practice, our Google Certified

Professional Data Engineer Exam dumps are highly relevant to what you actually need to get through the certifications tests. Moreover they impart you information in the format of Professional-Data-Engineer Questions and answers that is actually the format of your real certification test. Hence not only you get the required knowledge but also find the opportunity to practice real exam scenario. For consolidation of your learning, our Google Certified Professional Data Engineer Exam dumps PDF file also provide you sets of practice questions and answers. Doing them again and again, you enrich your knowledge and maximize chances of an outstanding exam success.

Google Certified Professional Data Engineer Exam Sample Questions (Q84-Q89):

NEW QUESTION # 84

You are administering a BigQuery on-demand environment. Your business intelligence tool is submitting hundreds of queries each day that aggregate a large (50 TB) sales history fact table at the day and month levels. These queries have a slow response time and are exceeding cost expectations. You need to decrease response time, lower query costs, and minimize maintenance. What should you do?

- A. Enable BI Engine and add your sales table as a preferred table.
- B. Create a scheduled query to build sales day and sales month aggregate tables on an hourly basis.
- C. Build authorized views on top of the sales table to aggregate data at the day and month level.
- **D. Build materialized views on top of the sales table to aggregate data at the day and month level.**

Answer: D

Explanation:

To improve response times and reduce costs for frequent queries aggregating a large sales history fact table, materialized views are a highly effective solution. Here's why option A is the best choice:

Materialized Views:

Materialized views store the results of a query physically and update them periodically, offering faster query responses for frequently accessed data.

They are designed to improve performance for repetitive and expensive aggregation queries by precomputing the results.

Efficiency and Cost Reduction:

By building materialized views at the day and month level, you significantly reduce the computation required for each query, leading to faster response times and lower query costs.

Materialized views also reduce the need for on-demand query execution, which can be costly when dealing with large datasets.

Minimized Maintenance:

Materialized views in BigQuery are managed automatically, with updates handled by the system, reducing the maintenance burden on your team.

Steps to Implement:

Identify Aggregation Queries:

Analyze the existing queries to identify common aggregation patterns at the day and month levels.

Create Materialized Views:

Create materialized views in BigQuery for the identified aggregation patterns. For example `CREATE MATERIALIZED VIEW`

`project.dataset.sales_daily_summary AS SELECT DATE(transaction_time) AS day, SUM(amount) AS total_sales FROM`

`project.dataset.sales GROUP BY day;` `CREATE MATERIALIZED VIEW` `project.dataset.sales_monthly_summary AS SELECT`

`EXTRACT(YEAR FROM transaction_time) AS year, EXTRACT(MONTH FROM transaction_time) AS month, SUM(amount)`

`AS total_sales FROM project.dataset.sales GROUP BY year, month;` Query Using Materialized Views:

Update existing queries to use the materialized views instead of directly querying the base table.

Reference Links:

[BigQuery Materialized Views](#)

[Optimizing Query Performance](#)

NEW QUESTION # 85

You are designing a data mesh on Google Cloud by using Dataplex to manage data in BigQuery and Cloud Storage. You want to simplify data asset permissions. You are creating a customer virtual lake with two user groups:

- * Data engineers, which require full data lake access
- * Analytic users, which require access to curated data

You need to assign access rights to these two groups. What should you do?

- A. 1. Grant the `dataplex.dataReader` role to the data engineer group on the customer data lake.

- 2. Grant the dataplex.dataOwner to the analytic user group on the customer curated zone.
- B. 1. Grant the bigquery.dataownex role on BigQuery datasets and the storage.objectcreator role on Cloud Storage buckets to data engineers.
2. Grant the bigquery.dataViewer role on BigQuery datasets and the storage.objectViewer role on Cloud Storage buckets to analytic users.
- C. 1. Grant the dataplex.dataOwner role to the data engineer group on the customer data lake.
2. Grant the dataplex.dataReader role to the analytic user group on the customer curated zone.
- D. 1. Grant the bigquery.dataViewer role on BigQuery datasets and the storage.objectviewer role on Cloud Storage buckets to data engineers.
2. Grant the bigquery.dataOwner role on BigQuery datasets and the storage.objectEditor role on Cloud Storage buckets to analytic users.

Answer: C

Explanation:

When designing a data mesh on Google Cloud using Dataplex to manage data in BigQuery and Cloud Storage, it is essential to simplify data asset permissions while ensuring that each user group has the appropriate access levels. Here's why option A is the best choice:

Data Engineer Group:

Data engineers require full access to the data lake to manage and operate data assets comprehensively. Granting the dataplex.dataOwner role to the data engineer group on the customer data lake ensures they have the necessary permissions to create, modify, and delete data assets within the lake.

Analytic User Group:

Analytic users need access to curated data but do not require full control over all data assets. Granting the dataplex.dataReader role to the analytic user group on the customer curated zone provides read-only access to the curated data, enabling them to analyze the data without the ability to modify or delete it.

Steps to Implement:

Grant Data Engineer Permissions:

Assign the dataplex.dataOwner role to the data engineer group on the customer data lake to ensure full access and management capabilities.

Grant Analytic User Permissions:

Assign the dataplex.dataReader role to the analytic user group on the customer curated zone to provide read-only access to curated data.

Reference:

Dataplex IAM Roles and Permissions

Managing Access in Dataplex

NEW QUESTION # 86

Which of the following statements about the Wide & Deep Learning model are true? (Select 2 answers.)

- A. A good use for the wide and deep model is a recommender system.
- B. A good use for the wide and deep model is a small-scale linear regression problem.
- C. The wide model is used for memorization, while the deep model is used for generalization.
- D. The wide model is used for generalization, while the deep model is used for memorization.

Answer: A,C

Explanation:

Can we teach computers to learn like humans do, by combining the power of memorization and generalization? It's not an easy question to answer, but by jointly training a wide linear model (for memorization) alongside a deep neural network (for generalization), one can combine the strengths of both to bring us one step closer. At Google, we call it Wide & Deep Learning. It's useful for generic large-scale regression and classification problems with sparse inputs (categorical features with a large number of possible feature values), such as recommender systems, search, and ranking problems.

Reference: <https://research.googleblog.com/2016/06/wide-deep-learning-better-together-with.html>

NEW QUESTION # 87

You are administering shared BigQuery datasets that contain views used by multiple teams in your organization. The marketing team is concerned about the variability of their monthly BigQuery analytics spend using the on-demand billing model. You need to help the marketing team establish a consistent BigQuery analytics spend each month. What should you do?

- A. Create a BigQuery Standard pay-as-you go reservation with a baseline of 0 slots and autoscaling set to 500 for the marketing team, and bill them back accordingly.
- **B. Create a BigQuery reservation with a baseline of 500 slots with no autoscaling for the marketing team, and bill them back accordingly.**
- C. Establish a BigQuery quota for the marketing team, and limit the maximum number of bytes scanned each day.
- D. Create a BigQuery Enterprise reservation with a baseline of 250 slots and autoscaling set to 500 for the marketing team, and bill them back accordingly.

Answer: B

Explanation:

To help the marketing team establish a consistent BigQuery analytics spend each month, you can use BigQuery reservations to allocate dedicated slots for their queries. This provides predictable costs by reserving a fixed amount of compute resources.

BigQuery Reservations:

BigQuery Reservations allow you to purchase dedicated query processing capacity in the form of slots.

By reserving slots, you can control costs and ensure that the marketing team has the necessary resources for their queries without unexpected increases in spending.

Baseline Slots:

Setting a baseline of 500 slots without autoscaling ensures a consistent allocation of resources.

This provides a predictable monthly cost, as the marketing team will be billed for the reserved slots regardless of actual usage.

Billing Back:

The marketing team's usage can be billed back based on the fixed reservation cost, ensuring budget predictability.

This approach avoids the variability associated with on-demand billing, where costs can fluctuate based on query volume and complexity.

No Autoscaling:

By not enabling autoscaling, you prevent additional costs from being incurred due to temporary increases in query demand.

This fixed reservation ensures that the marketing team only uses the allocated 500 slots, maintaining a consistent monthly spend.

Google Data Engineer Reference:

BigQuery Reservations Documentation

BigQuery Slot Reservations

Managing BigQuery Costs

Using a fixed reservation of 500 slots provides the marketing team with predictable costs and the necessary resources for their queries without unexpected billing variability.

NEW QUESTION # 88

You are integrating one of your internal IT applications and Google BigQuery, so users can query BigQuery from the application's interface. You do not want individual users to authenticate to BigQuery and you do not want to give them access to the dataset. You need to securely access BigQuery from your IT application.

What should you do?

- A. Integrate with a single sign-on (SSO) platform, and pass each user's credentials along with the query request
- B. Create groups for your users and give those groups access to the dataset
- **C. Create a service account and grant dataset access to that account. Use the service account's private key to access the dataset**
- D. Create a dummy user and grant dataset access to that user. Store the username and password for that user in a file on the files system, and use those credentials to access the BigQuery dataset

Answer: C

Explanation:

Explanation

NEW QUESTION # 89

.....

Actually, Professional-Data-Engineer exam really make you anxious. You may have been suffering from the complex study materials, why not try our Professional-Data-Engineer exam software of Pass4Leader to ease your burden. Our IT elite finally designs the best Professional-Data-Engineer exam study materials by collecting the complex questions and analyzing the focal points of the exam over

myportal.utt.edu.tt, myportal.utt.edu.tt, kianafciy745909.wikienlightenment.com, orlandomzf866766.wikisona.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, Disposable vapes

DOWNLOAD the newest Pass4Leader Professional-Data-Engineer PDF dumps from Cloud Storage for free:
https://drive.google.com/open?id=1TDIsNY_1et0a4fg2EejzC8OS7cbGJYUJ