

Exam Professional-Data-Engineer Registration & Related Professional-Data-Engineer Certifications



DOWNLOAD the newest PracticeDump Professional-Data-Engineer PDF dumps from Cloud Storage for free:
<https://drive.google.com/open?id=1-Cu6c7sOfjRBdg8D7zp2WwNw52yv1YEw>

We respect the private information of our customers. If you buy the Professional-Data-Engineer exam materials from us, your personal information will be protected well. Once the payment finished, we will not look the information of you, and we also won't send the junk mail to your email address. What's more, we offer you free update for 365 days for Professional-Data-Engineer Exam Dumps, so that you can get the recent information for the exam. The latest version will be automatically sent to you by our system, if you have any other questions, just contact us.

This course will show you how to manage big data including loading, extracting, cleaning, and validating data. At the end of the training, you can easily create machine learning and statistical models as well as visualizing query results. This program is a bit lengthy but you have to practice well to get the knowledge needed on the actual exam. These are the following modules covered in the course:

- Big Data Analytics with Cloud AI Platform Notebook
- Bigtable Streaming Features and High-Throughput BigQuery
- Handling Data Pipelines with Cloud Composer and Cloud Data Fusion
- Building a Data Warehouse
- Serverless Data Processing with Cloud Dataflow
- Introduction to Data Engineering
- Production ML Pipelines and use of Kubeflow
- Introduction to Processing Streaming Data
- Prebuilt ML Models APIs for Unsaturated Data

These modules involve everything the candidate requires for passing the Professional Data Engineer certification exam. Thus, you will not miss anything if you are taking this learning program keenly and apply the required knowledge in an appropriate way. You would end up getting a good score and achieving the Google Professional Data Engineer certification.

>> Exam Professional-Data-Engineer Registration <<

Trustable Exam Professional-Data-Engineer Registration Help You to Get

Acquainted with Real Professional-Data-Engineer Exam Simulation

With the development of information and communications technology, we are now living in a globalized world. Professional-Data-Engineer information technology learning is correspondingly popular all over the world. Modern technology has changed the way how we live and work. When it comes to the study materials selling in the market, qualities are patchy. But our Professional-Data-Engineer test material has been recognized by multitude of customers, which possess of the top-class quality, can help you pass exam successfully. On the other hand, our Professional-Data-Engineer Latest Dumps are designed by the most experienced experts, thus it can not only teach you knowledge, but also show you the method of learning in the most brief and efficient ways.

Google Certified Professional Data Engineer Exam Sample Questions (Q261-Q266):

NEW QUESTION # 261

Which SQL keyword can be used to reduce the number of columns processed by BigQuery?

- A. BETWEEN
- B. LIMIT
- C. WHERE
- D. SELECT

Answer: D

Explanation:

SELECT allows you to query specific columns rather than the whole table.

LIMIT, BETWEEN, and WHERE clauses will not reduce the number of columns processed by BigQuery.

Reference:

https://cloud.google.com/bigquery/launch-checklist#architecture_design_and_development_checklist

NEW QUESTION # 262

Your company is implementing a data warehouse using BigQuery and you have been tasked with designing the data model. You move your on-premises sales data warehouse with a star data schema to BigQuery but notice performance issues when querying the data of the past 30 days. Based on Google's recommended practices, what should you do to speed up the query without increasing storage costs?

- A. Denormalize the data
- B. Materialize the dimensional data in views
- C. Partition the data by transaction date
- D. Shard the data by customer ID

Answer: C

NEW QUESTION # 263

You created an analytics environment on Google Cloud so that your data scientist team can explore data without impacting the on-premises Apache Hadoop solution. The data in the on-premises Hadoop Distributed File System (HDFS) cluster is in Optimized Row Columnar (ORC) formatted files with multiple columns of Hive partitioning. The data scientist team needs to be able to explore the data in a similar way as they used the on-premises HDFS cluster with SQL on the Hive query engine. You need to choose the most cost-effective storage and processing solution. What should you do?

- A. Copy the ORC files on Cloud Storage, then create external BigQuery tables for the data scientist team.
- B. Import the ORC files to Bigtable tables for the data scientist team.
- C. Import the ORC files to BigQuery tables for the data scientist team.
- D. Copy the ORC files on Cloud Storage, then deploy a Dataproc cluster for the data scientist team.

Answer: A

NEW QUESTION # 264

Your car factory is pushing machine measurements as messages into a Pub/Sub topic in your Google Cloud project. A Dataflow streaming job, that you wrote with the Apache Beam SDK, reads these messages, sends acknowledgment to Pub/Sub, applies some custom business logic in a Dataflow instance, and writes the result to BigQuery. You want to ensure that if your business logic fails on a message, the message will be sent to a Pub

/Sub topic that you want to monitor for alerting purposes. What should you do?

- A. Enable retaining of acknowledged messages in your Pub/Sub pull subscription. Use Cloud Monitoring to monitor the subscription/num_retained_acked_messages metric on this subscription.
- B. Create a snapshot of your Pub/Sub pull subscription. Use Cloud Monitoring to monitor the snapshot /nummessages metric on this snapshot.
- C. Use an exception handling block in your Data Flow's Dataflow code to push the messages that failed to be transformed through a side output and to a new Pub/Sub topic. Use Cloud Monitoring to monitor the topic /num_jnacked_messages_by_region metric on this new topic.
- D. **Enable dead lettering in your Pub/Sub pull subscription, and specify a new Pub/Sub topic as the dead letter topic. Use Cloud Monitoring to monitor the subscription/dead_letter_message_count metric on your pull subscription.**

Answer: D

Explanation:

To ensure that messages failing to process in your Dataflow job are sent to a Pub/Sub topic for monitoring and alerting, the best approach is to use Pub/Sub's dead-letter topic feature. Here's why option C is the best choice:

* Dead-Letter Topic:

* Pub/Sub's dead-letter topic feature allows messages that fail to be processed successfully to be redirected to a specified topic. This ensures that these messages are not lost and can be reviewed for debugging and alerting purposes.

* Monitoring and Alerting:

* By specifying a new Pub/Sub topic as the dead-letter topic, you can use Cloud Monitoring to track metrics such as subscription/dead_letter_message_count, providing visibility into the number of failed messages.

* This allows you to set up alerts based on these metrics to notify the appropriate teams when failures occur.

Steps to Implement:

* Enable Dead-Letter Topic:

* Configure your Pub/Sub pull subscription to enable dead lettering and specify the new Pub/Sub topic for dead-letter messages.

* Set Up Monitoring:

* Use Cloud Monitoring to monitor the subscription/dead_letter_message_count metric on your pull subscription.

* Configure alerts based on this metric to notify the team of any processing failures.

Reference Links:

* Pub/Sub Dead Letter Policy

* Cloud Monitoring with Pub/Sub

NEW QUESTION # 265

MJTelco Case Study

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 relationships between data consumers and providers 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.

You need to compose visualization for operations teams with the following requirements:

- * Telemetry must include data from all 50,000 installations for the most recent 6 weeks (sampling once every minute)
- * The report must not be more than 3 hours delayed from live data.
- * The actionable report should only show suboptimal links.
- * Most suboptimal links should be sorted to the top.
- * Suboptimal links can be grouped and filtered by regional geography.
- * User response time to load the report must be <5 seconds.

You create a data source to store the last 6 weeks of data, and create visualizations that allow viewers to see multiple date ranges, distinct geographic regions, and unique installation types. You always show the latest data without any changes to your visualizations. You want to avoid creating and updating new visualizations each month. What should you do?

- A. Export the data to a spreadsheet, compose a series of charts and tables, one for each possible combination of criteria, and spread them across multiple tabs.
- B. Look through the current data and compose a small set of generalized charts and tables bound to criteria filters that allow value selection.
- C. Load the data into relational database tables, write a Google App Engine application that queries all rows, summarizes the data across each criteria, and then renders results using the Google Charts and visualization API.
- D. Look through the current data and compose a series of charts and tables, one for each possible combination of criteria.

Answer: B

NEW QUESTION # 266

.....

Don't let the Professional-Data-Engineer exam stress you out! Prepare with PracticeDump Professional-Data-Engineer exam dumps and boost your confidence in the real Professional-Data-Engineer exam. We ensure your road towards success without any mark of failure. Time is of the essence - don't wait to ace your Professional-Data-Engineer Certification Exam! Register yourself now.

Related Professional-Data-Engineer Certifications: https://www.practicdump.com/Professional-Data-Engineer_actualtests.html

- Valid Professional-Data-Engineer Exam Tutorial Professional-Data-Engineer Question Explanations Practice Professional-Data-Engineer Exam Fee Easily obtain Professional-Data-Engineer for free download through www.practicevce.com Professional-Data-Engineer Question Explanations
- Revolutionize Your Google Exam Preparation with Our Web-Based Professional-Data-Engineer Practice Test Software Search on www.pdfvce.com for Professional-Data-Engineer to obtain exam materials for free download Professional-Data-Engineer Valid Dumps Questions
- Professional-Data-Engineer Dump File Detailed Professional-Data-Engineer Study Plan Test Professional-Data-

Engineer Testking □ Search for 【 Professional-Data-Engineer 】 and download it for free on □ www.dumpsquestion.com
□ website □Professional-Data-Engineer Valid Braindumps

BTW, DOWNLOAD part of PracticeDump Professional-Data-Engineer dumps from Cloud Storage:

<https://drive.google.com/open?id=1-Cu6c7sOfjRBdg8D7zp2WwNw52yv1YEw>