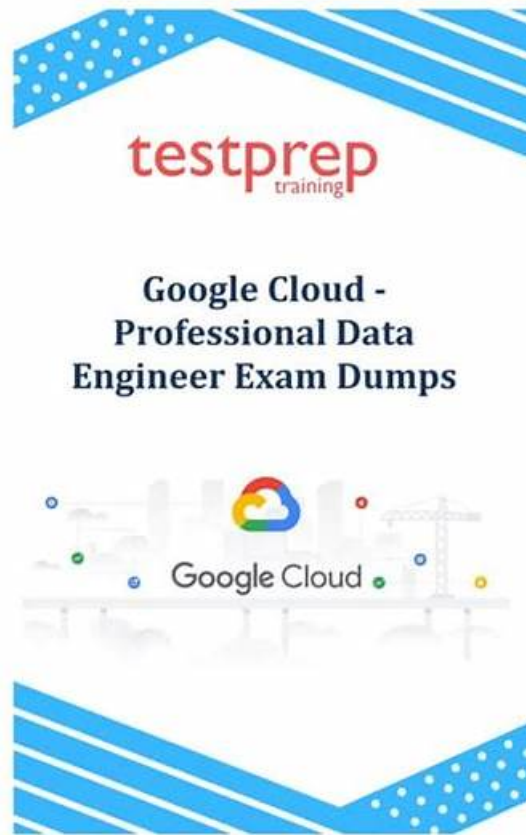


Google Professional-Data-Engineer Latest Dumps Files, Professional-Data-Engineer Latest Exam Format



BONUS!!! Download part of ExamTorrent Professional-Data-Engineer dumps for free: <https://drive.google.com/open?id=1WxFLG9sKu-K0NGSFPrnt7VSMerIhdgw>

In this rapid rhythm society, the competitions among talents are growing with each passing day, some job might ask more than one's academic knowledge it might also require the professional Professional-Data-Engineercertification and so on. It can't be denied that professional certification is an efficient way for employees to show their personal Google Certified Professional Data Engineer Exam abilities. In order to get more chances, more and more people tend to add shining points, for example a certification to their resumes. Passing exam won't be a problem anymore as long as you are familiar with our Professional-Data-Engineer Exam Material (only about 20 to 30 hours practice). High accuracy and high quality are the reasons why you should choose us.

Google Professional-Data-Engineer exam is designed to test the skills and knowledge of professionals who work with Google Cloud Platform. It is designed for those who have experience in designing, building, and managing data processing systems, and for those who work with machine learning models and data analysis tools. Professional-Data-Engineer Exam is designed to test the ability of professionals to use Google Cloud Platform tools to develop and implement solutions that meet the needs of businesses.

>> **Google Professional-Data-Engineer Latest Dumps Files** <<

Professional-Data-Engineer Latest Exam Format - Professional-Data-Engineer Reliable Test Voucher

With the qualification certificate, you are qualified to do this professional job. Therefore, getting the test Professional-Data-Engineer certification is of vital importance to our future employment. And the Professional-Data-Engineer study tool can provide a good

learning platform for users who want to get the test Professional-Data-Engineercertification in a short time. If you can choose to trust us, I believe you will have a good experience when you use the Professional-Data-Engineer study guide, and you can pass the exam and get a good grade in the test Professional-Data-Engineer certification.

Google Certified Professional Data Engineer Exam Sample Questions (Q239-Q244):

NEW QUESTION # 239

You used Cloud Dataprep to create a recipe on a sample of data in a BigQuery table. You want to reuse this recipe on a daily upload of data with the same schema, after the load job with variable execution time completes. What should you do?

- A. Create an App Engine cron job to schedule the execution of the Cloud Dataprep job.
- B. Export the Cloud Dataprep job as a Cloud Dataflow template, and incorporate it into a Cloud Composer job.
- C. Export the recipe as a Cloud Dataprep template, and create a job in Cloud Scheduler.
- D. Create a cron schedule in Cloud Dataprep.

Answer: C

NEW QUESTION # 240

Which of these is NOT a way to customize the software on Dataproc cluster instances?

- A. Set initialization actions
- B. Configure the cluster using Cloud Deployment Manager
- C. Modify configuration files using cluster properties
- D. Log into the master node and make changes from there

Answer: B

Explanation:

You can access the master node of the cluster by clicking the SSH button next to it in the Cloud Console.

You can easily use the `--properties` option of the `dataproc` command in the Google Cloud SDK to modify many common configuration files when creating a cluster.

When creating a Cloud Dataproc cluster, you can specify initialization actions in executables and/or scripts that Cloud Dataproc will run on all nodes in your Cloud Dataproc cluster immediately after the cluster is set up.

[<https://cloud.google.com/dataproc/docs/concepts/configuring-clusters/init-actions>] Reference:

<https://cloud.google.com/dataproc/docs/concepts/configuring-clusters/cluster-properties>

NEW QUESTION # 241

Which row keys are likely to cause a disproportionate number of reads and/or writes on a particular node in a Bigtable cluster (select 2 answers)?

- A. A timestamp followed by a stock symbol
- B. A non-sequential numeric ID
- C. A stock symbol followed by a timestamp
- D. A sequential numeric ID

Answer: A,D

Explanation:

using a timestamp as the first element of a row key can cause a variety of problems.

In brief, when a row key for a time series includes a timestamp, all of your writes will target a single node; fill that node; and then move onto the next node in the cluster, resulting in hotspotting.

Suppose your system assigns a numeric ID to each of your application's users. You might be tempted to use the user's numeric ID as the row key for your table. However, since new users are more likely to be active users, this approach is likely to push most of your traffic to a small number of nodes. [<https://cloud.google.com/bigtable/docs/schema-design>]

Reference: https://cloud.google.com/bigtable/docs/schema-design-time-series#ensure_that_your_row_key_avoids_hotspotting

NEW QUESTION # 242

You work for a large e-commerce company. You store your customers order data in Bigtable. You have a garbage collection policy set to delete the data after 30 days and the number of versions is set to 1. When the data analysts run a query to report total customer spending, the analysts sometimes see customer data that is older than 30 days. You need to ensure that the analysts do not see customer data older than 30 days while minimizing cost and overhead. What should you do?

- A. Set the expiring values of the column families to 30 days and set the number of versions to 2.
- B. Set the expiring values of the column families to 29 days and keep the number of versions to 1.
- C. Schedule a job daily to scan the data in the table and delete data older than 30 days.
- **D. Use a timestamp range filter in the query to fetch the customer's data for a specific range.**

Answer: D

Explanation:

By using a timestamp range filter in the query, you can ensure that the analysts only see the customer data that is within the desired time range, regardless of the garbage collection policy¹. This option is the most cost-effective and simple way to avoid fetching data that is marked for deletion by garbage collection, as it does not require changing the existing policy or creating additional jobs. You can use the Bigtable client libraries or the cbt CLI to apply a timestamp range filter to your read requests².

Option A is not effective, as it increases the number of versions to 2, which may cause more data to be retained and increase the storage costs. Option C is not reliable, as it reduces the expiring values to 29 days, which may not match the actual data arrival and usage patterns. Option D is not efficient, as it requires scheduling a job daily to scan and delete the data, which may incur additional overhead and complexity. Moreover, none of these options guarantee that the data older than 30 days will be immediately deleted, as garbage collection is an asynchronous process that can take up to a week to remove the data³. References:

* 1: Filters | Cloud Bigtable Documentation | Google Cloud

* 2: Read data | Cloud Bigtable Documentation | Google Cloud

* 3: Garbage collection overview | Cloud Bigtable Documentation | Google Cloud

NEW QUESTION # 243

You are migrating an application that tracks library books and information about each book, such as author or year published, from an on-premises data warehouse to BigQuery. In your current relational database, the author information is kept in a separate table and joined to the book information on a common key. Based on Google's recommended practice for schema design, how would you structure the data to ensure optimal speed of queries about the author of each book that has been borrowed?

- A. Create a table that is wide and includes a column for each attribute, including the author's first name, last name, date of birth, etc
- **B. Create a table that includes information about the books and authors, but nest the author fields inside the author column**
- C. Keep the schema the same, create a view that joins all of the tables, and always query the view
- D. Keep the schema the same, maintain the different tables for the book and each of the attributes, and query as you are doing today

Answer: B

NEW QUESTION # 244

.....

Practice tests (desktop and web-based) provide an Google Professional-Data-Engineer examination scenario so your preparation for the Google Certified Professional Data Engineer Exam (Professional-Data-Engineer) exam becomes quite easier. Since the real Professional-Data-Engineer examination costs a high penny, ExamTorrent provide a free demo of Google Professional-Data-Engineer Exam Dumps before your purchase. The free demo of the Google Certified Professional Data Engineer Exam (Professional-Data-Engineer) exam prep material is helpful to remove your doubts about it. The product is available in three versions which are PDF, Web-based practice test, and Desktop practice test software.

Professional-Data-Engineer Latest Exam Format: <https://www.examtorent.com/Professional-Data-Engineer-valid-vce-dumps.html>

- Professional-Data-Engineer Exam Latest Dumps Files- Latest Professional-Data-Engineer Latest Exam Format Pass Success
 Enter « www.examcollectionpass.com » and search for ► Professional-Data-Engineer to download for free

