

Professional-Data-Engineer Verified Answers | Professional-Data-Engineer Exam Answers



P.S. Free & New Professional-Data-Engineer dumps are available on Google Drive shared by PassLeaderVCE:
<https://drive.google.com/open?id=1p9kYChmlUW4VpG1BY2IabrDujVl5E-8P>

The system of Professional-Data-Engineer study materials is very smooth and you don't need to spend a lot of time installing it. We take into account all aspects on the Professional-Data-Engineer exam braindumps and save you as much time as possible. After the installation is complete, you can devote all of your time to studying Professional-Data-Engineer Exam Questions. And a lot of our worthy customers always praise the high-efficiency of our Professional-Data-Engineer learning guide. If you buy it, i guess you will love it as well.

Preparation Process

To perform well in the Google Professional Data Engineer certification exam, the candidates must be ready to devote ample time to preparation. There is a host of study materials available on the Internet, but if you want to be confident in the authenticity of the resources that you use, it is best to refer to the official platform. Google recommends that the applicants follow the Professional Data Engineer learning path, which is a comprehensive option involving in-person classes, online training, hands-on labs, and other resources from Google Cloud.

Besides that, it is recommended that the students use the official sample questions to familiarize themselves with the question formats that they will encounter during the actual exam. The official webpage also contains additional resources such as Google Cloud documentation and Google Cloud solutions. There is also an option of joining the subject-related webinar to get valuable preparation tips from the Google experts.

>> Professional-Data-Engineer Verified Answers <<

Authoritative Professional-Data-Engineer – 100% Free Verified Answers | Professional-Data-Engineer Exam Answers

The web-based practice test is similar to the desktop-based software, with all the same elements of the desktop practice exam. The mock exam can be accessed from any browser and does not require installation. The Professional-Data-Engineer questions in the

mock test are the same as those in the real exam. Candidates can take the web-based Google Certified Professional Data Engineer Exam (Professional-Data-Engineer) practice test immediately, regardless of the operating system and browser they are using.

Google Certified Professional Data Engineer Exam Sample Questions (Q207-Q212):

NEW QUESTION # 207

Your company's on-premises Apache Hadoop servers are approaching end-of-life, and IT has decided to migrate the cluster to Google Cloud Dataproc. A like-for-like migration of the cluster would require 50 TB of Google Persistent Disk per node. The CIO is concerned about the cost of using that much block storage.

You want to minimize the storage cost of the migration. What should you do?

- A. Use preemptible virtual machines (VMs) for the Cloud Dataproc cluster.
- B. Tune the Cloud Dataproc cluster so that there is just enough disk for all data.
- C. Put the data into Google Cloud Storage.
- D. Migrate some of the cold data into Google Cloud Storage, and keep only the hot data in Persistent Disk.

Answer: A

Explanation:

Explanation/Reference:

Reference: <https://cloud.google.com/dataproc/>

NEW QUESTION # 208

You want to use Google Stackdriver Logging to monitor Google BigQuery usage. You need an instant notification to be sent to your monitoring tool when new data is appended to a certain table using an insert job, but you do not want to receive notifications for other tables. What should you do?

- A. Make a call to the Stackdriver API to list all logs, and apply an advanced filter.
- B. In the Stackdriver logging admin interface, enable a log sink export to Google Cloud Pub/Sub, and subscribe to the topic from your monitoring tool.
- C. Using the Stackdriver API, create a project sink with advanced log filter to export to Pub/Sub, and subscribe to the topic from your monitoring tool.
- D. In the Stackdriver logging admin interface, and enable a log sink export to BigQuery.

Answer: D

Explanation:

Topic 1, 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 into. 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 parity 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 our 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.

NEW QUESTION # 209

You have spent a few days loading data from comma-separated values (CSV) files into the Google BigQuery table `CLICK_STREAM`. The column `DT` stores the epoch time of click events. For convenience, you chose a simple schema where every field is treated as the `STRING` type. Now, you want to compute web session durations of users who visit your site, and you want to change its data type to the `TIMESTAMP`. You want to minimize the migration effort without making future queries computationally expensive. What should you do?

- A. Construct a query to return every row of the table `CLICK_STREAM`, while using the built-in function to cast strings from the column `DT` into `TIMESTAMP` values. Run the query into a destination table `NEW_CLICK_STREAM`, in which the column `TS` is the `TIMESTAMP` type. Reference the table `NEW_CLICK_STREAM` instead of the table `CLICK_STREAM` from now on. In the future, new data is loaded into the table `NEW_CLICK_STREAM`.
- B. Add a column `TS` of the `TIMESTAMP` type to the table `CLICK_STREAM`, and populate the numeric values from the column `TS` for each row. Reference the column `TS` instead of the column `DT` from now on.
- C. Delete the table `CLICK_STREAM`, and then re-create it such that the column `DT` is of the `TIMESTAMP` type. Reload the data.
- D. Create a view `CLICK_STREAM_V`, where strings from the column `DT` are cast into `TIMESTAMP` values. Reference the view `CLICK_STREAM_V` instead of the table `CLICK_STREAM` from now on.
- E. Add two columns to the table `CLICK STREAM`: `TS` of the `TIMESTAMP` type and `IS_NEW` of the `BOOLEAN` type. Reload all data in append mode. For each appended row, set the value of `IS_NEW` to true. For future queries, reference the

column TS instead of the column DT, with the WHERE clause ensuring that the value of IS_NEW must be true.

Answer: E

NEW QUESTION # 210

Your analytics team wants to build a simple statistical model to determine which customers are most likely to work with your company again, based on a few different metrics. They want to run the model on Apache Spark, using data housed in Google Cloud Storage, and you have recommended using Google Cloud Dataproc to execute this job. Testing has shown that this workload can run in approximately 30 minutes on a 15-node cluster, outputting the results into Google BigQuery. The plan is to run this workload weekly.

How should you optimize the cluster for cost?

- A. Use SSDs on the worker nodes so that the job can run faster
- B. Migrate the workload to Google Cloud Dataflow
- C. Use a higher-memory node so that the job runs faster
- **D. Use pre-emptible virtual machines (VMs) for the cluster**

Answer: D

Explanation:

<https://cloud.google.com/dataproc/docs/concepts/compute/preemptible-vm>

NEW QUESTION # 211

You have historical data covering the last three years in BigQuery and a data pipeline that delivers new data to BigQuery daily. You have noticed that when the Data Science team runs a query filtered on a date column and limited to 30-90 days of data, the query scans the entire table. You also noticed that your bill is increasing more quickly than you expected. You want to resolve the issue as cost-effectively as possible while maintaining the ability to conduct SQL queries. What should you do?

- A. Recommend that the Data Science team export the table to a CSV file on Cloud Storage and use Cloud Datalab to explore the data by reading the files directly.
- **B. Modify your pipeline to maintain the last 30-90 days of data in one table and the longer history in a different table to minimize full table scans over the entire history.**
- C. Write an Apache Beam pipeline that creates a BigQuery table per day. Recommend that the Data Science team use wildcards on the table name suffixes to select the data they need.
- D. Re-create the tables using DDL. Partition the tables by a column containing a TIMESTAMP or DATE Type.

Answer: B

NEW QUESTION # 212

.....

Our Professional-Data-Engineer torrent prep can apply to any learner whether students or working staff, novices or practitioners with years of experience. To simplify complex concepts and add examples to explain anything that might be difficult to understand, studies on Professional-Data-Engineer exam questions can easily navigate learning and become the master of learning. Our Professional-Data-Engineer Exam Questions are committed to instill more important information with fewer questions and answers, so you can learn easily and efficiently in this process. Our Professional-Data-Engineer training guide will be your best choice.

Professional-Data-Engineer Exam Answers: <https://www.passleadervce.com/Google-Cloud-Certified/reliable-Professional-Data-Engineer-exam-learning-guide.html>

- Pass-Sure Professional-Data-Engineer Verified Answers - Leading Offer in Qualification Exams - Marvelous Professional-Data-Engineer: Google Certified Professional Data Engineer Exam Open www.examcollectionpass.com enter **【 Professional-Data-Engineer 】** and obtain a free download Clearer Professional-Data-Engineer Explanation
- Clearer Professional-Data-Engineer Explanation Professional-Data-Engineer Test Pass4sure New Professional-Data-Engineer Mock Test Search for \Rightarrow Professional-Data-Engineer \Leftarrow and download exam materials for free through { www.pdfvce.com } New Professional-Data-Engineer Mock Test
- New Professional-Data-Engineer Mock Test Test Professional-Data-Engineer Questions Vce Professional-Data-Engineer Training For Exam Easily obtain Professional-Data-Engineer for free download through **➡**

