

Valid Professional-Data-Engineer Exam Materials & Professional-Data-Engineer Valid Examcollection



BONUS!!! Download part of Prep4pass Professional-Data-Engineer dumps for free: <https://drive.google.com/open?id=1nldHuxfx1Ah5LglHWKAGVRxNTTCbcJk5>

In our study, we found that many people have the strongest ability to use knowledge for a period of time at the beginning of their knowledge. As time goes on, memory fades. Our Professional-Data-Engineer study materials are designed to help users consolidate what they have learned, will add to the instant of many training, the user can test their learning effect in time after finished the part of the learning content, have a special set of wrong topics in our Professional-Data-Engineer Study Materials, enable users to find their weak spot of knowledge in this function, iterate through constant practice, finally reach a high success rate.

We are here divide grieves with you to help you pass your Google Professional-Data-Engineer exam with ease. You can abandon the time-consuming thought from now on. You won't regret your decision of choosing our Google Professional-Data-Engineer study guide. In contrast, they will inspire your potential without obscure content to feel. After getting our Professional-Data-Engineer Exam Prep, you will not live under great stress during the Professional-Data-Engineer exam period.

>> Valid Professional-Data-Engineer Exam Materials <<

Valid Professional-Data-Engineer Exam Materials 100% Pass | Valid Google Certified Professional Data Engineer Exam Valid Examcollection Pass for

sure

There are lots of benefits of obtaining a certificate, it can help you enter a better company, have a high position in the company, improve your wages etc. Our Professional-Data-Engineer test materials will help you get the certificate successfully. We have channels to obtain the latest information about the exam, and we ensure you that you can get the latest information about the Professional-Data-Engineer Exam Dumps timely. Furthermore, you can get the downloading link and password for Professional-Data-Engineer test materials within ten minutes after purchasing.

Google Certified Professional Data Engineer Exam Sample Questions (Q219-Q224):

NEW QUESTION # 219

You have a table that contains millions of rows of sales data, partitioned by date. Various applications and users query this data many times a minute. The query requires aggregating values by using avg, max, and sum, and does not require joining to other tables. The required aggregations are only computed over the past year of data, though you need to retain full historical data in the base tables. You want to ensure that the query results always include the latest data from the tables, while also reducing computation cost, maintenance overhead, and duration. What should you do?

- A. Create a view to aggregate the base table data. Include a filter clause to specify the last year of partitions.
- **B. Create a materialized view to aggregate the base table data. Configure a partition expiration on the base table to retain only the last one year of partitions.**
- C. Create a materialized view to aggregate the base table data. Include a filter clause to specify the last one year of partitions.
- D. Create a new table that aggregates the base table data. Include a filter clause to specify the last year of partitions. Set up a scheduled query to recreate the new table every hour.

Answer: B

Explanation:

A materialized view is a database object that contains the results of a query, which can be updated periodically. It can improve the performance and efficiency of queries that involve aggregations, joins, or filters. By creating a materialized view to aggregate the base table data and include a filter clause to specify the last one year of partitions, you can ensure that the query results always include the latest data from the tables, while also reducing computation cost, maintenance overhead, and duration. The materialized view will automatically refresh when the base table data changes, and will only use the partitions that match the filter clause. Option A is incorrect because it will delete the historical data from the base table, which is not desired.

Option C is incorrect because it will create a redundant table that needs to be updated manually by a scheduled query, which is more complex and costly than using a materialized view. Option D is incorrect because a view does not store any data, but only references the base table data, which means it will not reduce the computation cost or duration of the query. References:

* Materialized views, ML models in data warehouse - Google Cloud

* Data Engineering with Google Cloud Platform - Packt Subscription

NEW QUESTION # 220

You are working on a sensitive project involving private user data. You have set up a project on Google Cloud Platform to house your work internally. An external consultant is going to assist with coding a complex transformation in a Google Cloud Dataflow pipeline for your project. How should you maintain users' privacy?

- **A. Create a service account and allow the consultant to log on with it.**
- B. Create an anonymized sample of the data for the consultant to work with in a different project.
- C. Grant the consultant the Cloud Dataflow Developer role on the project.
- D. Grant the consultant the Viewer role on the project.

Answer: A

NEW QUESTION # 221

You are migrating your on-premises data warehouse to BigQuery. One of the upstream data sources resides on a MySQL database that runs in your on-premises data center with no public IP addresses. You want to ensure that the data ingestion into BigQuery is done securely and does not go through the public internet.

What should you do?

- A. Use Datastream to replicate data from your on-premises MySQL database to BigQuery. Gather Datastream public IP addresses of the Google Cloud region that will be used to set up the stream. Add those IP addresses to the firewall allowlist of your on-premises data center. Use IP Allowlisting as the connectivity method and Server-only as the encryption type when setting up the connection profile in Datastream.
- B. Use Datastream to replicate data from your on-premises MySQL database to BigQuery. Set up Cloud Interconnect between your on-premises data center and Google Cloud. Use Private connectivity as the connectivity method and allocate an IP address range within your VPC network to the Datastream connectivity configuration. Use Server-only as the encryption type when setting up the connection profile in Datastream.
- C. Use Datastream to replicate data from your on-premises MySQL database to BigQuery. Use Forward-SSH tunnel as the connectivity method to establish a secure tunnel between Datastream and your on-premises MySQL database through a tunnel server in your on-premises data center. Use None as the encryption type when setting up the connection profile in Datastream.
- D. Update your existing on-premises ETL tool to write to BigQuery by using the BigQuery Open Database Connectivity (ODBC) driver. Set up the proxy parameter in the Simba. googlebigqueryodbc. ini file to point to your data center's NAT gateway.

Answer: B

Explanation:

To securely ingest data from an on-premises MySQL database into BigQuery without routing through the public internet, using Datastream with Private connectivity over Cloud Interconnect is the best approach.

Here's why:

* Datastream for Data Replication:

* Datastream provides a managed service for data replication from various sources, including on-premises databases, to Google Cloud services like BigQuery.

* Cloud Interconnect:

* Cloud Interconnect establishes a private connection between your on-premises data center and Google Cloud, ensuring that data transfer occurs over a secure, private network rather than the public internet.

* Private Connectivity:

* Using Private connectivity with Datastream leverages the established Cloud Interconnect to securely connect your on-premises MySQL database with Google Cloud. This method ensures that the data does not traverse the public internet.

* Encryption:

* Using Server-only encryption ensures that data is encrypted in transit between Datastream and BigQuery, adding an extra layer of security.

Steps to Implement:

* Set Up Cloud Interconnect:

* Establish a Cloud Interconnect between your on-premises data center and Google Cloud to create a private connection.

* Configure Datastream:

* Set up Datastream to use Private connectivity as the connection method and allocate an IP address range within your VPC network.

* Use Server-only encryption to ensure secure data transfer.

* Create Connection Profile:

* Create a connection profile in Datastream to define the connection parameters, including the use of Cloud Interconnect and Private connectivity.

Reference Links:

* Datastream Documentation

* Cloud Interconnect Documentation

* Setting Up Private Connectivity in Datastream

NEW QUESTION # 222

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

Answer: E

Explanation:

Topic 4, Main Questions Set B

NEW QUESTION # 223

Suppose you have a dataset of images that are each labeled as to whether or not they contain a human face. To create a neural network that recognizes human faces in images using this labeled dataset, what approach would likely be the most effective?

- A. Use K-means Clustering to detect faces in the pixels.
- **B. Use deep learning by creating a neural network with multiple hidden layers to automatically detect features of faces.**
- C. Use feature engineering to add features for eyes, noses, and mouths to the input data.
- D. Build a neural network with an input layer of pixels, a hidden layer, and an output layer with two categories.

Answer: B

Explanation:

Traditional machine learning relies on shallow nets, composed of one input and one output layer, and at most one hidden layer in between. More than three layers (including input and output) qualifies as "deep" learning. So deep is a strictly defined, technical term that means more than one hidden layer.

In deep-learning networks, each layer of nodes trains on a distinct set of features based on the previous layer's output. The further you advance into the neural net, the more complex the features your nodes can recognize, since they aggregate and recombine features from the previous layer.

A neural network with only one hidden layer would be unable to automatically recognize high-level features of faces, such as eyes, because it wouldn't be able to "build" these features using previous hidden layers that detect low-level features, such as lines.

Feature engineering is difficult to perform on raw image data.

K-means Clustering is an unsupervised learning method used to categorize unlabeled data.

NEW QUESTION # 224

.....

You can download the trial version of our Professional-Data-Engineer learning material for free. After using the trial version of our Professional-Data-Engineer study materials, I believe you will have a deeper understanding of the advantages of our Professional-Data-Engineer training engine. The development of society urges us to advance and use our Professional-Data-Engineer Study Materials to make us progress faster and become the leader of this era. The best you need is the best exam preparation materials. Our Professional-Data-Engineer exam simulation will accompany you to a better future.

Professional-Data-Engineer Valid Examcollection: https://www.prep4pass.com/Professional-Data-Engineer_exam-braindumps.html

We provide the well searched Professional-Data-Engineer exam dumps prepared by the Google Professional Certification Experts, In the end, you will build a clear knowledge structure of the Professional-Data-Engineer Valid Examcollection - Google Certified Professional Data Engineer Exam exam, Google Valid Professional-Data-Engineer Exam Materials In fact, there is no point in wasting much time on invalid input, If user wish to test the Google Professional-Data-Engineer study material before joining Prep4pass, they may do so with a free sample trial.

The approach is named The Power of Completion, and the tools have been honed by expert project managers through decades of experience, Furthermore, the Professional-Data-Engineer learning materials will help you pass exam easily and successfully, boost your Professional-Data-Engineer confidence to pursue your dream such as double your salary, get promotion and become senior management in your company.

Providing You High-quality Valid Professional-Data-Engineer Exam Materials with 100% Passing Guarantee

We provide the well searched Professional-Data-Engineer exam dumps prepared by the Google Professional Certification Experts, In the end, you will build a clear knowledge structure of the Google Certified Professional Data Engineer Exam exam.

In fact, there is no point in wasting much time on invalid input, If user wish to test the Google Professional-Data-Engineer study material before joining Prep4pass, they may do so with a free sample trial.

