![Pass2Lead logo](https://Pass2Lead.com)
PR OFESSIONAL-DATA-ENGINEER^{Q&As}

Professional Data Engineer on Google Cloud Platform

Pass Google PROFESSIONAL-DATA-ENGINEER Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass2lead.com/professional-data-engineer.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Google Official Exam Center

QUESTION 1

What are all of the BigQuery operations that Google charges for?

A. Storage, queries, and streaming inserts

B. Storage, queries, and loading data from a file

C. Storage, queries, and exporting data

D. Queries and streaming inserts

Correct Answer: A

Google charges for storage, queries, and streaming inserts. Loading data from a file and exporting data are free operations. Reference: https://cloud.google.com/bigquery/pricing

QUESTION 2

You are migrating your data warehouse to Google Cloud and decommissioning your on-premises data center Because this is a priority for your company, you know that bandwidth will be made available for the initial data load to the cloud.

The files being transferred are not large in number, but each file is 90 GB Additionally, you want your transactional systems to continually update the warehouse on Google Cloud in real time What tools should you use to migrate the data and ensure that it continues to write to your warehouse?

A. Storage Transfer Service for the migration, Pub/Sub and Cloud Data Fusion for the real-time updates

B. BigQuery Data Transfer Service lor the migration, Pub/Sub and Dataproc for the real-time updates

C. gsutil for the migration; Pub/Sub and Dataflow for the real-time updates

D. gsutil for both the migration and the real-time updates

Correct Answer: C

QUESTION 3

You need to deploy additional dependencies to all of a Cloud Dataproc cluster at startup using an existing initialization action. Company security policies require that Cloud Dataproc nodes do not have access to the Internet so public initialization actions cannot fetch resources. What should you do?

A. Deploy the Cloud SQL Proxy on the Cloud Dataproc master

B. Use an SSH tunnel to give the Cloud Dataproc cluster access to the Internet

C. Copy all dependencies to a Cloud Storage bucket within your VPC security perimeter

D. Use Resource Manager to add the service account used by the Cloud Dataproc cluster to the Network User role

Correct Answer: C

QUESTION 4

You need to choose a database for a new project that has the following requirements:

1.

Fully managed

2.

Able to automatically scale up

3.

Transactionally consistent

4.

Able to scale up to 6 TB

5.

Able to be queried using SQL Which database do you choose?

A. Cloud SQL

B. Cloud Bigtable

C. Cloud Spanner

D. Cloud Datastore

Correct Answer: A

QUESTION 5

You are operating a streaming Cloud Dataflow pipeline. Your engineers have a new version of the pipeline with a different windowing algorithm and triggering strategy. You want to update the running pipeline with the new version. You want to ensure that no data is lost during the update. What should you do?

A. Update the Cloud Dataflow pipeline inflight by passing the --update option with the --jobName set to the existing job name

B. Update the Cloud Dataflow pipeline inflight by passing the --update option with the --jobName set to a new unique job name

C. Stop the Cloud Dataflow pipeline with the Cancel option. Create a new Cloud Dataflow job with the updated code

D. Stop the Cloud Dataflow pipeline with the Drain option. Create a new Cloud Dataflow job with the updated code

Correct Answer: D

PROFESSIONAL-DATA-ENGINEER VCE Dumps

PROFESSIONAL-DATA-ENGINEER Practice Test

PROFESSIONAL-DATA-ENGINEER Exam Questions