Q&As

Professional Cloud Architect on Google Cloud Platform

Pass Google PROFESSIONAL-CLOUD-ARCHITECT Exam with 100% Guarantee

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

https://www.passapply.com/professional-cloud-architect.html

100% Passing Guarantee
100% Money Back Assurance

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

QUESTION 1

For this question, refer to the Helicopter Racing League (HRL) case study. HRL wants better prediction accuracy from their ML prediction models. They want you to use Google\\'s AI Platform so HRL can understand and interpret the predictions. What should you do?

A. Use Explainable AI.

B. Use Vision AI.

C. Use Google Cloud\\'s operations suite.

D. Use Jupyter Notebooks.

Correct Answer: A

Reference: https://cloud.google.com/ai-platform/prediction/docs/ai-explanations/preparing-metadata

---

QUESTION 2

You are analyzing and defining business processes to support your startup\\'s trial usage of GCP, and you don\\'t yet know what consumer demand for your product will be. Your manager requires you to minimize GCP service costs and adhere to Google best practices. What should you do?

A. Utilize free tier and sustained use discounts. Provision a staff position for service cost management.

B. Utilize free tier and sustained use discounts. Provide training to the team about service cost management.

C. Utilize free tier and committed use discounts. Provision a staff position for service cost management.

D. Utilize free tier and committed use discounts. Provide training to the team about service cost management.

Correct Answer: B

QUESTION 3

You are tasked with building an online analytical processing (OLAP) marketing analytics and reporting tool. This requires a relational database that can operate on hundreds of terabytes of data. What is the Google-recommended tool for such applications?

A. Cloud Spanner, because it is globally distributed

B. Cloud SQL, because it is a fully managed relational database

C. Cloud Firestore, because it offers real-time synchronization across devices

D. BigQuery, because it is designed for large-scale processing of tabular data

Correct Answer: D

Reference: https://cloud.google.com/files/BigQueryTechnicalWP.pdf

QUESTION 4

The application reliability team at your company this added a debug feature to their backend service to send all server events to Google Cloud Storage for eventual analysis. The event records are at least 50 KB and at most 15 MB and are expected to peak at 3,000 events per second. You want to minimize data loss.

Which process should you implement?

A. Append metadata to file body Compress individual files Name files with serverName Timestamp Create a new bucket if bucket is older than 1 hour and save individual files to the new bucket. Otherwise, save files to existing bucket.

B. Batch every 10,000 events with a single manifest file for metadata Compress event files and manifest file into a single archive file Name files using serverName EventSequence Create a new bucket if bucket is older than 1 day and save the single archive file to the new bucket. Otherwise, save the single archive file to existing bucket.

C. Compress individual files Name files with serverName EventSequence Save files to one bucket Set custom metadata headers for each object after saving

D. Append metadata to file body Compress individual files Name files with a random prefix pattern Save files to one bucket

Correct Answer: D

In order to maintain a high request rate, avoid using sequential names. Using completely random object names will give you the best load distribution. Randomness after a common prefix is effective under the prefix https://cloud.google.com/storage/docs/request-rate

QUESTION 5

For this question, refer to the TerramEarth case study. A new architecture that writes all incoming data to BigQuery has been introduced. You notice that the data is dirty, and want to ensure data quality on an automated daily basis while managing cost.

What should you do?

A. Set up a streaming Cloud Dataflow job, receiving data by the ingestion process. Clean the data in a Cloud Dataflow pipeline.

B. Create a Cloud Function that reads data from BigQuery and cleans it. Trigger it. Trigger the Cloud Function from a Compute Engine instance.

C. Create a SQL statement on the data in BigQuery, and save it as a view. Run the view daily, and save the result to a new table.

D. Use Cloud Dataprep and configure the BigQuery tables as the source. Schedule a daily job to clean the data.

Correct Answer: D

Questions