



## 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



VCE & PDF

PassApply.com

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

2024 Latest passapply PROFESSIONAL-CLOUD-ARCHITECT PDF and VCE dumps Download

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers





### QUESTION 1

Your company has an application running on Google Cloud that is collecting data from thousands of physical devices that are globally distributed. Data is published to Pub/Sub and streamed in real time into an SSD Cloud Bigtable cluster via a Dataflow pipeline. The operations team informs you that your Cloud Bigtable cluster has a hotspot, and queries are taking longer than expected. You need to resolve the problem and prevent it from happening in the future. What should you do?

- A. Advise your clients to use HBase APIs instead of NodeJS APIs.
- B. Delete records older than 30 days.
- C. Review your RowKey strategy and ensure that keys are evenly spread across the alphabet.
- D. Double the number of nodes you currently have.

Correct Answer: C

---

### QUESTION 2

Your company wants to track whether someone is present in a meeting room reserved for a scheduled meeting. There are 1000 meeting rooms across 5 offices on 3 continents. Each room is equipped with a motion sensor that reports its status every second. The data from the motion detector includes only a sensor ID and several different discrete items of information. Analysts will use this data, together with information about account owners and office locations.

Which database type should you use?

- A. Flat file
- B. NoSQL
- C. Relational
- D. Blobstore

Correct Answer: B

Relational databases were not designed to cope with the scale and agility challenges that face modern applications, nor were they built to take advantage of the commodity storage and processing power available today.

NoSQL fits well for:

Developers are working with applications that create massive volumes of new, rapidly changing data types -- structured, semi-structured, unstructured and polymorphic data.

Incorrect Answers:

D: The Blobstore API allows your application to serve data objects, called blobs, that are much larger than the size allowed for objects in the Datastore service. Blobs are useful for serving large files, such as video or image files, and for allowing users to upload large data files.

References: <https://www.mongodb.com/nosql-explained>

---



### QUESTION 3

Your company is designing its data lake on Google Cloud and wants to develop different ingestion pipelines to collect unstructured data from different sources. After the data is stored in Google Cloud, it will be processed in several data pipelines to build a recommendation engine for end users on the website. The structure of the data retrieved from the source systems can change at any time. The data must be stored exactly as it was retrieved for reprocessing purposes in case the data structure is incompatible with the current processing pipelines. You need to design an architecture to support the use case after you retrieve the data. What should you do?

- A. Send the data through the processing pipeline, and then store the processed data in a BigQuery table for reprocessing.
- B. Store the data in a BigQuery table. Design the processing pipelines to retrieve the data from the table.
- C. Send the data through the processing pipeline, and then store the processed data in a Cloud Storage bucket for reprocessing.
- D. Store the data in a Cloud Storage bucket. Design the processing pipelines to retrieve the data from the bucket.

Correct Answer: D

---

### QUESTION 4

For this question, refer to the Mountkirk Games case study. Mountkirk Games wants to migrate from their current analytics and statistics reporting model to one that meets their technical requirements on Google Cloud Platform. Which two steps should be part of their migration plan? (Choose two.)

- A. Evaluate the impact of migrating their current batch ETL code to Cloud Dataflow.
- B. Write a schema migration plan to denormalize data for better performance in BigQuery.
- C. Draw an architecture diagram that shows how to move from a single MySQL database to a MySQL cluster.
- D. Load 10 TB of analytics data from a previous game into a Cloud SQL instance, and run test queries against the full dataset to confirm that they complete successfully.
- E. Integrate Cloud Armor to defend against possible SQL injection attacks in analytics files uploaded to Cloud Storage.

Correct Answer: AB

---

### QUESTION 5

One of your primary business objectives is being able to trust the data stored in your application. You want to log all changes to the application data. How can you design your logging system to verify authenticity of your logs?

- A. Write the log concurrently in the cloud and on premises
- B. Use a SQL database and limit who can modify the log table
- C. Digitally sign each timestamp and log entry and store the signature



D. Create a JSON dump of each log entry and store it in Google Cloud Storage

Correct Answer: C

Write a log entry. If the log does not exist, it is created. You can specify a severity for the log entry, and you can write a structured log entry by specifying --payload-type=json and writing your message as a JSON string:

gcloud logging write LOG STRING gcloud logging write LOG JSON-STRING --payload-type=json References:

<https://cloud.google.com/logging/docs/reference/tools/gcloud-logging>

[Latest PROFESSIONAL-CLOUD-ARCHITECT Dumps](#)

[PROFESSIONAL-CLOUD-ARCHITECT Study Guide](#)

[PROFESSIONAL-CLOUD-ARCHITECT Exam Questions](#)