



VCE & PDF

PassApply.com

<https://www.passapply.com/cloud-digital-leader.html>

2024 Latest passapply CLOUD-DIGITAL-LEADER PDF and VCE dumps

Download

CLOUD-DIGITAL-LEADER^{Q&As}

Cloud Digital Leader

**Pass Google CLOUD-DIGITAL-LEADER Exam with
100% Guarantee**

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

<https://www.passapply.com/cloud-digital-leader.html>

100% Passing Guarantee
100% Money Back Assurance

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

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





QUESTION 1

Your organization is developing an application that will capture a large amount of data from millions of different sensor devices spread all around the world. Your organization needs a database that is suitable for worldwide, high-speed data storage of a large amount of unstructured data.

Which Google Cloud product should your organization choose?

- A. Firestore
- B. Cloud Data Fusion
- C. Cloud SQL
- D. Cloud Bigtable

Correct Answer: D

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

Cloud Bigtable is a sparsely populated table that can scale to billions of rows and thousands of columns, enabling you to store terabytes or even petabytes of data. A single value in each row is indexed; this value is known as the row key.

Bigtable is ideal for storing very large amounts of single-keyed data with very low latency. It supports high read and write throughput at low latency, and it is an ideal data source for MapReduce operations.

Bigtable is exposed to applications through multiple client libraries, including a supported extension to the Apache HBase library for Java. As a result, it integrates with the existing Apache ecosystem of open-source Big Data software.

Bigtable's powerful back-end servers offer several key advantages over a self-managed HBase installation:

Incredible scalability. Bigtable scales in direct proportion to the number of machines in your cluster. A self-managed HBase installation has a design bottleneck that limits the performance after a certain threshold is reached. Bigtable does not

have this bottleneck, so you can scale your cluster up to handle more reads and writes.

Simple administration. Bigtable handles upgrades and restarts transparently, and it automatically maintains high data durability. To replicate your data, simply add a second cluster to your instance, and replication starts automatically. No more

managing replicas or regions; just design your table schemas, and Bigtable will handle the rest for you.

Cluster resizing without downtime. You can increase the size of a Bigtable cluster for a few hours to handle a large load, then reduce the cluster's size again--all without any downtime. After you change a cluster's size, it typically takes just a

few minutes under load for Bigtable to balance performance across all of the nodes in your cluster.

QUESTION 2

An organization wants to use all available data to offer predictive suggestions on their website that improve over time. Which method should the organization use?



- A. Data automation
- B. Trends analysis
- C. Machine learning
- D. Multiple regression

Correct Answer: C

QUESTION 3

A startup is planning to create their entire suite of applications on Google Cloud. They are looking at various open source technologies to build applications. One of the consideration is about having a well integrated monitoring tool. They have to be able to constantly review load capacity and performance of their applications and virtual machines. What would you advise them to do?

- A. It is best to build a custom solution so that they know it integrates well with all their custom applications.
- B. Since they are using open source for applications, find another open source monitoring tool and integrate it, which could turn out to be very cheap.
- C. Use the Google Cloud Operations Suite which contains monitoring among other operations tools.
- D. Update the application code to regularly write to output logs. Export the logs to BigQuery to analyze them frequently.

Correct Answer: C

Operations Suite is well integrated into Google and it is the recommended option. References:
<https://cloud.google.com/products/operations>

QUESTION 4

An organization needs a platform to create custom end-to-end artificial intelligence models. Which Google Cloud product or service should the organization use?

- A. Dataproc
- B. Compute Engine
- C. Recommendations AI
- D. Vertex AI

Correct Answer: D

Recommendations AI enables you to build an end-to-end personalized recommendation system based on state-of-the-art deep learning ML models, without a need for expertise in ML or recommendation systems. With Vertex AI, both AutoML training and custom training are available options. Whichever option you choose for training, you can save models, deploy models, and request predictions with Vertex AI. <https://cloud.google.com/vertex-ai>

**QUESTION 5**

An organization needs frequent access to only a subset of their data. They want to reduce costs by depositing the rest of their data across Nearline Coldline and Archive repositories Which Google Cloud product should the organization use?

- A. Filestore
- B. Cloud Spanner
- C. Data Catalog
- D. Cloud Storage

Correct Answer: D

Per Google docs, specifically for GCP Cloud Storage there exists four types of storage with one of them, standard storage, being described as "storage for data that is frequently accessed ("hot" data) and/or stored for only brief periods of time." <https://cloud.google.com/storage>

[CLOUD-DIGITAL-LEADER
PDF Dumps](#)

[CLOUD-DIGITAL-LEADER
VCE Dumps](#)

[CLOUD-DIGITAL-LEADER
Exam Questions](#)