



**Free Questions for Professional-Cloud-Architect by certsdeals**

**Shared by Beasley on 07-06-2022**

**For More Free Questions and Preparation Resources**

**Check the Links on Last Page**

# Question 1

---

## Question Type: MultipleChoice

---

Your company is running a stateless application on a Compute Engine instance. The application is used heavily during regular business hours and lightly outside of business hours. Users are reporting that the application is slow during peak hours. You need to optimize the application's performance. What should you do?

### Options:

---

- A)** Create a snapshot of the existing disk. Create an instance template from the snapshot. Create an autoscaled managed instance group from the instance template.
- B)** Create a snapshot of the existing disk. Create a custom image from the snapshot. Create an autoscaled managed instance group from the custom image.
- C)** Create a custom image from the existing disk. Create an instance template from the custom image. Create an autoscaled managed instance group from the instance template.
- D)** Create an instance template from the existing disk. Create a custom image from the instance template. Create an autoscaled managed instance group from the custom image.

### Answer:

---

B

**Explanation:**

---

<https://cloud.google.com/compute/docs/instance-templates/create-instance-templates>

## Question 2

---

**Question Type: MultipleChoice**

---

Your marketing department wants to send out a promotional email campaign. The development team wants to minimize direct operation management. They project a wide range of possible customer responses, from 100 to 500,000 click-throughs per day. The link leads to a simple website that explains the promotion and collects user information and preferences. Which infrastructure should you recommend?

**Options:**

---

- A) Use Google App Engine to serve the website and Google Cloud Datastore to store user data.
- B) Use a Google Container Engine cluster to serve the website and store data to persistent disk.

- C) Use a managed instance group to serve the website and Google Cloud Bigtable to store user data.
- D) Use a single compute Engine virtual machine (VM) to host a web server, backed by Google Cloud SQL.

**Answer:**

---

A

**Explanation:**

---



References: <https://cloud.google.com/storage-options/>

## Question 3

---

**Question Type:** MultipleChoice

---

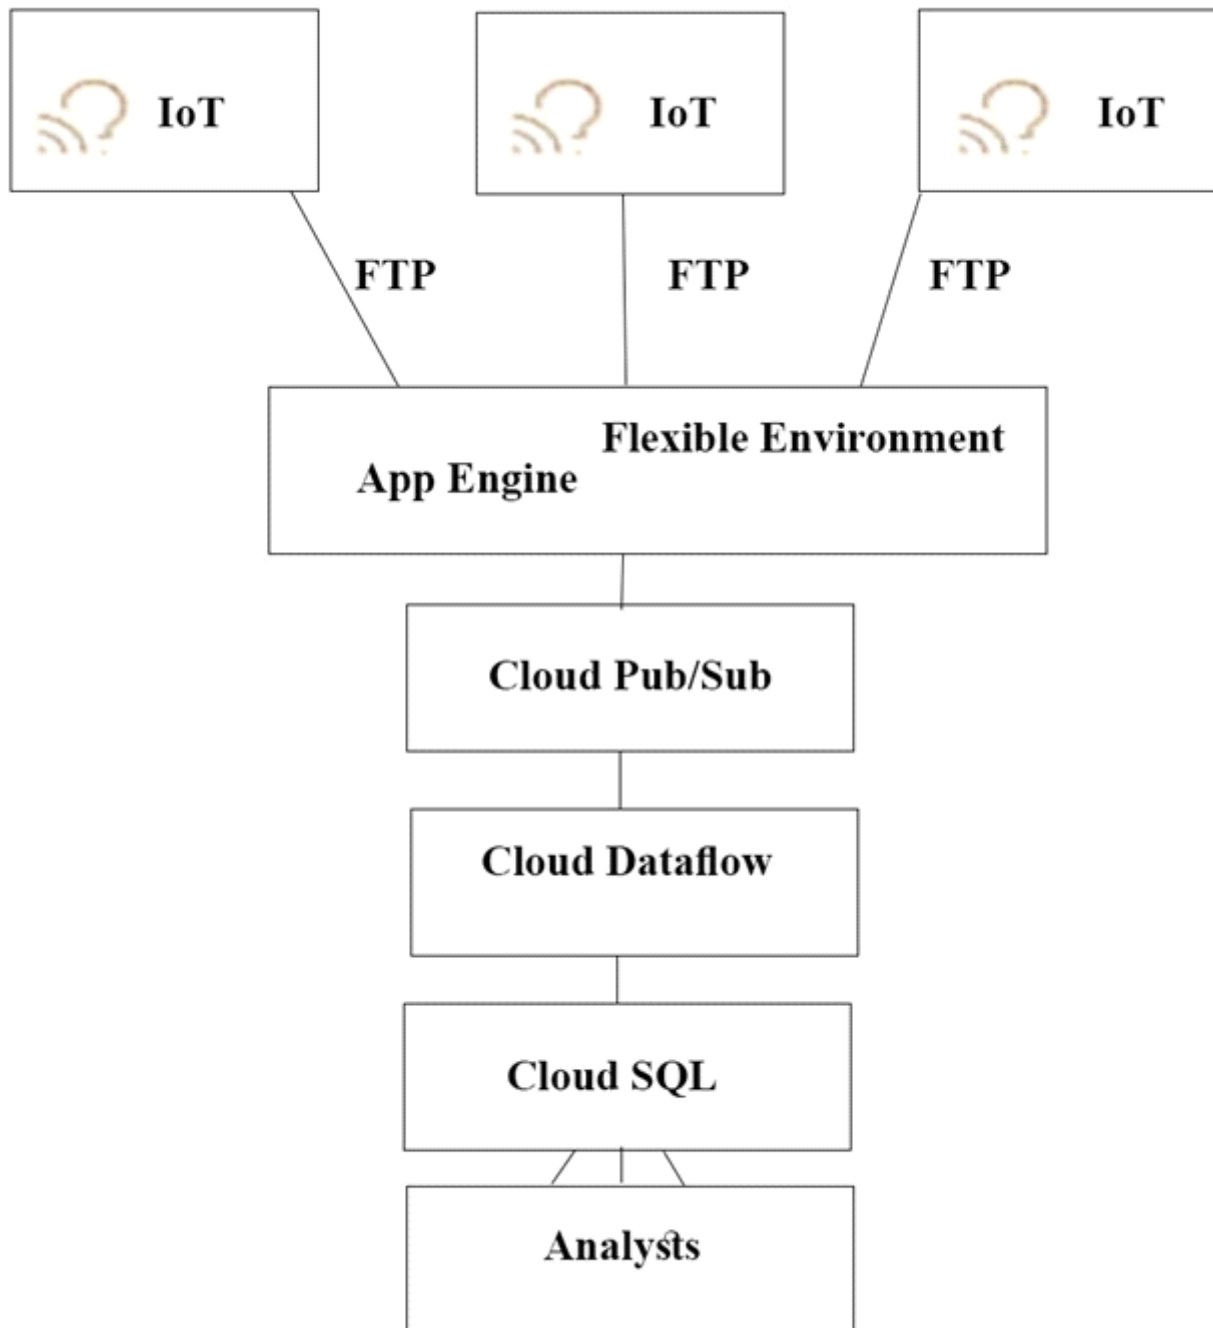
For this question, refer to the TerramEarth case study.

TerramEarth's CTO wants to use the raw data from connected vehicles to help identify approximately when a vehicle in the development team to focus their failure. You want to allow analysts to centrally query the vehicle data. Which architecture should you recommend?

**Options:**


---

A) Option




**B)** Option






C) Option



D) Option



**Answer:**

---

A

**Explanation:**

---

<https://cloud.google.com/solutions/iot/>

<https://cloud.google.com/solutions/designing-connected-vehicle-platform>

[https://cloud.google.com/solutions/designing-connected-vehicle-platform#data\\_ingestion](https://cloud.google.com/solutions/designing-connected-vehicle-platform#data_ingestion)

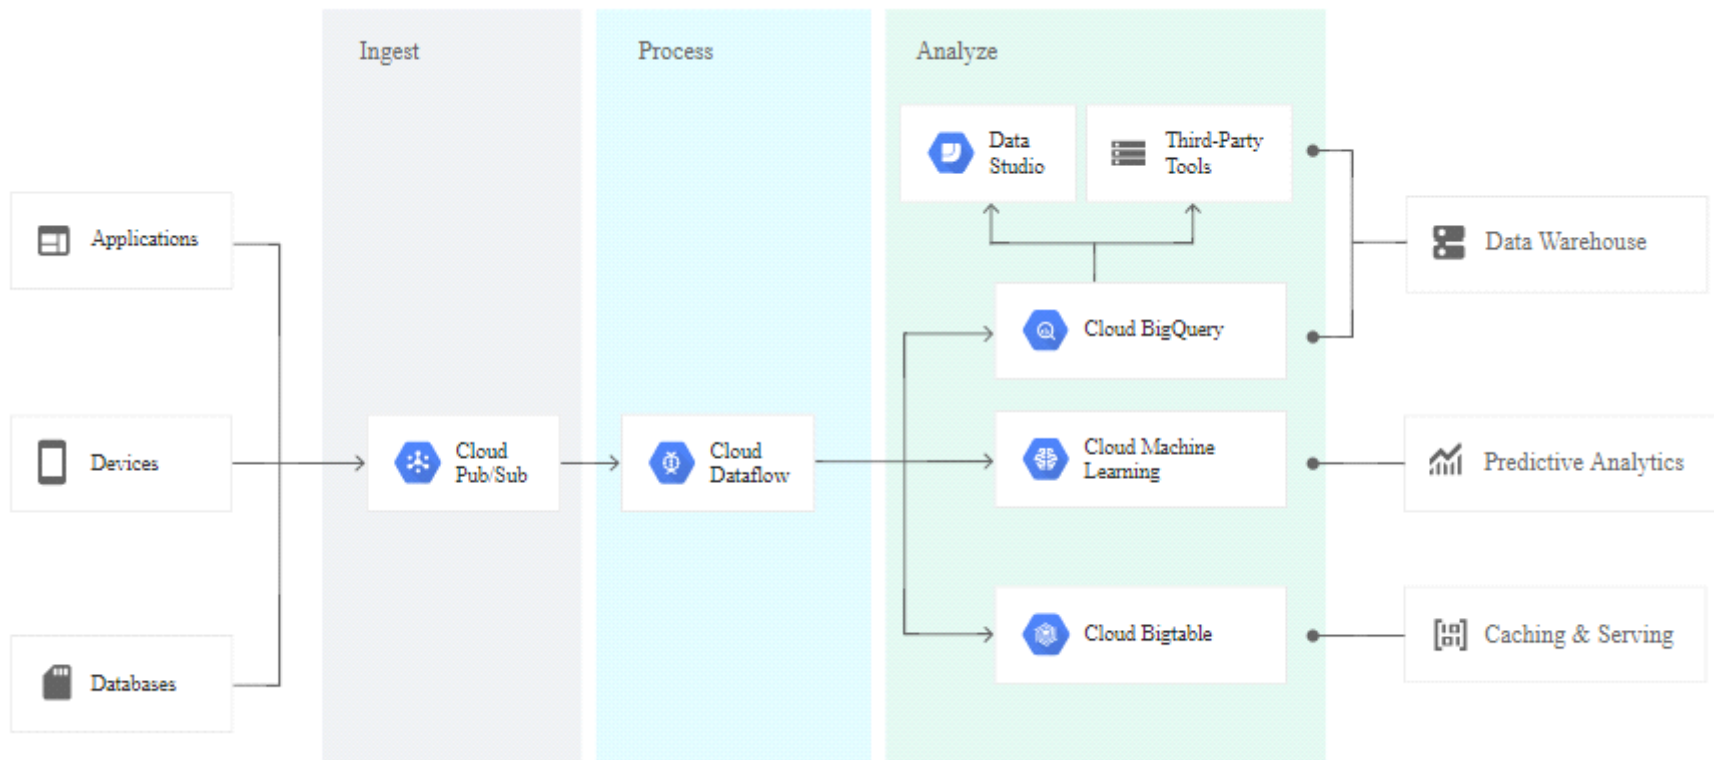
<http://www.eweek.com/big-data-and-analytics/google-touts-value-of-cloud-iot-core-for-analyzing-connected-car-data>

<https://cloud.google.com/solutions/iot/>

The push endpoint can be a load balancer.

A container cluster can be used.

Cloud Pub/Sub for Stream Analytics



References: <https://cloud.google.com/pubsub/>

<https://cloud.google.com/solutions/iot/>

<https://cloud.google.com/solutions/designing-connected-vehicle-platform>

[https://cloud.google.com/solutions/designing-connected-vehicle-platform#data\\_ingestion](https://cloud.google.com/solutions/designing-connected-vehicle-platform#data_ingestion)

<http://www.eweek.com/big-data-and-analytics/google-touts-value-of-cloud-iot-core-for-analyzing-connected-car-data>

## Question 4

---

**Question Type:** MultipleChoice

---

Y and ensure that you don't run out of storage, maintain 75% CPU usage cores, and keep replication lag below 60 seconds. What are the correct steps to meet your requirements?

### Options:

---

- A)** 1) Enable automatic storage increase for the instance.  
2) Create a Stackdriver alert when CPU usage exceeds 75%, and change the instance type to reduce CPU usage.  
3) Create a Stackdriver alert for replication lag, and shard the database to reduce replication time.
- B)** 1) Enable automatic storage increase for the instance.  
2) Change the instance type to a 32-core machine type to keep CPU usage below 75%.  
3) Create a Stackdriver alert for replication lag, and shard the database to reduce replication time.
- C)** 1) Create a Stackdriver alert when storage exceeds 75%, and increase the available storage on the instance to create more space.  
2) Deploy memcached to reduce CPU load.



3) Change the instance type to a 32-core machine type to reduce replication lag.

**D)** 1) Create a Stackdriver alert when storage exceeds 75%, and increase the available storage on the instance to create more space.

2) Deploy memcached to reduce CPU load.

3) Create a Stackdriver alert for replication lag, and change the instance type to a 32-core machine type to reduce replication lag.

**Answer:**

---

A

## Question 5

---

**Question Type:** MultipleChoice

---

For this question, refer to the Mountkirk Games case study. Which managed storage option meets Mountkirk's technical requirement for storing game activity in a time series database service?

**Options:**

---

**A)** Cloud Bigtable

**B)** Cloud Spanner

C) BigQuery

D) Cloud Datastore

**Answer:**

---

A

## Question 6

---

**Question Type: MultipleChoice**

---

For this question, refer to the Mountkirk Games case study. Mountkirk Games wants you to design a way to test the analytics platform's resilience to changes in mobile network latency. What should you do?

**Options:**

---

- A) Deploy failure injection software to the game analytics platform that can inject additional latency to mobile client analytics traffic.
- B) Build a test client that can be run from a mobile phone emulator on a Compute Engine virtual machine, and run multiple copies in Google Cloud Platform regions all over the world to generate realistic traffic.
- C) Add the ability to introduce a random amount of delay before beginning to process analytics files uploaded from mobile devices.
- D) Create an opt-in beta of the game that runs on players' mobile devices and collects response times from analytics endpoints running

in Google Cloud Platform regions all over the world.

**Answer:**

---

C

## Question 7

---

**Question Type: MultipleChoice**

---

For this question, refer to the Dress4Win case study. Which of the compute services should be migrated as –is and would still be an optimized architecture for performance in the cloud?

**Options:**

---

- A) Web applications deployed using App Engine standard environment
- B) RabbitMQ deployed using an unmanaged instance group
- C) Hadoop/Spark deployed using Cloud Dataproc Regional in High Availability mode
- D) Jenkins, monitoring, bastion hosts, security scanners services deployed on custom machine types

## Answer:

---

A

## Question 8

---

### Question Type: MultipleChoice

---

You are using Cloud SQL as the database backend for a large CRM deployment. You want to scale as usage increases and ensure that you don't run out of storage, maintain 75% CPU usage cores, and keep replication

lag below 60 seconds. What are the correct steps to meet your requirements?

### Options:

---

- A)** 1. Enable automatic storage increase for the instance.  
2. Create a Stackdriver alert when CPU usage exceeds 75%, and change the instance type to reduce CPU usage.  
3. Create a Stackdriver alert for replication lag, and shard the database to reduce replication time.
- B)** 1. Enable automatic storage increase for the instance.  
2. Change the instance type to a 32-core machine type to keep CPU usage below 75%.  
3. Create a Stackdriver alert for replication lag, and shard the database to reduce replication time.
- C)** 1. Create a Stackdriver alert when storage exceeds 75%, and increase the available storage on the

instance to create more space.

2. Deploy memcached to reduce CPU load.

3. Change the instance type to a 32-core machine type to reduce replication lag.

**D)** 1. Create a Stackdriver alert when storage exceeds 75%, and increase the available storage on the instance to create more space.

2. Deploy memcached to reduce CPU load.

3. Create a Stackdriver alert for replication lag, and change the instance type to a 32-core machine type to reduce replication lag.

**Answer:**

---

A

## Question 9

---

**Question Type: MultipleChoice**

---

For this question, refer to the Mountkirk Games case study.

Mountkirk Games wants to set up a continuous delivery pipeline. Their architecture includes many small services that they want to be able to update and roll back quickly. Mountkirk Games has the following requirements:

- Services are deployed redundantly across multiple regions in the US and Europe.
- Only frontend services are exposed on the public internet.

- They can provide a single frontend IP for their fleet of services.
- Deployment artifacts are immutable.

Which set of products should they use?

**Options:**

---

- A)** Google Cloud Storage, Google Cloud Dataflow, Google Compute Engine
- B)** Google Cloud Storage, Google App Engine, Google Network Load Balancer
- C)** Google Container Registry, Google Container Engine, Google HTTP(s) Load Balancer
- D)** Google Cloud Functions, Google Cloud Pub/Sub, Google Cloud Deployment Manager

**Answer:**

---

D

**To Get Premium Files for Professional-Cloud-Architect Visit**

**<https://www.p2pexams.com/products/professional-cloud-architect>**

**For More Free Questions Visit**

**<https://www.p2pexams.com/google/pdf/professional-cloud-architect>**

