C.

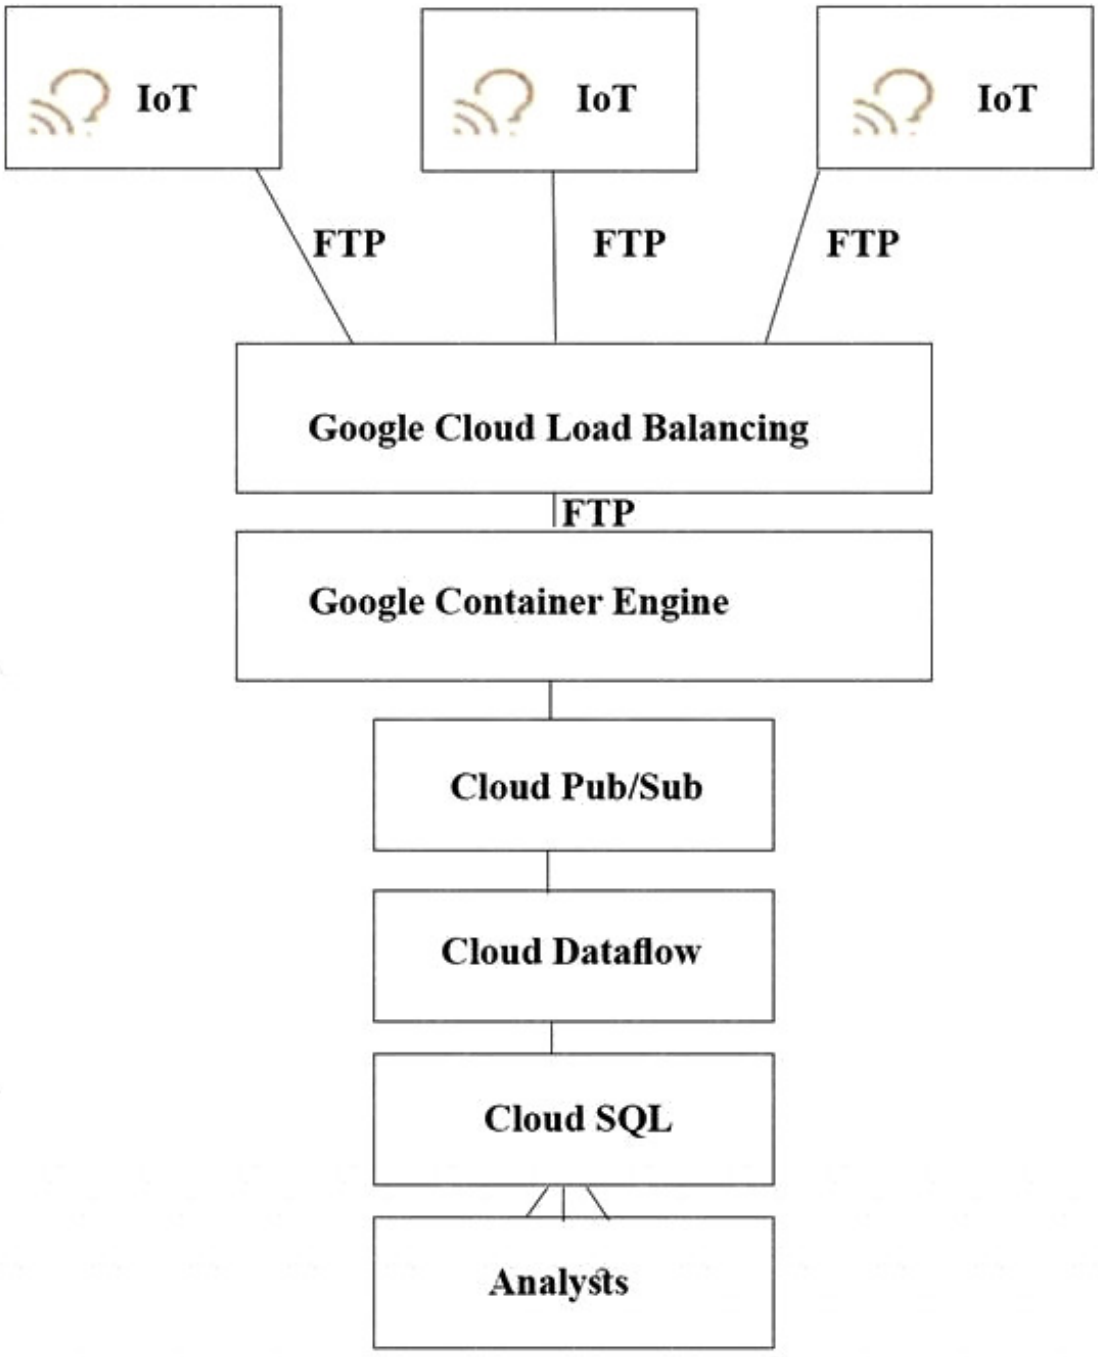
IoT
IoT
IoT
FTP
FTP
FTP
Google Cloud Load Balancing
FTP
Google Container Engine
Cloud Pub/Sub
Cloud Dataflow
Cloud SQL
Analysts

D.

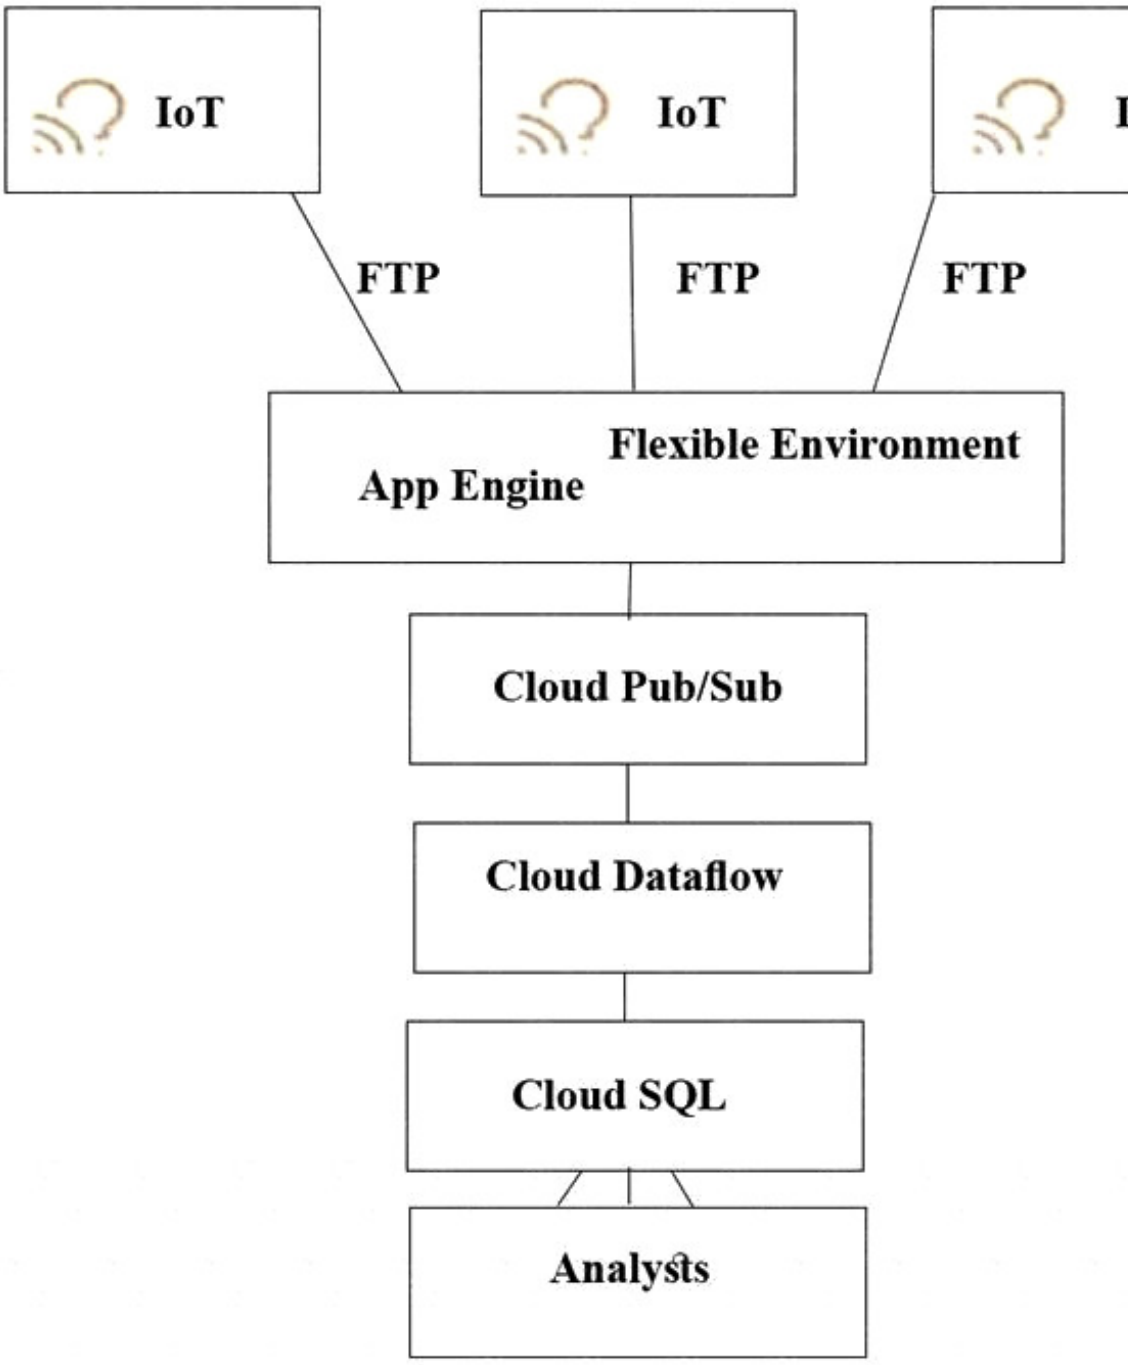

Correct Answer: A
Explanation:
https://cloud.google.com/solutions/iot/
https://cloud.google.com/solutions/designing-connected-vehicle-platform
https://cloud.google.com/solutions/designing-connected-vehicle-platform#data_ingestion
http://www.eweek.com/big-data-and-analytics/google-touts-value-of-cloud-iot-core-for-analyzing-

conne cted-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

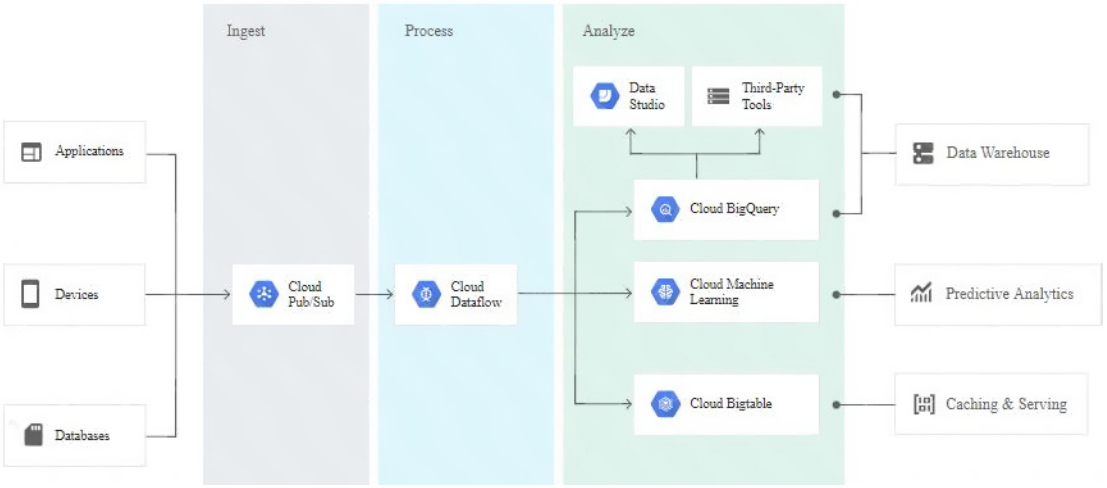

References:
https://cloud.google.com/pubsub/
https://cloud.google.com/solutions/iot/
https://cloud.google.com/solutions/designing-connected-vehicle-platform
https://cloud.google.com/solutions/designing-connected-vehicle-platform#data_ingestion
http://www.eweek.com/big-data-and-analytics/google-touts-value-of-cloud-iot-core-for-analyzing-connected-car-data
https://cloud.google.com/solutions/iot/

QUESTION 5
For this question refer to the TerramEarth case study. Which of TerramEarth's legacy enterprise processes will experience significant change as a result of increased Google Cloud Platform adoption.

A. Opex/capex allocation, LAN changes, capacity planning
B. Capacity planning, TCO calculations, opex/capex allocation
C. Capacity planning, utilization measurement, data center expansion
D. Data Center expansion, TCO calculations, utilization measurement

Correct Answer: B
Explanation:
Capacity planning, TCO calculations, opex/capex allocation From the case study, it can conclude that Management (CXO) all concern rapid provision of resources (infrastructure) for growing as well as cost management, such as Cost optimization in Infrastructure, trade up front capital expenditures (Capex) for ongoing operating expenditures (Opex), and Total cost of ownership (TCO)

QUESTION 6
For this question, refer to the TerramEarth case study. The TerramEarth development team wants to create an API to meet the company's business requirements. You want the development team to focus their development effort on business value versus creating a custom framework. Which method should they use?

A. Use Google App Engine with Google Cloud Endpoints. Focus on an API for dealers and partners.
B. Use Google App Engine with a JAX-RS Jersey Java-based framework. Focus on an API for the public.
C. Use Google App Engine with the Swagger (open API Specification) framework. Focus on an API for the public.
D. Use Google Container Engine with a Django Python container. Focus on an API for the public.
E. Use Google Container Engine with a Tomcat container with the Swagger (Open API Specification) framework. Focus on an API for dealers and partners.

Correct Answer: A
Explanation:
https://cloud.google.com/endpoints/docs/openapi/about-cloud-endpoints?hl=en_US&_ga=2.21787131.-1712523161.1522785064
https://cloud.google.com/endpoints/docs/openapi/architecture-overview
https://cloud.google.com/storage/docs/gsutil/commands/test

Develop, deploy, protect and monitor your APIs with Google Cloud Endpoints. Using an Open API Specification or one of our API frameworks, Cloud Endpoints gives you the tools you need for every phase of API development.

From scenario:
Business Requirements
Decrease unplanned vehicle downtime to less than 1 week, without increasing the cost of carrying surplus inventory

Support the dealer network with more data on how their customers use their equipment to better position new products and services

Have the ability to partner with different companies ?especially with seed and fertilizer suppliers in the fast-growing agricultural business ?to create compelling joint offerings for their customers.

Reference:
https://cloud.google.com/certification/guides/cloud-architect/casestudy-terramearth

QUESTION 7
For this question, refer to the TerramEarth case study. TerramEarth's 20 million vehicles are scattered around the world. Based on the vehicle's location its telemetry data is stored in a Google Cloud Storage (GCS) regional bucket (US. Europe, or Asia). The CTO has asked you to run a report on the raw telemetry data to determine why vehicles are breaking down after 100 K miles. You want to run this job on all the data. What is the most cost-effective way to run this job?

A. Move all the data into 1 zone, then launch a Cloud Dataproc cluster to run the job.
B. Move all the data into 1 region, then launch a Google Cloud Dataproc cluster to run the job.
C. Launch a cluster in each region to preprocess and compress the raw data, then move the data into a multi region bucket and use a Dataproc cluster to finish the job.
D. Launch a cluster in each region to preprocess and compress the raw data, then move the data into a region bucket and use a Cloud Dataproc cluster to finish the jo

Correct Answer: D
Explanation:
Storageguarantees 2 replicates which are geo diverse (100 miles apart) which can get better remote latency and availability.

More importantly, is that multiregional heavily leverages Edge caching and CDNs to provide the content to the end users.

All this redundancy and caching means that Multiregional comes with overhead to sync and ensure consistency between geo-diverse areas. As such, it's much better for write-once-read-many scenarios. This means frequently accessed (e.g. "hot" objects) around the world, such as website content, streaming videos, gaming or mobile applications.

References:
https://medium.com/google-cloud/google-cloud-storage-what-bucket-class-for-the-best-performance-5c847ac8f9f2

QUESTION 8
For this question, refer to the TerramEarth case study. TerramEarth has equipped unconnected trucks with servers and sensors to collet telemetry data. Next year they want to use the data to train machine learning models. They want to store this data in the cloud while reducing costs. What should they do?

A. Have the vehicle' computer compress the data in hourly snapshots, and store it in a Google Cloud storage (GCS) Nearline bucket.
B. Push the telemetry data in Real-time to a streaming dataflow job that compresses the data, and store it in Google BigQuery.
C. Push the telemetry data in real-time to a streaming dataflow job that compresses the data, and store it in Cloud Bigtable.
D. Have the vehicle's computer compress the data in hourly snapshots, a Store it in a GCS Coldline bucket.

Correct Answer: D
Explanation:
Coldline Storage is the best choice for data that you plan to access at most once a year, due to its slightly lower availability, 90-day minimum storage duration, costs for data access, and higher per-operation costs. For example:

Cold Data Storage - Infrequently accessed data, such as data stored for legal or regulatory reasons, can be stored at low cost as Coldline Storage, and be available when you need it.

Disaster recovery - In the event of a disaster recovery event, recovery time is key. Cloud Storage provides low latency access to data stored as Coldline Storage.

References:
https://cloud.google.com/storage/docs/storage-classes

QUESTION 9
For this question refer to the TerramEarth case study. Operational parameters such as oil pressure are adjustable on each of TerramEarth's vehicles to increase their efficiency, depending on their environmental conditions. Your primary goal is to increase the operating efficiency of all