

Google Cloud Architect Professional

1. For this question, refer to the Mountkirk Games case study. Mountkirk Games wants you to design their new testing strategy. How should the test coverage differ from their existing backends on the other platforms?

A. Tests should scale well beyond the prior approaches.

B. Unit tests are no longer required, only end-to-end tests.

C. Tests should be applied after the release is in the production environment.

D. Tests should include directly testing the Google Cloud Platform (GCP) infrastructure.

Answer(s): A

2. For this question, refer to the Mountkirk Games case study. Mountkirk Games has deployed their new backend on Google Cloud Platform (GCP). You want to create a thorough testing process for new versions of the backend before they are released to the public. You want the testing environment to scale in an economical way. How should you design the process?

A. Create a scalable environment in GCP for simulating production load.

B. Use the existing infrastructure to test the GCP-based backend at scale.

C. Build stress tests into each component of your application using resources internal to GCP to simulate load.

D. Create a set of static environments in GCP to test different levels of load -- for example, high, medium, and low.

Answer(s): A

3. 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?

A. Google Cloud Storage, Google Cloud Dataflow, Google Compute Engine

B. Google Cloud Storage, Google App Engine, Google Network Load Balancer

C. Google Kubernetes Registry, Google Container Engine, Google HTTP(S) Load Balancer

D. Google Cloud Functions, Google Cloud Pub/Sub, Google Cloud Deployment Manager

Answer(s): C

4. For this question, refer to the Mountkirk Games case study. Mountkirk Games' gaming servers are not automatically scaling properly. Last month, they rolled out a new feature, which suddenly became very popular. A record number of users are trying to use the service, but many of them are getting 503 errors and very slow response times.

What should they investigate first?

A. Verify that the database is online.

B. Verify that the project quota hasn't been exceeded.

C. Verify that the new feature code did not introduce any performance bugs.

D. Verify that the load-testing team is not running their tool against production.

Answer(s): B

5. For this question, refer to the Mountkirk Games case study. Mountkirk Games needs to create a repeatable and configurable mechanism for deploying isolated application environments. Developers and testers can access each other's environments and resources, but they cannot access staging or production resources. The staging environment needs access to some services from production.

What should you do to isolate development environments from staging and production?

A. Create a project for development and test and another for staging and production.

B. Create a network for development and test and another for staging and production.

C. Create one subnetwork for development and another for staging and production.

D. Create one project for development, a second for staging and a third for production.

Answer(s): D

6. For this question, refer to the Mountkirk Games case study. Mountkirk Games wants to set up a real-time analytics platform for their new game. The new platform must meet their technical requirements.

Which combination of Google technologies will meet all of their requirements?

A. Container Engine, Cloud Pub/Sub, and Cloud SQL

B. Cloud Dataflow, Cloud Storage, Cloud Pub/Sub, and BigQuery

C. Cloud SQL, Cloud Storage, Cloud Pub/Sub, and Cloud Dataflow

D. Cloud Dataproc, Cloud Pub/Sub, Cloud SQL, and Cloud Dataflow

E. Cloud Pub/Sub, Compute Engine, Cloud Storage, and Cloud Dataproc

Answer(s): B

7. 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

A. Which architecture should you recommend?A)B)C)D)

B. Option A

C. Option B

D. Option C

E. Option D

Answer(s): A

8. 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.

Answer(s): A

9. For this question, refer to the TerramEarth case study

Your development team has created a structured API to retrieve vehicle data

A. They want to allow third parties to develop tools for dealerships that use this vehicle event data. You want to support delegated authorization against this data. What should you do?

B. Build or leverage an OAuth-compatible access control system.

C. Build SAML 2.0 SSO compatibility into your authentication system.

D. Restrict data access based on the source IP address of the partner systems.

E. Create secondary credentials for each dealer that can be given to the trusted third party.

Answer(s): A

10. For this question, refer to the TerramEarth case study.

TerramEarth plans to connect all 20 million vehicles in the field to the cloud. This increases the volume to 20 million 600 byte records a second for 40 TB an hour. How should you design the data ingestion?

A. Vehicles write data directly to GCS.

B. Vehicles write data directly to Google Cloud Pub/Sub.

C. Vehicles stream data directly to Google BigQuery.

D. Vehicles continue to write data using the existing system (FTP).

Answer(s): B

11. For this question, refer to the TerramEarth case study

You analyzed TerramEarth's business requirement to reduce downtime, and found that they can achieve a majority of time saving by reducing customers' wait time for parts. You decided to focus on reduction of the 3 weeks aggregate reporting time. Which modifications to the company's processes should you recommend?

A. Migrate from CSV to binary format, migrate from FTP to SFTP transport, and develop machine learning analysis of metrics.

B. Migrate from FTP to streaming transport, migrate from CSV to binary format, and develop machine learning analysis of metrics.

C. Increase fleet cellular connectivity to 80%, migrate from FTP to streaming transport, and develop machine learning analysis of metrics.

D. Migrate from FTP to SFTP transport, develop machine learning analysis of metrics, and increase dealer local inventory by a fixed factor.

Answer(s): C

12. 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

Answer(s): B

13. For this question, refer to the TerramEarth case study.

To speed up data retrieval, more vehicles will be upgraded to cellular connections and be able to transmit data to the ETL process. The current FTP process is error-prone and restarts the data transfer from the start of the file when connections fail, which happens often. You want to improve the reliability of the solution and minimize data transfer time on the cellular connections.

What should you do?

A. Use one Google Container Engine cluster of FTP servers. Save the data to a Multi-Regional bucket. Run the ETL process using data in the bucket.

B. Use multiple Google Container Engine clusters running FTP servers located in different regions. Save the data to Multi-Regional buckets in us, eu, and asia. Run the ETL process using the data in the bucket.

C. Directly transfer the files to different Google Cloud Multi-Regional Storage bucket locations in us, eu, and asia using Google APIs over HTTP(S). Run the ETL process using the data in the bucket.

D. Directly transfer the files to a different Google Cloud Regional Storage bucket location in us, eu, and asia using Google APIs over HTTP(S). Run the ETL process to retrieve the data from each Regional bucket.

Answer(s): D

14. 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 dat

A. What is the most cost-effective way to run this job?

B. Move all the data into 1 zone, then launch a Cloud Dataproc cluster to run the job.

C. Move all the data into 1 region, then launch a Google Cloud Dataproc cluster to run the job.

D. 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.

E. 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

Answer(s): D

15. For this question, refer to the TerramEarth case study.

TerramEarth has equipped unconnected trucks with servers and sensors to collect telemetry dat

A. 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?

B. Have the vehicle's computer compress the data in hourly snapshots, and store it in a Google Cloud storage (GCS) Nearline bucket.

C. Push the telemetry data in Real-time to a streaming dataflow job that compresses the data, and store it in Google BigQuery.

D. Push the telemetry data in real-time to a streaming dataflow job that compresses the data, and store it in Cloud Bigtable.

E. Have the vehicle's computer compress the data in hourly snapshots, and store it in a GCS Coldline bucket.

Answer(s): D

16. 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 20 million cellular and unconnected vehicles in the field. How can you accomplish this goal?

A. Have your engineers inspect the data for patterns, and then create an algorithm with rules that make operational adjustments automatically.

B. Capture all operating data, train machine learning models that identify ideal operations, and run locally to make operational adjustments automatically.

C. Implement a Google Cloud Dataflow streaming job with a sliding window, and use Google Cloud Messaging (GCM) to make operational adjustments automatically.

D. Capture all operating data, train machine learning models that identify ideal operations, and host in Google Cloud Machine Learning (ML) Platform to make operational adjustments automatically.

Answer(s): B

17. Your agricultural division is experimenting with fully autonomous vehicles. You want your architecture to promote strong security during vehicle operation.

Which two architecture should you consider?

Choose 2 answers:

A. Treat every micro service call between modules on the vehicle as untrusted.

B. Require IPv6 for connectivity to ensure a secure address space.

C. Use a trusted platform module (TPM) and verify firmware and binaries on boot.

D. Use a functional programming language to isolate code execution cycles.

E. Use multiple connectivity subsystems for redundancy.

F. Enclose the vehicle's drive electronics in a Faraday cage to isolate chips.

Answer(s): A C

18. For this question, refer to the JencoMart case study.

The JencoMart security team requires that all Google Cloud Platform infrastructure is deployed using a least privilege model with separation of duties for administration between production and development resources.

What Google domain and project structure should you recommend?

A. Create two G Suite accounts to manage users: one for development/test/staging and one for production. Each account should contain one project for every application.

B. Create two G Suite accounts to manage users: one with a single project for all development applications and one with a single project for all production applications.

C. Create a single G Suite account to manage users with each stage of each application in its own project.

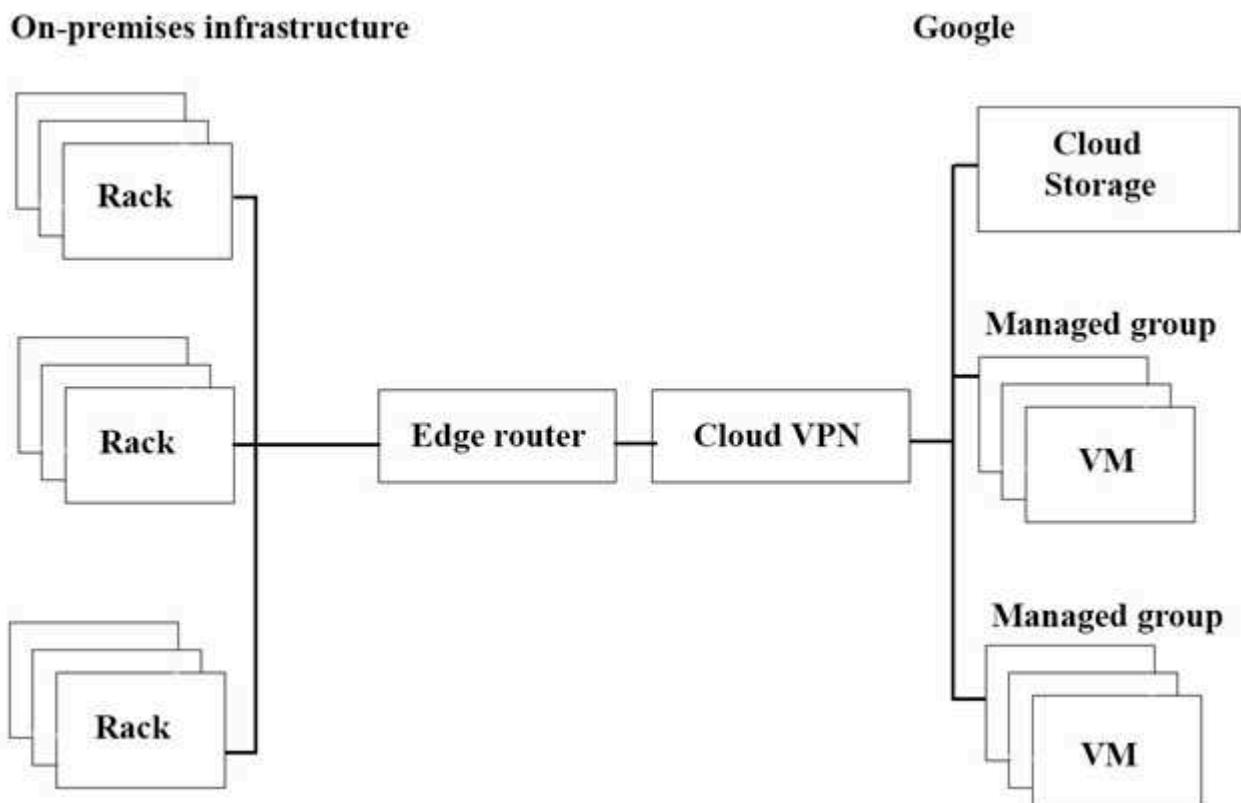
D. Create a single G Suite account to manage users with one project for the development/test/staging environment and one project for the production environment.

Answer(s): D

19. For this question, refer to the JencoMart case study.

The migration of JencoMart's application to Google Cloud Platform (GCP) is progressing too slowly. The infrastructure is shown in the diagram. You want to maximize throughput.

What are three potential bottlenecks? (Choose 3 answers.)



- A. A single VPN tunnel, which limits throughput
- B. A tier of Google Cloud Storage that is not suited for this task
- C. A copy command that is not suited to operate over long distances
- D. Fewer virtual machines (VMs) in GCP than on-premises machines
- E. A separate storage layer outside the VMs, which is not suited for this task
- F. Complicated internet connectivity between the on-premises infrastructure and GCP

Answer(s): A D F

20. For this question, refer to the JencoMart case study

A few days after JencoMart migrates the user credentials database to Google Cloud Platform and shuts down the old server, the new database server stops responding to SSH connections. It is still serving database requests to the application servers correctly.

What three steps should you take to diagnose the problem? Choose 3 answers

- A. Delete the virtual machine (VM) and disks and create a new one.
- B. Delete the instance, attach the disk to a new VM, and investigate.
- C. Take a snapshot of the disk and connect to a new machine to investigate.
- D. Check inbound firewall rules for the network the machine is connected to.
- E. Connect the machine to another network with very simple firewall rules and investigate.
- F. Print the Serial Console output for the instance for troubleshooting, activate the interactive console, and investigate.

Answer(s): C D F
