

100% Money Back
Guarantee

Vendor:Google

Exam Code:PROFESSIONAL-CLOUD-ARCHITECT

Exam Name:Professional Cloud Architect on Google
Cloud Platform

Version:Demo

QUESTION 1

For this question, refer to the TerramEarth case study. You have broken down a legacy monolithic application into a few containerized RESTful microservices. You want to run those microservices on Cloud Run. You also want to make sure the services are highly available with low latency to your customers. What should you do?

- A. Deploy Cloud Run services to multiple availability zones. Create Cloud Endpoints that point to the services. Create a global HTTP(S) Load Balancing instance and attach the Cloud Endpoints to its backend.
- B. Deploy Cloud Run services to multiple regions. Create serverless network endpoint groups pointing to the services. Add the serverless NEGs to a backend service that is used by a global HTTP(S) Load Balancing instance.
- C. Deploy Cloud Run services to multiple regions. In Cloud DNS, create a latency-based DNS name that points to the services.
- D. Deploy Cloud Run services to multiple availability zones. Create a TCP/IP global load balancer. Add the Cloud Run Endpoints to its backend service.

Correct Answer: B

QUESTION 2

Your company has an application deployed on Anthos clusters (formerly Anthos GKE) that is running multiple microservices. The cluster has both Anthos Service Mesh and Anthos Config Management configured. End users inform you that the application is responding very slowly. You want to identify the microservice that is causing the delay. What should you do?

- A. Use the Service Mesh visualization in the Cloud Console to inspect the telemetry between the microservices.
- B. Use Anthos Config Management to create a ClusterSelector selecting the relevant cluster. On the Google Cloud Console page for Google Kubernetes Engine, view the Workloads and filter on the cluster. Inspect the configurations of the filtered workloads.
- C. Use Anthos Config Management to create a namespaceSelector selecting the relevant cluster namespace. On the Google Cloud Console page for Google Kubernetes Engine, visit the workloads and filter on the namespace. Inspect the configurations of the filtered workloads.
- D. Reinstall istio using the default istio profile in order to collect request latency. Evaluate the telemetry between the microservices in the Cloud Console.

Correct Answer: A

The Anthos Service Mesh pages in the Google Cloud Console provide both summary and in-depth metrics, charts, and graphs that enable you to observe service behavior. You can monitor the overall health of your services, or drill down on a specific service to set a service level objective (SLO) or troubleshoot an issue. <https://cloud.google.com/service-mesh/docs/observability/explore-dashboard> <https://cloud.google.com/anthos/service-mesh>

QUESTION 3

Your BigQuery project has several users. For audit purposes, you need to see how many queries each user ran in the last month. What should you do?

- A. Connect Google Data Studio to BigQuery. Create a dimension for the users and a metric for the amount of queries per user.
- B. In the BigQuery interface, execute a query on the JOBS table to get the required information.
- C. Use ``bq show`` to list all jobs. Per job, use ``bq ls`` to list job information and get the required information.
- D. Use Cloud Audit Logging to view Cloud Audit Logs, and create a filter on the query operation to get the required information.

Correct Answer: D

Reference <https://cloud.google.com/bigquery/docs/access-control>

QUESTION 4

You are deploying a PHP App Engine Standard service with SQL as the backend. You want to minimize the number of queries to the database. What should you do?

- A. Set the memcache service level to dedicated. Create a key from the hash of the query, and return database values from memcache before issuing a query to Cloud SQL.
- B. Set the memcache service level to dedicated. Create a cron task that runs every minute to populate the cache with keys containing query results.
- C. Set the memcache service level to shared. Create a cron task that runs every minute to save all expected queries to a key called "cached-queries".
- D. Set the memcache service level to shared. Create a key called "cached-queries", and return database values from the key before using a query to Cloud SQL.

Correct Answer: A

<https://cloud.google.com/appengine/docs/standard/php/memcache/using>

QUESTION 5

Your company's test suite is a custom C++ application that runs tests throughout each day on Linux virtual machines. The full test suite takes several hours to complete, running on a limited number of on-premises servers reserved for testing. Your company wants to move the testing infrastructure to the cloud, to reduce the amount of time it takes to fully test a change to the system, while changing the tests as little as possible.

Which cloud infrastructure should you recommend?

- A. Google Compute Engine unmanaged instance groups and Network Load Balancer
- B. Google Compute Engine managed instance groups with auto-scaling
- C. Google Cloud Dataproc to run Apache Hadoop jobs to process each test

D. Google App Engine with Google StackDriver for logging

Correct Answer: B

QUESTION 6

Your customer wants to do resilience testing of their authentication layer. This consists of a regional managed instance group serving a public REST API that reads from and writes to a Cloud SQL instance. What should you do?

- A. Engage with a security company to run web scrapes that look your users' authentication data on malicious websites and notify you if any if found.
- B. Deploy intrusion detection software to your virtual machines to detect and log unauthorized access.
- C. Schedule a disaster simulation exercise during which you can shut off all VMs in a zone to see how your application behaves.
- D. Configure a read replica for your Cloud SQL instance in a different zone than the master, and then manually trigger a failover while monitoring KPIs for our REST API.

Correct Answer: C

QUESTION 7

As part of implementing their disaster recovery plan, your company is trying to replicate their production MySQL database from their private data center to their GCP project using a Google Cloud VPN connection. They are experiencing latency issues and a small amount of packet loss that is disrupting the replication. What should they do?

- A. Configure their replication to use UDP.
- B. Configure a Google Cloud Dedicated Interconnect.
- C. Restore their database daily using Google Cloud SQL.
- D. Add additional VPN connections and load balance them.
- E. Send the replicated transaction to Google Cloud Pub/Sub.

Correct Answer: B

QUESTION 8

For this question, refer to the TerraEarth case study. You are building a microservice-based application for TerraEarth. The application is based on Docker containers. You want to follow Google-recommended practices to build the application continuously and store the build artifacts. What should you do?

- A. Configure a trigger in Cloud Build for new source changes. Invoke Cloud Build to build container images for each microservice, and tag them using the code commit hash. Push the images to the Container Registry.

B. Configure a trigger in Cloud Build for new source changes. The trigger invokes build jobs and build container images for the microservices. Tag the images with a version number, and push them to Cloud Storage.

C. Create a Scheduler job to check the repo every minute. For any new change, invoke Cloud Build to build container images for the microservices. Tag the images using the current timestamp, and push them to the Container Registry.

D. Configure a trigger in Cloud Build for new source changes. Invoke Cloud Build to build one container image, and tag the image with the label `latest`. Push the image to the Container Registry.

Correct Answer: A

QUESTION 9

Your company is building a new architecture to support its data-centric business focus. You are responsible for setting up the network. Your company's mobile and web-facing applications will be deployed on-premises, and all data analysis will be conducted in GCP. The plan is to process and load 7 years of archived .csv files totaling 900 TB of data and then continue loading 10 TB of data daily. You currently have an existing 100-MB internet connection.

What actions will meet your company's needs?

A. Compress and upload both archived files and files uploaded daily using the `gsutil -m` option.

B. Lease a Transfer Appliance, upload archived files to it, and send it, and send it to Google to transfer archived data to Cloud Storage. Establish a connection with Google using a Dedicated Interconnect or Direct Peering connection and use it to upload files daily.

C. Lease a Transfer Appliance, upload archived files to it, and send it, and send it to Google to transfer archived data to Cloud Storage. Establish one Cloud VPN Tunnel to VPC networks over the public internet, and compress and upload files daily using the `gsutil -m` option.

D. Lease a Transfer Appliance, upload archived files to it, and send it to Google to transfer archived data to Cloud Storage. Establish a Cloud VPN Tunnel to VPC networks over the public internet, and compress and upload files daily.

Correct Answer: B

<https://cloud.google.com/interconnect/docs/how-to/direct-peering>

QUESTION 10

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

Correct Answer: D

References: <https://cloud.google.com/appengine/docs/standard/go/creating-separate-dev-environments>

QUESTION 11

For this question, refer to the Mountkirk Games case study. You need to analyze and define the technical architecture for the database workloads for your company, Mountkirk Games. Considering the business and technical requirements, what should you do?

- A. Use Cloud SQL for time series data, and use Cloud Bigtable for historical data queries.
- B. Use Cloud SQL to replace MySQL, and use Cloud Spanner for historical data queries.
- C. Use Cloud Bigtable to replace MySQL, and use BigQuery for historical data queries.
- D. Use Cloud Bigtable for time series data, use Cloud Spanner for transactional data, and use BigQuery for historical data queries.

Correct Answer: D

QUESTION 12

For this question, refer to the Helicopter Racing League (HRL) case study. The HRL development team releases a new version of their predictive capability application every Tuesday evening at 3 a.m. UTC to a repository. The security team at HRL has developed an in-house penetration test Cloud Function called Airwolf. The security team wants to run Airwolf against the predictive capability application as soon as it is released every Tuesday. You need to set up Airwolf to run at the recurring weekly cadence. What should you do?

- A. Set up Cloud Tasks and a Cloud Storage bucket that triggers a Cloud Function.
- B. Set up a Cloud Logging sink and a Cloud Storage bucket that triggers a Cloud Function.
- C. Configure the deployment job to notify a Pub/Sub queue that triggers a Cloud Function.
- D. Set up Identity and Access Management (IAM) and Confidential Computing to trigger a Cloud Function.

Correct Answer: C