

Practice Exam Questions



Google Cloud

Cloud Developer



Google Certified Professional



EXAMKILLER

Help Pass Your Exam At First Try

Google

Exam GCP-Cloud-Developer

Google Cloud Certified Professional - Cloud Developer

Version: 10.0

[Total Questions: 227]

Topic break down

Topic	No. of Questions
Topic 1: HipLocal Case StudyCompany Overview	19
Topic 2: Misc Questions	208

Topic 1, HipLocal Case StudyCompany Overview

HipLocal is a community application designed to facilitate communication between people in close proximity. It is used for event planning and organizing sporting events, and for businesses to connect with their local communities. HipLocal launched recently in a few neighborhoods in Dallas and is rapidly growing into a global phenomenon. Its unique style of hyper-local community communication and business outreach is in demand around the world.

Executive statement

We are the number one local community app; it's time to take our local community services global. Our venture capital investors want to see rapid growth and the same great experience for new local and virtual communities that come online, whether their members are 10 or 10,000 miles away from each other.

Solution concept

HipLocal wants to expand their existing service, with updated functionality, in new regions to better serve their global customers. They want to hire and train a new team to support these regions in their time zones. They will need to ensure that the application scales smoothly and provides clear uptime data.

Existing technical environment

HipLocal's environment is a mix of on-premises hardware and infrastructure running in Google Cloud Platform. The HipLocal team understands their application well but has limited experience in global scale applications. Their existing technical environment is as follows:

 uk.co.certification.simulator.questionpool.PList@e3218b0

Business requirements

HipLocal's investors want to expand their footprint and support the increase in demand they are seeing. Their requirements are:

 uk.co.certification.simulator.questionpool.PList@e3219c0

Technical requirements

 uk.co.certification.simulator.questionpool.PList@e321ae0

Question No : 1 - (Topic 1)

In order for HipLocal to store application state and meet their stated business requirements, which database service should they migrate to?

- A. Cloud Spanner
- B. Cloud Datastore
- C. Cloud Memorystore as a cache
- D. Separate Cloud SQL clusters for each region

Answer: D

Question No : 2 - (Topic 1)

For this question, refer to the HipLocal case study.

HipLocal is expanding into new locations. They must capture additional data each time the application is launched in a new European country. This is causing delays in the development process due to constant schema changes and a lack of environments for conducting testing on the application changes. How should they resolve the issue while meeting the business requirements?

- A.** Create new Cloud SQL instances in Europe and North America for testing and deployment. Provide developers with local MySQL instances to conduct testing on the application changes.
- B.** Migrate data to Bigtable. Instruct the development teams to use the Cloud SDK to emulate a local Bigtable development environment.
- C.** Move from Cloud SQL to MySQL hosted on Compute Engine. Replicate hosts across regions in the Americas and Europe. Provide developers with local MySQL instances to conduct testing on the application changes.
- D.** Migrate data to Firestore in Native mode and set up instan

Answer: B

Question No : 3 - (Topic 1)

HipLocal has connected their Hadoop infrastructure to GCP using Cloud Interconnect in order to query data stored on persistent disks.

Which IP strategy should they use?

- A.** Create manual subnets.
- B.** Create an auto mode subnet.
- C.** Create multiple peered VPCs.
- D.** Provision a single instance for NAT.

Answer: A

Question No : 4 - (Topic 1)

For this question, refer to the HipLocal case study.

How should HipLocal redesign their architecture to ensure that the application scales to support a large increase in users?

- A.** Use Google Kubernetes Engine (GKE) to run the application as a microservice. Run the MySQL database on a dedicated GKE node.
- B.** Use multiple Compute Engine instances to run MySQL to store state information. Use a Google Cloud-managed load balancer to distribute the load between instances. Use managed instance groups for scaling.
- C.** Use Memorystore to store session information and CloudSQL to store state information. Use a Google Cloud-managed load balancer to distribute the load between instances. Use managed instance groups for scaling.
- D.** Use a Cloud Storage bucket to serve the application as a static website, and use another Cloud Storage bucket to store user state information.

Answer: D

Question No : 5 - (Topic 1)

HipLocal's .net-based auth service fails under intermittent load.

What should they do?

- A.** Use App Engine for autoscaling.
- B.** Use Cloud Functions for autoscaling.
- C.** Use a Compute Engine cluster for the service.
- D.** Use a dedicated Compute Engine virtual machine instance for the service.

Answer: D

Reference: <https://www.qwiklabs.com/focuses/611?parent=catalog>

Question No : 6 - (Topic 1)

HipLocal's data science team wants to analyze user reviews.

How should they prepare the data?

- A.** Use the Cloud Data Loss Prevention API for redaction of the review dataset.
- B.** Use the Cloud Data Loss Prevention API for de-identification of the review dataset.

- C. Use the Cloud Natural Language Processing API for redaction of the review dataset.
- D. Use the Cloud Natural Language Processing API for de-identification of the review dataset.

Answer: B

Explanation: <https://cloud.google.com/dlp/docs/deidentify-sensitive-data>

Question No : 7 - (Topic 1)

Which service should HipLocal use to enable access to internal apps?

- A. Cloud VPN
- B. Cloud Armor
- C. Virtual Private Cloud
- D. Cloud Identity-Aware Proxy

Answer: D

Reference: <https://cloud.google.com/iap/docs/cloud-iap-for-on-prem-apps-overview>

Question No : 8 - (Topic 1)

In order to meet their business requirements, how should HipLocal store their application state?

- A. Use local SSDs to store state.
- B. Put a memcache layer in front of MySQL.
- C. Move the state storage to Cloud Spanner.
- D. Replace the MySQL instance with Cloud SQL.

Answer: B

Question No : 9 - (Topic 1)

For this question, refer to the HipLocal case study.

Which Google Cloud product addresses HipLocal's business requirements for service level indicators and objectives?

- A. Cloud Profiler
- B. Cloud Monitoring
- C. Cloud Trace
- D. Cloud Logging

Answer: B

Explanation: <https://cloud.google.com/stackdriver/docs/solutions/slo-monitoring#defn-sli>

Question No : 10 - (Topic 1)

Which service should HipLocal use for their public APIs?

- A. Cloud Armor
- B. Cloud Functions
- C. Cloud Endpoints
- D. Shielded Virtual Machines

Answer: D

Question No : 11 - (Topic 1)

HipLocal is configuring their access controls.

Which firewall configuration should they implement?

- A. Block all traffic on port 443.
- B. Allow all traffic into the network.
- C. Allow traffic on port 443 for a specific tag.
- D. Allow all traffic on port 443 into the network.

Answer: D

Question No : 12 - (Topic 1)

HipLocal's APIs are showing occasional failures, but they cannot find a pattern. They want to collect some

metrics to help them troubleshoot.

What should they do?

- A. Take frequent snapshots of all of the VMs.
- B. Install the Stackdriver Logging agent on the VMs.
- C. Install the Stackdriver Monitoring agent on the VMs.
- D. Use Stackdriver Trace to look for performance bottlenecks.

Answer: C

Question No : 13 - (Topic 1)

Which database should HipLocal use for storing user activity?

- A. BigQuery
- B. Cloud SQL
- C. Cloud Spanner
- D. Cloud Datastore

Answer: A

Question No : 14 - (Topic 1)

For this question refer to the HipLocal case study.

HipLocal wants to reduce the latency of their services for users in global locations. They have created read replicas of their database in locations where their users reside and configured their service to read traffic using those replicas. How should they further reduce latency for all database interactions with the least amount of effort?

- A. Migrate the database to Bigtable and use it to serve all global user traffic.
- B. Migrate the database to Cloud Spanner and use it to serve all global user traffic.

- C. Migrate the database to Firestore in Datastore mode and use it to serve all global user traffic.
- D. Migrate the services to Google Kubernetes Engine and use a load balancer service to better scale the application.

Answer: D

Question No : 15 - (Topic 1)

HipLocal wants to reduce the number of on-call engineers and eliminate manual scaling.

Which two services should they choose? (Choose two.)

- A. Use Google App Engine services.
- B. Use serverless Google Cloud Functions.
- C. Use Knative to build and deploy serverless applications.
- D. Use Google Kubernetes Engine for automated deployments.
- E. Use a large Google Compute Engine cluster for deployments.

Answer: B,C

Question No : 16 - (Topic 1)

For this question, refer to the HipLocal case study.

How should HipLocal increase their API development speed while continuing to provide the QA team with a stable testing environment that meets feature requirements?

- A. Include unit tests in their code, and prevent deployments to QA until all tests have a passing status.
- B. Include performance tests in their code, and prevent deployments to QA until all tests have a passing status.
- C. Create health checks for the QA environment, and redeploy the APIs at a later time if the environment is unhealthy.
- D. Redeploy the APIs to App Engine using Traffic Splitting. Do not move QA traffic to the new versions if errors are found.

Answer: B

Question No : 17 - (Topic 1)

For this question, refer to the HipLocal case study.

HipLocal's application uses Cloud Client Libraries to interact with Google Cloud. HipLocal needs to configure authentication and authorization in the Cloud Client Libraries to implement least privileged access for the application. What should they do?

- A.** Create an API key. Use the API key to interact with Google Cloud.
- B.** Use the default compute service account to interact with Google Cloud.
- C.** Create a service account for the application. Export and deploy the private key for the application. Use the service account to interact with Google Cloud.
- D.** Create a service account for the application and for each Google Cloud API used by the application. Export and deploy the private keys used by the application. Use the service account with one Google Cloud API to interact with Google Cloud.

Answer: A

Question No : 18 - (Topic 1)

HipLocal wants to improve the resilience of their MySQL deployment, while also meeting their business and technical requirements.

Which configuration should they choose?

- A.** Use the current single instance MySQL on Compute Engine and several read-only MySQL servers on Compute Engine.
- B.** Use the current single instance MySQL on Compute Engine, and replicate the data to Cloud SQL in an external master configuration.
- C.** Replace the current single instance MySQL instance with Cloud SQL, and configure high availability.
- D.** Replace the current single instance MySQL instance with Cloud SQL, and Google provides redundancy without further configuration.

Answer: B

Question No : 19 - (Topic 1)

For this question, refer to the HipLocal case study.

A recent security audit discovers that HipLocal's database credentials for their Compute Engine-hosted MySQL databases are stored in plain text on persistent disks. HipLocal needs to reduce the risk of these credentials being stolen. What should they do?

- A.** Create a service account and download its key. Use the key to authenticate to Cloud Key Management Service (KMS) to obtain the database credentials.
- B.** Create a service account and download its key. Use the key to authenticate to Cloud Key Management Service (KMS) to obtain a key used to decrypt the database credentials.
- C.** Create a service account and grant it the roles/iam.serviceAccountUser role. Impersonate as this account and authenticate using the Cloud SQL Proxy.
- D.** Grant the roles/secretmanager.secretAccessor role to the Compute Engine service account. Store and access the database credentials with the Secret Manager API.

Answer: D

Explanation: <https://cloud.google.com/secret-manager/docs/overview>

Topic 2, Misc Questions

Question No : 20 - (Topic 2)

You have written a Cloud Function that accesses other Google Cloud resources. You want to secure the environment using the principle of least privilege. What should you do?

- A.** Create a new service account that has Editor authority to access the resources. The deployer is given permission to get the access token.
- B.** Create a new service account that has a custom IAM role to access the resources. The deployer is given permission to get the access token.
- C.** Create a new service account that has Editor authority to access the resources. The deployer is given permission to act as the new service account.
- D.** Create a new service account that has a custom IAM role to access the resources. The deployer is given permission to act as the new service account.

Answer: D

Reference: <https://cloud.google.com/blog/products/application-development/least-privilege-for-cloud-functions-using-cloud-iam>

Question No : 21 - (Topic 2)

You are deploying a microservices application to Google Kubernetes Engine (GKE). The application will receive daily updates. You expect to deploy a large number of distinct containers that will run on the Linux operating system (OS). You want to be alerted to any known OS vulnerabilities in the new containers. You want to follow Google-recommended best practices. What should you do?

- A.** Use the gcloud CLI to call Container Analysis to scan new container images. Review the vulnerability results before each deployment.
- B.** Enable Container Analysis, and upload new container images to Artifact Registry. Review the vulnerability results before each deployment.
- C.** Enable Container Analysis, and upload new container images to Artifact Registry. Review the critical vulnerability results before each deployment.
- D.** Use the Container Analysis REST API to call Container Analysis to scan new container images. Review the vulnerability results before each deployment.

Answer: B

Explanation:

<https://cloud.google.com/container-analysis/docs/automated-scanning-howto>

<https://cloud.google.com/container-analysis/docs/os-overview> says: The Container Scanning API allows you to automate OS vulnerability detection, scanning each time you push an image to Container Registry or Artifact Registry. Enabling this API also triggers language package scans for Go and Java vulnerabilities (Preview).

Question No : 22 - (Topic 2)

Your company has a data warehouse that keeps your application information in BigQuery.

The BigQuery data warehouse keeps 2 PBs of user data. Recently, your company expanded your user base to include EU users and needs to comply with these requirements:

Your company must be able to delete all user account information upon user request.

All EU user data must be stored in a single region specifically for EU users.

Which two actions should you take? (Choose two.)

- A. Use BigQuery federated queries to query data from Cloud Storage.
- B. Create a dataset in the EU region that will keep information about EU users only.
- C. Create a Cloud Storage bucket in the EU region to store information for EU users only.
- D. Re-upload your data using to a Cloud Dataflow pipeline by filtering your user records out.
- E. Use DML statements in BigQuery to update/delete user records based on their requests.

Answer: C,E

Reference: <https://cloud.google.com/solutions/bigquery-data-warehouse>

Question No : 23 - (Topic 2)

You are running a web application on Google Kubernetes Engine that you inherited. You want to determine whether the application is using libraries with known vulnerabilities or is vulnerable to XSS attacks. Which service should you use?

- A. Google Cloud Armor
- B. Debugger
- C. Web Security Scanner
- D. Error Reporting

Answer: C

Explanation:

<https://cloud.google.com/security-command-center/docs/concepts-web-security-scanner-overview>

Web Security Scanner identifies security vulnerabilities in your App Engine, Google Kubernetes Engine (GKE), and Compute Engine web applications. It crawls your application, following all links within the scope of your starting URLs, and attempts to

exercise as many user inputs and event handlers as possible.

Question No : 24 - (Topic 2)

You are developing a web application that contains private images and videos stored in a Cloud Storage bucket. Your users are anonymous and do not have Google Accounts. You want to use your application-specific logic to control access to the images and videos. How should you configure access?

- A.** Cache each web application user's IP address to create a named IP table using Google Cloud Armor. Create a Google Cloud Armor security policy that allows users to access the backend bucket.
- B.** Grant the Storage Object Viewer IAM role to allUsers. Allow users to access the bucket after authenticating through your web application.
- C.** Configure Identity-Aware Proxy (IAP) to authenticate users into the web application. Allow users to access the bucket after authenticating through IAP.
- D.** Generate a signed URL that grants read access to the bucket. Allow users to access the URL after authenticating through your web application.

Answer: D

Explanation:

<https://cloud.google.com/storage/docs/access-control/signed-urls#should-you-use>

In some scenarios, you might not want to require your users to have a Google account in order to access Cloud Storage, but you still want to control access using your application-specific logic. The typical way to address this use case is to provide a signed URL to a user, which gives the user read, write, or delete access to that resource for a limited time. You specify an expiration time when you create the signed URL. Anyone who knows the URL can access the resource until the expiration time for the URL is reached or the key used to sign the URL is rotated.

Question No : 25 - (Topic 2)

You want to create “fully baked” or “golden” Compute Engine images for your application. You need to bootstrap your application to connect to the appropriate database according to the environment the application is running on (test, staging, production). What should you do?

- A.** Embed the appropriate database connection string in the image. Create a different image for each environment.
- B.** When creating the Compute Engine instance, add a tag with the name of the database to be connected. In your application, query the Compute Engine API to pull the tags for the current instance, and use the tag to construct the appropriate database connection string.
- C.** When creating the Compute Engine instance, create a metadata item with a key of “DATABASE” and a value for the appropriate database connection string. In your application, read the “DATABASE” environment variable, and use the value to connect to the appropriate database.
- D.** When creating the Compute Engine instance, create a metadata item with a key of “DATABASE” and a value for the appropriate database connection string. In your application, query the metadata server for the “DATABASE” value, and use the value to connect to the appropriate database.

Answer: C

Question No : 26 - (Topic 2)

You have an application deployed in production. When a new version is deployed, some issues don't arise until the application receives traffic from users in production. You want to reduce both the impact and the number of users affected.

Which deployment strategy should you use?

- A.** Blue/green deployment
- B.** Canary deployment
- C.** Rolling deployment
- D.** Recreate deployment

Answer: A

Reference: <https://thenewstack.io/deployment-strategies/>

Question No : 27 - (Topic 2)

Your company has created an application that uploads a report to a Cloud Storage bucket. When the report is uploaded to the bucket, you want to publish a message to a Cloud Pub/Sub topic. You want to implement a solution that will take a small amount of effort to implement. What should you do?

- A.** Configure the Cloud Storage bucket to trigger Cloud Pub/Sub notifications when objects are modified.
- B.** Create an App Engine application to receive the file; when it is received, publish a message to the Cloud Pub/Sub topic.
- C.** Create a Cloud Function that is triggered by the Cloud Storage bucket. In the Cloud Function, publish a message to the Cloud Pub/Sub topic.
- D.** Create an application deployed in a Google Kubernetes Engine cluster to receive the file; when it is received, publish a message to the Cloud Pub/Sub topic.

Answer: C

Explanation: <https://cloud.google.com/storage/docs/pubsub-notifications>

Question No : 28 - (Topic 2)

You are developing an application using different microservices that should remain internal to the cluster. You want to be able to configure each microservice with a specific number of replicas. You also want to be able to address a specific microservice from any other microservice in a uniform way, regardless of the number of replicas the microservice scales to. You need to implement this solution on Google Kubernetes Engine. What should you do?

- A.** Deploy each microservice as a Deployment. Expose the Deployment in the cluster using a Service, and use the Service DNS name to address it from other microservices within the cluster.
- B.** Deploy each microservice as a Deployment. Expose the Deployment in the cluster using an Ingress, and use the Ingress IP address to address the Deployment from other microservices within the cluster.
- C.** Deploy each microservice as a Pod. Expose the Pod in the cluster using a Service, and use the Service DNS name to address the microservice from other microservices within the cluster.
- D.** Deploy each microservice as a Pod. Expose the Pod in the cluster using an Ingress, and use the Ingress IP address name to address the Pod from other microservices within the cluster.

Answer: A

Question No : 29 - (Topic 2)

You recently joined a new team that has a Cloud Spanner database instance running in production. Your manager has asked you to optimize the Spanner instance to reduce cost while maintaining high reliability and availability of the database. What should you do?

- A.** Use Cloud Logging to check for error logs, and reduce Spanner processing units by small increments until you find the minimum capacity required.
- B.** Use Cloud Trace to monitor the requests per sec of incoming requests to Spanner, and reduce Spanner processing units by small increments until you find the minimum capacity required.
- C.** Use Cloud Monitoring to monitor the CPU utilization, and reduce Spanner processing units by small increments until you find the minimum capacity required.
- D.** Use Snapshot Debugger to check for application errors, and reduce Spanner processing units by small increments until you find the minimum capacity required.

Answer: C

Explanation:

https://cloud.google.com/spanner/docs/compute-capacity#increasing_and_decreasing_compute_capacity

Question No : 30 - (Topic 2)

Your company has a BigQuery dataset named "Master" that keeps information about employee travel and

expenses. This information is organized by employee department. That means employees should only be able

to view information for their department. You want to apply a security framework to enforce this requirement

with the minimum number of steps.

What should you do?

- A.** Create a separate dataset for each department. Create a view with an appropriate WHERE clause to select records from a particular dataset for the specific department. Authorize this view to access records from your Master dataset. Give employees the permission to this department-specific dataset.
- B.** Create a separate dataset for each department. Create a data pipeline for each department to copy appropriate information from the Master dataset to the specific dataset for the department. Give employees the permission to this department-specific dataset.
- C.** Create a dataset named Master dataset. Create a separate view for each department in the Master dataset. Give employees access to the specific view for their department.
- D.** Create a dataset named Master dataset. Create a separate table for each department in the Master dataset. Give employees access to the specific table for their department.

Answer: B

Question No : 31 - (Topic 2)

Your team manages a Google Kubernetes Engine (GKE) cluster where an application is running. A different team is planning to integrate with this application. Before they start the integration, you need to ensure that the other team cannot make changes to your application, but they can deploy the integration on GKE. What should you do?

- A.** Using Identity and Access Management (IAM), grant the Viewer IAM role on the cluster project to the other team.
- B.** Create a new GKE cluster. Using Identity and Access Management (IAM), grant the Editor role on the cluster project to the other team.
- C.** Create a new namespace in the existing cluster. Using Identity and Access Management (IAM), grant the Editor role on the cluster project to the other team.
- D.** Create a new namespace in the existing cluster. Using Kubernetes role-based access control (RBAC), grant the Admin role on the new namespace to the other team.

Answer: D

Question No : 32 - (Topic 2)

You are planning to migrate a MySQL database to the managed Cloud SQL database for Google Cloud. You have Compute Engine virtual machine instances that will connect with this Cloud SQL instance. You do not want to whitelist IPs for the Compute Engine instances to be able to access Cloud SQL.

What should you do?

- A. Enable private IP for the Cloud SQL instance.
- B. Whitelist a project to access Cloud SQL, and add Compute Engine instances in the whitelisted project.
- C. Create a role in Cloud SQL that allows access to the database from external instances, and assign the Compute Engine instances to that role.
- D. Create a CloudSQL instance on one project. Create Compute engine instances in a different project. Create a VPN between these two projects to allow internal access to CloudSQL.

Answer: C

Reference: <https://cloud.google.com/sql/docs/mysql/connect-external-app>

Question No : 33 - (Topic 2)

You recently developed a new application. You want to deploy the application on Cloud Run without a Dockerfile. Your organization requires that all container images are pushed to a centrally managed container repository. How should you build your container using Google Cloud services? (Choose two.)

- A. Push your source code to Artifact Registry.
- B. Submit a Cloud Build job to push the image.
- C. Use the pack build command with pack CLI.
- D. Include the --source flag with the gcloud run deploy CLI command.
- E. Include the --platform=kubernetes flag with the gcloud run deploy CLI command.

Answer: A,C

Explanation: <https://cloud.google.com/run/docs/deploying#images>

<https://cloud.google.com/blog/products/containers-kubernetes/google-cloud-now-supports-buildpacks>