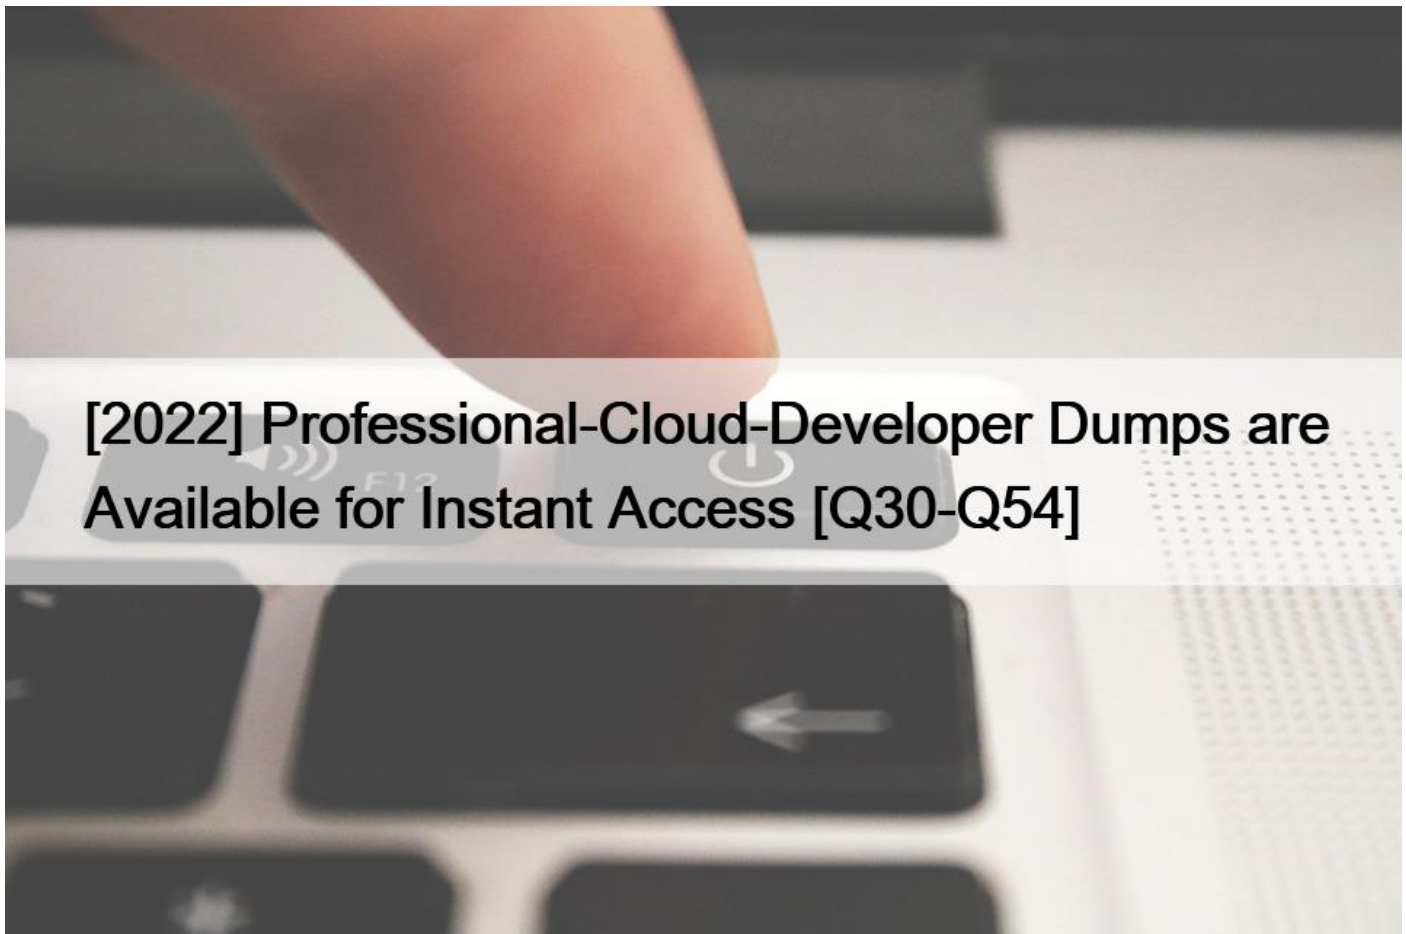


[2022 Professional-Cloud-Developer Dumps are Available for Instant Access [Q30-Q54]



[2022 Professional-Cloud-Developer Dumps are Available for Instant Access Valid Professional-Cloud-Developer Dumps for Helping Passing Professional-Cloud-Developer Exam!]

Google Professional-Cloud-Developer Exam Syllabus Topics:

TopicDetailsTopic 1- Implementing Appropriate Deployment Strategies Based On The Target Compute Environment- Creating A Load Balancer For Compute Engine InstancesTopic 2- Integrating An Application With Data And Storage Services- Writing An SQL Query To Retrieve Data From Relational DatabasesTopic 3- Reading And Updating An Entity In A Cloud Datastore Transaction From An Application- Using Apis To Read- Write To Data ServicesTopic 4- Publishing And Consuming From Data Ingestion Sources- Authenticating Users By Using Oauth2 Web Flow And Identity Aware ProxyTopic 5- Launching A Compute Instance Using GCP Console And Cloud SDK- Creating An Autoscaled Managed Instance Group Using An Instance TemplateTopic 6- Setting Up Your Development Environment, Considerations- Building And Testing ApplicationsTopic 7- Reviewing Test Results Of Continuous Integration Pipeline- Developing Unit Tests For All Code WrittenTopic 8- Configuring Compute Services Network Settings- Configuring A Cloud Pub- Sub Push Subscription To Call An EndpointTopic 9- Operating System Versions And Base Runtimes Of Services- Ooogle-Recommended Practices And DocumentationTopic 10 - Security Mechanisms That Protect Services And Resources- Choosing Data Storage Options Based On Use Case ConsiderationsTopic 11- Defining A Key Structure For High Write Applications Using Cloud Storage- Using Cloud Storage To Run A Static WebsiteTopic 12- Deploying Applications And Services On Google Kubernetes Engine- Deploying

Applications And Services On Google Kubernetes Engine
Topic 13- Defining Database Schemas For Google-Managed Databases- Re-Architecting Applications From Local Services To Google Cloud Platform
Topic 14- Developing An Integration Pipeline Using Services- Emulating GCP Services For Local Application Development
Topic 15- Google-Recommended Practices And Documentation- Deploying And Securing An API With Cloud Endpoints
Topic 16- Deploying Applications And Services On Compute Engine- Deploying An Application To App Engine
Topic 17- Designing Highly Scalable, Available, And Reliable Cloud-Native Applications- Geographic Distribution Of Google Cloud Services

Google Professional Cloud Developer Certification Path

The Google Professional Cloud Developer Certification is the highest level of certification mainly focussing to the Google Professional Cloud Developer.

There is no prerequisite for this exam but still it would be best to follow some sequence in order to prove immense knowledge as a Google Professional Cloud Developer.

You can complete Google Associate Certifications then approach for the professional certification. For more information related to Google cloud certification track [Google-certification-path](#)

The content of the qualifying test for the Google Professional Cloud Developer certification comprises of 5 topics covering specific knowledge and skills. The candidates need to thoroughly study a detailed exam guide available on the official website before attending the test. The highlights of the topics that constitute the structure of the exam are enumerated below: **Section 1: Designing Highly Scalable, Available, and Reliable**

Cloud-Native Apps

Within this subject area, the examinees need to demonstrate their proficiency in designing high-performing applications & APIs; designing secure applications; managing application data; executing application modernization.

NO.30 Which of the following customer statements would alert you to a safety issue? (Choose two.)

- * My iPhone flashed and sparked when I tried to charge it.
- * The corner of my iPad is badly bent.
- * My iPhone has fluctuating sound levels. Sometimes it is deafening.
- * The screen is too bright. It hurts my eyes.
- * My new Apple Watch makes me itchy and my wrist is red and irritated.
- * The home button on my iPhone seems to have sunk.

NO.31 Please refer to the following information to answer the questions on the right.

Rachel is starting a repair on a three-year-old MacBook Pro.

After opening the device, she takes some time to visually inspect the top case assembly with battery.

During an embedded battery inspection which of the following issues should Rachel look for? (Choose two.)

- * Updated battery firmware
- * Dot imprints
- * Battery is the correct color
- * Scratches
- * Battery-compliance shipping label

NO.32 You are using Cloud Build to create a new Docker image on each source code commit to a Cloud Source Repositories repository. Your application is built on every commit to the master branch. You want to release specific commits made to the master branch in an automated method. What should you do?

- * Manually trigger the build for new releases.

- * Create a build trigger on a Git tag pattern. Use a Git tag convention for new releases.
- * Create a build trigger on a Git branch name pattern. Use a Git branch naming convention for new releases.
- * Commit your source code to a second Cloud Source Repositories repository with a second Cloud Build trigger. Use this repository for new releases only.

Reference:

<https://docs.docker.com/docker-hub/builds/>

NO.33 Your application performs well when tested locally, but it runs significantly slower when you deploy it to App Engine standard environment. You want to diagnose the problem. What should you do?

- * File a ticket with Cloud Support indicating that the application performs faster locally.
- * Use Stackdriver Debugger Snapshots to look at a point-in-time execution of the application.
- * Use Stackdriver Trace to determine which functions within the application have higher latency.
- * Add logging commands to the application and use Stackdriver Logging to check where the latency problem occurs.

Topic 1, HipLocal Case Study

Company 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:

- * Existing APIs run on Compute Engine virtual machine instances hosted in GCP
- * State is stored in a single instance MySQL database in GCP
- * Data is exported to an on-premises Teradata/Vertica data warehouse
- * Data analytics is performed in an on-premises Hadoop environment
- * The application has no logging

* There are basic indicators of uptime; alerts are frequently fired when the APIs are unresponsive
Business Requirements
HipLocal's investors want to expand their footprint and support the increase in demand they are seeing. Their requirements are:

- * Expand availability of the application to new regions
- * Increase the number of concurrent users that can be supported
- * Ensure a consistent experience for users when they travel to different regions
- * Obtain user activity metrics to better understand how to monetize their product
- * Ensure compliance with regulations in the new regions (for example, GDPR)
- * Reduce infrastructure management time and cost
- * Adopt the Google-recommended practices for cloud computing

Technical Requirements

- * The application and backend must provide usage metrics and monitoring
- * APIs require strong authentication and authorization
- * Logging must be increased, and data should be stored in a cloud analytics platform
- * Move to serverless architecture to facilitate elastic scaling
- * Provide authorized access to internal apps in a secure manner

NO.34 Which service should HipLocal use to enable access to internal apps?

- * Cloud VPN
- * Cloud Armor
- * Virtual Private Cloud
- * Cloud Identity-Aware Proxy

Reference:

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

NO.35 Your company's development teams want to use Cloud Build in their projects to build and push Docker images to Container Registry. The operations team requires all Docker images to be published to a centralized, securely managed Docker registry that the operations team manages.

What should you do?

- * Use Container Registry to create a registry in each development team's project. Configure the Cloud Build build to push the Docker image to the project's registry. Grant the operations team access to each development team's registry.
- * Create a separate project for the operations team that has Container Registry configured. Assign appropriate permissions to the Cloud Build service account in each developer team's project to allow access to the operation team's registry.
- * Create a separate project for the operations team that has Container Registry configured. Create a Service Account for each

development team and assign the appropriate permissions to allow it access to the operations team's registry. Store the service account key file in the source code repository and use it to authenticate against the operations team's registry.

* Create a separate project for the operations team that has the open source Docker Registry deployed on a Compute Engine virtual machine instance. Create a username and password for each development team. Store the username and password in the source code repository and use it to authenticate against the operations team's Docker registry.

NO.36 You are a SaaS provider deploying dedicated blogging software to customers in your Google Kubernetes Engine (GKE) cluster. You want to configure a secure multi-tenant platform to ensure that each customer has access to only their own blog and can't affect the workloads of other customers. What should you do?

- * Enable Application-layer Secrets on the GKE cluster to protect the cluster.
- * Deploy a namespace per tenant and use Network Policies in each blog deployment.
- * Use GKE Audit Logging to identify malicious containers and delete them on discovery.
- * Build a custom image of the blogging software and use Binary Authorization to prevent untrusted image deployments.

NO.37 Case study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Company 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 10000 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:

- * Existing APIs run on Compute Engine virtual machine instances hosted in GCP.
- * State is stored in a single instance MySQL database in GCP.
- * Data is exported to an on-premises Teradata/Vertica data warehouse.
- * Data analytics is performed in an on-premises Hadoop environment.
- * The application has no logging.
- * There are basic indicators of uptime; alerts are frequently fired when the APIs are unresponsive.

Business Requirements

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

- * Expand availability of the application to new regions.
- * Increase the number of concurrent users that can be supported.
- * Ensure a consistent experience for users when they travel to different regions.
- * Obtain user activity metrics to better understand how to monetize their product.
- * Ensure compliance with regulations in the new regions (for example, GDPR).
- * Reduce infrastructure management time and cost.
- * Adopt the Google-recommended practices for cloud computing.

Technical Requirements

- * The application and backend must provide usage metrics and monitoring.
- * APIs require strong authentication and authorization.
- * Logging must be increased, and data should be stored in a cloud analytics platform.

- * Move to serverless architecture to facilitate elastic scaling.
- * Provide authorized access to internal apps in a secure manner.

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

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

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

NO.38 Your company is planning to migrate their on-premises Hadoop environment to the cloud. Increasing storage cost and maintenance of data stored in HDFS is a major concern for your company. You also want to make minimal changes to existing data analytics jobs and existing architecture. How should you proceed with the migration?

- * Migrate your data stored in Hadoop to BigQuery. Change your jobs to source their information from BigQuery instead of the on-premises Hadoop environment.
- * Create Compute Engine instances with HDD instead of SSD to save costs. Then perform a full migration of your existing environment into the new one in Compute Engine instances.
- * Create a Cloud Dataproc cluster on Google Cloud Platform, and then migrate your Hadoop environment to the new Cloud Dataproc cluster. Move your HDFS data into larger HDD disks to save on storage costs.
- * Create a Cloud Dataproc cluster on Google Cloud Platform, and then migrate your Hadoop code objects to the new cluster. Move your data to Cloud Storage and leverage the Cloud Dataproc connector to run jobs on that data.

NO.39 Your service adds text to images that it reads from Cloud Storage. During busy times of the year, requests to Cloud Storage fail with an HTTP 429 Too Many Requests status code.

How should you handle this error?

- * Add a cache-control header to the objects.
- * Request a quota increase from the GCP Console.
- * Retry the request with a truncated exponential backoff strategy.
- * Change the storage class of the Cloud Storage bucket to Multi-regional.

Explanation/Reference: <https://developers.google.com/gmail/api/v1/reference/quota>

NO.40 You migrated your applications to Google Cloud Platform and kept your existing monitoring platform. You now find that your notification system is too slow for time critical problems.

What should you do?

- * Replace your entire monitoring platform with Stackdriver.
- * Install the Stackdriver agents on your Compute Engine instances.
- * Use Stackdriver to capture and alert on logs, then ship them to your existing platform.
- * Migrate some traffic back to your old platform and perform AB testing on the two platforms concurrently.

Reference:

<https://cloud.google.com/monitoring/>

NO.41 Leigh states that her MacBook Pro (Retina, 15-inch, Mid 2015) does not recognize the SD card she brought with her. You have her reproduce the issue and discover that she is properly inserting the card. However, it is not recognized by Photos or Image

Capture. What question should you ask next to isolate the issue to hardware?

- * What kind of files or images are on the SD card?
- * Are you running the latest version of Photos?
- * Have you tried resetting your SMC and NVRAM?
- * Have you had this issue with all SD cards or just this one?

NO.42 You want to view the memory usage of your application deployed on Compute Engine. What should you do?

- * Install the Stackdriver Client Library.
- * Install the Stackdriver Monitoring Agent.
- * Use the Stackdriver Metrics Explorer.
- * Use the Google Cloud Platform Console.

Reference:

<https://stackoverflow.com/questions/43991246/google-cloud-platform-how-to-monitor-memory-usage-of-vm-instances>

NO.43 You recently developed an application. You need to call the Cloud Storage API from a Compute Engine instance that doesn't have a public IP address. What should you do?

- * Use Carrier Peering
- * Use VPC Network Peering
- * Use Shared VPC networks
- * Use Private Google Access

<https://cloud.google.com/vpc/docs/private-google-access>

NO.44 Your application takes an input from a user and publishes it to the user's contacts. This input is stored in a table in Cloud Spanner. Your application is more sensitive to latency and less sensitive to consistency.

How should you perform reads from Cloud Spanner for this application?

- * Perform Read-Only transactions.
- * Perform stale reads using single-read methods.
- * Perform strong reads using single-read methods.
- * Perform stale reads using read-write transactions.

NO.45 Case Study

Company 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 10000 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:

- * Existing APIs run on Compute Engine virtual machine instances hosted in GCP.
- * State is stored in a single instance MySQL database in GCP.
- * Data is exported to an on-premises Teradata/Vertica data warehouse.
- * Data analytics is performed in an on-premises Hadoop environment.
- * The application has no logging.
- * There are basic indicators of uptime; alerts are frequently fired when the APIs are unresponsive.

Business Requirements

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

- * Expand availability of the application to new regions.
- * Increase the number of concurrent users that can be supported.
- * Ensure a consistent experience for users when they travel to different regions.
- * Obtain user activity metrics to better understand how to monetize their product.
- * Ensure compliance with regulations in the new regions (for example, GDPR).
- * Reduce infrastructure management time and cost.
- * Adopt the Google-recommended practices for cloud computing.

Technical Requirements

- * The application and backend must provide usage metrics and monitoring.
- * APIs require strong authentication and authorization.
- * Logging must be increased, and data should be stored in a cloud analytics platform.

- * Move to serverless architecture to facilitate elastic scaling.
- * Provide authorized access to internal apps in a secure manner.

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

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

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

Explanation/Reference:

NO.46 Your team is developing a new application using a PostgreSQL database and Cloud Run. You are responsible for ensuring that all traffic is kept private on Google Cloud. You want to use managed services and follow Google-recommended best practices. What should you do?

- * 1. Enable Cloud SQL and Cloud Run in the same project.
- 2. Configure a private IP address for Cloud SQL. Enable private services access.
- 3. Create a Serverless VPC Access connector.
- 4. Configure Cloud Run to use the connector to connect to Cloud SQL.
- * 1. Install PostgreSQL on a Compute Engine virtual machine (VM), and enable Cloud Run in the same project.
- 2. Configure a private IP address for the VM. Enable private services access.
- 3. Create a Serverless VPC Access connector.
- 4. Configure Cloud Run to use the connector to connect to the VM hosting PostgreSQL.
- * 1. Use Cloud SQL and Cloud Run in different projects.
- 2. Configure a private IP address for Cloud SQL. Enable private services access.
- 3. Create a Serverless VPC Access connector.
- 4. Set up a VPN connection between the two projects. Configure Cloud Run to use the connector to connect to Cloud SQL.
- * 1. Install PostgreSQL on a Compute Engine VM, and enable Cloud Run in different projects.
- 2. Configure a private IP address for the VM. Enable private services access.
- 3. Create a Serverless VPC Access connector.
- 4. Set up a VPN connection between the two projects. Configure Cloud Run to use the connector to access the VM hosting PostgreSQL

<https://cloud.google.com/sql/docs/postgres/connect-run#private-ip>

NO.47 Your code is running on Cloud Functions in projectA . It is supposed to write an object in a Cloud Storage bucket owned by project B.

However, the write call is failing with the error `403 Forbidden`;

What should you do to correct the problem?

- * Grant your user account the roles/storage.objectCreator role for the Cloud Storage bucket.
- * Grant your user account the roles/iam.serviceAccountUser role for the service-PROJECTA@gcf-adminrobot.

iam.gserviceaccount.com service account.

- * Grant the service-PROJECTA@gcf-admin-robot.iam.gserviceaccount.com service account the roles/ storage.objectCreator role for the Cloud Storage bucket.
- * Enable the Cloud Storage API in project B.

NO.48 You are using Cloud Build build to promote a Docker image to Development, Test, and Production environments. You need to ensure that the same Docker image is deployed to each of these environments.

How should you identify the Docker image in your build?

- * Use the latest Docker image tag.
- * Use a unique Docker image name.
- * Use the digest of the Docker image.
- * Use a semantic version Docker image tag.

NO.49 Case Study

Company 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 10000 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:

- * Existing APIs run on Compute Engine virtual machine instances hosted in GCP.
- * State is stored in a single instance MySQL database in GCP.
- * Data is exported to an on-premises Teradata/Vertica data warehouse.
- * Data analytics is performed in an on-premises Hadoop environment.
- * The application has no logging.
- * There are basic indicators of uptime; alerts are frequently fired when the APIs are unresponsive.

Business Requirements

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

- * Expand availability of the application to new regions.
- * Increase the number of concurrent users that can be supported.
- * Ensure a consistent experience for users when they travel to different regions.
- * Obtain user activity metrics to better understand how to monetize their product.
- * Ensure compliance with regulations in the new regions (for example, GDPR).
- * Reduce infrastructure management time and cost.
- * Adopt the Google-recommended practices for cloud computing.

Technical Requirements

- * The application and backend must provide usage metrics and monitoring.
- * APIs require strong authentication and authorization.
- * Logging must be increased, and data should be stored in a cloud analytics platform.
- * Move to serverless architecture to facilitate elastic scaling.
- * Provide authorized access to internal apps in a secure manner.

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

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

Explanation/Reference:

NO.50 HipLocal's data science team wants to analyze user reviews.

How should they prepare the data?

- * Use the Cloud Data Loss Prevention API for redaction of the review dataset.
- * Use the Cloud Data Loss Prevention API for de-identification of the review dataset.
- * Use the Cloud Natural Language Processing API for redaction of the review dataset.
- * Use the Cloud Natural Language Processing API for de-identification of the review dataset.

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

NO.51 You are planning to deploy your application in a Google Kubernetes Engine (GKE) cluster. The application exposes an HTTP-based health check at /healthz. You want to use this health check endpoint to determine whether traffic should be routed to the pod by the load balancer.

Which code snippet should you include in your Podconfiguration?

*

```
livenessProbe:  
  httpGet:  
    path: /healthz  
    port: 80
```

*

```
readinessProbe:  
  httpGet:  
    path: /healthz  
    port: 80
```

*

```
loadbalancerHealthCheck:  
  httpGet:  
    path: /healthz  
    port: 80
```

*

```
healthCheck:  
  httpGet:  
    path: /healthz  
    port: 80
```

For the GKE ingress controller to use your readinessProbes as health checks, the Pods for an Ingress must exist at the time of Ingress creation. If your replicas are scaled to 0, the default health check will apply.

NO.52 You are planning to deploy your application in a Google Kubernetes Engine (GKE) cluster. Your application can scale horizontally, and each instance of your application needs to have a stable network identity and its own persistent disk.

Which GKE object should you use?

- * Deployment
- * StatefulSet
- * ReplicaSet
- * ReplicaController

NO.53 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?

- * Enable private IP for the Cloud SQL instance.
- * Whitelist a project to access Cloud SQL, and add Compute Engine instances in the whitelisted project.
- * Create a role in Cloud SQL that allows access to the database from external instances, and assign the Compute Engine instances to that role.
- * 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.

NO.54 HipLocal's net-based auth service fails under intermittent load.

What should they do?

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

Updated Professional-Cloud-Developer Dumps Questions For Google Exam:

<https://www.dumpsmaterials.com/Professional-Cloud-Developer-real-torrent.html>