


Professional-Cloud-Developer Practice Test Questions Answers Updated 140 Questions [Q52-Q68]



Professional-Cloud-Developer Practice Test Questions Answers Updated 140 Questions
Professional-Cloud-Developer dumps & Cloud Developer Sure Practice with 140 Questions

Q52. Your website is deployed on Compute Engine. Your marketing team wants to test conversion rates between 3 different website designs.

Which approach should you use?

- * Deploy the website on App Engine and use traffic splitting.
- * Deploy the website on App Engine as three separate services.
- * Deploy the website on Cloud Functions and use traffic splitting.
- * Deploy the website on Cloud Functions as three separate functions.

Q53. You recently deployed your application in Google Kubernetes Engine, and now need to release a new version of your application. You need the ability to instantly roll back to the previous version in case there are issues with the new version. Which deployment model should you use?

- * Perform a rolling deployment, and test your new application after the deployment is complete.
- * Perform A/B testing, and test your application periodically after the new tests are implemented.
- * Perform a blue/green deployment, and test your new application after the deployment is complete.
- * Perform a canary deployment, and test your new application periodically after the new version is deployed.

Q54. Your data is stored in Cloud Storage buckets. Fellow developers have reported that data downloaded from Cloud Storage is resulting in slow API performance. You want to research the issue to provide details to the GCP support team.

Which command should you run?

- * `gsutil test -o output.json gs://my-bucket`
- * `gsutil perfdiag -o output.json gs://my-bucket`
- * `gcloud compute scp example-instance:~/test-data -o output.json gs://my-bucket`
- * `gcloud services test -o output.json gs://my-bucket`

Explanation/Reference: <https://groups.google.com/forum/#!topic/gce-discussion/xB19Jq5HDsY>

Q55. Your App Engine standard configuration is as follows:

```
service: production
```

```
instance_class: B1
```

You want to limit the application to 5 instances. Which code snippet should you include in your configuration?

- * `manual_scaling:instances: 5min_pending_latency: 30ms`
- * `manual_scaling:max_instances: 5idle_timeout: 10m`
- * `basic_scaling:instances: 5min_pending_latency: 30ms`
- * `basic_scaling:max_instances: 5idle_timeout: 10m`

Q56. You are developing a JPEG image-resizing API hosted on Google Kubernetes Engine (GKE). Callers of the service will exist within the same GKE cluster. You want clients to be able to get the IP address of the service.

What should you do?

- * Define a GKE Service. Clients should use the name of the A record in Cloud DNS to find the service's cluster IP address.
- * Define a GKE Service. Clients should use the service name in the URL to connect to the service.
- * Define a GKE Endpoint. Clients should get the endpoint name from the appropriate environment variable in the client container.
- * Define a GKE Endpoint. Clients should get the endpoint name from Cloud DNS.

Q57. 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.

Q58. 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?

- * 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.
- * 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.
- * 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.
- * 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.

Q59. You are running an application on App Engine that you inherited. You want to find out whether the application is using insecure binaries or is vulnerable to XSS attacks. Which service should you use?

- * Cloud Armor
- * Stackdriver Debugger
- * Cloud Security Scanner
- * Stackdriver Error Reporting

Q60. 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.

Reference:

<https://cloud.google.com/solutions/best-practices-cloud-spanner-gaming-database>

Q61. In the systematic troubleshooting approach, which of the following statements is true about isolating an issue?

- * Asking the customer to reproduce an issue can help determine if the issue is with the device.
- * If an issue cannot be reproduced, it is likely a hardware issue.
- * Replacing an internal component will determine if the issue is related to environment.
- * Basing a conclusion on past experience with similar issues is a proven troubleshooting method.

Q62. You have deployed an HTTP(s) Load Balancer with the gcloud commands shown below.

```
export NAME=load-balancer

# create network
gcloud compute networks create ${NAME}

# add instance
gcloud compute instances create ${NAME}-backend-instance-1 --subnet ${NAME} --no-address

# create the instance group
gcloud compute instance-groups unmanaged create ${NAME}-i
gcloud compute instance-groups unmanaged set-named-ports ${NAME}-i --named-ports http:80
gcloud compute instance-groups unmanaged add-instances ${NAME}-i --instances ${NAME}-instance-1

# configure health checks
gcloud compute health-checks create http ${NAME}-http-hc --port 80

# create backend service
gcloud compute backend-services create ${NAME}-http-bes --health-checks ${NAME}-http-hc --protocol HTTP
--global
gcloud compute backend-services add-backend ${NAME}-http-bes --instance-group ${NAME}-i --balancing-mode
100000 --capacity-scaler 1.0 --global --instance-group-zone us-east1-d

# create url maps and forwarding rule
gcloud compute url-maps create ${NAME}-http-urlmap --default-service ${NAME}-http-bes
gcloud compute target-http-proxies create ${NAME}-http-proxy --url-map ${NAME}-http-urlmap
gcloud compute forwarding-rules create ${NAME}-http-fw --global --ip-protocol ICP --target-http-proxy
--ports 80
```

Health checks to port 80 on the Compute Engine virtual machine instance are failing and no traffic is sent to your instances. You want to resolve the problem.

Which commands should you run?

- * gcloud compute instances add-access-config \${NAME}-backend-instance-1
- * gcloud compute instances add-tags \${NAME}-backend-instance-1 --tags http-server

- * gcloud compute firewall-rules create allow-lb --network load-balancer --allow tcp --source-ranges 130.211.0.0/22,35.191.0.0/16 --direction INGRESS
- * gcloud compute firewall-rules create allow-lb --network load-balancer --allow tcp --destination-ranges 130.211.0.0/22,35.191.0.0/16 --direction EGRESS

Q63. Your code is running on Cloud Functions in project A.

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.

Q64. You have an application in production. It is deployed on Compute Engine virtual machine instances controlled by a managed instance group. Traffic is routed to the instances via a HTTP(s) load balancer. Your users are unable to access your application. You want to implement a monitoring technique to alert you when the application is unavailable.

Which technique should you choose?

- * Smoke tests
- * Stackdriver uptime checks
- * Cloud Load Balancing --health checks
- * Managed instance group --health checks

Reference:

<https://medium.com/google-cloud/stackdriver-monitoring-automation-part-3-uptime-checks-476b8507f59c>

Q65. Your App Engine standard configuration is as follows:

```
service: production
```

```
instance_class: B1
```

You want to limit the application to 5 instances.

Which code snippet should you include in your configuration?

```
manual_scaling:
```

- * instances: 5

```
min_pending_latency: 30ms
```

manual_scaling:

* max_instances: 5

idle_timeout: 10m

basic_scaling:

* instances: 5

min_pending_latency: 30ms

basic_scaling:

* max_instances: 5

idle_timeout: 10m

Q66. Your application requires service accounts to be authenticated to GCP products via credentials stored on its host Compute Engine virtual machine instances. You want to distribute these credentials to the host instances as securely as possible.

What should you do?

- * Use HTTP signed URLs to securely provide access to the required resources.
- * Use the instance's service account Application Default Credentials to authenticate to the required resources.
- * Generate a P12 file from the GCP Console after the instance is deployed, and copy the credentials to the host instance before starting the application.
- * Commit the credential JSON file into your application's source repository, and have your CI/CD process package it with the software that is deployed to the instance.

Explanation/Reference: <https://cloud.google.com/compute/docs/api/how-tos/authorization>

Q67. 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'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?

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

Explanation/Reference:

Q68. Your application is logging to Stackdriver. You want to get the count of all requests on all /api/alpha/* endpoints.

What should you do?

- * Add a Stackdriver counter metric for path:/api/alpha/.
- * Add a Stackdriver counter metric for endpoint:/api/alpha/*.
- * Export the logs to Cloud Storage and count lines matching /api/alpha.
- * Export the logs to Cloud Pub/Sub and count lines matching /api/alpha.

Explanation/Reference:

Revision Books Let's now focus on the details of the revision books that will supplement the above training programs well: Cloud Native DevOps with Kubernetes: Building, Deploying, and Scaling Modern Applications in the Cloud (1st Edition)

This book is perfectly designed to cover everything there's to know about Kubernetes. It is a valuable study material that addresses vital details starting with what Kubernetes is, its origin, and what the future holds for this platform. You will build a high-level

knowledge of containers including how they work, how to manage them, and most importantly, the tips for designing cloud-native infrastructure and similar services. This guide suits both experienced IT professionals and beginner-level individuals who are only getting started with Kubernetes. Through this resource, candidates will gain important hands-on skills related to the writing and deployment of Kubernetes apps, configuration and operation of Kubernetes clusters, and cloud infrastructure automation using popular tools such as Helm. What's more, this book will also help you build fundamental knowledge of Kubernetes security, including Role Based Access Control (RBAC). When opting to study using this book, be sure to get the official resource that's written by John Arundel alongside Justin Domingus. **Building Secure & Reliable Systems: Best Practices for Designing, Implementing, and Maintaining Systems (1st Edition)**The authors of this book (Heather Adkins, Piotr Lewandoski, Betsy Beyer, Ana Oprea, and others) created this comprehensive self-study tool with an aim to help candidates learn how to accommodate security and reliability solutions in the lifecycle of systems and software. This material gives explicit knowledge of highly secure systems through coding, debugging, and testing practices, design techniques, and strategies required to address security incidents. Check it out from Amazon and get easy access to one of the most valuable study resources you will ever find for the Google Professional Cloud Developer exam preparation.

New Professional-Cloud-Developer Exam Questions| Real Professional-Cloud-Developer Dumps:
https://www.braindumpsit.com/Professional-Cloud-Developer_real-exam.html