

Google.Professional-Cloud-Developer.v2022-09-22.q112

Exam Code:	Professional-Cloud-Developer
Exam Name:	Google Certified Professional - Cloud Developer
Certification Provider:	Google
Free Question Number:	112
Version:	v2022-09-22
# of views:	1985
# of Questions views:	1110
https://www.freeqas.com/qa/Google/Professional-Cloud-Developer/Google.Professional-Cloud-Developer.v2022-09-22.q112.html	

NEW QUESTION: 1

Your application is running on Compute Engine and is showing sustained failures for a small number of requests. You have narrowed the cause down to a single Compute Engine instance, but the instance is unresponsive to SSH.

What should you do next?

- A. Take a snapshot of the disk and attach it to a new machine.
- B. Delete the machine and create a new one.
- C. Reboot the machine.
- D. Enable and check the serial port output.

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 2

How would you define a "positive no" to a co-worker?

- A. Lead with a positive then switch to a negative, such as, "Definitely not."
- B. It's leading with "no" and not engaging in any further conversation with the customer.
- C. It is a way to correct a customer by informing them of what you can do.
- D. You say no with a positive voice inflection.

Answer: C ([LEAVE A REPLY](#))

Explanation/Reference:

NEW QUESTION: 3

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?

- A. `gsutil test -o output.json gs://my-bucket`

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

Answer: B (LEAVE A REPLY)

Explanation/Reference: <https://groups.google.com/forum/#!topic/gce-discussion/xBI9Jq5HDsY>

NEW QUESTION: 4

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?

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

Answer: (SHOW ANSWER)

NEW QUESTION: 5

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 6

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.)

- 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 ([LEAVE A REPLY](#))

Explanation/Reference:

NEW QUESTION: 7

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 .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 (LEAVE A REPLY)

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

NEW QUESTION: 8

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?

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

Answer: (SHOW ANSWER)

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

NEW QUESTION: 9

Your existing application keeps user state information in a single MySQL database. This state information is very user-specific and depends heavily on how long a user has been using an application. The MySQL database is causing challenges to maintain and enhance the schema for various users.

Which storage option should you choose?

- A. Cloud SQL
- B. Cloud Storage
- C. Cloud Spanner
- D. Cloud Datastore/Firestore

Answer: A (LEAVE A REPLY)

Explanation/Reference: <https://cloud.google.com/solutions/migrating-mysql-to-cloudsql-concept>

NEW QUESTION: 10

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?

- A. Perform stale reads using single-read methods.
- B. Perform stale reads using read-write transactions.
- C. Perform Read-Only transactions.
- D. Perform strong reads using single-read methods.

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 11

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?

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

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 12

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: ([SHOW ANSWER](#))

NEW QUESTION: 13

You are building an API that will be used by Android and iOS apps. The API must:

- * Support HTTPs
- * Minimize bandwidth cost
- * Integrate easily with mobile apps

Which API architecture should you use?

- A. RESTful APIs
- B. MQTT for APIs
- C. gRPC-based APIs
- D. SOAP-based APIs

Answer: A ([LEAVE A REPLY](#))

Reference:

<https://www.devteam.space/blog/how-to-build-restful-api-for-your-mobile-app/>

NEW QUESTION: 14

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 ([LEAVE A REPLY](#))

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

NEW QUESTION: 15

You are working on a social media application. You plan to add a feature that allows users to upload images.

These images will be 2 MB - 1 GB in size. You want to minimize their infrastructure operations overhead for this feature.

What should you do?

- A.** Change the application to accept images directly and store them in the database that stores other user information.
- B.** Change the application to create signed URLs for Cloud Storage. Transfer these signed URLs to the client application to upload images to Cloud Storage.
- C.** Set up a web server on GCP to accept user images and create a file store to keep uploaded files. Change the application to retrieve images from the file store.
- D.** Create a separate bucket for each user in Cloud Storage. Assign a separate service account to allow write access on each bucket. Transfer service account credentials to the client application based on user information. The application uses this service account to upload images to Cloud Storage.

Answer: B ([LEAVE A REPLY](#))

Explanation/Reference: <https://cloud.google.com/blog/products/storage-data-transfer/uploading-images-directly-to-cloud-storage-by-using-signed-url>

NEW QUESTION: 16

You are load testing your server application. During the first 30 seconds, you observe that a previously inactive Cloud Storage bucket is now servicing 2000 write requests per second and 7500 read requests per second.

Your application is now receiving intermittent 5xx and 429 HTTP responses from the Cloud Storage JSON API as the demand escalates. You want to decrease the failed responses from the Cloud Storage API.

What should you do?

- A. Distribute the uploads across a large number of individual storage buckets.
- B. Use the XML API instead of the JSON API for interfacing with Cloud Storage.
- C. Pass the HTTP response codes back to clients that are invoking the uploads from your application.
- D. Limit the upload rate from your application clients so that the dormant bucket's peak request rate is reached more gradually.

Answer: (SHOW ANSWER)

Explanation/Reference: <https://cloud.google.com/storage/docs/request-rate>

Valid Professional-Cloud-Developer Dumps shared by PrepPdf.com for Helping Passing Professional-Cloud-Developer Exam! PrepPdf.com now offer the **newest Professional-Cloud-Developer exam dumps**, the PrepPdf.com Professional-Cloud-Developer exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com Professional-Cloud-Developer dumps with Test Engine here:

<https://www.preppdf.com/Google/Professional-Cloud-Developer-prepaway-exam-dumps.html>

(400 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 17

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?

- A. Deploy the website on App Engine and use traffic splitting.
- B. Deploy the website on App Engine as three separate services.
- C. Deploy the website on Cloud Functions and use traffic splitting.
- D. Deploy the website on Cloud Functions as three separate functions.

Answer: A (LEAVE A REPLY)

Explanation/Reference: <https://cloud.google.com/appengine/docs/standard/python/splitting-traffic>

NEW QUESTION: 18

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:

A. instances: 5

min_pending_latency: 30ms

basic_scaling:

B. max_instances: 5

idle_timeout: 10m

basic_scaling:

C. max_instances: 5

idle_timeout: 10m

D. instances: 5

min_pending_latency: 30ms

manual_scaling:

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 19

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?

A. Manually trigger the build for new releases.

B. Create a build trigger on a Git tag pattern. Use a Git tag convention for new releases.

C. Create a build trigger on a Git branch name pattern. Use a Git branch naming convention for new releases.

D. Commit your source code to a second Cloud Source Repositories repository with a second Cloud Build trigger. Use this repository for new releases only.

Answer: C ([LEAVE A REPLY](#))

Reference:

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

NEW QUESTION: 20

You are developing an ecommerce application that stores customer, order, and inventory data as relational tables inside Cloud Spanner. During a recent load test, you discover that Spanner performance is not scaling linearly as expected. Which of the following is the cause?

A. The use of the STRING data type for arbitrary-precision values.

B. The use of LIKE instead of STARTS_WITH keyword for parameterized SQL queries.

C. The use of 64-bit numeric types for 32-bit numbers.

D. The use of Version 1 UUIDs as primary keys that increase monotonically.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 21

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.

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: A ([LEAVE A REPLY](#))

Explanation/Reference:

NEW QUESTION: 22

You are using Cloud Build to build a Docker image. You need to modify the build to execute unit and run integration tests. When there is a failure, you want the build history to clearly display the stage at which the build failed.

What should you do?

- A. Create a Cloud Build build config file with a single build step to compile unit and integration tests.
- B. Create a Cloud Build build config file with separate cloud builder steps to compile and execute unit and integration tests.

- C. Add RUN commands in the Dockerfile to execute unit and integration tests.
- D. Create a Cloud Build build config file that will spawn a separate cloud build pipeline for unit and integration tests.

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 23

A computer service technician says, "I don't use ESD precautions and have never had a problem." Which of the following is the correct response to this statement?

- A. ESD damage may not appear immediately.
- B. ESD happens only to inexperienced technicians.
- C. ESD only occurs in very rare circumstances.
- D. ESD damage is really not as bad as everyone thinks.

Answer: C ([LEAVE A REPLY](#))

Explanation/Reference: <https://electronics.stackexchange.com/questions/2945/should-i-be-worried-about-electro-static-discharge>

NEW QUESTION: 24

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?

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

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 25

Your teammate has asked you to review the code below, which is adding a credit to an account balance in Cloud Datastore. Which improvement should you suggest your teammate make?

```
public Entity creditAccount(long accountId, long
creditAmount) {
    Entity account = datastore.get
    (keyFactory.newKey(accountId));
    account = Entity.builder(account).set(
        "balance", account.getLong("balance")
+ creditAmount).build();
    datastore.put(account);
    return account;
}
```

- A. Use a strongly consistent transactional database.
- B. Get the entity with an ancestor query.
- C. Get and put the entity in a transaction.
- D. Don't return the account entity from the function.

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 26

You configured your Compute Engine instance group to scale automatically according to overall CPU usage.

However, your application's response latency increases sharply before the cluster has finished adding up instances. You want to provide a more consistent latency experience for your end users by changing the configuration of the instance group autoscaler.

Which two configuration changes should you make? (Choose two.)

- A.** Increase the target CPU usage for the instance group autoscaler.
- B.** Decrease the cool-down period for instances added to the group.
- C.** Add the label "AUTOSCALE" to the instance group template.
- D.** Decrease the target CPU usage for the instance group autoscaler.
- E.** Remove the health-check for individual VMs in the instance group.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 27

Your company wants to expand their users outside the United States for their popular application. The company wants to ensure 99.999% availability of the database for their application and also wants to minimize the read latency for their users across the globe.

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

- A.** Create a Cloud Dataflow pipeline to replicate data across different databases.
- B.** Create a cluster with at least 1 Spanner node.
- C.** Create a multi-regional Cloud Spanner instance with "nam3" configuration.
- D.** Create a cluster with at least 3 Spanner nodes.
- E.** Create a minimum of two Cloud Spanner instances in separate regions with at least one node.
- F.** Create a multi-regional Cloud Spanner instance with "nam-asia-eur1" configuration.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 28

You have an application controlled by a managed instance group. When you deploy a new version of the application, costs should be minimized and the number of instances should not increase. You want to ensure that, when each new instance is created, the deployment only continues if the new instance is healthy. What should you do?

- A.** Perform a rolling-action with maxSurge set to 1, maxUnavailable set to 0.
- B.** Perform a rolling-action with maxSurge set to 0, maxUnavailable set to 1
- C.** Perform a rolling-action with maxHealthy set to 1, maxUnhealthy set to 0.
- D.** Perform a rolling-action with maxHealthy set to 0, maxUnhealthy set to 1.

Answer: A ([LEAVE A REPLY](#))

Reference:

<https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instance-groups>

NEW QUESTION: 29

You have two tables in an ANSI-SQL compliant database with identical columns that you need to quickly combine into a single table, removing duplicate rows from the result set.

What should you do?

- A. Use the JOIN operator in SQL to combine the tables.
- B. Use nested WITH statements to combine the tables.
- C. Use the UNION operator in SQL to combine the tables.
- D. Use the UNION ALL operator in SQL to combine the tables.

Answer: ([SHOW ANSWER](#))

Reference:

https://www.techonthenet.com/sql/union_all.php

NEW QUESTION: 30

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: ([SHOW ANSWER](#))

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

NEW QUESTION: 31

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?

- A. Smoke tests
- B. Stackdriver uptime checks
- C. Cloud Load Balancing - health checks
- D. Managed instance group - health checks

Answer: B ([LEAVE A REPLY](#))

Reference:

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

Valid Professional-Cloud-Developer Dumps shared by PrepPdf.com for Helping Passing Professional-Cloud-Developer Exam! PrepPdf.com now offer the **newest Professional-Cloud-Developer exam dumps**, the PrepPdf.com Professional-Cloud-Developer exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com Professional-Cloud-Developer dumps with Test Engine here:
<https://www.preppdf.com/Google/Professional-Cloud-Developer-prepaway-exam-dumps.html>
(400 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 32

You are developing an ecommerce web application that uses App Engine standard environment and Memorystore for Redis. When a user logs into the app, the application caches the user's information (e.g., session, name, address, preferences), which is stored for quick retrieval during checkout.

While testing your application in a browser, you get a 502 Bad Gateway error. You have determined that the application is not connecting to Memorystore. What is the reason for this error?

- A.** You configured your Serverless VPC Access connector in a different region than your App Engine instance.
- B.** You configured your application to use a Serverless VPC Access connector on a different subnet in a different availability zone than your App Engine instance.
- C.** The firewall rule allowing a connection between App Engine and Memorystore was removed during an infrastructure update by the DevOps team.
- D.** Your Memorystore for Redis instance was deployed without a public IP address.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 33

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 ([LEAVE A REPLY](#))

Reference:

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

NEW QUESTION: 34

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.)

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

Answer: B,C ([LEAVE A REPLY](#))

NEW QUESTION: 35

Your existing application keeps user state information in a single MySQL database. This state information is very user-specific and depends heavily on how long a user has been using an application. The MySQL database is causing challenges to maintain and enhance the schema for various users.

Which storage option should you choose?

- A. Cloud SQL
- B. Cloud Storage
- C. Cloud Spanner
- D. Cloud Datastore/Firestore

Answer: ([SHOW ANSWER](#))

Reference:

<https://cloud.google.com/solutions/migrating-mysql-to-cloudsql-concept>

NEW QUESTION: 36

You support an application that uses the Cloud Storage API. You review the logs and discover multiple HTTP 503 Service Unavailable error responses from the API. Your application logs the error and does not take any further action. You want to implement Google-recommended retry logic to improve success rates. Which approach should you take?

- A. Retry the failures in batch after a set number of failures is logged.
- B. Retry each failure at a set time interval up to a maximum number of times.
- C. Retry each failure at increasing time intervals up to a maximum number of tries.
- D. Retry each failure at decreasing time intervals up to a maximum number of tries.

Answer: C ([LEAVE A REPLY](#))

<https://cloud.google.com/storage/docs/retry-strategy>

NEW QUESTION: 37

You are deploying your application to a Compute Engine virtual machine instance with the Stackdriver Monitoring Agent installed. Your application is a unix process on the instance. You want to be alerted if the unix process has not run for at least 5 minutes. You are not able to change the application to generate metrics or logs.

Which alert condition should you configure?

- A. Uptime check
- B. Process health
- C. Metric absence
- D. Metric threshold

Answer: ([SHOW ANSWER](#))

Explanation/Reference: <https://cloud.google.com/monitoring/alerts/concepts-indepth>

NEW QUESTION: 38

Which of the following customer statements would alert you to a safety issue? (Choose two.)

- A. My iPhone flashed and sparked when I tried to charge it.
- B. My iPhone has fluctuating sound levels. Sometimes it is deafening.
- C. The corner of my iPad is badly bent.
- D. The screen is too bright. It hurts my eyes.
- E. The home button on my iPhone seems to have sunk.

F. My new Apple Watch makes me itchy and my wrist is red and irritated.

Answer: A,E ([LEAVE A REPLY](#))

NEW QUESTION: 39

Your teammate has asked you to review the code below. Its purpose is to efficiently add a large number of small rows to a BigQuery table.

```
BigQuery service = BigQueryOptions.newBuilder().build().getService();

public void writeToBigQuery(Collection<Map<String, String>> rows){
    for(Map<String, String> row : rows){
        InsertAllRequest insertRequest = InsertAllRequest.newBuilder(
            "datasetId", "tableId",
            InsertAllRequest.RowToInsert.of(row)).build();
        service.insertAll(insertRequest);
    }
}
```

Which improvement should you suggest your teammate make?

- A. Write each row to a Cloud Storage object, then load into BigQuery.
- B. Include multiple rows with each request.
- C. Perform the inserts in parallel by creating multiple threads.
- D. Write each row to a Cloud Storage object in parallel, then load into BigQuery.

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 40

Your application is running on Compute Engine and is showing sustained failures for a small number of requests. You have narrowed the cause down to a single Compute Engine instance, but the instance is unresponsive to SSH. What should you do next?

- A. Take a snapshot of the disk and attach it to a new machine.
- B. Delete the machine and create a new one.
- C. Enable and check the serial port output.
- D. Reboot the machine.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 41

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 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.

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 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.

D. 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.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 42

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?

A. Perform Read-Only transactions.

B. Perform stale reads using single-read methods.

C. Perform strong reads using single-read methods.

D. Perform stale reads using read-write transactions.

Answer: D ([LEAVE A REPLY](#))

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

NEW QUESTION: 43

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 --port-name http
--global
gcloud compute backend-services add-backend ${NAME}-http-bes --instance-group ${NAME}-i --balancing-mode RATE --max-rate
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 ${NAME}-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?

- A. `gcloud compute instances add-access-config ${NAME}-backend-instance-1`
- B. `gcloud compute instances add-tags ${NAME}-backend-instance-1 --tags http-server`
- C. `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`
- D. `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`

Answer: C ([LEAVE A REPLY](#))

Reference:

<https://cloud.google.com/vpc/docs/special-configurations>

NEW QUESTION: 44

You are deploying your application on a Compute Engine instance that communicates with Cloud SQL. You will use Cloud SQL Proxy to allow your application to communicate to the database using the service account associated with the application's instance. You want to follow the Google-recommended best practice of providing minimum access for the role assigned to the service account. What should you do?

- A. Assign the Project Owner role.
- B. Assign the Cloud SQL Client role.
- C. Assign the Cloud SQL Editor role.
- D. Assign the Project Editor role.

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 45

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: ([SHOW ANSWER](#)**)**

Reference:

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

NEW QUESTION: 46

Please refer to the following information to answer the questions on the right.

Hannah recently picked up her iMac after a repair. The sound from the speakers was distorted and unclear and a technician determined that a repair would resolve the issue.

Once Hannah returned home with the iMac, the sound issue occurred again.

Hannah has returned. She is angry and she is cursing at the technician.

From the following, which are part of the 5-step conflict resolution model? (Choose two.)

- A. "I can prove to you exactly how you are wrong about this issue."
- B. "Calm down!"
- C. "Though your frustration is understandable, we are in a family-friendly environment so I am going to have to ask you to be mindful of the language you are using."
- D. "If I can ask some questions about the last repair and what has occurred since then, I can figure out what the appropriate next step is. I might be able to suggest a solution. Does that sound okay, Hannah?"
- E. "Are you sure your children didn't do something to it?"
- F. "Is this really the same issue? Might it be a different issue?"

Answer: B,D (LEAVE A REPLY)

Valid Professional-Cloud-Developer Dumps shared by PrepPdf.com for Helping Passing Professional-Cloud-Developer Exam! PrepPdf.com now offer the **newest Professional-Cloud-Developer exam dumps**, the PrepPdf.com Professional-Cloud-Developer exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com Professional-Cloud-Developer dumps with Test Engine here:
<https://www.preppdf.com/Google/Professional-Cloud-Developer-prepaway-exam-dumps.html>
(400 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 47

Your application is deployed in a Google Kubernetes Engine (GKE) cluster. When a new version of your application is released, your CI/CD tool updates the `spec.template.spec.containers[0].image` value to reference the Docker image of your new application version. When the Deployment object applies the change, you want to deploy at least 1 replica of the new version and maintain the previous replicas until the new replica is healthy. Which change should you make to the GKE Deployment object shown below?


```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: ecommerce-frontend-deployment
spec:
  replicas: 3
  selector:
    matchLabels:
      app: ecommerce-frontend
  template:
    metadata:
      labels:
        app: ecommerce-frontend
    spec:
      containers:
        - name: ecommerce-frontend-webapp
          image: ecommerce-frontend-webapp:1.7.9
          ports:
            - containerPort: 80
```

- A. Set the Deployment strategy to Recreate with maxSurge set to 1, maxUnavailable set to 0.
- B. Set the Deployment strategy to RollingUpdate with maxSurge set to 1, maxUnavailable set to 0.
- C. Set the Deployment strategy to Recreate with maxSurge set to 0, maxUnavailable set to 1.
- D. Set the Deployment strategy to RollingUpdate with maxSurge set to 0, maxUnavailable set to 1.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 48

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

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 49

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?

- A. Deploy the website on App Engine and use traffic splitting.
- B. Deploy the website on App Engine as three separate services.
- C. Deploy the website on Cloud Functions and use traffic splitting.
- D. Deploy the website on Cloud Functions as three separate functions.

Answer: A ([LEAVE A REPLY](#))

Reference:

<https://cloud.google.com/appengine/docs/standard/python/splitting-traffic>

NEW QUESTION: 50

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 ([LEAVE A REPLY](#))

Reference:

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

NEW QUESTION: 51

You are using Cloud Build to build and test application source code stored in Cloud Source Repositories. The build process requires a build tool not available in the Cloud Build environment.

What should you do?

- A. Download the binary from the internet during the build process.
- B. Build a custom cloud builder image and reference the image in your build steps.
- C. Include the binary in your Cloud Source Repositories repository and reference it in your build scripts.
- D. Ask to have the binary added to the Cloud Build environment by filing a feature request against the Cloud

Answer: B ([LEAVE A REPLY](#))

Build public Issue Tracker.

NEW QUESTION: 52

You have been tasked with planning the migration of your company's application from on-premises to Google Cloud. Your company's monolithic application is an ecommerce website. The application will be migrated to microservices deployed on Google Cloud in stages. The majority of your company's revenue is generated through online sales, so it is important to minimize risk during the migration. You need to prioritize features and select the first functionality to migrate.

What should you do?

- A. Migrate Order fulfillment, which has integrations to the order database, inventory system, and third-party shipping vendor.

- B.** Migrate the Shopping cart, which has integrations to the frontend, cart database, inventory system, and payment processing system.
- C.** Migrate Payment processing, which has integrations to the frontend, order database, and third-party payment vendor.
- D.** Migrate the Product catalog, which has integrations to the frontend and product database.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 53

You migrated some of your applications to Google Cloud. You are using a legacy monitoring platform deployed on-premises for both on-premises and cloud-deployed applications. You discover that your notification system is responding slowly to time-critical problems in the cloud applications. What should you do?

- A.** Replace your monitoring platform with Cloud Monitoring.
- B.** Migrate some traffic back to your old platform. Perform A/B testing on the two platforms concurrently.
- C.** Install the Cloud Monitoring agent on your Compute Engine instances.
- D.** Use Cloud Logging and Cloud Monitoring to capture logs, monitor, and send alerts. Send them to your existing platform.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 54

You are deploying your application to a Compute Engine virtual machine instance. Your application is configured to write its log files to disk. You want to view the logs in Stackdriver Logging without changing the application code.

What should you do?

- A.** Use a Stackdriver Logging Library to log directly from the application to Stackdriver Logging.
- B.** Provide the log file folder path in the metadata of the instance to configure it to send the application logs.
- C.** Change the application to log to /var/log so that its logs are automatically sent to Stackdriver Logging.
- D.** Install the Stackdriver Logging Agent and configure it to send the application logs.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 55

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.

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

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 56

When evaluating an issue, which of the following is an empathetic and supportive closed question about updating software on an iOS device?

- A. How to update your software is in the iPhone User's guide, why didn't you read it before using this phone?
- B. You do know you should always update your software, correct?
- C. Updating your software is extremely easy, why haven't you done so?
- D. I notice that your iPhone software is not up to date; would you like me to walk you through updating the software?

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 57

Which of the following types of data are backed up to iCloud? (Choose two.)

- A. Apple Pay information and settings
- B. iCloud Music Library
- C. Messages
- D. Device Settings
- E. Touch ID Settings

Answer: C,D ([LEAVE A REPLY](#))

NEW QUESTION: 58

Your analytics system executes queries against a BigQuery dataset. The SQL query is executed in batch and passes the contents of a SQL file to the BigQuery CLI. Then it redirects the BigQuery CLI output to another process. However, you are getting a permission error from the BigQuery CLI when the queries are executed.


You want to resolve the issue. What should you do?

- A. Create a new dataset in BigQuery, and copy the source table to the new dataset Query the new dataset and table from the CLI.
- B. Create a view in BigQuery from the SQL query and SELECT*from the view in the CLI.
- C. Grant the service account BigQuery Data Viewer and BigQuery Job User roles.
- D. Grant the service account BigQuery Data Editor and BigQuery Data Viewer roles.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 59

Your team is developing an application in Google Cloud that executes with user identities maintained by Cloud Identity. Each of your application's users will have an associated Pub/Sub topic to which messages are published, and a Pub/Sub subscription where the same user will retrieve published messages. You need to ensure that only authorized users can publish and subscribe to their own specific Pub/Sub topic and subscription. What should you do?



- A. Bind the user identity to the pubsub.publisher and pubsub.subscriber roles at the resource level.
- B. Grant the user identity a custom role that contains the pubsub.topics.create and pubsub.subscriptions.create permissions.
- C. Grant the user identity the pubsub.publisher and pubsub.subscriber roles at the project level.

D. Configure the application to run as a service account that has the pubsub.publisher and pubsub.subscriber roles.

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 60

HipLocal is configuring their access controls.

Which firewall configuration should they implement?

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

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 61

You are creating an App Engine application that writes a file to any user's Google Drive.

How should the application authenticate to the Google Drive API?

- A. With an OAuth Client ID that uses the <https://www.googleapis.com/auth/drive.file> scope to obtain an access token for each user.
- B. With an OAuth Client ID with delegated domain-wide authority.
- C. With the App Engine service account and <https://www.googleapis.com/auth/drive.file> scope that generates a signed JWT.
- D. With the App Engine service account with delegated domain-wide authority.

Answer: B ([LEAVE A REPLY](#))

Reference:

<https://developers.google.com/drive/api/v3/about-auth>

Valid Professional-Cloud-Developer Dumps shared by PrepPdf.com for Helping Passing Professional-Cloud-Developer Exam! PrepPdf.com now offer the **newest Professional-Cloud-Developer exam dumps**, the PrepPdf.com Professional-Cloud-Developer exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com Professional-Cloud-Developer dumps with Test Engine here:
<https://www.preppdf.com/Google/Professional-Cloud-Developer-prepaway-exam-dumps.html>
(400 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 62

You have an application controlled by a managed instance group. When you deploy a new version of the application, costs should be minimized and the number of instances should not increase. You want to ensure that, when each new instance is created, the deployment only continues if the new instance is healthy. What should you do?

- A. Perform a rolling-action with maxSurge set to 0, maxUnavailable set to 1
- B. Perform a rolling-action with maxSurge set to 1, maxUnavailable set to 0.
- C. Perform a rolling-action with maxHealthy set to 1, maxUnhealthy set to 0.
- D. Perform a rolling-action with maxHealthy set to 0, maxUnhealthy set to 1.

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 63

Your development team is using Cloud Build to promote a Node.js application built on App Engine from your staging environment to production. The application relies on several directories of photos stored in a Cloud Storage bucket named webphotos-staging in the staging environment. After the promotion, these photos must be available in a Cloud Storage bucket named webphotos-prod in the production environment. You want to automate the process where possible. What should you do?

A)

Manually copy the photos to webphotos-prod.

B)

Add a startup script in the application's app.yaml file to move the photos from webphotos-staging to webphotos-prod.

C)

Add a build step in the cloudbuild.yaml file before the promotion step with the arguments:

```
- name: gcr.io/cloud-builders/gsutil
  args: ['cp', '-r', 'gs://webphotos-staging',
        'gs://webphotos-prod']
  waitFor: ['-']
```

D)

Add a build step in the cloudbuild.yaml file before the promotion step with the arguments:

```
- name: gcr.io/cloud-builders/gcloud
  args: ['cp', '-A', 'gs://webphotos-staging',
        'gs://webphotos-prod']
  waitFor: ['-']
```

- A. Option C
- B. Option B
- C. Option D
- D. Option A

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 64

You need to copy directory local-scripts and all of its contents from your local workstation to a Compute Engine virtual machine instance.

Which command should you use?

- A. gsutil cp --project "my-gcp-project" -r ~/local-scripts/ gcp-instance-name:~/ server-scripts/ --zone "us-east1-b"

- B.** `gcloud compute mv --project "my-gcp-project" --recurse ~/local-scripts/ gcpinstance-name:~/server-scripts/ --zone "us-east1-b"`
- C.** `gsutil cp --project "my-gcp-project" -R ~/local-scripts/ gcp-instance-name:~/ server-scripts/ --zone "us-east1-b"`
- D.** `gcloud compute scp --project "my-gcp-project" --recurse ~/local-scripts/ gcpinstance-name:~/server-scripts/ --zone "us-east1-b"`

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 65

Select the statement that appropriately sets customer expectations when fixing an issue.

- A.** "We currently do not have that part in stock. I can call you as soon as the part arrives."
- B.** "I promise to have the product ready by midday tomorrow."
- C.** "Repairing this component is easy. We never fail to do it in less than four hours."
- D.** "Our courier always delivers on time."

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 66

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?

- A.** Cloud Amor
- B.** Stackdriver Debugger
- C.** Stackdriver Error Reporting
- D.** Cloud Security Scanner

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 67

Your team develops stateless services that run on Google Kubernetes Engine (GKE). You need to deploy a new service that will only be accessed by other services running in the GKE cluster. The service will need to scale as quickly as possible to respond to changing load. What should you do?

- A.** Use a Vertical Pod Autoscaler to scale the containers, and expose them via a ClusterIP Service.
- B.** Use a Vertical Pod Autoscaler to scale the containers, and expose them via a NodePort Service.
- C.** Use a Horizontal Pod Autoscaler to scale the containers, and expose them via a ClusterIP Service.
- D.** Use a Horizontal Pod Autoscaler to scale the containers, and expose them via a NodePort Service.

Answer: C ([LEAVE A REPLY](#))

<https://cloud.google.com/kubernetes-engine/docs/concepts/service>

NEW QUESTION: 68

Which of the following statements is true about an actively swelling embedded battery?

- A. Actively swelling batteries are not considered a safety concern.
- B. Actively swelling batteries are likely caused by large amounts of liquid inside the batteries.
- C. Actively swelling batteries are unlikely to lead to thermal runaway.
- D. Actively swelling batteries are more likely to rupture and release gasses.

Answer: D ([LEAVE A REPLY](#))

Explanation/Reference: <https://www.ifixit.com/News/11211/what-to-do-with-a-swollen-iphone-battery>

NEW QUESTION: 69

Your existing application keeps user state information in a single MySQL database. This state information is very user-specific and depends heavily on how long a user has been using an application. The MySQL database is causing challenges to maintain and enhance the schema for various users.

Which storage option should you choose?

- A. Cloud SQL
- B. Cloud Storage
- C. Cloud Spanner
- D. Cloud Datastore/Firestore

Answer: A ([LEAVE A REPLY](#))

Explanation/Reference:

Reference: <https://cloud.google.com/solutions/migrating-mysql-to-cloudsql-concept>

NEW QUESTION: 70

You have a mixture of packaged and internally developed applications hosted on a Compute Engine instance that is running Linux. These applications write log records as text in local files. You want the logs to be written to Cloud Logging. What should you do?

- A. Install a Google version of fluentd on the Compute Engine instance.
- B. Using cron, schedule a job to copy the log files to Cloud Storage once a day.
- C. Pipe the content of the files to the Linux Syslog daemon.
- D. Install a Google version of collectd on the Compute Engine instance.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 71

Laura is performing a visual inspection of an embedded battery. What kinds of cosmetic issues should she look for? (Choose two.)

- A. UPC Code
- B. Apple product serial number
- C. IATA complyce label

- D. Scratches
- E. QR Code
- F. Apple logo
- G. Dents or dot imprints

Answer: D,G ([LEAVE A REPLY](#))

Explanation/Reference: <https://docplayer.net/50514766-Embedded-battery-safety.html>

NEW QUESTION: 72

Your existing application keeps user state information in a single MySQL database. This state information is very user-specific and depends heavily on how long a user has been using an application. The MySQL database is causing challenges to maintain and enhance the schema for various users.

Which storage option should you choose?

- A. Cloud Storage
- B. Cloud Datastore/Firestore
- C. Cloud Spanner
- D. Cloud SQL

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 73

You are parsing a log file that contains three columns: a timestamp, an account number (a string), and a transaction amount (a number). You want to calculate the sum of all transaction amounts for each unique account number efficiently.

Which data structure should you use?

- A. A hash table
- B. A two-dimensional array
- C. A linked list
- D. A comma-delimited string

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 74

You are deploying a single website on App Engine that needs to be accessible via the URL <http://www.altostrat.com/>. What should you do?

- A. Verify domain ownership with Webmaster Central. Create a DNS CNAME record to point to the App Engine canonical name ghs.googlehosted.com.
- B. Verify domain ownership with Webmaster Central. Define an A record pointing to the single global App Engine IP address.
- C. Define a mapping in `dispatch.yaml` to point the domain www.altostrat.com to your App Engine service. Create a DNS CNAME record to point to the App Engine canonical name ghs.googlehosted.com.

D. Define a mapping in dispatch.yaml to point the domain www.altostrat.com to your App Engine service. Define an A record pointing to the single global App Engine IP address.

Answer: A ([LEAVE A REPLY](#))

Reference:

<https://cloud.google.com/appengine/docs/flexible/dotnet/mapping-custom-domains?hl=fa>

NEW QUESTION: 75

You are developing a corporate tool on Compute Engine for the finance department, which needs to authenticate users and verify that they are in the finance department. All company employees use G Suite.

What should you do?

- A.** Configure Cloud Armor Security Policies to restrict access to only corporate IP address ranges. Issue client side certificates to everybody in the finance team and verify the certificates in the application.
- B.** Configure Cloud Armor Security Policies to restrict access to only corporate IP address ranges. Verify the provided JSON Web Token within the application.
- C.** Enable Cloud Identity-Aware Proxy on the HTTP(s) load balancer and restrict access to a Google Group containing users in the finance department. Issue client-side certificates to everybody in the finance team and verify the certificates in the application.
- D.** Enable Cloud Identity-Aware Proxy on the HTTP(s) load balancer and restrict access to a Google Group containing users in the finance department. Verify the provided JSON Web Token within the application.

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 76

Lew is not certain that his Apple ID is working correctly in macOS. Which of the following methods would allow Lew to verify that he is using his correct Apple ID in macOS?

- A.** Open System Preferences, click iCloud, and verify that the account is signed in.
- B.** Click the Apple menu, select Apple ID, and ensure that Lew's Apple ID is listed.
- C.** Open the Apple ID pane in System Preferences.
- D.** Go to the My Apple ID page (appleid.apple.com) and verify the account is signed in.

Answer: A ([LEAVE A REPLY](#))

Explanation/Reference: <https://appletoolbox.com/how-to-sign-in-and-sign-out-of-icloud-on-macos-catalina/>

Valid Professional-Cloud-Developer Dumps shared by PrepPdf.com for Helping Passing Professional-Cloud-Developer Exam! PrepPdf.com now offer the **newest Professional-Cloud-Developer exam dumps**, the PrepPdf.com Professional-Cloud-Developer exam questions have been updated and answers have been corrected get the **newest**

PrepPdf.com Professional-Cloud-Developer dumps with Test Engine here:

<https://www.preppdf.com/Google/Professional-Cloud-Developer-prepaway-exam-dumps.html>

(400 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 77

You are developing an application that will allow clients to download a file from your website for a specific period of time. How should you design the application to complete this task while following Google-recommended best practices?

- A. Generate and assign a Cloud Storage-signed URL for the file. Make the URL available for the client to download.
- B. Configure the application to send the file to the client as an email attachment.
- C. Create a temporary Cloud Storage bucket with time expiration specified, and give download permissions to the bucket. Copy the file, and send it to the client.
- D. Generate the HTTP cookies with time expiration specified. If the time is valid, copy the file from the Cloud Storage bucket, and make the file available for the client to download.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 78

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?

- A. ReplicaController
- B. Deployment
- C. StatefulSet
- D. ReplicaSet

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 79

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 --port-name http
--global
gcloud compute backend-services add-backend ${NAME}-http-bes --instance-group ${NAME}-i --balancing-mode RATE --max-rate
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 ${NAME}-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?

- A. `gcloud compute instances add-tags ${NAME}-backend-instance-1 --tags http-server`
- B. `gcloud compute instances add-access-config ${NAME}-backend-instance-1`
- C. `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`
- D. `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`

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 80

Your security team is auditing all deployed applications running in Google Kubernetes Engine. After completing the audit, your team discovers that some of the applications send traffic within the cluster in clear text. You need to ensure that all application traffic is encrypted as quickly as possible while minimizing changes to your applications and maintaining support from Google. What should you do?

- A. Install Istio, enable proxy injection on your application namespace, and then enable mTLS.
- B. Define Trusted Network ranges within the application, and configure the applications to allow traffic only from those networks.
- C. Use an automated process to request SSL Certificates for your applications from Let's Encrypt and add them to your applications.
- D. Use Network Policies to block traffic between applications.

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 81

Your teammate has asked you to review the code below. Its purpose is to efficiently add a large number of small rows to a BigQuery table.


```

BigQuery service = BigQueryOptions.newBuilder().build().getService();

public void writeToBigQuery(Collection<Map<String, String>> rows){
    for(Map<String, String> row : rows){
        InsertAllRequest insertRequest = InsertAllRequest.newBuilder(
            "datasetId", "tableId",
            InsertAllRequest.RowToInsert.of(row)).build();
        service.insertAll(insertRequest);
    }
}

```

Which improvement should you suggest your teammate make?

- A. Perform the inserts in parallel by creating multiple threads.
- B. Include multiple rows with each request.
- C. Write each row to a Cloud Storage object in parallel, then load into BigQuery.
- D. Write each row to a Cloud Storage object, then load into BigQuery.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 82

From the following, select the best description of cosmetic condition that would be included in good case notes.

- A. Significant dent in lower right corner suggests impact
- B. Case is dented like it was dropped
- C. Badly dinged in addition to crack
- D. Big ding; probably dropped or thrown

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 83

You are working on a social media application. You plan to add a feature that allows users to upload images. These images will be 2 MB - 1 GB in size. You want to minimize their infrastructure operations overhead for this feature. What should you do?

- A. Change the application to accept images directly and store them in the database that stores other user information.
- B. Change the application to create signed URLs for Cloud Storage. Transfer these signed URLs to the client application to upload images to Cloud Storage.
- C. Set up a web server on GCP to accept user images and create a file store to keep uploaded files. Change the application to retrieve images from the file store.
- D. Create a separate bucket for each user in Cloud Storage. Assign a separate service account to allow write access on each bucket. Transfer service account credentials to the client application based on user information. The application uses this service account to upload images to Cloud Storage.

Answer: B ([LEAVE A REPLY](#))

Reference:

<https://cloud.google.com/blog/products/storage-data-transfer/uploading-images-directly-to-cloud-storage-by-using-signed-url>

NEW QUESTION: 84

You are creating an App Engine application that writes a file to any user's Google Drive. How should the application authenticate to the Google Drive API?

- A.** With an OAuth Client ID that uses the <https://www.googleapis.com/auth/drive.filescope> to obtain an access token for each user.
- B.** With an OAuth Client ID with delegated domain-wide authority.
- C.** With the App Engine service account and <https://www.googleapis.com/auth/drive.filescope> that generates a signed JWT.
- D.** With the App Engine service account with delegated domain-wide authority.

Answer: ([SHOW ANSWER](#))

Explanation/Reference: <https://developers.google.com/drive/api/v3/about-auth>

NEW QUESTION: 85

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** ([LEAVE A REPLY](#))

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

NEW QUESTION: 86

You are developing a web application that will be accessible over both HTTP and HTTPS and will run on Compute Engine instances. On occasion, you will need to SSH from your remote laptop into one of the Compute Engine instances to conduct maintenance on the app. How should you configure the instances while following Google-recommended best practices?

- A.** Configure Cloud Identity-Aware Proxy API for SSH access. Then configure the Compute Engine servers with private IP addresses behind an HTTP(s) load balancer for the application web traffic.

- B.** Set up a backend with Compute Engine web server instances with a private IP address behind an HTTP(S) load balancer. Set up a bastion host with a public IP address and open firewall ports. Connect to the web instances using the bastion host.
- C.** Set up a backend with Compute Engine web server instances with a private IP address behind a TCP proxy load balancer.
- D.** Configure the firewall rules to allow all ingress traffic to connect to the Compute Engine web servers, with each server having a unique external IP address.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 87

In which of the following circumstances should a technician ensure that he or she is grounded to minimize ESD?

- A.** The technician is working on a device that is plugged in to AC power.
- B.** The technician is cleaning the display on a device after reassembly.
- C.** The technician is conducting a current test on a device.
- D.** The technician is replacing a circuit board inside a device.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 88

You are designing a schema for a Cloud Spanner customer database. You want to store a phone number array field in a customer table. You also want to allow users to search customers by phone number. How should you design this schema?

- A.** Create a table named Customers. Add an Array field in a table that will hold phone numbers for the customer.
- B.** Create a table named Customers. Create a table named Phones. Add a CustomerId field in the Phones table to find the CustomerId from a phone number.
- C.** Create a table named Customers. Add an Array field in a table that will hold phone numbers for the customer. Create a secondary index on the Array field.
- D.** Create a table named Customers as a parent table. Create a table named Phones, and interleave this table into the Customer table. Create an index on the phone number field in the Phones table.

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 89

You are creating and running containers across different projects in Google Cloud. The application you are developing needs to access Google Cloud services from within Google Kubernetes Engine (GKE).

What should you do?

- A.** Assign a Google service account to the GKE nodes.
- B.** Use a Google service account to run the Pod with Workload Identity.
- C.** Store the Google service account credentials as a Kubernetes Secret.

D. Use a Google service account with GKE role-based access control (RBAC).

Answer: ([SHOW ANSWER](#))

<https://cloud.google.com/kubernetes-engine/docs/concepts/workload-identity>

NEW QUESTION: 90

You need to migrate an internal file upload API with an enforced 500-MB file size limit to App Engine.

What should you do?

- A.** Use FTP to upload files.
- B.** Use CPanel to upload files.
- C.** Use signed URLs to upload files.
- D.** Change the API to be a multipart file upload API.

Answer: **C** ([LEAVE A REPLY](#))

Reference:

https://wiki.christophchamp.com/index.php?title=Google_Cloud_Platform

NEW QUESTION: 91

Your application is deployed in a Google Kubernetes Engine (GKE) cluster. You want to expose this application publicly behind a Cloud Load Balancing HTTP(S) load balancer. What should you do?

- A.** Configure a GKE Ingress resource.
- B.** Configure a GKE Service resource.
- C.** Configure a GKE Ingress resource with type: LoadBalancer.
- D.** Configure a GKE Service resource with type: LoadBalancer.

Answer: **A** ([LEAVE A REPLY](#))

Reference:

<https://cloud.google.com/kubernetes-engine/docs/concepts/ingress>

Valid Professional-Cloud-Developer Dumps shared by PrepPdf.com for Helping Passing Professional-Cloud-Developer Exam! PrepPdf.com now offer the **newest Professional-Cloud-Developer exam dumps**, the PrepPdf.com Professional-Cloud-Developer exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com Professional-Cloud-Developer dumps with Test Engine here:

<https://www.preppdf.com/Google/Professional-Cloud-Developer-prepaway-exam-dumps.html>

(400 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 92

You want to view the memory usage of your application deployed on Compute Engine. What should you do?

- A. Use the Stackdriver Metrics Explorer.
- B. Install the Stackdriver Monitoring Agent.
- C. Use the Google Cloud Platform Console.
- D. Install the Stackdriver Client Library.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 93

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?

- A. Migrate your data stored in Hadoop to BigQuery. Change your jobs to source their information from BigQuery instead of the on-premises Hadoop environment.
- B. 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.
- C. 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.
- D. 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.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 94

Please refer to the following information to answer the questions on the right.

Debra has brought her iPad Pro to you for troubleshooting. She has been experiencing very slow charging of the battery and says the only way she can get a full charge is to leave the iPad alone for more than 12 hours.

She brought in an Apple Power adapter and Lightning to USB cable. When you examine the power adapter, you notice that it is an iPhone power adapter not an iPad power adapter. You also noticed she was using an iPhone 6s while waiting to see you. Debra thinks her iPad is defective because her iPhone has no issue charging with this power adapter and cable.

Based on the information Debra provided, in which problem space does this issue fit?

- A. User Data
- B. Software
- C. Environmental
- D. Educational opportunity

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 95

You are creating a Google Kubernetes Engine (GKE) cluster and run this command:

```
> gcloud container clusters create large-cluster --num-nodes 200
```

The command fails with the error:

```
insufficient regional quota to satisfy request: resource "CPUS": request requires '200.0' and is short '176.0'. project has a quota of '24.0' with '24.0' available
```

You want to resolve the issue. What should you do?

- A. Request additional GKE quota in the GCP Console.
- B. Request additional Compute Engine quota in the GCP Console.
- C. Open a support case to request additional GKE quota.
- D. Decouple services in the cluster, and rewrite new clusters to function with fewer cores.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 96

What are the features of using iCloud as a backup method for an iOS device?

- A. Always encrypts backups.
- B. Backups can be encrypted or un-encrypted.
- C. Offers up to 200 TB of storage.
- D. The customer can make backups over cellular networks.

Answer: ([SHOW ANSWER](#))

Explanation/Reference: <https://www.pcworld.com/article/2154792/backups-via-icloud-or-itunes-the-crucial-differences-that-affect-your-data.html>

NEW QUESTION: 97

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 ([LEAVE A REPLY](#))

Reference:

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

NEW QUESTION: 98

Your teammate has asked you to review the code below, which is adding a credit to an account balance in Cloud Datastore.

Which improvement should you suggest your teammate make?

```
public Entity creditAccount(long accountId, long
creditAmount) {
    Entity account = datastore.get
(keyFactory.newKey(accountId));
    account = Entity.builder(account).set(
        "balance", account.getLong("balance")
+ creditAmount).build();
    datastore.put(account);
    return account;
}
```

- A. Get the entity with an ancestor query.
- B. Get and put the entity in a transaction.
- C. Use a strongly consistent transactional database.
- D. Don't return the account entity from the function.

Answer: ([SHOW ANSWER](#))

Explanation/Reference:

NEW QUESTION: 99

You are building a new API. You want to minimize the cost of storing and reduce the latency of serving images.

Which architecture should you use?

- A. App Engine backed by Cloud Storage
- B. Cloud Content Delivery Network (CDN) backed by Cloud Storage
- C. Compute Engine backed by Persistent Disk
- D. Transfer Appliance backed by Cloud Filestore

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 100

You recently developed a new service on Cloud Run. The new service authenticates using a custom service and then writes transactional information to a Cloud Spanner database. You need to verify that your application can support up to 5,000 read and 1,000 write transactions per second while identifying any bottlenecks that occur. Your test infrastructure must be able to autoscale. What should you do?

- A. Build a test harness to generate requests and deploy it to Cloud Run. Analyze the VPC Flow Logs using Cloud Logging.
- B. Create a Google Kubernetes Engine cluster running the Locust or JMeter images to dynamically generate load tests. Analyze the results using Cloud Trace.
- C. Create a Cloud Task to generate a test load. Use Cloud Scheduler to run 60,000 Cloud Task transactions per minute for 10 minutes. Analyze the results using Cloud Monitoring.

D. Create a Compute Engine instance that uses a LAMP stack image from the Marketplace, and use Apache Bench to generate load tests against the service. Analyze the results using Cloud Trace.

Answer: B ([LEAVE A REPLY](#))

<https://cloud.google.com/architecture/distributed-load-testing-using-gke>

NEW QUESTION: 101

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?

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

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 102

Which equipment is used to check if an AC power outlet is properly grounded?

- A. A conductive workbench mat
- B. A nylon probe tool (black stick)
- C. A grounding cord with alligator clips
- D. A ground polarity tester

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 103

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 ([LEAVE A REPLY](#))

Topic 1, HipLocal

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

NEW QUESTION: 104

You deployed a new application to Google Kubernetes Engine and are experiencing some performance degradation. Your logs are being written to Cloud Logging, and you are using a Prometheus sidecar model for capturing metrics. You need to correlate the metrics and data from the logs to troubleshoot the performance issue and send real-time alerts while minimizing costs. What should you do?

- A.** Export the Cloud Logging logs and stream the Prometheus metrics to BigQuery. Run a recurring query to join the results, and send notifications using Cloud Tasks.
- B.** Create custom metrics from the Cloud Logging logs, and use Prometheus to import the results using the Cloud Monitoring REST API.
- C.** Export the Prometheus metrics and use Cloud Monitoring to view them as external metrics. Configure Cloud Monitoring to create log-based metrics from the logs, and correlate them with the Prometheus data.
- D.** Export the Cloud Logging logs and the Prometheus metrics to Cloud Bigtable. Run a query to join the results, and analyze in Google Data Studio.

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 105

Your API backend is running on multiple cloud providers. You want to generate reports for the network latency of your API.

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

- A.** Use Stackdriver Profiler to generate report.
- B.** Use Zipkin collector to gather data.
- C.** Use Stackdriver Debugger to generate report.
- D.** Use Stackdriver Trace to generate reports.
- E.** Use Fluentd agent to gather data.

Answer: A,D ([LEAVE A REPLY](#))

NEW QUESTION: 106

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?

- A.** 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.

B. 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.

C. 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.

D. 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.

Answer: A (LEAVE A REPLY)

Explanation/Reference: <https://cloud.google.com/container-registry/>

Valid Professional-Cloud-Developer Dumps shared by PrepPdf.com for Helping Passing Professional-Cloud-Developer Exam! PrepPdf.com now offer the **newest Professional-Cloud-Developer exam dumps**, the PrepPdf.com Professional-Cloud-Developer exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com Professional-Cloud-Developer dumps with Test Engine here:
<https://www.preppdf.com/Google/Professional-Cloud-Developer-prepaway-exam-dumps.html>
(400 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 107

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

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 108

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?

- A.** Migrate your data stored in Hadoop to BigQuery. Change your jobs to source their information from BigQuery instead of the on-premises Hadoop environment.
- B.** 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.
- C.** 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.
- D.** 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.

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 109

You plan to make a simple HTML application available on the internet. This site keeps information about FAQs for your application. The application is static and contains images, HTML, CSS, and Javascript. You want to make this application available on the internet with as few steps as possible.

What should you do?

- A.** Upload your application to Cloud Storage.
- B.** Upload your application to an App Engine environment.
- C.** Create a Compute Engine instance with Apache web server installed. Configure Apache web server to host the application.
- D.** Containerize your application first. Deploy this container to Google Kubernetes Engine (GKE) and assign an external IP address to the GKE pod hosting the application.

Answer: ([SHOW ANSWER](#)**)**

Reference:

<https://cloud.google.com/storage/docs/hosting-static-website>

NEW QUESTION: 110

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.

After gaining agreement with the customer about repair time and cost, Rachel begins the replacement of the top case assembly with battery. If during the repair the battery cells begin to actively swell, how should Rachel respond?

- A.** Leave the room for fifteen minutes.

- B. Find a fire extinguisher and keep it nearby.
- C. Pour sand over the battery immediately.
- D. Carefully continue the repair while monitoring the battery.

Answer: C (LEAVE A REPLY)

NEW QUESTION: 111

Tony would like to create a new Apple ID to use with iTunes. It is his preference not to give his credit card information.

Which of the following is the correct statement about his situation?

- A. Tony will have to purchase an iTunes gift card before he can setup his Apple ID.
- B. Tony will have to use a bank account and routing number to setup his Apple ID.
- C. Tony can create an Apple ID in iTunes without a credit card or other payment method.
- D. Tony will not be able to use an Apple ID with iTunes unless he is willing to give credit card information.

Answer: C (LEAVE A REPLY)

Explanation/Reference: <https://support.apple.com/en-us/HT204034>

Valid Professional-Cloud-Developer Dumps shared by PrepPdf.com for Helping Passing Professional-Cloud-Developer Exam! PrepPdf.com now offer the **newest Professional-Cloud-Developer exam dumps**, the PrepPdf.com Professional-Cloud-Developer exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com Professional-Cloud-Developer dumps with Test Engine here:
<https://www.preppdf.com/Google/Professional-Cloud-Developer-prepaway-exam-dumps.html>
(400 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)