

Google.Professional-Cloud-Developer.v2023-06-10.q116

Exam Code:	Professional-Cloud-Developer
Exam Name:	Google Certified Professional - Cloud Developer
Certification Provider:	Google
Free Question Number:	116
Version:	v2023-06-10
# of views:	1426
# of Questions views:	1150
https://www.freeqas.com/qa/Google/Professional-Cloud-Developer/Google.Professional-Cloud-Developer.v2023-06-10.q116.html	

NEW QUESTION: 1

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 2

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

What should you do?

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

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

Explanation

NEW QUESTION: 3

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](#))

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

NEW QUESTION: 4

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

What should you do?

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

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

Explanation/Reference:

NEW QUESTION: 5

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

Which IP strategy should they use?

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

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

NEW QUESTION: 6

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

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

NEW QUESTION: 7

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

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

NEW QUESTION: 8

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: A (LEAVE A REPLY)

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

NEW QUESTION: 9

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.** Replace your entire monitoring platform with Stackdriver.
- 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.** Migrate some traffic back to your old platform and perform AB testing on the two platforms concurrently.

Answer: B (LEAVE A REPLY)

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

NEW QUESTION: 10

Case Study

Company Overview

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

Executive Statement

We are the number one local community app; it's time to take our local community services global. Our venture capital investors want to see rapid growth and the same great experience for new local and virtual communities that come online, whether their members are 10 or 10000 miles away from each other.

Solution Concept

HipLocal wants to expand their existing service, with updated functionality, in new regions to better serve their global customers. They want to hire and train a new team to support

these regions in their time zones. They will need to ensure that the application scales smoothly and provides clear uptime data.

Existing Technical Environment

HipLocal's environment is a mix of on-premises hardware and infrastructure running in Google Cloud Platform.

The HipLocal team understands their application well, but has limited experience in global scale applications.

Their existing technical environment is as follows:

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

Business Requirements

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

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

Technical Requirements

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

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

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 11

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: B (LEAVE A REPLY)

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

NEW QUESTION: 12

When would a technician need to use a positive 'no' with a customer? (Choose three.)

- A. Apple introduces new product.
- B. The device is an obsolete device that we no longer service.
- C. The customer is misinformed about service options.
- D. The customer wants to purchase a new device.
- E. The device is out of warranty.
- F. The technician does not feel like helping the customer.
- G. The device is eligible for repair.

Answer: B,C,E (LEAVE A REPLY)

NEW QUESTION: 13

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 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.
- C. Create a dataset named Master dataset. Create a separate view for each department in the Master dataset. Give employees access to the specific view for their department.
- D. Create a 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.

Answer: D (LEAVE A REPLY)

NEW QUESTION: 14

Leigh states that her MacBook Pro (Retina, 15-inch, Mid 2015) does not recognize the SD card she brought with her. You have her reproduce the issue and discover that she is properly inserting the card. However, it is not recognized by Photos or Image Capture. What question should you ask next to isolate the issue to hardware?

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

Answer: A (LEAVE A REPLY)

NEW QUESTION: 15

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>

NEW QUESTION: 16

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)

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

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> (265 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 17

You manage your company's ecommerce platform's payment system, which runs on Google Cloud. Your company must retain user logs for 1 year for internal auditing purposes and for 3 years to meet compliance requirements. You need to store new user logs on Google Cloud to minimize on-premises storage usage and ensure that they are easily searchable. You want to minimize effort while ensuring that the logs are stored correctly. What should you do?

- A. Store the logs in a Cloud Storage bucket with bucket lock turned on.
- B. Store the logs in a Cloud Storage bucket with a 3-year retention period.
- C. Store the logs in Cloud Logging as custom logs with a custom retention period.
- D. Store the logs in a Cloud Storage bucket with a 1-year retention period. After 1 year, move the logs to another bucket with a 2-year retention period.

Answer: (SHOW ANSWER)

Explanation

<https://cloud.google.com/logging/docs/buckets#custom-retention>

NEW QUESTION: 18

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. Create a Compute Engine instance with Apache web server installed. Configure Apache web server to host the application.
- C. 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.
- D. Upload your application to an App Engine environment.

Answer: (SHOW ANSWER)

NEW QUESTION: 19

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

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

NEW QUESTION: 20

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

Explanation/Reference:

NEW QUESTION: 22

You are developing an internal application that will allow employees to organize community events within your company. You deployed your application on a single Compute Engine instance. Your company uses Google Workspace (formerly G Suite), and you need to ensure that the company employees can authenticate to the application from anywhere. What should you do?

- A.** Add a public IP address to your instance, and restrict access to the instance using firewall rules. Allow your company's proxy as the only source IP address.
- B.** Add an HTTP(S) load balancer in front of the instance, and set up Identity-Aware Proxy (IAP).

Configure the IAP settings to allow your company domain to access the website.

- C.** Set up a VPN tunnel between your company network and your instance's VPC location on Google Cloud. Configure the required firewall rules and routing information to both the on-premises and Google Cloud networks.
- D.** Add a public IP address to your instance, and allow traffic from the internet. Generate a random hash, and create a subdomain that includes this hash and points to your instance. Distribute this DNS address to your company's employees.

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

Explanation

<https://cloud.google.com/blog/topics/developers-practitioners/control-access-your-web-sites-identity-aware-prox>

NEW QUESTION: 23

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.

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

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

Answer: D (LEAVE A REPLY)

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

NEW QUESTION: 24

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: 25

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:

D)

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

A. Option B

B. Option D

C. Option A

D. Option C

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 26

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:/

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: 27

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. Remove the health-check for individual VMs in the instance group.

B. Decrease the cool-down period for instances added to the group.

C. Decrease the target CPU usage for the instance group autoscaler.

D. Increase the target CPU usage for the instance group autoscaler.

E. Add the label "AUTOSCALE" to the instance group template.

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

NEW QUESTION: 28

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

What should you do?

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

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

NEW QUESTION: 29

Examine the image.

Which feature in macOS is opened by clicking on this dock icon?

- A. System Preferences
- B. Launchpad
- C. Finder
- D. Spotlight

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

Explanation/Reference: <https://support.apple.com/en-us/HT202635>

NEW QUESTION: 30

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: 31

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?

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

Answer: A ([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> (265 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 32

You are a developer working with the CI/CD team to troubleshoot a new feature that your team introduced. The CI/CD team used HashiCorp Packer to create a new Compute Engine image from your development branch. The image was successfully built, but is not booting up. You need to investigate the issue with the CI/CD team. What should you do?

- A.** Create a new feature branch, and ask the build team to rebuild the image.
- B.** Shut down the deployed virtual machine, export the disk, and then mount the disk locally to access the boot logs.
- C.** Install Packer locally, build the Compute Engine image locally, and then run it in your personal Google Cloud project.
- D.** Check Compute Engine OS logs using the serial port, and check the Cloud Logging logs to confirm access to the serial port.

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

<https://cloud.google.com/compute/docs/troubleshooting/troubleshooting-using-serial-console>

NEW QUESTION: 33

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

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

NEW QUESTION: 34

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: (SHOW ANSWER)

Reference:

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

NEW QUESTION: 35

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.** Set up a backend with Compute Engine web server instances with a private IP address behind a TCP proxy load balancer.
- B.** 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.

C. 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.

D. 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.

Answer: C (LEAVE A REPLY)

Reference: https://cloud.google.com/compute/docs/instances/connecting-advanced#cloud_iap

https://cloud.google.com/solutions/connecting-securely#storing_host_keys_by_enabling_guest_attributes

NEW QUESTION: 36

Your company has a BigQuery dataset named "Master" that keeps information about employee travel and expenses. This information is organized by employee department. That means employees should only be able to view information for their department. You want to apply a security framework to enforce this requirement with the minimum number of steps.

What should you do?

A. Create a separate dataset for each department. Create a view with an appropriate WHERE clause to select records from a particular dataset for the specific department. Authorize this view to access records from your Master dataset. Give employees the permission to this department-specific dataset.

B. Create a separate dataset for each department. Create a data pipeline for each department to copy appropriate information from the Master dataset to the specific dataset for the department. Give employees the permission to this department-specific dataset.

C. Create a dataset named Master dataset. Create a separate view for each department in the Master dataset. Give employees access to the specific view for their department.

D. Create a dataset named Master dataset. Create a separate table for each department in the Master dataset. Give employees access to the specific table for their department.

Answer: B (LEAVE A REPLY)

Explanation

NEW QUESTION: 37

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 Armor

B. Stackdriver Debugger

C. Cloud Security Scanner

D. Stackdriver Error Reporting

Answer: (SHOW ANSWER)

Reference:

<https://cloud.google.com/security-scanner>

NEW QUESTION: 38

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

Answer: B (LEAVE A REPLY)

NEW QUESTION: 39

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)

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

NEW QUESTION: 40

This architectural diagram depicts a system that streams data from thousands of devices. You want to ingest data into a pipeline, store the data, and analyze the data using SQL statements. Which Google Cloud services should you use for steps 1, 2, 3, and 4?

- A.** 1) Pub/Sub
2) Dataflow
3) Firestore
4) BigQuery
- B.** 1) Pub/Sub

- 2) Dataflow
- 3) BigQuery
- 4) Firestore
- C.** 1) App Engine
- 2) Pub/Sub
- 3) BigQuery
- 4) Firestore
- D.** 1) Dataflow
- 2) Pub/Sub
- 3) Firestore
- 4) BigQuery

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 41

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

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

NEW QUESTION: 42

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:

476b8507f59c

NEW QUESTION: 43

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](#))

Explanation

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

NEW QUESTION: 44

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: 45

You have a container deployed on Google Kubernetes Engine. The container can sometimes be slow to launch, so you have implemented a liveness probe. You notice that the liveness probe occasionally fails on launch. What should you do?

A. Add a startup probe.

B. Increase the initial delay for the liveness probe.

C. Increase the CPU limit for the container.

D. Add a readiness probe.

Answer: (SHOW ANSWER)

Explanation

<https://kubernetes.io/docs/tasks/configure-pod-container/configure-liveness-readiness-startup-probes/#configure->

NEW QUESTION: 46

You are in the final stage of migrating an on-premises data center to Google Cloud. You are quickly approaching your deadline, and discover that a web API is running on a server slated for decommissioning.

You need to recommend a solution to modernize this API while migrating to Google Cloud.

The modernized web API must meet the following requirements:

- * Autoscales during high traffic periods at the end of each month

- * Written in Python 3.x

- * Developers must be able to rapidly deploy new versions in response to frequent code changes

You want to minimize cost, effort, and operational overhead of this migration.

What should you do?

A. Modernize and deploy the code on App Engine flexible environment.

B. Modernize and deploy the code on App Engine standard environment.

C. Deploy the modernized application to an n1-standard-1 Compute Engine instance.

D. Ask the development team to re-write the application to run as a Docker container on Google Kubernetes Engine.

Answer: B (LEAVE A REPLY)

Explanation

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

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> (265 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 47

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. Set up a backend with Compute Engine web server instances with a private IP address behind a TCP proxy load balancer.
- B. 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.
- C. 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.
- 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: B ([LEAVE A REPLY](#))

NEW QUESTION: 48

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

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

NEW QUESTION: 49

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

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

Answer: ([SHOW ANSWER](#))

Reference:

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

NEW QUESTION: 50

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

The command fails with the error:

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: 51

Your service adds text to images that it reads from Cloud Storage. During busy times of the year, requests to Cloud Storage fail with an HTTP 429 "Too Many Requests" status code.

How should you handle this error?

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

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

Reference:

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

NEW QUESTION: 52

You are developing an application that needs to store files belonging to users in Cloud Storage. You want each user to have their own subdirectory in Cloud Storage. When a new user is created, the corresponding empty subdirectory should also be created. What should you do?

- A. Create an object with the name of the subdirectory ending with a trailing slash ('/') that is zero bytes in length.

- B.** Create an object with the name of the subdirectory, and then immediately delete the object within that subdirectory.
- C.** Create an object with the name of the subdirectory that is zero bytes in length and has WRITER access control list permission.
- D.** Create an object with the name of the subdirectory that is zero bytes in length. Set the Content-Type metadata to CLOUDSTORAGE_FOLDER.

Answer: (SHOW ANSWER)

Explanation

<https://cloud.google.com/storage/docs/folders>

If you create an empty folder using the Google Cloud console, Cloud Storage creates a zero-byte object as a placeholder. For example, if you create a folder called folder in a bucket called my-bucket, a zero-byte object called gs://my-bucket/folder/ is created. This placeholder is discoverable by other tools when listing the objects in the bucket, for example when using the gsutil ls command.

NEW QUESTION: 53

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: B (LEAVE A REPLY)

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

NEW QUESTION: 54

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: 55

Your team develops services that run on Google Cloud. You want to process messages sent to a Pub/Sub topic, and then store them. Each message must be processed exactly once to avoid duplication of data and any data conflicts. You need to use the cheapest and most simple solution. What should you do?

- A.** Process the messages with a Dataproc job, and write the output to storage.
- B.** Process the messages with a Dataflow streaming pipeline using Apache Beam's PubSubIO package, and write the output to storage.
- C.** Process the messages with a Cloud Function, and write the results to a BigQuery location where you can run a job to deduplicate the data.
- D.** Retrieve the messages with a Dataflow streaming pipeline, store them in Cloud Bigtable, and use another Dataflow streaming pipeline to deduplicate messages.

Answer: B (LEAVE A REPLY)

Explanation

<https://cloud.google.com/dataflow/docs/concepts/streaming-with-cloud-pubsub>

NEW QUESTION: 56

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

Answer: A (LEAVE A REPLY)

NEW QUESTION: 57

You are using Cloud Run to host a web application. You need to securely obtain the application project ID and region where the application is running and display this information to users. You want to use the most performant approach. What should you do?

- A.** Make an API call to the Cloud Asset Inventory API from the application and format the request to include instance metadata.
- B.** Use HTTP requests to query the available metadata server at the `http://metadata.google.internal/` endpoint with the Metadata-Flavor: Google header.

C. In the Google Cloud console, navigate to the Project Dashboard and gather configuration details.

Navigate to the Cloud Run "Variables & Secrets" tab, and add the desired environment variables in Key:Value format.

D. In the Google Cloud console, navigate to the Project Dashboard and gather configuration details. Write the application configuration information to Cloud Run's in-memory container filesystem.

Answer: C (LEAVE A REPLY)

NEW QUESTION: 58

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: A (LEAVE A REPLY)

Reference:

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

NEW QUESTION: 59

You need to migrate a standalone Java application running in an on-premises Linux virtual machine (VM) to Google Cloud in a cost-effective manner. You decide not to take the lift-and-shift approach, and instead you plan to modernize the application by converting it to a container. How should you accomplish this task?

A. Use Migrate for Anthos to migrate the VM to your Google Kubernetes Engine (GKE) cluster as a container.

B. Export the VM as a raw disk and import it as an image. Create a Compute Engine instance from the Imported image.

C. Use Migrate for Compute Engine to migrate the VM to a Compute Engine instance, and use Cloud Build to convert it to a container.

D. Use Jib to build a Docker image from your source code, and upload it to Artifact Registry. Deploy the application in a GKE cluster, and test the application.

Answer: D (LEAVE A REPLY)

Explanation

<https://cloud.google.com/blog/products/application-development/introducing-jib-build-java-docker-images-bette>

NEW QUESTION: 60

Your development team has been tasked with maintaining a .NET legacy application. The application incurs occasional changes and was recently updated. Your goal is to ensure that the application provides consistent results while moving through the CI/CD pipeline from environment to environment. You want to minimize the cost of deployment while making sure that external factors and dependencies between hosting environments are not problematic. Containers are not yet approved in your organization. What should you do?

- A.** Rewrite the application using .NET Core, and deploy to Cloud Run. Use revisions to separate the environments.
- B.** Use Cloud Build to deploy the application as a new Compute Engine image for each build. Use this image in each environment.
- C.** Deploy the application using MS Web Deploy, and make sure to always use the latest, patched MS Windows Server base image in Compute Engine.
- D.** Use Cloud Build to package the application, and deploy to a Google Kubernetes Engine cluster. Use namespaces to separate the environments.

Answer: B (LEAVE A REPLY)

Explanation

https://cloud.google.com/architecture/modernization-path-dotnet-applications-google-cloud#phase_1_rehost_in_

<https://cloud.google.com/architecture/modernization-path-dotnet-applications-google-cloud>

NEW QUESTION: 61

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>

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> (265 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 62

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. Metric absence
- C. Process health
- D. Metric threshold

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

NEW QUESTION: 63

You want to migrate an on-premises container running in Knative to Google Cloud. You need to make sure that the migration doesn't affect your application's deployment strategy, and you want to use a fully managed service. Which Google Cloud service should you use to deploy your container?

- A. Cloud Run
- B. Compute Engine
- C. Google Kubernetes Engine
- D. App Engine flexible environment

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

Explanation

<https://cloud.google.com/blog/products/serverless/knative-based-cloud-run-services-are-ga>

NEW QUESTION: 64

You are using Cloud Build build to promote a Docker image to Development, Test, and Production environments. You need to ensure that the same Docker image is deployed to each of these environments.

How should you identify the Docker image in your build?

- A. Use the latest Docker image tag.
- B. Use a unique Docker image name.
- C. Use the digest of the Docker image.
- D. Use a semantic version Docker image tag.

Answer: D (LEAVE A REPLY)

Explanation/Reference:

NEW QUESTION: 65

You are configuring a continuous integration pipeline using Cloud Build to automate the deployment of new container images to Google Kubernetes Engine (GKE). The pipeline builds the application from its source code, runs unit and integration tests in separate steps, and pushes the container to Container Registry. The application runs on a Python web server.

The Dockerfile is as follows:

```
FROM python:3.7-alpine -
COPY . /app -
WORKDIR /app -
RUN pip install -r requirements.txt
CMD [ "unicorn", "-w 4", "main:app" ]
```

You notice that Cloud Build runs are taking longer than expected to complete. You want to decrease the build time. What should you do? (Choose two.)

- A. Select a virtual machine (VM) size with higher CPU for Cloud Build runs.
- B. Deploy a Container Registry on a Compute Engine VM in a VPC, and use it to store the final images.
- C. Cache the Docker image for subsequent builds using the -- cache-from argument in your build config file.
- D. Change the base image in the Dockerfile to ubuntu:latest, and install Python 3.7 using a package manager utility.
- E. Store application source code on Cloud Storage, and configure the pipeline to use gsutil to download the source code.

Answer: (SHOW ANSWER)

Explanation

<https://cloud.google.com/build/docs/optimize-builds/increase-vcpu-for-builds> By default, Cloud Build runs your builds on a standard virtual machine (VM). In addition to the standard VM, Cloud Build provides several high-CPU VM types to run builds. To increase the speed of your build, select a machine with a higher vCPU to run builds. Keep in mind that although selecting a high vCPU machine increases your build speed, it may also

increase the startup time of your build as Cloud Build only starts non-standard machines on demand.

[https://cloud.google.com/build/docs/optimize-builds/speeding-up-](https://cloud.google.com/build/docs/optimize-builds/speeding-up-builds#using_a_cached_docker_image)

[builds#using_a_cached_docker_image](https://cloud.google.com/build/docs/optimize-builds/speeding-up-builds#using_a_cached_docker_image) The easiest way to increase the speed of your Docker image build is by specifying a cached image that can be used for subsequent builds. You can specify the cached image by adding the `--cache-from` argument in your build config file, which will instruct Docker to build using that image as a cache source.

NEW QUESTION: 66

You want to re-architect a monolithic application so that it follows a microservices model. You want to accomplish this efficiently while minimizing the impact of this change to the business.

Which approach should you take?

- A. Deploy the application to Compute Engine and turn on autoscaling.
- B. Replace the application's features with appropriate microservices in phases.
- C. Refactor the monolithic application with appropriate microservices in a single effort and deploy it.
- D. Build a new application with the appropriate microservices separate from the monolith and replace it when it is complete.

Answer: C (LEAVE A REPLY)

Reference: <https://cloud.google.com/solutions/migrating-a-monolithic-app-to-microservices-gke>

NEW QUESTION: 67

You recently developed an application. You need to call the Cloud Storage API from a Compute Engine instance that doesn't have a public IP address. What should you do?

- A. Use Carrier Peering
- B. Use VPC Network Peering
- C. Use Shared VPC networks
- D. Use Private Google Access

Answer: D (LEAVE A REPLY)

Explanation

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

NEW QUESTION: 68

Your team develops services that run on Google Kubernetes Engine. You need to standardize their log data using Google-recommended practices and make the data more useful in the fewest number of steps. What should you do? (Choose two.)

- A. Create aggregated exports on application logs to BigQuery to facilitate log analytics.
- B. Create aggregated exports on application logs to Cloud Storage to facilitate log analytics.

- C.** Write log output to standard output (stdout) as single-line JSON to be ingested into Cloud Logging as structured logs.
- D.** Mandate the use of the Logging API in the application code to write structured logs to Cloud Logging.
- E.** Mandate the use of the Pub/Sub API to write structured data to Pub/Sub and create a Dataflow streaming pipeline to normalize logs and write them to BigQuery for analytics.

Answer: A,C (LEAVE A REPLY)

Explanation

https://cloud.google.com/stackdriver/docs/solutions/gke/managing-logs#best_practices

NEW QUESTION: 69

You are developing a microservice-based application that will be deployed on a Google Kubernetes Engine cluster. The application needs to read and write to a Spanner database. You want to follow security best practices while minimizing code changes. How should you configure your application to retrieve Spanner credentials?

- A.** Configure the appropriate service accounts, and use Workload Identity to run the pods.
- B.** Store the application credentials as Kubernetes Secrets, and expose them as environment variables.
- C.** Configure the appropriate routing rules, and use a VPC-native cluster to directly connect to the database.
- D.** Store the application credentials using Cloud Key Management Service, and retrieve them whenever a database connection is made.

Answer: A (LEAVE A REPLY)

Explanation

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

NEW QUESTION: 70

For this question, refer to the HipLocal case study.

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

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

Answer: D (LEAVE A REPLY)

Explanation

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

NEW QUESTION: 71

You want to upload files from an on-premises virtual machine to Google Cloud Storage as part of a data migration. These files will be consumed by Cloud DataProc Hadoop cluster in a GCP environment.

Which command should you use?

- A. `gsutil cp [LOCAL_OBJECT] gs://[DESTINATION_BUCKET_NAME]/`
- B. `gcloud cp [LOCAL_OBJECT] gs://[DESTINATION_BUCKET_NAME]/`
- C. `hadoop fs cp [LOCAL_OBJECT] gs://[DESTINATION_BUCKET_NAME]/`
- D. `gcloud dataproc cp [LOCAL_OBJECT] gs://[DESTINATION_BUCKET_NAME]/`

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

Explanation

The `gsutil cp` command allows you to copy data between your local file. storage. boto files generated by running "gsutil config"

NEW QUESTION: 72

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: ([SHOW ANSWER](#))

Explanation

NEW QUESTION: 73

You want to view the memory usage of your application deployed on Compute Engine.

What should you do?

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 74

Elizabeth would rather not answer phone calls using her iMac. Where in macOS can Elizabeth turn off iPhone Cellular Calls?

- A. Turn off iPhone Cellular Calls in iCloud preferences.
- B. Turn off iPhone Cellular Calls in System Preferences.
- C. Turn off iPhone Cellular Calls in FaceTime preferences.
- D. Turn off iPhone Cellular Calls in Messages preferences.

Answer: (SHOW ANSWER)

Explanation/Reference: <https://discussions.apple.com/thread/6836476>

NEW QUESTION: 75

What requirements are needed for iCloud to automatically back up an iOS device each day? (Choose two.)

- A. The device has two-factor authentication turned on.
- B. The device is powered off.
- C. The device is updated to latest iOS.
- D. The device is connected to a power source.
- E. The device's screen is locked.

Answer: D,E (LEAVE A REPLY)

Explanation/Reference:

Reference: <https://support.apple.com/en-us/HT203977>

NEW QUESTION: 76

Your company needs a database solution that stores customer purchase history and meets the following requirements:

Customers can query their purchase immediately after submission.

Purchases can be sorted on a variety of fields.

Distinct record formats can be stored at the same time.

Which storage option satisfies these requirements?

- A. Firestore in Native mode
- B. Cloud SQL using a SQL SELECT statement
- C. Cloud Storage using an object read
- D. Firestore in Datastore mode using a global query

Answer: (SHOW ANSWER)

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> (265 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 77

April states she would like to use the cellular network from her iPhone to access the Internet for free on her Mac. What true statement can you give April?

- A. "You must enable Personal Hotspot on your iPhone before the feature can work for free."
- B. "Personal Hotspot only works on CDMA networks."
- C. "Personal Hotspot is a great way to access the Internet for free!"
- D. "Personal Hotspot can come with additional charges. You should contact your carrier."

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

NEW QUESTION: 78

You are planning to deploy hundreds of microservices in your Google Kubernetes Engine (GKE) cluster. How should you secure communication between the microservices on GKE using a managed service?

- A. Use global HTTP(S) Load Balancing with managed SSL certificates to protect your services
- B. Deploy open source Istio in your GKE cluster, and enable mTLS in your Service Mesh
- C. Install cert-manager on GKE to automatically renew the SSL certificates.
- D. Install Anthos Service Mesh, and enable mTLS in your Service Mesh.

Answer: (SHOW ANSWER)

Explanation

https://cloud.google.com/service-mesh/docs/overview#security_benefits

- Ensures encryption in transit. Using mTLS for authentication also ensures that all TCP communications are encrypted in transit.

NEW QUESTION: 79

Your team develops services that run on Google Kubernetes Engine. Your team's code is stored in Cloud Source Repositories. You need to quickly identify bugs in the code before it is deployed to production. You want to invest in automation to improve developer feedback and make the process as efficient as possible.

What should you do?

- A. Use Spinnaker to automate building container images from code based on Git tags.
- B. Use Cloud Build to automate building container images from code based on Git tags.
- C. Use Spinnaker to automate deploying container images to the production environment.
- D. Use Cloud Build to automate building container images from code based on forked versions.

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

Reference: <https://spinnaker.io/docs/guides/tutorials/codelabs/kubernetes-v2-source-to-prod/>

NEW QUESTION: 80

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: 81

You work for an organization that manages an ecommerce site. Your application is deployed behind a global HTTP(S) load balancer. You need to test a new product recommendation algorithm. You plan to use A/B testing to determine the new algorithm's effect on sales in a randomized way. How should you test this feature?

- A. Split traffic between versions using weights.
- B. Enable the new recommendation feature flag on a single instance.
- C. Mirror traffic to the new version of your application.
- D. Use HTTP header-based routing.

Answer: ([SHOW ANSWER](#))

Explanation

https://cloud.google.com/load-balancing/docs/https/traffic-management-global#traffic_actions_weight-based_tra

Deploying a new version of an existing production service generally incurs some risk. Even if your tests pass in staging, you probably don't want to subject 100% of your users to the new version immediately. With traffic management, you can define percentage-based traffic splits across multiple backend services.

For example, you can send 95% of the traffic to the previous version of your service and 5% to the new version of your service. After you've validated that the new production version works as expected, you can gradually shift the percentages until 100% of the traffic reaches the new version of your service. Traffic splitting is typically used for deploying new versions, A/B testing, service migration, and similar processes.

https://cloud.google.com/traffic-director/docs/advanced-traffic-management#weight-based_traffic_splitting_for_

https://cloud.google.com/architecture/implementing-deployment-and-testing-strategies-on-gke#split_the_traffic_
https://cloud.google.com/load-balancing/docs/https/traffic-management-global#traffic_actions_weight-based_tra

NEW QUESTION: 82

You are writing a Compute Engine hosted application in project A that needs to securely authenticate to a Cloud Pub/Sub topic in project B.

What should you do?

- A.** Configure the instances with a service account owned by project B. Add the service account as a Cloud Pub/Sub publisher to project A.
- B.** Configure Application Default Credentials to use the private key of a service account owned by project B.

Add the service account as a Cloud Pub/Sub publisher to project A.

- C.** Configure the instances with a service account owned by project A. Add the service account as a publisher on the topic.
- D.** Configure Application Default Credentials to use the private key of a service account owned by project A.

Add the service account as a publisher on the topic

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 83

You want to view the memory usage of your application deployed on Compute Engine.

What should you do?

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

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

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

NEW QUESTION: 84

Your development team has been asked to refactor an existing monolithic application into a set of composable microservices. Which design aspects should you implement for the new application? (Choose two.)

- A.** Create an API contract agreement between the microservice implementation and microservice caller.
- B.** Require asynchronous communications between all microservice implementations and microservice callers.

- C. Develop the microservice code in the same programming language used by the microservice caller.
- D. Implement a versioning scheme to permit future changes that could be incompatible with the current interface.
- E. Ensure that sufficient instances of the microservice are running to accommodate the performance requirements.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 85

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

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 86

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

Which configuration should they choose?

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

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

NEW QUESTION: 87

You manage an ecommerce application that processes purchases from customers who can subsequently cancel or change those purchases. You discover that order volumes are highly variable and the backend order-processing system can only process one request at a time. You want to ensure seamless performance for customers regardless of usage

volume. It is crucial that customers' order update requests are performed in the sequence in which they were generated. What should you do?

- A. Send the purchase and change requests over WebSockets to the backend.
- B. Send the purchase and change requests as REST requests to the backend.
- C. Use a Pub/Sub subscriber in pull mode and use a data store to manage ordering.
- D. Use a Pub/Sub subscriber in push mode and use a data store to manage ordering.

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

Explanation

<https://cloud.google.com/pubsub/docs/pull>

NEW QUESTION: 88

You are designing a deployment technique for your new applications on Google Cloud. As part of your deployment planning, you want to use live traffic to gather performance metrics for both new and existing applications. You need to test against the full production load prior to launch. What should you do?

- A. Use A/B testing with traffic mirroring during deployment
- B. Use rolling updates deployment
- C. Use blue/green deployment
- D. Use canary deployment

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

NEW QUESTION: 89

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: (SHOW ANSWER)

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

NEW QUESTION: 90

Your code is running on Cloud Functions in project A.

It is supposed to write an object in a Cloud Storage bucket owned by project B.

However, the write call is failing with the error "403 Forbidden".

What should you do to correct the problem?

- A. Enable the Cloud Storage API in project B.

B. Grant your user account the roles/storage.objectCreator role for the Cloud Storage bucket.

C. Grant your user account the roles/iam.serviceAccountUser role for the service-PROJECTA@gcf-adminrobot.

iam.gserviceaccount.com service account.

D. Grant the service-PROJECTA@gcf-admin-robot.iam.gserviceaccount.com service account the roles/ storage.objectCreator role for the Cloud Storage bucket.

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

NEW QUESTION: 91

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 cluster with at least 1 Spanner node.

B. Create a cluster with at least 3 Spanner nodes.

C. Create a minimum of two Cloud Spanner instances in separate regions with at least one node.

D. Create a Cloud Dataflow pipeline to replicate data across different databases.

E. Create a multi-regional Cloud Spanner instance with "nam-asia-eur1" configuration.

F. Create a multi-regional Cloud Spanner instance with "nam3" configuration.

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

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> (265 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 92

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

What should you do?

A. Use HTTP signed URLs to securely provide access to the required resources.

B. Use the instance's service account Application Default Credentials to authenticate to the required resources.

- C.** Generate a P12 file from the GCP Console after the instance is deployed, and copy the credentials to the host instance before starting the application.
- D.** Commit the credential JSON file into your application's source repository, and have your CI/CD process package it with the software that is deployed to the instance.

Answer: B (LEAVE A REPLY)

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

NEW QUESTION: 93

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.** "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?"
- B.** "Calm down!"
- C.** "I can prove to you exactly how you are wrong about this issue."
- D.** "Are you sure your children didn't do something to it?"
- E.** "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."
- F.** "Is this really the same issue? Might it be a different issue?"

Answer: A,B (LEAVE A REPLY)

NEW QUESTION: 94

Your operations team has asked you to create a script that lists the Cloud Bigtable, Memorystore, and Cloud SQL databases running within a project. The script should allow users to submit a filter expression to limit the results presented. How should you retrieve the data?

- A.** Use the HBase API, Redis API, and MySQL connection to retrieve database lists. Combine the results, and then apply the filter to display the results
- B.** Use the HBase API, Redis API, and MySQL connection to retrieve database lists. Filter the results individually, and then combine them to display the results
- C.** Run `gcloud bigtable instances list`, `gcloud redis instances list`, and `gcloud sql databases list`. Use a filter within the application, and then display the results
- D.** Run `gcloud bigtable instances list`, `gcloud redis instances list`, and `gcloud sql databases list`. Use `--filter` flag with each command, and then display the results

Answer: D (LEAVE A REPLY)

Explanation

<https://cloud.google.com/sdk/gcloud/reference/topic/filters>

Most gcloud commands return a list of resources on success. By default they are pretty-printed on the standard output. The --format=NAME[ATTRIBUTES](PROJECTION) and --filter=EXPRESSION flags along with projections can be used to format and change the default output to a more meaningful result. Use the --format flag to change the default output format of a command. For details run \$ gcloud topic formats.

NEW QUESTION: 95

You have an application deployed in Google Kubernetes Engine (GKE) that reads and processes Pub/Sub messages. Each Pod handles a fixed number of messages per minute. The rate at which messages are published to the Pub/Sub topic varies considerably throughout the day and week, including occasional large batches of messages published at a single moment.

You want to scale your GKE Deployment to be able to process messages in a timely manner. What GKE feature should you use to automatically adapt your workload?

- A. Vertical Pod Autoscaler in Auto mode
- B. Vertical Pod Autoscaler in Recommendation mode
- C. Horizontal Pod Autoscaler based on an external metric
- D. Horizontal Pod Autoscaler based on resources utilization

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

Explanation

<https://kubernetes.io/docs/tasks/run-application/horizontal-pod-autoscale/>

NEW QUESTION: 96

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 compliance 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: 97

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. Transfer Appliance backed by Cloud Filestore

- C. Cloud Content Delivery Network (CDN) backed by Cloud Storage
- D. Compute Engine backed by Persistent Disk

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

NEW QUESTION: 98

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

Answer: ([SHOW ANSWER](#))

Explanation

NEW QUESTION: 99

You are developing an application hosted on Google Cloud that uses a MySQL relational database schema.

The application will have a large volume of reads and writes to the database and will require backups and ongoing capacity planning. Your team does not have time to fully manage the database but can take on small administrative tasks. How should you host the database?

- A. Configure Cloud SQL to host the database, and import the schema into Cloud SQL.
- B. Deploy MySQL from the Google Cloud Marketplace to the database using a client, and import the schema.
- C. Configure Bigtable to host the database, and import the data into Bigtable.
- D. Configure Cloud Spanner to host the database, and import the schema into Cloud Spanner.
- E. Configure Firestore to host the database, and import the data into Firestore.

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

Explanation

<https://cloud.google.com/spanner/docs/migrating-mysql-to-spanner#migration-process>

Cloud SQL: Cloud SQL is a web service that allows you to create, configure, and use relational databases that live in Google's cloud. It is a fully-managed service that maintains, manages, and administers your databases, allowing you to focus on your applications and services.

<https://cloud.google.com/sql/docs/mysql> Cloud SQL for MySQL is a fully-managed database service that helps you set up, maintain, manage, and administer your MySQL relational databases on Google Cloud Platform.

NEW QUESTION: 100

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: 101

Which of the following expresses empathy rather than sympathy?

- A.** I feel so sorry for you. I bet you feel disappointed.
- B.** I am sorry this happened to you. I can see why you may be disappointed
- C.** Oh, what a pity! This must make you feel disappointed.
- D.** Poor you! You must feel awful.

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

NEW QUESTION: 102

You are developing a new application that has the following design requirements:

Creation and changes to the application infrastructure are versioned and auditable.

The application and deployment infrastructure uses Google-managed services as much as possible.

The application runs on a serverless compute platform.

How should you design the application's architecture?

- A.** 1. Store the application and infrastructure source code in a Git repository.
2. Use Cloud Build to deploy the application infrastructure with Terraform.

3. Deploy the application to a Cloud Function as a pipeline step.

B. 1. Deploy Jenkins from the Google Cloud Marketplace, and define a continuous integration pipeline in Jenkins.

2. Configure a pipeline step to pull the application source code from a Git repository.

3. Deploy the application source code to App Engine as a pipeline step.

C. 1. Create a continuous integration pipeline on Cloud Build, and configure the pipeline to deploy the application infrastructure using Deployment Manager templates.

2. Configure a pipeline step to create a container with the latest application source code.

3. Deploy the container to a Compute Engine instance as a pipeline step.

D. 1. Deploy the application infrastructure using gcloud commands.

2. Use Cloud Build to define a continuous integration pipeline for changes to the application source code.

3. Configure a pipeline step to pull the application source code from a Git repository, and create a containerized application.

4. Deploy the new container on Cloud Run as a pipeline step.

Answer: D (LEAVE A REPLY)

Reference: <https://cloud.google.com/docs/ci-cd>

NEW QUESTION: 103

Your company's development teams want to use various open source operating systems in their Docker builds.

When images are created in published containers in your company's environment, you need to scan them for Common Vulnerabilities and Exposures (CVEs). The scanning process must not impact software development agility. You want to use managed services where possible. What should you do?

A. Enable the Vulnerability scanning setting in the Container Registry.

B. Create a Cloud Function that is triggered on a code check-in and scan the code for CVEs.

C. Disallow the use of non-commercially supported base images in your development environment.

D. Use Cloud Monitoring to review the output of Cloud Build to determine whether a vulnerable version has been used.

Answer: A (LEAVE A REPLY)

Explanation

<https://cloud.google.com/container-analysis/docs/os-overview>

NEW QUESTION: 104

Your team is responsible for maintaining an application that aggregates news articles from many different sources. Your monitoring dashboard contains publicly accessible real-time reports and runs on a Compute Engine instance as a web application. External

stakeholders and analysts need to access these reports via a secure channel without authentication. How should you configure this secure channel?

- A.** Add a public IP address to the instance. Use the service account key of the instance to encrypt the traffic.
- B.** Use Cloud Scheduler to trigger Cloud Build every hour to create an export from the reports. Store the reports in a public Cloud Storage bucket.
- C.** Add an HTTP(S) load balancer in front of the monitoring dashboard. Configure Identity-Aware Proxy to secure the communication channel.
- D.** Add an HTTP(S) load balancer in front of the monitoring dashboard. Set up a Google-managed SSL certificate on the load balancer for traffic encryption.

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

Explanation

<https://cloud.google.com/load-balancing/docs/ssl-certificates/google-managed-certs>

NEW QUESTION: 105

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

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

NEW QUESTION: 106

You want to view the memory usage of your application deployed on Compute Engine.

What should you do?

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

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

Reference:

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

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> (265 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 107

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: 108

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](#))

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

NEW QUESTION: 109

You are using Cloud Build build to promote a Docker image to Development, Test, and Production environments. You need to ensure that the same Docker image is deployed to each of these environments.

How should you identify the Docker image in your build?

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

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

NEW QUESTION: 110

You recently developed an application. You need to call the Cloud Storage API from a Compute Engine instance that doesn't have a public IP address. What should you do?

- A.** Use Carrier Peering
- B.** Use VPC Network Peering
- C.** Use Shared VPC networks
- D.** Use Private Google Access

Answer: (SHOW ANSWER)

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

NEW QUESTION: 111

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

Answer: (SHOW ANSWER)

NEW QUESTION: 112

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

Which configuration should they choose?

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

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

NEW QUESTION: 113

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.** Replace your entire monitoring platform with Stackdriver.
- 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.** Migrate some traffic back to your old platform and perform AB testing on the two platforms concurrently.

Answer: ([SHOW ANSWER](#))

Reference:

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

NEW QUESTION: 114

You have an application deployed in production. When a new version is deployed, you want to ensure that all production traffic is routed to the new version of your application.

You also want to keep the previous version deployed so that you can revert to it if there is an issue with the new version.

Which deployment strategy should you use?

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

D. Rolling deployment

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

NEW QUESTION: 115

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>

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> (265 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)