Actual exam question from Google's Professional Cloud Developer

Question #: 1

Topic #: 1

[All Professional Cloud Developer Questions]

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]/

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 2
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 3
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 4

Topic #: 1

[All Professional Cloud Developer Questions]

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

```
export NAME-load-balancer

# create network
gcloud compute networks create ${NAME}

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

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

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

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

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

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

Which commands should you run?

A. gcloud compute instances add-access-config ${NAME}-backend-instance-1

B. gcloud compute instances add-tags ${NAME}-backend-instance-1 --tags http-server

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

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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 5

Topic #: 1

[All Professional Cloud Developer Questions]

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

Which approach should you use?

A. Deploy the website on App Engine and use traffic splitting.

B. Deploy the website on App Engine as three separate services.

C. Deploy the website on Cloud Functions and use traffic splitting.

D. Deploy the website on Cloud Functions as three separate functions.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 6

Topic #: 1

[All Professional Cloud Developer Questions]

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

Which command should you use?

A. gsutil cp --project ɹ€my-gcp-projectɹ€ -r ~/local-scripts/ gcp-instance-name:~/server-scripts/ --zone ɹ€us-east1-bɹ€

B. gsutil cp --project ɹ€my-gcp-projectɹ€ -R ~/local-scripts/ gcp-instance-name:~/server-scripts/ --zone ɹ€us-east1-bɹ€

C. gcloud compute scp --project ɹ€my-gcp-projectɹ€ --recurse ~/local-scripts/ gcp-instance-name:~/server-scripts/ --zone ɹ€us-east1-bɹ€

D. gcloud compute mv --project ɹ€my-gcp-projectɹ€ --recurse ~/local-scripts/ gcp-instance-name:~/server-scripts/ --zone ɹ€us-east1-bɹ€

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 7

Topic #: 1

[All Professional Cloud Developer Questions]

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

A. Uptime check

B. Process health

C. Metric absence

D. Metric threshold

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 8
Topic #: 1
[All Professional Cloud Developer Questions]

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

A. Use the JOIN operator in SQL to combine the tables.

B. Use nested WITH statements to combine the tables.

C. Use the UNION operator in SQL to combine the tables.

D. Use the UNION ALL operator in SQL to combine the tables.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 9
Topic #: 1
[All Professional Cloud Developer Questions]

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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 10
Topic #: 1
[All Professional Cloud Developer Questions]

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

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

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

D. Create a cluster with at least 1 Spanner node.

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

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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 11

Topic #: 1

[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 12

Topic #: 1

[All Professional Cloud Developer Questions]

You are planning to deploy your application in a Google Kubernetes Engine (GKE) cluster. The application exposes an HTTP-based health check at /healthz. You want to use this health check endpoint to determine whether traffic should be routed to the pod by the load balancer.

Which code snippet should you include in your Pod configuration?

A.

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

B.

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

C.

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

D.

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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 13

Topic #: 1

[All Professional Cloud Developer Questions]

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

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

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

Which improvement should you suggest your teammate make?

A. Include multiple rows with each request.

B. Perform the inserts in parallel by creating multiple threads.

C. Write each row to a Cloud Storage object, then load into BigQuery.

D. Write each row to a Cloud Storage object in parallel, then load into BigQuery.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 14
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 15

Topic #: 1

[All Professional Cloud Developer Questions]

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 Build public Issue Tracker.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 16

Topic #: 1

[All Professional Cloud Developer Questions]

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

What should you do?

A. Install the Stackdriver Logging Agent and configure it to send the application logs.

B. Use a Stackdriver Logging Library to log directly from the application to Stackdriver Logging.

C. Provide the log file folder path in the metadata of the instance to configure it to send the application logs.

D. Change the application to log to /var/log so that its logs are automatically sent to Stackdriver Logging.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 17
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 18

Topic #: 1

[All Professional Cloud Developer Questions]

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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 19
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 20

Topic #: 1

[All Professional Cloud Developer Questions]

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

Which change should you make to the GKE Deployment object shown below?

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

A. Set the Deployment strategy to RollingUpdate with maxSurge set to 0, maxUnavailable set to 1.

B. Set the Deployment strategy to RollingUpdate with maxSurge set to 1, maxUnavailable set to 0.

C. Set the Deployment strategy to Recreate with maxSurge set to 0, maxUnavailable set to 1.

D. Set the Deployment strategy to Recreate with maxSurge set to 1, maxUnavailable set to 0.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 21
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 22
Topic #: 1
[All Professional Cloud Developer Questions]

Your company has deployed a new API to App Engine Standard environment. During testing, the API is not behaving as expected. You want to monitor the application over time to diagnose the problem within the application code without redeploying the application.
Which tool should you use?

A. Stackdriver Trace

B. Stackdriver Monitoring

C. Stackdriver Debug Snapshots

D. Stackdriver Debug Logpoints

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 23

Topic #: 1

[All Professional Cloud Developer Questions]

You want to use the Stackdriver Logging Agent to send an application's log file to Stackdriver from a Compute Engine virtual machine instance.
After installing the Stackdriver Logging Agent, what should you do first?

A. Enable the Error Reporting API on the project.

B. Grant the instance full access to all Cloud APIs.

C. Configure the application log file as a custom source.

D. Create a Stackdriver Logs Export Sink with a filter that matches the application's log entries.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 24
Topic #: 1
[All Professional Cloud Developer Questions]

Your company has a BigQuery data mart that provides analytics information to hundreds of employees. One user of wants to run jobs without interrupting important workloads. This user isn't concerned about the time it takes to run these jobs. You want to fulfill this request while minimizing cost to the company and the effort required on your part.
What should you do?

A. Ask the user to run the jobs as batch jobs.

B. Create a separate project for the user to run jobs.

C. Add the user as a job.user role in the existing project.

D. Allow the user to run jobs when important workloads are not running.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 25

Topic #: 1

[All Professional Cloud Developer Questions]

You want to notify on-call engineers about a service degradation in production while minimizing development time.

What should you do?

A. Use Cloud Function to monitor resources and raise alerts.

B. Use Cloud Pub/Sub to monitor resources and raise alerts.

C. Use Stackdriver Error Reporting to capture errors and raise alerts.

D. Use Stackdriver Monitoring to monitor resources and raise alerts.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 26
Topic #: 1
[All Professional Cloud Developer Questions]

You are writing a single-page web application with a user-interface that communicates with a third-party API for content using XMLHttpRequest. The data displayed on the UI by the API results is less critical than other data displayed on the same web page, so it is acceptable for some requests to not have the API data displayed in the UI. However, calls made to the API should not delay rendering of other parts of the user interface. You want your application to perform well when the API response is an error or a timeout.
What should you do?

A. Set the asynchronous option for your requests to the API to false and omit the widget displaying the API results when a timeout or error is encountered.

B. Set the asynchronous option for your request to the API to true and omit the widget displaying the API results when a timeout or error is encountered.

C. Catch timeout or error exceptions from the API call and keep trying with exponential backoff until the API response is successful.

D. Catch timeout or error exceptions from the API call and display the error response in the UI widget.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 27

Topic #: 1

[All Professional Cloud Developer Questions]

You are creating a web application that runs in a Compute Engine instance and writes a file to any user's Google Drive. You need to configure the application to authenticate to the Google Drive API. What should you do?

A. Use an OAuth Client ID that uses the https://www.googleapis.com/auth/drive.file scope to obtain an access token for each user.

B. Use an OAuth Client ID with delegated domain-wide authority.

C. Use the App Engine service account and https://www.googleapis.com/auth/drive.file scope to generate a signed JSON Web Token (JWT).

D. Use the App Engine service account with delegated domain-wide authority.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 28

Topic #: 1

[All Professional Cloud Developer Questions]

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

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

The command fails with the error:

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

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

A. Request additional GKE quota in the GCP Console.

B. Request additional Compute Engine quota in the GCP Console.

C. Open a support case to request additional GKE quota.

D. Decouple services in the cluster, and rewrite new clusters to function with fewer cores.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 29
Topic #: 1
[All Professional Cloud Developer Questions]

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

A. A linked list

B. A hash table

C. A two-dimensional array

D. A comma-delimited string

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 30
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 31
Topic #: 1
[All Professional Cloud Developer Questions]

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 - heath checks

D. Managed instance group - heath checks

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 32
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 33

Topic #: 1

[All Professional Cloud Developer Questions]

Your application is controlled by a managed instance group. You want to share a large read-only data set between all the instances in the managed instance group. You want to ensure that each instance can start quickly and can access the data set via its filesystem with very low latency. You also want to minimize the total cost of the solution.

What should you do?

A. Move the data to a Cloud Storage bucket, and mount the bucket on the filesystem using Cloud Storage FUSE.

B. Move the data to a Cloud Storage bucket, and copy the data to the boot disk of the instance via a startup script.

C. Move the data to a Compute Engine persistent disk, and attach the disk in read-only mode to multiple Compute Engine virtual machine instances.

D. Move the data to a Compute Engine persistent disk, take a snapshot, create multiple disks from the snapshot, and attach each disk to its own instance.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 34
Topic #: 1
[All Professional Cloud Developer Questions]

You are developing an HTTP API hosted on a Compute Engine virtual machine instance that needs to be invoked by multiple clients within the same Virtual Private Cloud (VPC). You want clients to be able to get the IP address of the service.
What should you do?

A. Reserve a static external IP address and assign it to an HTTP(S) load balancing service's forwarding rule. Clients should use this IP address to connect to the service.

B. Reserve a static external IP address and assign it to an HTTP(S) load balancing service's forwarding rule. Then, define an A record in Cloud DNS. Clients should use the name of the A record to connect to the service.

C. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url https://[INSTANCE_NAME].[ZONE].c.[PROJECT_ID].internal/.

D. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url https://[API_NAME]/[API_VERSION]/.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 35

Topic #: 1

[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 36
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 37
Topic #: 1
[All Professional Cloud Developer Questions]

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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 38

Topic #: 1

[All Professional Cloud Developer Questions]

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. Compute Engine backed by Persistent Disk

C. Transfer Appliance backed by Cloud Filestore

D. Cloud Content Delivery Network (CDN) backed by Cloud Storage

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 39
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 40
Topic #: 1
[All Professional Cloud Developer Questions]

You are planning to deploy your application in a Google Kubernetes Engine (GKE) cluster. Your application can scale horizontally, and each instance of your application needs to have a stable network identity and its own persistent disk.
Which GKE object should you use?

A. Deployment

B. StatefulSet

C. ReplicaSet

D. ReplicaController

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 41
Topic #: 1
[All Professional Cloud Developer Questions]

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

A. Add RUN commands in the Dockerfile to execute unit and integration tests.

B. Create a Cloud Build build config file with a single build step to compile unit and integration tests.

C. Create a Cloud Build build config file that will spawn a separate cloud build pipeline for unit and integration tests.

D. Create a Cloud Build build config file with separate cloud builder steps to compile and execute unit and integration tests.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 42

Topic #: 1

[All Professional Cloud Developer Questions]

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

B. Grant your user account the roles/iam.serviceAccountUser role for the service-PROJECTA@gcf-admin-robot.iam.gserviceaccount.com service account.

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

D. Enable the Cloud Storage API in project B.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 43
Topic #: 1
[All Professional Cloud Developer Questions]

Case study -
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

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

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

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

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

Existing Technical Environment -
HipLocal's environment is a mix of on-premises hardware and infrastructure running in Google Cloud Platform. The HipLocal team understands their application well, but has limited experience in global scale applications. Their existing technical environment is as follows:
* Existing APIs run on Compute Engine virtual machine instances hosted in GCP.
* State is stored in a single instance MySQL database in GCP.
* Data is exported to an on-premises Teradata/Vertica data warehouse.
* Data analytics is performed in an on-premises Hadoop environment.
* The application has no logging.
* There are basic indicators of uptime; alerts are frequently fired when the APIs are unresponsive.

Business Requirements -
HipLocal's investors want to expand their footprint and support the increase in demand they are seeing. Their requirements are:
* Expand availability of the application to new regions.
* Increase the number of concurrent users that can be supported.
* Ensure a consistent experience for users when they travel to different regions.
* Obtain user activity metrics to better understand how to monetize their product.
* Ensure compliance with regulations in the new regions (for example, GDPR).
* Reduce infrastructure management time and cost.
* Adopt the Google-recommended practices for cloud computing.

Technical Requirements -
* The application and backend must provide usage metrics and monitoring.
* APIs require strong authentication and authorization.
* Logging must be increased, and data should be stored in a cloud analytics platform.
* Move to serverless architecture to facilitate elastic scaling.
* Provide authorized access to internal apps in a secure manner.
HipLocal's .net-based auth service fails under intermittent load.
What should they do?

A. Use App Engine for autoscaling.

B. Use Cloud Functions for autoscaling.

C. Use a Compute Engine cluster for the service.

D. Use a dedicated Compute Engine virtual machine instance for the service.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 44
Topic #: 1
[All Professional Cloud Developer Questions]

Case study -
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

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

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

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

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

Existing Technical Environment -
HipLocal's environment is a mix of on-premises hardware and infrastructure running in Google Cloud Platform. The HipLocal team understands their application well, but has limited experience in global scale applications. Their existing technical environment is as follows:
* Existing APIs run on Compute Engine virtual machine instances hosted in GCP.
* State is stored in a single instance MySQL database in GCP.
* Data is exported to an on-premises Teradata/Vertica data warehouse.
* Data analytics is performed in an on-premises Hadoop environment.
* The application has no logging.
* There are basic indicators of uptime; alerts are frequently fired when the APIs are unresponsive.

Business Requirements -
HipLocal's investors want to expand their footprint and support the increase in demand they are seeing. Their requirements are:
* Expand availability of the application to new regions.
* Increase the number of concurrent users that can be supported.
* Ensure a consistent experience for users when they travel to different regions.
* Obtain user activity metrics to better understand how to monetize their product.
* Ensure compliance with regulations in the new regions (for example, GDPR).
* Reduce infrastructure management time and cost.
* Adopt the Google-recommended practices for cloud computing.

Technical Requirements -
* The application and backend must provide usage metrics and monitoring.
* APIs require strong authentication and authorization.
* Logging must be increased, and data should be stored in a cloud analytics platform.
* Move to serverless architecture to facilitate elastic scaling.
* Provide authorized access to internal apps in a secure manner.
HipLocal's APIs are having occasional application failures. They want to collect application information specifically to troubleshoot the issue. What should they do?

A. Take frequent snapshots of the virtual machines.

B. Install the Cloud Logging agent on the virtual machines.

C. Install the Cloud Monitoring agent on the virtual machines.

D. Use Cloud Trace to look for performance bottlenecks.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 45
Topic #: 1
[All Professional Cloud Developer Questions]

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 has connected their Hadoop infrastructure to GCP using Cloud Interconnect in order to query data stored on persistent disks.
Which IP strategy should they use?

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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 46
Topic #: 1
[All Professional Cloud Developer Questions]

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.
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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 47
Topic #: 1
[All Professional Cloud Developer Questions]

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 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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 48
Topic #: 1
[All Professional Cloud Developer Questions]

Case study -
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

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

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

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

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

Existing Technical Environment -
HipLocal's environment is a mix of on-premises hardware and infrastructure running in Google Cloud Platform. The HipLocal team understands their application well, but has limited experience in global scale applications. Their existing technical environment is as follows:
* Existing APIs run on Compute Engine virtual machine instances hosted in GCP.
* State is stored in a single instance MySQL database in GCP.
* Data is exported to an on-premises Teradata/Vertica data warehouse.
* Data analytics is performed in an on-premises Hadoop environment.
* The application has no logging.
* There are basic indicators of uptime; alerts are frequently fired when the APIs are unresponsive.

Business Requirements -
HipLocal's investors want to expand their footprint and support the increase in demand they are seeing. Their requirements are:
* Expand availability of the application to new regions.
* Increase the number of concurrent users that can be supported.
* Ensure a consistent experience for users when they travel to different regions.
* Obtain user activity metrics to better understand how to monetize their product.
* Ensure compliance with regulations in the new regions (for example, GDPR).
* Reduce infrastructure management time and cost.
* Adopt the Google-recommended practices for cloud computing.

Technical Requirements -
* The application and backend must provide usage metrics and monitoring.
* APIs require strong authentication and authorization.
* Logging must be increased, and data should be stored in a cloud analytics platform.
* Move to serverless architecture to facilitate elastic scaling.
* Provide authorized access to internal apps in a secure manner.
In order to meet their business requirements, how should HipLocal store their application state?

A. Use local SSDs to store state.

B. Put a memcache layer in front of MySQL.

C. Move the state storage to Cloud Spanner.

D. Replace the MySQL instance with Cloud SQL.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 49
Topic #: 1
[All Professional Cloud Developer Questions]

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.
Which service should HipLocal use for their public APIs?

A. Cloud Armor

B. Cloud Functions

C. Cloud Endpoints

D. Shielded Virtual Machines

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 50
Topic #: 1
[All Professional Cloud Developer Questions]

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 improve the resilience of their MySQL deployment, while also meeting their business and technical requirements.
Which configuration should they choose?

A. Use the current single instance MySQL on Compute Engine and several read-only MySQL servers on Compute Engine.

B. Use the current single instance MySQL on Compute Engine, and replicate the data to Cloud SQL in an external master configuration.

C. Replace the current single instance MySQL instance with Cloud SQL, and configure high availability.

D. Replace the current single instance MySQL instance with Cloud SQL, and Google provides redundancy without further configuration.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 51
Topic #: 1
[All Professional Cloud Developer Questions]

Your application is running in multiple Google Kubernetes Engine clusters. It is managed by a Deployment in each cluster. The Deployment has created multiple replicas of your Pod in each cluster. You want to view the logs sent to stdout for all of the replicas in your Deployment in all clusters.
Which command should you use?

A. kubectl logs [PARAM]

B. gcloud logging read [PARAM]

C. kubectl exec ɹ€"it [PARAM] journalctl

D. gcloud compute ssh [PARAM] ɹ€"-command= ɹ€sudo journalctlɹ€

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 52
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 53

Topic #: 1

[All Professional Cloud Developer Questions]

You are designing a schema for a table that will be moved from MySQL to Cloud Bigtable. The MySQL table is as follows:

```
AccountActivity
(
Account_id int,
Event_timestamp datetime,
Transaction_type string,
Amount numeric(18, 4)
) primary key (Account_id, Event_timestamp)
```

How should you design a row key for Cloud Bigtable for this table?

A. Set Account_id as a key.

B. Set Account_id_Event_timestamp as a key.

C. Set Event_timestamp_Account_id as a key.

D. Set Event_timestamp as a key.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 54

Topic #: 1

[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 55
Topic #: 1
[All Professional Cloud Developer Questions]

You have an analytics application that runs hundreds of queries on BigQuery every few minutes using BigQuery API. You want to find out how much time these queries take to execute.
What should you do?

A. Use Stackdriver Monitoring to plot slot usage.

B. Use Stackdriver Trace to plot API execution time.

C. Use Stackdriver Trace to plot query execution time.

D. Use Stackdriver Monitoring to plot query execution times.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 56
Topic #: 1
[All Professional Cloud Developer Questions]

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

A. Create a table named Customers. Add an Array field in a table that will hold phone numbers for the customer.

B. Create a table named Customers. Create a table named Phones. Add a CustomerId field in the Phones table to find the CustomerId from a phone number.

C. Create a table named Customers. Add an Array field in a table that will hold phone numbers for the customer. Create a secondary index on the Array field.

D. Create a table named Customers as a parent table. Create a table named Phones, and interleave this table into the Customer table. Create an index on the phone number field in the Phones table.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 57
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 58

Topic #: 1

[All Professional Cloud Developer Questions]

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

A. Cloud Amor

B. Stackdriver Debugger

C. Cloud Security Scanner

D. Stackdriver Error Reporting

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 59
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 60
Topic #: 1
[All Professional Cloud Developer Questions]

Your application is built as a custom machine image. You have multiple unique deployments of the machine image. Each deployment is a separate managed instance group with its own template. Each deployment requires a unique set of configuration values. You want to provide these unique values to each deployment but use the same custom machine image in all deployments. You want to use out-of-the-box features of Compute Engine.
What should you do?

A. Place the unique configuration values in the persistent disk.

B. Place the unique configuration values in a Cloud Bigtable table.

C. Place the unique configuration values in the instance template startup script.

D. Place the unique configuration values in the instance template instance metadata.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 61
Topic #: 1
[All Professional Cloud Developer Questions]

Your application performs well when tested locally, but it runs significantly slower after you deploy it to a Compute Engine instance. You need to diagnose the problem. What should you do?
What should you do?

A. File a ticket with Cloud Support indicating that the application performs faster locally.

B. Use Cloud Debugger snapshots to look at a point-in-time execution of the application.

C. Use Cloud Profiler to determine which functions within the application take the longest amount of time.

D. Add logging commands to the application and use Cloud Logging to check where the latency problem occurs.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 62

Topic #: 1

[All Professional Cloud Developer Questions]

You have an application running in App Engine. Your application is instrumented with Stackdriver Trace. The /product-details request reports details about four known unique products at /sku-details as shown below. You want to reduce the time it takes for the request to complete.

What should you do?

A. Increase the size of the instance class.

B. Change the Persistent Disk type to SSD.

C. Change /product-details to perform the requests in parallel.

D. Store the /sku-details information in a database, and replace the webservice call with a database query.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 63
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 64
Topic #: 1
[All Professional Cloud Developer Questions]

Your App Engine standard configuration is as follows:
service: production
instance_class: B1
You want to limit the application to 5 instances.
Which code snippet should you include in your configuration?

A. manual_scaling: instances: 5 min_pending_latency: 30ms

B. manual_scaling: max_instances: 5 idle_timeout: 10m

C. basic_scaling: instances: 5 min_pending_latency: 30ms

D. basic_scaling: max_instances: 5 idle_timeout: 10m

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 65

Topic #: 1

[All Professional Cloud Developer Questions]

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

A. Grant the service account BigQuery Data Viewer and BigQuery Job User roles.

B. Grant the service account BigQuery Data Editor and BigQuery Data Viewer roles.

C. Create a view in BigQuery from the SQL query and SELECT* from the view in the CLI.

D. Create a new dataset in BigQuery, and copy the source table to the new dataset Query the new dataset and table from the CLI.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 66
Topic #: 1
[All Professional Cloud Developer Questions]

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. Reboot the machine.

B. Enable and check the serial port output.

C. Delete the machine and create a new one.

D. Take a snapshot of the disk and attach it to a new machine.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 67
Topic #: 1
[All Professional Cloud Developer Questions]

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. Add the label ג€AUTOSCALEג€ to the instance group template.

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

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

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

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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 68
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 69
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 70
Topic #: 1
[All Professional Cloud Developer Questions]

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

A. Configure a GKE Ingress resource.

B. Configure a GKE Service resource.

C. Configure a GKE Ingress resource with type: LoadBalancer.

D. Configure a GKE Service resource with type: LoadBalancer.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 71
Topic #: 1
[All Professional Cloud Developer Questions]

Your company is planning to migrate their on-premises Hadoop environment to the cloud. Increasing storage cost and maintenance of data stored in HDFS is a major concern for your company. You also want to make minimal changes to existing data analytics jobs and existing architecture.
How should you proceed with the migration?

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

B. Create Compute Engine instances with HDD instead of SSD to save costs. Then perform a full migration of your existing environment into the new one in Compute Engine instances.

C. Create a Cloud Dataproc cluster on Google Cloud Platform, and then migrate your Hadoop environment to the new Cloud Dataproc cluster. Move your HDFS data into larger HDD disks to save on storage costs.

D. Create a Cloud Dataproc cluster on Google Cloud Platform, and then migrate your Hadoop code objects to the new cluster. Move your data to Cloud Storage and leverage the Cloud Dataproc connector to run jobs on that data.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 72

Topic #: 1

[All Professional Cloud Developer Questions]

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 ɪ€"o output.json gs://my-bucket

B. gsutil perfdiag ɪ€"o output.json gs://my-bucket

C. gcloud compute scp example-instance:~/test-data ɪ€"o output.json gs://my-bucket

D. gcloud services test ɪ€"o output.json gs://my-bucket

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 73
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 74
Topic #: 1
[All Professional Cloud Developer Questions]

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 to 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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 75

Topic #: 1

[All Professional Cloud Developer Questions]

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

Which improvement should you suggest your teammate make?

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

A. Get the entity with an ancestor query.

B. Get and put the entity in a transaction.

C. Use a strongly consistent transactional database.

D. Don't return the account entity from the function.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 76
Topic #: 1
[All Professional Cloud Developer Questions]

Your company stores their source code in a Cloud Source Repositories repository. Your company wants to build and test their code on each source code commit to the repository and requires a solution that is managed and has minimal operations overhead.
Which method should they use?

A. Use Cloud Build with a trigger configured for each source code commit.

B. Use Jenkins deployed via the Google Cloud Platform Marketplace, configured to watch for source code commits.

C. Use a Compute Engine virtual machine instance with an open source continuous integration tool, configured to watch for source code commits.

D. Use a source code commit trigger to push a message to a Cloud Pub/Sub topic that triggers an App Engine service to build the source code.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 77
Topic #: 1
[All Professional Cloud Developer Questions]

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 the instances with a service account owned by project A. Add the service account as a publisher on the topic.

C. 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.

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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 78
Topic #: 1
[All Professional Cloud Developer Questions]

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

A. Enable Cloud Identity-Aware Proxy on the HTTP(s) load balancer and restrict access to a Google Group containing users in the finance department. Verify the provided JSON Web Token within the application.

B. Enable Cloud Identity-Aware Proxy on the HTTP(s) load balancer and restrict access to a Google Group containing users in the finance department. Issue client-side certificates to everybody in the finance team and verify the certificates in the application.

C. Configure Cloud Armor Security Policies to restrict access to only corporate IP address ranges. Verify the provided JSON Web Token within the application.

D. Configure Cloud Armor Security Policies to restrict access to only corporate IP address ranges. Issue client side certificates to everybody in the finance team and verify the certificates in the application.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 79

Topic #: 1

[All Professional Cloud Developer Questions]

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

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

A. Use Zipkin collector to gather data.

B. Use Fluentd agent to gather data.

C. Use Stackdriver Trace to generate reports.

D. Use Stackdriver Debugger to generate report.

E. Use Stackdriver Profiler to generate report.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 80
Topic #: 1
[All Professional Cloud Developer Questions]

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.
Which database should HipLocal use for storing user activity?

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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 81
Topic #: 1
[All Professional Cloud Developer Questions]

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 is configuring their access controls.
Which firewall configuration should they implement?

A. Block all traffic on port 443.

B. Allow all traffic into the network.

C. Allow traffic on port 443 for a specific tag.

D. Allow all traffic on port 443 into the network.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 82
Topic #: 1
[All Professional Cloud Developer Questions]

Case study -
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

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

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

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

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

Existing Technical Environment -
HipLocal's environment is a mix of on-premises hardware and infrastructure running in Google Cloud Platform. The HipLocal team understands their application well, but has limited experience in global scale applications. Their existing technical environment is as follows:
* Existing APIs run on Compute Engine virtual machine instances hosted in GCP.
* State is stored in a single instance MySQL database in GCP.
* Data is exported to an on-premises Teradata/Vertica data warehouse.
* Data analytics is performed in an on-premises Hadoop environment.
* The application has no logging.
* There are basic indicators of uptime; alerts are frequently fired when the APIs are unresponsive.

Business Requirements -
HipLocal's investors want to expand their footprint and support the increase in demand they are seeing. Their requirements are:
* Expand availability of the application to new regions.
* Increase the number of concurrent users that can be supported.
* Ensure a consistent experience for users when they travel to different regions.
* Obtain user activity metrics to better understand how to monetize their product.
* Ensure compliance with regulations in the new regions (for example, GDPR).
* Reduce infrastructure management time and cost.
* Adopt the Google-recommended practices for cloud computing.

Technical Requirements -
* The application and backend must provide usage metrics and monitoring.
* APIs require strong authentication and authorization.
* Logging must be increased, and data should be stored in a cloud analytics platform.
* Move to serverless architecture to facilitate elastic scaling.
* Provide authorized access to internal apps in a secure manner.
HipLocal's data science team wants to analyze user reviews.
How should they prepare the data?

A. Use the Cloud Data Loss Prevention API for redaction of the review dataset.

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

C. Use the Cloud Natural Language Processing API for redaction of the review dataset.

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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 83
Topic #: 1
[All Professional Cloud Developer Questions]

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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 84
Topic #: 1
[All Professional Cloud Developer Questions]

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. Rolling deployment

D. Recreate deployment

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 85

Topic #: 1

[All Professional Cloud Developer Questions]

You are porting an existing Apache/MySQL/PHP application stack from a single machine to Google Kubernetes Engine. You need to determine how to containerize the application. Your approach should follow Google-recommended best practices for availability. What should you do?

A. Package each component in a separate container. Implement readiness and liveness probes.

B. Package the application in a single container. Use a process management tool to manage each component.

C. Package each component in a separate container. Use a script to orchestrate the launch of the components.

D. Package the application in a single container. Use a bash script as an entrypoint to the container, and then spawn each component as a background job.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 86

Topic #: 1

[All Professional Cloud Developer Questions]

You are developing an application that will be launched on Compute Engine instances into multiple distinct projects, each corresponding to the environments in your software development process (development, QA, staging, and production). The instances in each project have the same application code but a different configuration. During deployment, each instance should receive the application's configuration based on the environment it serves. You want to minimize the number of steps to configure this flow. What should you do?

A. When creating your instances, configure a startup script using the gcloud command to determine the project name that indicates the correct environment.

B. In each project, configure a metadata key ɹ€environmentɹ€ whose value is the environment it serves. Use your deployment tool to query the instance metadata and configure the application based on the ɹ€environmentɹ€ value.

C. Deploy your chosen deployment tool on an instance in each project. Use a deployment job to retrieve the appropriate configuration file from your version control system, and apply the configuration when deploying the application on each instance.

D. During each instance launch, configure an instance custom-metadata key named ɹ€environmentɹ€ whose value is the environment the instance serves. Use your deployment tool to query the instance metadata, and configure the application based on the ɹ€environmentɹ€ value.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 87

Topic #: 1

[All Professional Cloud Developer Questions]

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

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

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

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

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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 88

Topic #: 1

[All Professional Cloud Developer Questions]

You are developing an application that reads credit card data from a Pub/Sub subscription. You have written code and completed unit testing. You need to test the Pub/Sub integration before deploying to Google Cloud. What should you do?

A. Create a service to publish messages, and deploy the Pub/Sub emulator. Generate random content in the publishing service, and publish to the emulator.

B. Create a service to publish messages to your application. Collect the messages from Pub/Sub in production, and replay them through the publishing service.

C. Create a service to publish messages, and deploy the Pub/Sub emulator. Collect the messages from Pub/Sub in production, and publish them to the emulator.

D. Create a service to publish messages, and deploy the Pub/Sub emulator. Publish a standard set of testing messages from the publishing service to the emulator.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 89
Topic #: 1
[All Professional Cloud Developer Questions]

You are designing an application that will subscribe to and receive messages from a single Pub/Sub topic and insert corresponding rows into a database. Your application runs on Linux and leverages preemptible virtual machines to reduce costs. You need to create a shutdown script that will initiate a graceful shutdown. What should you do?

A. Write a shutdown script that uses inter-process signals to notify the application process to disconnect from the database.

B. Write a shutdown script that broadcasts a message to all signed-in users that the Compute Engine instance is going down and instructs them to save current work and sign out.

C. Write a shutdown script that writes a file in a location that is being polled by the application once every five minutes. After the file is read, the application disconnects from the database.

D. Write a shutdown script that publishes a message to the Pub/Sub topic announcing that a shutdown is in progress. After the application reads the message, it disconnects from the database.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 90

Topic #: 1

[All Professional Cloud Developer Questions]

You work for a web development team at a small startup. Your team is developing a Node.js application using Google Cloud services, including Cloud Storage and Cloud Build. The team uses a Git repository for version control. Your manager calls you over the weekend and instructs you to make an emergency update to one of the company's websites, and you're the only developer available. You need to access Google Cloud to make the update, but you don't have your work laptop. You are not allowed to store source code locally on a non-corporate computer. How should you set up your developer environment?

A. Use a text editor and the Git command line to send your source code updates as pull requests from a public computer.

B. Use a text editor and the Git command line to send your source code updates as pull requests from a virtual machine running on a public computer.

C. Use Cloud Shell and the built-in code editor for development. Send your source code updates as pull requests.

D. Use a Cloud Storage bucket to store the source code that you need to edit. Mount the bucket to a public computer as a drive, and use a code editor to update the code. Turn on versioning for the bucket, and point it to the team's Git repository.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 91
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 92

Topic #: 1

[All Professional Cloud Developer Questions]

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 canary deployment

B. Use blue/green deployment

C. Use rolling updates deployment

D. Use A/B testing with traffic mirroring during deployment

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 93

Topic #: 1

[All Professional Cloud Developer Questions]

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

A. Retry the failures in batch after a set number of failures is logged.

B. Retry each failure at a set time interval up to a maximum number of times.

C. Retry each failure at increasing time intervals up to a maximum number of tries.

D. Retry each failure at decreasing time intervals up to a maximum number of tries.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 94

Topic #: 1

[All Professional Cloud Developer Questions]

You need to redesign the ingestion of audit events from your authentication service to allow it to handle a large increase in traffic. Currently, the audit service and the authentication system run in the same Compute Engine virtual machine. You plan to use the following Google Cloud tools in the new architecture:

- ✑ Multiple Compute Engine machines, each running an instance of the authentication service
- ✑ Multiple Compute Engine machines, each running an instance of the audit service
- ✑ Pub/Sub to send the events from the authentication services.

How should you set up the topics and subscriptions to ensure that the system can handle a large volume of messages and can scale efficiently?

A. Create one Pub/Sub topic. Create one pull subscription to allow the audit services to share the messages.

B. Create one Pub/Sub topic. Create one pull subscription per audit service instance to allow the services to share the messages.

C. Create one Pub/Sub topic. Create one push subscription with the endpoint pointing to a load balancer in front of the audit services.

D. Create one Pub/Sub topic per authentication service. Create one pull subscription per topic to be used by one audit service.

E. Create one Pub/Sub topic per authentication service. Create one push subscription per topic, with the endpoint pointing to one audit service.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 95

Topic #: 1

[All Professional Cloud Developer Questions]

You are developing a marquee stateless web application that will run on Google Cloud. The rate of the incoming user traffic is expected to be unpredictable, with no traffic on some days and large spikes on other days. You need the application to automatically scale up and down, and you need to minimize the cost associated with running the application. What should you do?

A. Build the application in Python with Firestore as the database. Deploy the application to Cloud Run.

B. Build the application in C# with Firestore as the database. Deploy the application to App Engine flexible environment.

C. Build the application in Python with CloudSQL as the database. Deploy the application to App Engine standard environment.

D. Build the application in Python with Firestore as the database. Deploy the application to a Compute Engine managed instance group with autoscaling.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 96

Topic #: 1

[All Professional Cloud Developer Questions]

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

A. Create a new service account that has Editor authority to access the resources. The deployer is given permission to get the access token.

B. Create a new service account that has a custom IAM role to access the resources. The deployer is given permission to get the access token.

C. Create a new service account that has Editor authority to access the resources. The deployer is given permission to act as the new service account.

D. Create a new service account that has a custom IAM role to access the resources. The deployer is given permission to act as the new service account.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 97

Topic #: 1

[All Professional Cloud Developer Questions]

You are a SaaS provider deploying dedicated blogging software to customers in your Google Kubernetes Engine (GKE) cluster. You want to configure a secure multi-tenant platform to ensure that each customer has access to only their own blog and can't affect the workloads of other customers. What should you do?

A. Enable Application-layer Secrets on the GKE cluster to protect the cluster.

B. Deploy a namespace per tenant and use Network Policies in each blog deployment.

C. Use GKE Audit Logging to identify malicious containers and delete them on discovery.

D. Build a custom image of the blogging software and use Binary Authorization to prevent untrusted image deployments.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 98

Topic #: 1

[All Professional Cloud Developer Questions]

You have decided to migrate your Compute Engine application to Google Kubernetes Engine. You need to build a container image and push it to Artifact Registry using Cloud Build. What should you do? (Choose two.)

A. Run gcloud builds submit in the directory that contains the application source code.

B. Run gcloud run deploy app-name --image gcr.io/$PROJECT_ID/app-name in the directory that contains the application source code.

C. Run gcloud container images add-tag gcr.io/$PROJECT_ID/app-name gcr.io/$PROJECT_ID/app-name:latest in the directory that contains the application source code.

D. In the application source directory, create a file named cloudbuild.yaml that contains the following contents:

```
steps:
- name: 'gcr.io/cloud-builders/docker'
  args: ['build', '-t', 'gcr.io/$PROJECT_ID/app-name', '.']
- name: 'gcr.io/cloud-builders/docker'
  args: ['push', 'gcr.io/$PROJECT_ID/app-name']
```

E. In the application source directory, create a file named cloudbuild.yaml that contains the following contents:

```
steps:
- name: 'gcr.io/cloud-builders/gcloud'
  args: ['app', 'deploy']
  timeout: '1600s'
```

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 99
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 100

Topic #: 1

[All Professional Cloud Developer Questions]

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.yami file to move the photos from webphotos-staging to webphotos-prod.

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

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

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

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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 101

Topic #: 1

[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 102

Topic #: 1

[All Professional Cloud Developer Questions]

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. Pipe the content of the files to the Linux Syslog daemon.

B. Install a Google version of fluentd on the Compute Engine instance.

C. Install a Google version of collectd on the Compute Engine instance.

D. Using cron, schedule a job to copy the log files to Cloud Storage once a day.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 103

Topic #: 1

[All Professional Cloud Developer Questions]

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

A. Embed the appropriate database connection string in the image. Create a different image for each environment.

B. When creating the Compute Engine instance, add a tag with the name of the database to be connected. In your application, query the Compute Engine API to pull the tags for the current instance, and use the tag to construct the appropriate database connection string.

C. When creating the Compute Engine instance, create a metadata item with a key of ɪ€DATABASEɪ€ and a value for the appropriate database connection string. In your application, read the ɪ€DATABASEɪ€ environment variable, and use the value to connect to the appropriate database.

D. When creating the Compute Engine instance, create a metadata item with a key of ɪ€DATABASEɪ€ and a value for the appropriate database connection string. In your application, query the metadata server for the ɪ€DATABASEɪ€ value, and use the value to connect to the appropriate database.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 104

Topic #: 1

[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 105
Topic #: 1
[All Professional Cloud Developer Questions]

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

A. Assign the Project Editor role.

B. Assign the Project Owner role.

C. Assign the Cloud SQL Client role.

D. Assign the Cloud SQL Editor role.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 106

Topic #: 1

[All Professional Cloud Developer Questions]

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

A. Use a Vertical Pod Autoscaler to scale the containers, and expose them via a ClusterIP Service.

B. Use a Vertical Pod Autoscaler to scale the containers, and expose them via a NodePort Service.

C. Use a Horizontal Pod Autoscaler to scale the containers, and expose them via a ClusterIP Service.

D. Use a Horizontal Pod Autoscaler to scale the containers, and expose them via a NodePort Service.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 107
Topic #: 1
[All Professional Cloud Developer Questions]

You recently migrated a monolithic application to Google Cloud by breaking it down into microservices. One of the microservices is deployed using Cloud Functions. As you modernize the application, you make a change to the API of the service that is backward-incompatible. You need to support both existing callers who use the original API and new callers who use the new API. What should you do?

A. Leave the original Cloud Function as-is and deploy a second Cloud Function with the new API. Use a load balancer to distribute calls between the versions.

B. Leave the original Cloud Function as-is and deploy a second Cloud Function that includes only the changed API. Calls are automatically routed to the correct function.

C. Leave the original Cloud Function as-is and deploy a second Cloud Function with the new API. Use Cloud Endpoints to provide an API gateway that exposes a versioned API.

D. Re-deploy the Cloud Function after making code changes to support the new API. Requests for both versions of the API are fulfilled based on a version identifier included in the call.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 108

Topic #: 1

[All Professional Cloud Developer Questions]

You are developing an application that will allow users to read and post comments on news articles. You want to configure your application to store and display user-submitted comments using Firestore. How should you design the schema to support an unknown number of comments and articles?

A. Store each comment in a subcollection of the article.

B. Add each comment to an array property on the article.

C. Store each comment in a document, and add the comment's key to an array property on the article.

D. Store each comment in a document, and add the comment's key to an array property on the user profile.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 109

Topic #: 1

[All Professional Cloud Developer Questions]

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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 110

Topic #: 1

[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 111

Topic #: 1

[All Professional Cloud Developer Questions]

You manage an application that runs in a Compute Engine instance. You also have multiple backend services executing in stand-alone Docker containers running in Compute Engine instances. The Compute Engine instances supporting the backend services are scaled by managed instance groups in multiple regions. You want your calling application to be loosely coupled. You need to be able to invoke distinct service implementations that are chosen based on the value of an HTTP header found in the request. Which Google Cloud feature should you use to invoke the backend services?

A. Traffic Director

B. Service Directory

C. Anthos Service Mesh

D. Internal HTTP(S) Load Balancing

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 112

Topic #: 1

[All Professional Cloud Developer Questions]

Your team is developing an ecommerce platform for your company. Users will log in to the website and add items to their shopping cart. Users will be automatically logged out after 30 minutes of inactivity. When users log back in, their shopping cart should be saved. How should you store users' session and shopping cart information while following Google-recommended best practices?

A. Store the session information in Pub/Sub, and store the shopping cart information in Cloud SQL.

B. Store the shopping cart information in a file on Cloud Storage where the filename is the SESSION ID.

C. Store the session and shopping cart information in a MySQL database running on multiple Compute Engine instances.

D. Store the session information in Memorystore for Redis or Memorystore for Memcached, and store the shopping cart information in Firestore.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 113
Topic #: 1
[All Professional Cloud Developer Questions]

You are designing a resource-sharing policy for applications used by different teams in a Google Kubernetes Engine cluster. You need to ensure that all applications can access the resources needed to run. What should you do? (Choose two.)

A. Specify the resource limits and requests in the object specifications.

B. Create a namespace for each team, and attach resource quotas to each namespace.

C. Create a LimitRange to specify the default compute resource requirements for each namespace.

D. Create a Kubernetes service account (KSA) for each application, and assign each KSA to the namespace.

E. Use the Anthos Policy Controller to enforce label annotations on all namespaces. Use taints and tolerations to allow resource sharing for namespaces.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 114

Topic #: 1

[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 115

Topic #: 1

[All Professional Cloud Developer Questions]

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).

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 116

Topic #: 1

[All Professional Cloud Developer Questions]

You have containerized a legacy application that stores its configuration on an NFS share. You need to deploy this application to Google Kubernetes Engine (GKE) and do not want the application serving traffic until after the configuration has been retrieved. What should you do?

A. Use the gsutil utility to copy files from within the Docker container at startup, and start the service using an ENTRYPOINT script.

B. Create a PersistentVolumeClaim on the GKE cluster. Access the configuration files from the volume, and start the service using an ENTRYPOINT script.

C. Use the COPY statement in the Dockerfile to load the configuration into the container image. Verify that the configuration is available, and start the service using an ENTRYPOINT script.

D. Add a startup script to the GKE instance group to mount the NFS share at node startup. Copy the configuration files into the container, and start the service using an ENTRYPOINT script.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 117

Topic #: 1

[All Professional Cloud Developer Questions]

Your team is developing a new application using a PostgreSQL database and Cloud Run. You are responsible for ensuring that all traffic is kept private on Google Cloud. You want to use managed services and follow Google-recommended best practices. What should you do?

A. 1. Enable Cloud SQL and Cloud Run in the same project. 2. Configure a private IP address for Cloud SQL. Enable private services access. 3. Create a Serverless VPC Access connector. 4. Configure Cloud Run to use the connector to connect to Cloud SQL.

B. 1. Install PostgreSQL on a Compute Engine virtual machine (VM), and enable Cloud Run in the same project. 2. Configure a private IP address for the VM. Enable private services access. 3. Create a Serverless VPC Access connector. 4. Configure Cloud Run to use the connector to connect to the VM hosting PostgreSQL.

C. 1. Use Cloud SQL and Cloud Run in different projects. 2. Configure a private IP address for Cloud SQL. Enable private services access. 3. Create a Serverless VPC Access connector. 4. Set up a VPN connection between the two projects. Configure Cloud Run to use the connector to connect to Cloud SQL.

D. 1. Install PostgreSQL on a Compute Engine VM, and enable Cloud Run in different projects. 2. Configure a private IP address for the VM. Enable private services access. 3. Create a Serverless VPC Access connector. 4. Set up a VPN connection between the two projects. Configure Cloud Run to use the connector to access the VM hosting PostgreSQL

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 118

Topic #: 1

[All Professional Cloud Developer Questions]

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

A. Configure the application to send the file to the client as an email attachment.

B. Generate and assign a Cloud Storage-signed URL for the file. Make the URL available for the client to download.

C. Create a temporary Cloud Storage bucket with time expiration specified, and give download permissions to the bucket. Copy the file, and send it to the client.

D. Generate the HTTP cookies with time expiration specified. If the time is valid, copy the file from the Cloud Storage bucket, and make the file available for the client to download.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 119

Topic #: 1

[All Professional Cloud Developer Questions]

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. Develop the microservice code in the same programming language used by the microservice caller.

B. Create an API contract agreement between the microservice implementation and microservice caller.

C. Require asynchronous communications between all microservice implementations and microservice callers.

D. Ensure that sufficient instances of the microservice are running to accommodate the performance requirements.

E. Implement a versioning scheme to permit future changes that could be incompatible with the current interface.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 120

Topic #: 1

[All Professional Cloud Developer Questions]

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

A. Create custom metrics from the Cloud Logging logs, and use Prometheus to import the results using the Cloud Monitoring REST API.

B. Export the Cloud Logging logs and the Prometheus metrics to Cloud Bigtable. Run a query to join the results, and analyze in Google Data Studio.

C. Export the Cloud Logging logs and stream the Prometheus metrics to BigQuery. Run a recurring query to join the results, and send notifications using Cloud Tasks.

D. Export the Prometheus metrics and use Cloud Monitoring to view them as external metrics. Configure Cloud Monitoring to create log-based metrics from the logs, and correlate them with the Prometheus data.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 121

Topic #: 1

[All Professional Cloud Developer Questions]

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

A. Migrate the Product catalog, which has integrations to the frontend and product database.

B. Migrate Payment processing, which has integrations to the frontend, order database, and third-party payment vendor.

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

D. Migrate the Shopping cart, which has integrations to the frontend, cart database, inventory system, and payment processing system.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 122

Topic #: 1

[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 123

Topic #: 1

[All Professional Cloud Developer Questions]

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

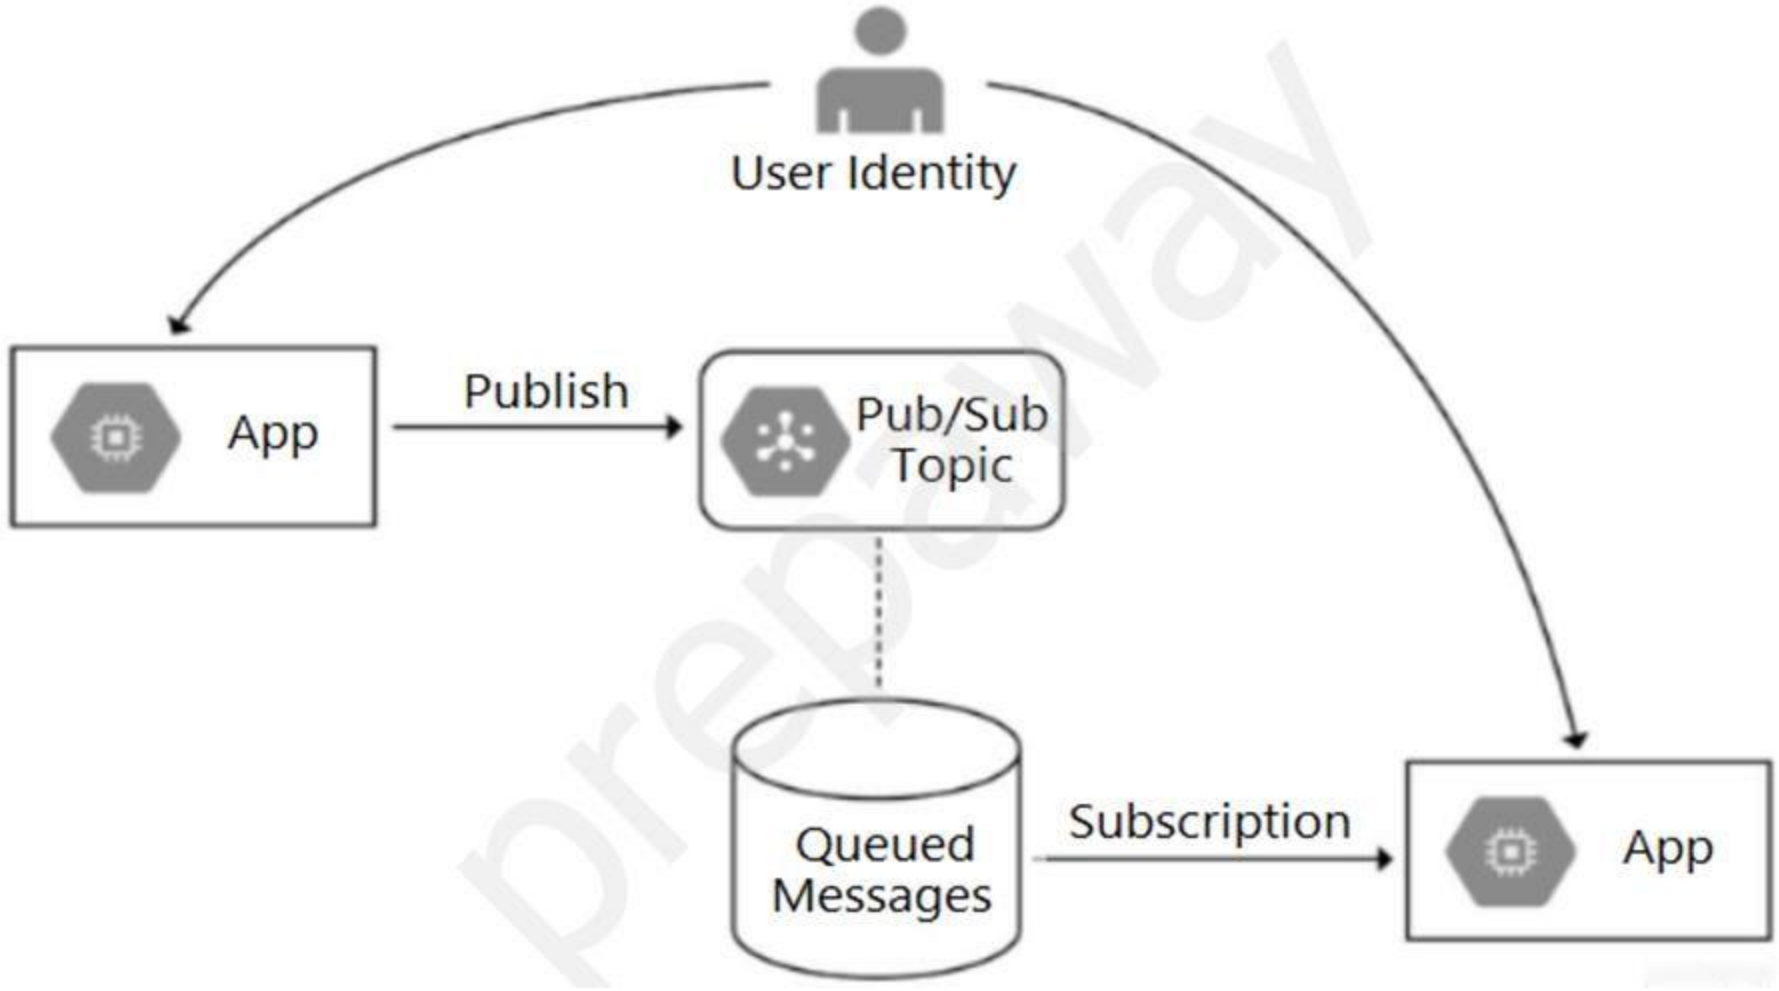

A. Bind the user identity to the pubsub.publisher and pubsub.subscriber roles at the resource level.

B. Grant the user identity the pubsub.publisher and pubsub.subscriber roles at the project level.

C. Grant the user identity a custom role that contains the pubsub.topics.create and pubsub.subscriptions.create permissions.

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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 124

Topic #: 1

[All Professional Cloud Developer Questions]

You are evaluating developer tools to help drive Google Kubernetes Engine adoption and integration with your development environment, which includes VS Code and IntelliJ. What should you do?

A. Use Cloud Code to develop applications.

B. Use the Cloud Shell integrated Code Editor to edit code and configuration files.

C. Use a Cloud Notebook instance to ingest and process data and deploy models.

D. Use Cloud Shell to manage your infrastructure and applications from the command line.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 125
Topic #: 1
[All Professional Cloud Developer Questions]

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 Serverless VPC Access connector in a different region than your App Engine instance.

C. The firewall rule allowing a connection between App Engine and Memorystore was removed during an infrastructure update by the DevOps team.

D. 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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 126

Topic #: 1

[All Professional Cloud Developer Questions]

Your team develops services that run on Google Cloud. You need to build a data processing service and will use Cloud Functions. The data to be processed by the function is sensitive. You need to ensure that invocations can only happen from authorized services and follow Google-recommended best practices for securing functions. What should you do?

A. Enable Identity-Aware Proxy in your project. Secure function access using its permissions.

B. Create a service account with the Cloud Functions Viewer role. Use that service account to invoke the function.

C. Create a service account with the Cloud Functions Invoker role. Use that service account to invoke the function.

D. Create an OAuth 2.0 client ID for your calling service in the same project as the function you want to secure. Use those credentials to invoke the function.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 127

Topic #: 1

[All Professional Cloud Developer Questions]

You are deploying your applications on Compute Engine. One of your Compute Engine instances failed to launch. What should you do? (Choose two.)

A. Determine whether your file system is corrupted.

B. Access Compute Engine as a different SSH user.

C. Troubleshoot firewall rules or routes on an instance.

D. Check whether your instance boot disk is completely full.

E. Check whether network traffic to or from your instance is being dropped.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 128

Topic #: 1

[All Professional Cloud Developer Questions]

Your web application is deployed to the corporate intranet. You need to migrate the web application to Google Cloud. The web application must be available only to company employees and accessible to employees as they travel. You need to ensure the security and accessibility of the web application while minimizing application changes. What should you do?

A. Configure the application to check authentication credentials for each HTTP(S) request to the application.

B. Configure Identity-Aware Proxy to allow employees to access the application through its public IP address.

C. Configure a Compute Engine instance that requests users to log in to their corporate account. Change the web application DNS to point to the proxy Compute Engine instance. After authenticating, the Compute Engine instance forwards requests to and from the web application.

D. Configure a Compute Engine instance that requests users to log in to their corporate account. Change the web application DNS to point to the proxy Compute Engine instance. After authenticating, the Compute Engine issues an HTTP redirect to a public IP address hosting the web application.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 129
Topic #: 1
[All Professional Cloud Developer Questions]

You have an application that uses an HTTP Cloud Function to process user activity from both desktop browser and mobile application clients. This function will serve as the endpoint for all metric submissions using HTTP POST.
Due to legacy restrictions, the function must be mapped to a domain that is separate from the domain requested by users on web or mobile sessions. The domain for the Cloud Function is https://fn.example.com. Desktop and mobile clients use the domain https://www.example.com. You need to add a header to the function's HTTP response so that only those browser and mobile sessions can submit metrics to the Cloud Function. Which response header should you add?

A. Access-Control-Allow-Origin: *

B. Access-Control-Allow-Origin: https://*.example.com

C. Access-Control-Allow-Origin: https://fn.example.com

D. Access-Control-Allow-origin: https://www.example.com

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 130
Topic #: 1
[All Professional Cloud Developer Questions]

You have an HTTP Cloud Function that is called via POST. Each submission's request body has a flat, unnested JSON structure containing numeric and text data. After the Cloud Function completes, the collected data should be immediately available for ongoing and complex analytics by many users in parallel. How should you persist the submissions?

A. Directly persist each POST request's JSON data into Datastore.

B. Transform the POST request's JSON data, and stream it into BigQuery.

C. Transform the POST request's JSON data, and store it in a regional Cloud SQL cluster.

D. Persist each POST request's JSON data as an individual file within Cloud Storage, with the file name containing the request identifier.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 131

Topic #: 1

[All Professional Cloud Developer Questions]

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. Install Istio, enable proxy injection on your application namespace, and then enable mTLS.

C. Define Trusted Network ranges within the application, and configure the applications to allow traffic only from those networks.

D. Use an automated process to request SSL Certificates for your applications from Let's Encrypt and add them to your applications.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 132

Topic #: 1

[All Professional Cloud Developer Questions]

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

A. Replace your monitoring platform with Cloud Monitoring.

B. Install the Cloud Monitoring agent on your Compute Engine instances.

C. Migrate some traffic back to your old platform. Perform A/B testing on the two platforms concurrently.

D. Use Cloud Logging and Cloud Monitoring to capture logs, monitor, and send alerts. Send them to your existing platform.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 133

Topic #: 1

[All Professional Cloud Developer Questions]

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

A. Perform a rolling deployment, and test your new application after the deployment is complete.

B. Perform A/B testing, and test your application periodically after the new tests are implemented.

C. Perform a blue/green deployment, and test your new application after the deployment is. complete.

D. Perform a canary deployment, and test your new application periodically after the new version is deployed.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 134

Topic #: 1

[All Professional Cloud Developer Questions]

You developed a JavaScript web application that needs to access Google Drive's API and obtain permission from users to store files in their Google Drives. You need to select an authorization approach for your application. What should you do?

A. Create an API key.

B. Create a SAML token.

C. Create a service account.

D. Create an OAuth Client ID.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 135

Topic #: 1

[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 136

Topic #: 1

[All Professional Cloud Developer Questions]

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 Storage using an object read

C. Cloud SQL using a SQL SELECT statement

D. Firestore in Datastore mode using a global query

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 137

Topic #: 1

[All Professional Cloud Developer Questions]

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

A. Build a test harness to generate requests and deploy it to Cloud Run. Analyze the VPC Flow Logs using Cloud Logging.

B. Create a Google Kubernetes Engine cluster running the Locust or JMeter images to dynamically generate load tests. Analyze the results using Cloud Trace.

C. Create a Cloud Task to generate a test load. Use Cloud Scheduler to run 60,000 Cloud Task transactions per minute for 10 minutes. Analyze the results using Cloud Monitoring.

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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 138

Topic #: 1

[All Professional Cloud Developer Questions]

You are using Cloud Build for your CI/CD pipeline to complete several tasks, including copying certain files to Compute Engine virtual machines. Your pipeline requires a flat file that is generated in one builder in the pipeline to be accessible by subsequent builders in the same pipeline. How should you store the file so that all the builders in the pipeline can access it?

A. Store and retrieve the file contents using Compute Engine instance metadata.

B. Output the file contents to a file in /workspace. Read from the same /workspace file in the subsequent build step.

C. Use gsutil to output the file contents to a Cloud Storage object. Read from the same object in the subsequent build step.

D. Add a build argument that runs an HTTP POST via curl to a separate web server to persist the value in one builder. Use an HTTP GET via curl from the subsequent build step to read the value.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 139

Topic #: 1

[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 140
Topic #: 1
[All Professional Cloud Developer Questions]

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 [ "gunicorn", "-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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 141
Topic #: 1
[All Professional Cloud Developer Questions]

You are building a CI/CD pipeline that consists of a version control system, Cloud Build, and Container Registry. Each time a new tag is pushed to the repository, a Cloud Build job is triggered, which runs unit tests on the new code builds a new Docker container image, and pushes it into Container Registry. The last step of your pipeline should deploy the new container to your production Google Kubernetes Engine (GKE) cluster. You need to select a tool and deployment strategy that meets the following requirements:

- Zero downtime is incurred
- Testing is fully automated
- Allows for testing before being rolled out to users
- Can quickly rollback if needed

What should you do?

A. Trigger a Spinnaker pipeline configured as an A/B test of your new code and, if it is successful, deploy the container to production.

B. Trigger a Spinnaker pipeline configured as a canary test of your new code and, if it is successful, deploy the container to production.

C. Trigger another Cloud Build job that uses the Kubernetes CLI tools to deploy your new container to your GKE cluster, where you can perform a canary test.

D. Trigger another Cloud Build job that uses the Kubernetes CLI tools to deploy your new container to your GKE cluster, where you can perform a shadow test.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 142

Topic #: 1

[All Professional Cloud Developer Questions]

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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 143

Topic #: 1

[All Professional Cloud Developer Questions]

You need to deploy a new European version of a website hosted on Google Kubernetes Engine. The current and new websites must be accessed via the same HTTP(S) load balancer's external IP address, but have different domain names. What should you do?

A. Define a new Ingress resource with a host rule matching the new domain

B. Modify the existing Ingress resource with a host rule matching the new domain

C. Create a new Service of type LoadBalancer specifying the existing IP address as the loadBalancerIP

D. Generate a new Ingress resource and specify the existing IP address as the kubernetes.io/ingress.global-static-ip-name annotation value

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 144
Topic #: 1
[All Professional Cloud Developer Questions]

You are developing a single-player mobile game backend that has unpredictable traffic patterns as users interact with the game throughout the day and night. You want to optimize costs by ensuring that you have enough resources to handle requests, but minimize over-provisioning. You also want the system to handle traffic spikes efficiently. Which compute platform should you use?

A. Cloud Run

B. Compute Engine with managed instance groups

C. Compute Engine with unmanaged instance groups

D. Google Kubernetes Engine using cluster autoscaling

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 145

Topic #: 1

[All Professional Cloud Developer Questions]

The development teams in your company want to manage resources from their local environments. You have been asked to enable developer access to each team's Google Cloud projects. You want to maximize efficiency while following Google-recommended best practices. What should you do?

A. Add the users to their projects, assign the relevant roles to the users, and then provide the users with each relevant Project ID.

B. Add the users to their projects, assign the relevant roles to the users, and then provide the users with each relevant Project Number.

C. Create groups, add the users to their groups, assign the relevant roles to the groups, and then provide the users with each relevant Project ID.

D. Create groups, add the users to their groups, assign the relevant roles to the groups, and then provide the users with each relevant Project Number.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 146

Topic #: 1

[All Professional Cloud Developer Questions]

Your company's product team has a new requirement based on customer demand to autoscale your stateless and distributed service running in a Google Kubernetes Engine (GKE) duster. You want to find a solution that minimizes changes because this feature will go live in two weeks. What should you do?

A. Deploy a Vertical Pod Autoscaler, and scale based on the CPU load.

B. Deploy a Vertical Pod Autoscaler, and scale based on a custom metric.

C. Deploy a Horizontal Pod Autoscaler, and scale based on the CPU toad.

D. Deploy a Horizontal Pod Autoscaler, and scale based on a custom metric.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 147

Topic #: 1

[All Professional Cloud Developer Questions]

Your application is composed of a set of loosely coupled services orchestrated by code executed on Compute Engine. You want your application to easily bring up new Compute Engine instances that find and use a specific version of a service. How should this be configured?

A. Define your service endpoint information as metadata that is retrieved at runtime and used to connect to the desired service.

B. Define your service endpoint information as label data that is retrieved at runtime and used to connect to the desired service.

C. Define your service endpoint information to be retrieved from an environment variable at runtime and used to connect to the desired service.

D. Define your service to use a fixed hostname and port to connect to the desired service. Replace the service at the endpoint with your new version.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 148

Topic #: 1

[All Professional Cloud Developer Questions]

You are developing a microservice-based application that will run on Google Kubernetes Engine (GKE). Some of the services need to access different Google Cloud APIs. How should you set up authentication of these services in the cluster following Google-recommended best practices? (Choose two.)

A. Use the service account attached to the GKE node.

B. Enable Workload Identity in the cluster via the gcloud command-line tool.

C. Access the Google service account keys from a secret management service.

D. Store the Google service account keys in a central secret management service.

E. Use gcloud to bind the Kubernetes service account and the Google service account using roles/iam.workloadIdentity.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 149

Topic #: 1

[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 150
Topic #: 1
[All Professional Cloud Developer Questions]

The new version of your containerized application has been tested and is ready to deploy to production on Google Kubernetes Engine. You were not able to fully load-test the new version in pre-production environments, and you need to make sure that it does not have performance problems once deployed. Your deployment must be automated. What should you do?

A. Use Cloud Load Balancing to slowly ramp up traffic between versions. Use Cloud Monitoring to look for performance issues.

B. Deploy the application via a continuous delivery pipeline using canary deployments. Use Cloud Monitoring to look for performance issues. and ramp up traffic as the metrics support it.

C. Deploy the application via a continuous delivery pipeline using blue/green deployments. Use Cloud Monitoring to look for performance issues, and launch fully when the metrics support it.

D. Deploy the application using kubectl and set the spec.updateStrategv.type to RollingUpdate. Use Cloud Monitoring to look for performance issues, and run the kubectl rollback command if there are any issues.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 151
Topic #: 1
[All Professional Cloud Developer Questions]

Users are complaining that your Cloud Run-hosted website responds too slowly during traffic spikes. You want to provide a better user experience during traffic peaks. What should you do?

A. Read application configuration and static data from the database on application startup.

B. Package application configuration and static data into the application image during build time.

C. Perform as much work as possible in the background after the response has been returned to the user.

D. Ensure that timeout exceptions and errors cause the Cloud Run instance to exit quickly so a replacement instance can be started.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 152
Topic #: 1
[All Professional Cloud Developer Questions]

You are a developer working on an internal application for payroll processing. You are building a component of the application that allows an employee to submit a timesheet, which then initiates several steps:

• An email is sent to the employee and manager, notifying them that the timesheet was submitted.
• A timesheet is sent to payroll processing for the vendor's API.
• A timesheet is sent to the data warehouse for headcount planning.

These steps are not dependent on each other and can be completed in any order. New steps are being considered and will be implemented by different development teams. Each development team will implement the error handling specific to their step. What should you do?

A. Deploy a Cloud Function for each step that calls the corresponding downstream system to complete the required action.

B. Create a Pub/Sub topic for each step. Create a subscription for each downstream development team to subscribe to their step's topic.

C. Create a Pub/Sub topic for timesheet submissions. Create a subscription for each downstream development team to subscribe to the topic.

D. Create a timesheet microservice deployed to Google Kubernetes Engine. The microservice calls each downstream step and waits for a successful response before calling the next step.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 153

Topic #: 1

[All Professional Cloud Developer Questions]

You are designing an application that uses a microservices architecture. You are planning to deploy the application in the cloud and on-premises. You want to make sure the application can scale up on demand and also use managed services as much as possible. What should you do?

A. Deploy open source Istio in a multi-cluster deployment on multiple Google Kubernetes Engine (GKE) clusters managed by Anthos.

B. Create a GKE cluster in each environment with Anthos, and use Cloud Run for Anthos to deploy your application to each cluster.

C. Install a GKE cluster in each environment with Anthos, and use Cloud Build to create a Deployment for your application in each cluster.

D. Create a GKE cluster in the cloud and install open-source Kubernetes on-premises. Use an external load balancer service to distribute traffic across the two environments.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 154

Topic #: 1

[All Professional Cloud Developer Questions]

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

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 155

Topic #: 1

[All Professional Cloud Developer Questions]

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?

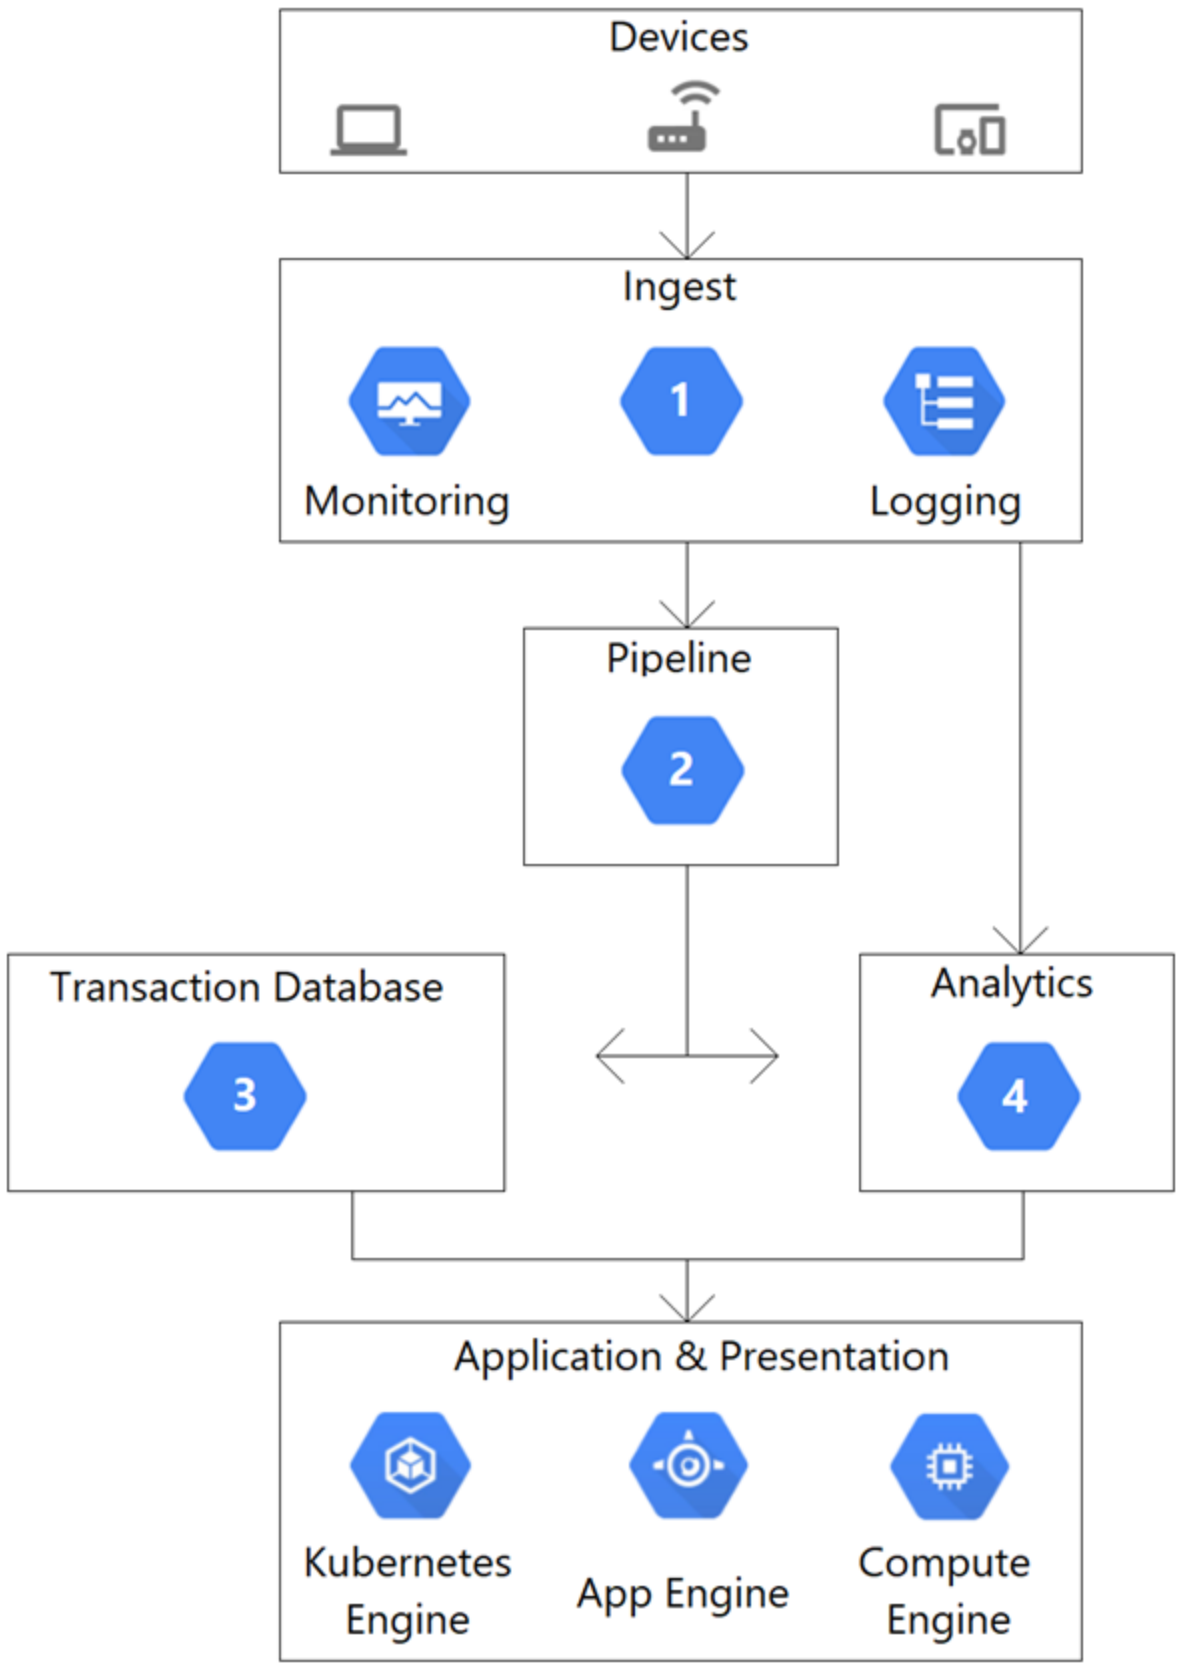

A. 1. App Engine
2. Pub/Sub
3. BigQuery
4. Firestore

B. 1. Dataflow
2. Pub/Sub
3. Firestore
4. BigQuery

C. 1. Pub/Sub
2. Dataflow
3. BigQuery
4. Firestore

D. 1. Pub/Sub
2. Dataflow
3. Firestore
4. BigQuery

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 156

Topic #: 1

[All Professional Cloud Developer Questions]

Your company just experienced a Google Kubernetes Engine (GKE) API outage due to a zone failure. You want to deploy a highly available GKE architecture that minimizes service interruption to users in the event of a future zone failure. What should you do?

A. Deploy Zonal clusters

B. Deploy Regional clusters

C. Deploy Multi-Zone clusters

D. Deploy GKE on-premises clusters

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 157

Topic #: 1

[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 158

Topic #: 1

[All Professional Cloud Developer Questions]

You are running a containerized application on Google Kubernetes Engine. Your container images are stored in Container Registry. Your team uses CI/CD practices. You need to prevent the deployment of containers with known critical vulnerabilities. What should you do?

A. • Use Web Security Scanner to automatically crawl your application
• Review your application logs for scan results, and provide an attestation that the container is free of known critical vulnerabilities
• Use Binary Authorization to implement a policy that forces the attestation to be provided before the container is deployed

B. • Use Web Security Scanner to automatically crawl your application
• Review the scan results in the scan details page in the Cloud Console, and provide an attestation that the container is free of known critical vulnerabilities
• Use Binary Authorization to implement a policy that forces the attestation to be provided before the container is deployed

C. • Enable the Container Scanning API to perform vulnerability scanning
• Review vulnerability reporting in Container Registry in the Cloud Console, and provide an attestation that the container is free of known critical vulnerabilities
• Use Binary Authorization to implement a policy that forces the attestation to be provided before the container is deployed

D. • Enable the Container Scanning API to perform vulnerability scanning
• Programmatically review vulnerability reporting through the Container Scanning API, and provide an attestation that the container is free of known critical vulnerabilities
• Use Binary Authorization to implement a policy that forces the attestation to be provided before the container is deployed

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 159

Topic #: 1

[All Professional Cloud Developer Questions]

You have an on-premises application that authenticates to the Cloud Storage API using a user-managed service account with a user-managed key. The application connects to Cloud Storage using Private Google Access over a Dedicated Interconnect link. You discover that requests from the application to access objects in the Cloud Storage bucket are failing with a 403 Permission Denied error code. What is the likely cause of this issue?

A. The folder structure inside the bucket and object paths have changed.

B. The permissions of the service account's predefined role have changed.

C. The service account key has been rotated but not updated on the application server.

D. The Interconnect link from the on-premises data center to Google Cloud is experiencing a temporary outage.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 160

Topic #: 1

[All Professional Cloud Developer Questions]

You are using the Cloud Client Library to upload an image in your application to Cloud Storage. Users of the application report that occasionally the upload does not complete and the client library reports an HTTP 504 Gateway Timeout error. You want to make the application more resilient to errors. What changes to the application should you make?

A. Write an exponential backoff process around the client library call.

B. Write a one-second wait time backoff process around the client library call.

C. Design a retry button in the application and ask users to click if the error occurs.

D. Create a queue for the object and inform the users that the application will try again in 10 minutes.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 161

Topic #: 1

[All Professional Cloud Developer Questions]

You are building a mobile application that will store hierarchical data structures in a database. The application will enable users working offline to sync changes when they are back online. A backend service will enrich the data in the database using a service account. The application is expected to be very popular and needs to scale seamlessly and securely. Which database and IAM role should you use?

A. Use Cloud SQL, and assign the roles/cloudsql.editor role to the service account.

B. Use Bigtable, and assign the roles/bigtable.viewer role to the service account.

C. Use Firestore in Native mode and assign the roles/datastore.user role to the service account.

D. Use Firestore in Datastore mode and assign the roles/datastore.viewer role to the service account.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 162

Topic #: 1

[All Professional Cloud Developer Questions]

Your application is deployed on hundreds of Compute Engine instances in a managed instance group (MIG) in multiple zones. You need to deploy a new instance template to fix a critical vulnerability immediately but must avoid impact to your service. What setting should be made to the MIG after updating the instance template?

A. Set the Max Surge to 100%.

B. Set the Update mode to Opportunistic.

C. Set the Maximum Unavailable to 100%.

D. Set the Minimum Wait time to 0 seconds.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 163

Topic #: 1

[All Professional Cloud Developer Questions]

You made a typo in a low-level Linux configuration file that prevents your Compute Engine instance from booting to a normal run level. You just created the Compute Engine instance today and have done no other maintenance on it, other than tweaking files. How should you correct this error?

A. Download the file using scp, change the file, and then upload the modified version

B. Configure and log in to the Compute Engine instance through SSH, and change the file

C. Configure and log in to the Compute Engine instance through the serial port, and change the file

D. Configure and log in to the Compute Engine instance using a remote desktop client, and change the file

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 164

Topic #: 1

[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 165
Topic #: 1
[All Professional Cloud Developer Questions]

Your company's corporate policy states that there must be a copyright comment at the very beginning of all source files. You want to write a custom step in Cloud Build that is triggered by each source commit. You need the trigger to validate that the source contains a copyright and add one for subsequent steps if not there. What should you do?

A. Build a new Docker container that examines the files in /workspace and then checks and adds a copyright for each source file. Changed files are explicitly committed back to the source repository.

B. Build a new Docker container that examines the files in /workspace and then checks and adds a copyright for each source file. Changed files do not need to be committed back to the source repository.

C. Build a new Docker container that examines the files in a Cloud Storage bucket and then checks and adds a copyright for each source file. Changed files are written back to the Cloud Storage bucket.

D. Build a new Docker container that examines the files in a Cloud Storage bucket and then checks and adds a copyright for each source file. Changed files are explicitly committed back to the source repository.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 166

Topic #: 1

[All Professional Cloud Developer Questions]

One of your deployed applications in Google Kubernetes Engine (GKE) is having intermittent performance issues. Your team uses a third-party logging solution. You want to install this solution on each node in your GKE cluster so you can view the logs. What should you do?

A. Deploy the third-party solution as a DaemonSet

B. Modify your container image to include the monitoring software

C. Use SSH to connect to the GKE node, and install the software manually

D. Deploy the third-party solution using Terraform and deploy the logging Pod as a Kubernetes Deployment

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 167
Topic #: 1
[All Professional Cloud Developer Questions]

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, and that they analyze and respond to any issues that occur.

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.
• Release cycles include development freezes to allow for QA testing.
• The application has no logging.
• Applications are manually deployed by infrastructure engineers during periods of slow traffic on weekday evenings.
• 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.
• Support 10x as many concurrent users.
• 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.
○ Develop standardized workflows and processes around application lifecycle management.
○ Define service level indicators (SLIs) and service level objectives (SLOs).

Technical Requirements -
• Provide secure communications between the on-premises data center and cloud-hosted applications and infrastructure.
• The application must provide usage metrics and monitoring.
• APIs require authentication and authorization.
• Implement faster and more accurate validation of new features.
• Logging and performance metrics must provide actionable information to be able to provide debugging information and alerts.
• Must scale to meet user demand.

For this question, refer to the HipLocal case study.

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

A. Use Google Kubernetes Engine (GKE) to run the application as a microservice. Run the MySQL database on a dedicated GKE node.

B. Use multiple Compute Engine instances to run MySQL to store state information. Use a Google Cloud-managed load balancer to distribute the load between instances. Use managed instance groups for scaling.

C. Use Memorystore to store session information and CloudSQL to store state information. Use a Google Cloud-managed load balancer to distribute the load between instances. Use managed instance groups for scaling.

D. Use a Cloud Storage bucket to serve the application as a static website, and use another Cloud Storage bucket to store user state information.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 168
Topic #: 1
[All Professional Cloud Developer Questions]

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, and that they analyze and respond to any issues that occur.

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.
• Release cycles include development freezes to allow for QA testing.
• The application has no logging.
• Applications are manually deployed by infrastructure engineers during periods of slow traffic on weekday evenings.
• 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.
• Support 10x as many concurrent users.
• 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.
○ Develop standardized workflows and processes around application lifecycle management.
○ Define service level indicators (SLIs) and service level objectives (SLOs).

Technical Requirements -
• Provide secure communications between the on-premises data center and cloud-hosted applications and infrastructure.
• The application must provide usage metrics and monitoring.
• APIs require authentication and authorization.
• Implement faster and more accurate validation of new features.
• Logging and performance metrics must provide actionable information to be able to provide debugging information and alerts.
• Must scale to meet user demand.

For this question, refer to the HipLocal case study.

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

A. Include unit tests in their code, and prevent deployments to QA until all tests have a passing status.

B. Include performance tests in their code, and prevent deployments to QA until all tests have a passing status.

C. Create health checks for the QA environment, and redeploy the APIs at a later time if the environment is unhealthy.

D. Redeploy the APIs to App Engine using Traffic Splitting. Do not move QA traffic to the new versions if errors are found.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 169
Topic #: 1
[All Professional Cloud Developer Questions]

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, and that they analyze and respond to any issues that occur.

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.
• Release cycles include development freezes to allow for QA testing.
• The application has no logging.
• Applications are manually deployed by infrastructure engineers during periods of slow traffic on weekday evenings.
• 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.
• Support 10x as many concurrent users.
• 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.
○ Develop standardized workflows and processes around application lifecycle management.
○ Define service level indicators (SLIs) and service level objectives (SLOs).

Technical Requirements -
• Provide secure communications between the on-premises data center and cloud-hosted applications and infrastructure.
• The application must provide usage metrics and monitoring.
• APIs require authentication and authorization.
• Implement faster and more accurate validation of new features.
• Logging and performance metrics must provide actionable information to be able to provide debugging information and alerts.
• Must scale to meet user demand.

For this question, refer to the HipLocal case study.

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

A. Create an API key. Use the API key to interact with Google Cloud.

B. Use the default compute service account to interact with Google Cloud.

C. Create a service account for the application. Export and deploy the private key for the application. Use the service account to interact with Google Cloud.

D. Create a service account for the application and for each Google Cloud API used by the application. Export and deploy the private keys used by the application. Use the service account with one Google Cloud API to interact with Google Cloud.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 170
Topic #: 1
[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 171

Topic #: 1

[All Professional Cloud Developer Questions]

You are developing an application that consists of several microservices running in a Google Kubernetes Engine cluster. One microservice needs to connect to a third-party database running on-premises. You need to store credentials to the database and ensure that these credentials can be rotated while following security best practices. What should you do?

A. Store the credentials in a sidecar container proxy, and use it to connect to the third-party database.

B. Configure a service mesh to allow or restrict traffic from the Pods in your microservice to the database.

C. Store the credentials in an encrypted volume mount, and associate a Persistent Volume Claim with the client Pod.

D. Store the credentials as a Kubernetes Secret, and use the Cloud Key Management Service plugin to handle encryption and decryption.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 172

Topic #: 1

[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 173
Topic #: 1
[All Professional Cloud Developer Questions]

Your company has a new security initiative that requires all data stored in Google Cloud to be encrypted by customer-managed encryption keys. You plan to use Cloud Key Management Service (KMS) to configure access to the keys. You need to follow the "separation of duties" principle and Google-recommended best practices. What should you do? (Choose two.)

A. Provision Cloud KMS in its own project.

B. Do not assign an owner to the Cloud KMS project.

C. Provision Cloud KMS in the project where the keys are being used.

D. Grant the roles/cloudkms.admin role to the owner of the project where the keys from Cloud KMS are being used.

E. Grant an owner role for the Cloud KMS project to a different user than the owner of the project where the keys from Cloud KMS are being used.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 174

Topic #: 1

[All Professional Cloud Developer Questions]

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.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 175

Topic #: 1

[All Professional Cloud Developer Questions]

Your organization has recently begun an initiative to replatform their legacy applications onto Google Kubernetes Engine. You need to decompose a monolithic application into microservices. Multiple instances have read and write access to a configuration file, which is stored on a shared file system. You want to minimize the effort required to manage this transition, and you want to avoid rewriting the application code. What should you do?

A. Create a new Cloud Storage bucket, and mount it via FUSE in the container.

B. Create a new persistent disk, and mount the volume as a shared PersistentVolume.

C. Create a new Filestore instance, and mount the volume as an NFS PersistentVolume.

D. Create a new ConfigMap and volumeMount to store the contents of the configuration file.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 176
Topic #: 1
[All Professional Cloud Developer Questions]

Your development team has built several Cloud Functions using Java along with corresponding integration and service tests. You are building and deploying the functions and launching the tests using Cloud Build. Your Cloud Build job is reporting deployment failures immediately after successfully validating the code. What should you do?

A. Check the maximum number of Cloud Function instances.

B. Verify that your Cloud Build trigger has the correct build parameters.

C. Retry the tests using the truncated exponential backoff polling strategy.

D. Verify that the Cloud Build service account is assigned the Cloud Functions Developer role.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 177

Topic #: 1

[All Professional Cloud Developer Questions]

You manage a microservices application on Google Kubernetes Engine (GKE) using Istio. You secure the communication channels between your microservices by implementing an Istio AuthorizationPolicy, a Kubernetes NetworkPolicy, and mTLS on your GKE cluster. You discover that HTTP requests between two Pods to specific URLs fail, while other requests to other URLs succeed. What is the cause of the connection issue?

A. A Kubernetes NetworkPolicy resource is blocking HTTP traffic between the Pods.

B. The Pod initiating the HTTP requests is attempting to connect to the target Pod via an incorrect TCP port.

C. The Authorization Policy of your cluster is blocking HTTP requests for specific paths within your application.

D. The cluster has mTLS configured in permissive mode, but the Pod's sidecar proxy is sending unencrypted traffic in plain text.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 178
Topic #: 1
[All Professional Cloud Developer Questions]

You recently migrated an on-premises monolithic application to a microservices application on Google Kubernetes Engine (GKE). The application has dependencies on backend services on-premises, including a CRM system and a MySQL database that contains personally identifiable information (PII). The backend services must remain on-premises to meet regulatory requirements.

You established a Cloud VPN connection between your on-premises data center and Google Cloud. You notice that some requests from your microservices application on GKE to the backend services are failing due to latency issues caused by fluctuating bandwidth, which is causing the application to crash. How should you address the latency issues?

A. Use Memorystore to cache frequently accessed PII data from the on-premises MySQL database

B. Use Istio to create a service mesh that includes the microservices on GKE and the on-premises services

C. Increase the number of Cloud VPN tunnels for the connection between Google Cloud and the on-premises services

D. Decrease the network layer packet size by decreasing the Maximum Transmission Unit (MTU) value from its default value on Cloud VPN

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 179

Topic #: 1

[All Professional Cloud Developer Questions]

Your company has deployed a new API to a Compute Engine instance. During testing, the API is not behaving as expected. You want to monitor the application over 12 hours to diagnose the problem within the application code without redeploying the application. Which tool should you use?

A. Cloud Trace

B. Cloud Monitoring

C. Cloud Debugger logpoints

D. Cloud Debugger snapshots

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 180

Topic #: 1

[All Professional Cloud Developer Questions]

You are designing an application that consists of several microservices. Each microservice has its own RESTful API and will be deployed as a separate Kubernetes Service. You want to ensure that the consumers of these APIs aren't impacted when there is a change to your API, and also ensure that third-party systems aren't interrupted when new versions of the API are released. How should you configure the connection to the application following Google-recommended best practices?

A. Use an Ingress that uses the API's URL to route requests to the appropriate backend.

B. Leverage a Service Discovery system, and connect to the backend specified by the request.

C. Use multiple clusters, and use DNS entries to route requests to separate versioned backends.

D. Combine multiple versions in the same service, and then specify the API version in the POST request.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 181

Topic #: 1

[All Professional Cloud Developer Questions]

Your team is building an application for a financial institution. The application's frontend runs on Compute Engine, and the data resides in Cloud SQL and one Cloud Storage bucket. The application will collect data containing PII, which will be stored in the Cloud SQL database and the Cloud Storage bucket. You need to secure the PII data. What should you do?

A. 1. Create the relevant firewall rules to allow only the frontend to communicate with the Cloud SQL database
2. Using IAM, allow only the frontend service account to access the Cloud Storage bucket

B. 1. Create the relevant firewall rules to allow only the frontend to communicate with the Cloud SQL database
2. Enable private access to allow the frontend to access the Cloud Storage bucket privately

C. 1. Configure a private IP address for Cloud SQL
2. Use VPC-SC to create a service perimeter
3. Add the Cloud SQL database and the Cloud Storage bucket to the same service perimeter

D. 1. Configure a private IP address for Cloud SQL
2. Use VPC-SC to create a service perimeter
3. Add the Cloud SQL database and the Cloud Storage bucket to different service perimeters

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 182

Topic #: 1

[All Professional Cloud Developer Questions]

You are designing a chat room application that will host multiple rooms and retain the message history for each room. You have selected Firestore as your database. How should you represent the data in Firestore?

A. Create a collection for the rooms. For each room, create a document that lists the contents of the messages

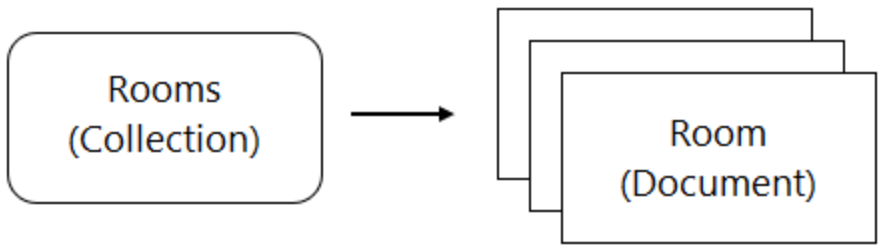

B. Create a collection for the rooms. For each room, create a collection that contains a document for each message

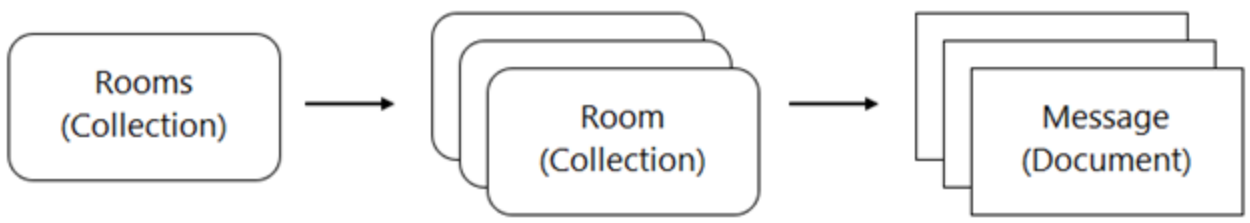

C. Create a collection for the rooms. For each room, create a document that contains a collection for documents, each of which contains a message.

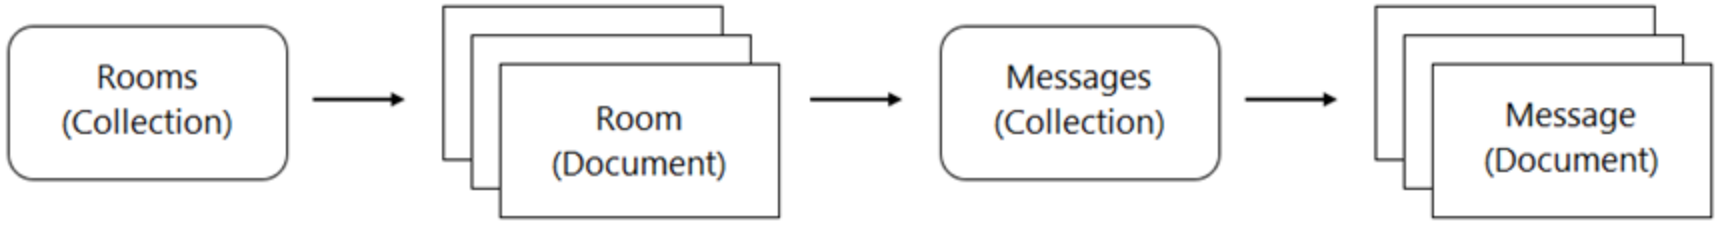

D. Create a collection for the rooms, and create a document for each room. Create a separate collection for messages, with one document per message. Each room's document contains a list of references to the messages.

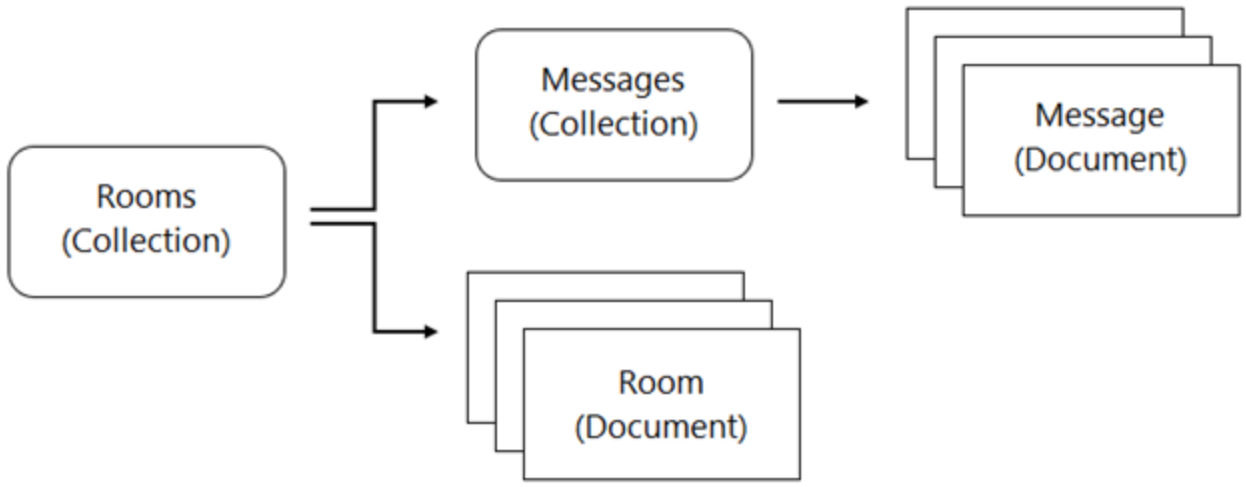

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 183

Topic #: 1

[All Professional Cloud Developer Questions]

You are developing an application that will handle requests from end users. You need to secure a Cloud Function called by the application to allow authorized end users to authenticate to the function via the application while restricting access to unauthorized users. You will integrate Google Sign-In as part of the solution and want to follow Google-recommended best practices. What should you do?

A. Deploy from a source code repository and grant users the roles/cloudfunctions.viewer role.

B. Deploy from a source code repository and grant users the roles/cloudfunctions.invoker role

C. Deploy from your local machine using gcloud and grant users the roles/cloudfunctions.admin role

D. Deploy from your local machine using gcloud and grant users the roles/cloudfunctions.developer role

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 184

Topic #: 1

[All Professional Cloud Developer Questions]

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

A. Google Cloud Armor

B. Debugger

C. Web Security Scanner

D. Error Reporting

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer

Question #: 185

Topic #: 1

[All Professional Cloud Developer Questions]

You are building a highly available and globally accessible application that will serve static content to users. You need to configure the storage and serving components. You want to minimize management overhead and latency while maximizing reliability for users. What should you do?

A. 1. Create a managed instance group. Replicate the static content across the virtual machines (VMs)
2. Create an external HTTP(S) load balancer.
3. Enable Cloud CDN, and send traffic to the managed instance group.

B. 1. Create an unmanaged instance group. Replicate the static content across the VMs.
2. Create an external HTTP(S) load balancer
3. Enable Cloud CDN, and send traffic to the unmanaged instance group.

C. 1. Create a Standard storage class, regional Cloud Storage bucket. Put the static content in the bucket
2. Reserve an external IP address, and create an external HTTP(S) load balancer
3. Enable Cloud CDN, and send traffic to your backend bucket

D. 1. Create a Standard storage class, multi-regional Cloud Storage bucket. Put the static content in the bucket.
2. Reserve an external IP address, and create an external HTTP(S) load balancer.
3. Enable Cloud CDN, and send traffic to your backend bucket.

Show Suggested Answer

Actual exam question from Google's Professional Cloud Developer
Question #: 186
Topic #: 1
[All Professional Cloud Developer Questions]

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, and that they analyze and respond to any issues that occur.

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.
• Release cycles include development freezes to allow for QA testing.
• The application has no logging.
• Applications are manually deployed by infrastructure engineers during periods of slow traffic on weekday evenings.
• 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.
• Support 10x as many concurrent users.
• 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.
○ Develop standardized workflows and processes around application lifecycle management.
○ Define service level indicators (SLIs) and service level objectives (SLOs).

Technical Requirements -
• Provide secure communications between the on-premises data center and cloud-hosted applications and infrastructure.
• The application must provide usage metrics and monitoring.
• APIs require authentication and authorization.
• Implement faster and more accurate validation of new features.
• Logging and performance metrics must provide actionable information to be able to provide debugging information and alerts.
• Must scale to meet user demand.

For this question refer to the HipLocal case study.

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

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

Show Suggested Answer