

Google Certified Professional - Cloud Developer.VCEplus.premium.exam.50q

Passing Score: 800
Time Limit: 120 min
File Version: 1.0



Website: <https://vceplus.com>

VCE to PDF Converter: <https://vceplus.com/vce-to-pdf/>

Facebook: <https://www.facebook.com/VCE.For.All.VN/>

Twitter : https://twitter.com/VCE_Plus

Google Certified Professional - Cloud Developer

Version 1.0



Exam A

QUESTION 1

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]/`

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

QUESTION 2 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.



Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 3

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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 4

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

```
export NAME=load-balancer

# create network
gcloud compute networks create ${NAME}

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

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

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

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

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

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

Which commands should you run?

- A. `gcloud compute instances add-access-config ${NAME}-backend-instance-1`
- B. `gcloud compute instances add-tags ${NAME}-backend-instance-1 --tags http-server`
- C. `gcloud compute firewall-rules create allow-lb --network load-balancer --allow tcp --source-ranges 130.211.0.0/22,35.191.0.0/16 --direction INGRESS`
- D. `gcloud compute firewall-rules create allow-lb --network load-balancer --allow tcp --destination-ranges 130.211.0.0/22,35.191.0.0/16 --direction EGRESS`

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 5

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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 6

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

Which command should you use?

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/sdk/gcloud/reference/compute/copy-files>

QUESTION 7

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 8 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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 9 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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 10

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.



Correct Answer: BF

Section: (none)

Explanation

Explanation/Reference:

QUESTION 11

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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 12

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

readinessProbe:
  httpGet:
    path: /healthz
    port: 80

loadBalancerHealthCheck:
  httpGet:
    path: /healthz
    port: 80

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



B.

C.

D.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

QUESTION 13

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.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 14 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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 15 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.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 16 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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:



QUESTION 17 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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 18 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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 19 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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 20

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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 21

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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 22

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://rominirani.com/gcp-stackdriver-tutorial-debug-snapshots-traces-logging-and-logpoints-1ba49e4780e6>

QUESTION 23

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.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:**QUESTION 24**

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.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:**QUESTION 25**

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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 26

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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:



QUESTION 27 You are creating an App Engine application that writes a file to any user's Google Drive.

How should the application authenticate to the Google Drive API?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 28

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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 29 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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 30

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.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 31

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

Which technique should you choose?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 32

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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 33

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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 34

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]/`.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 35

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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 36

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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 37

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 38

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 39

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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 40 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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://livebook.manning.com/book/kubernetes-in-action/chapter-10/46>

QUESTION 41 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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 42 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.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:



QUESTION 43

Case Study

Company Overview

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

Executive Statement

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

Solution Concept

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

Existing Technical Environment

HipLocal's environment is a mix of on-premises hardware and infrastructure running in Google Cloud Platform. The HipLocal team understands their application well, but has limited experience in global scale applications. Their existing technical environment is as follows:

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

Business Requirements

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

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

Technical Requirements

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

HipLocal's .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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

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



QUESTION 44

Case Study

Company Overview

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

Executive Statement

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

Solution Concept

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

Existing Technical Environment

HipLocal's environment is a mix of on-premises hardware and infrastructure running in Google Cloud Platform. The HipLocal team understands their application well, but has limited experience in global scale applications. Their existing technical environment is as follows:

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

Business Requirements

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

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

Technical Requirements

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

HipLocal's APIs are showing occasional failures, but they cannot find a pattern. They want to collect some metrics to help them troubleshoot.

What should they do?

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:



QUESTION 45

Case Study

Company Overview

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

Executive Statement

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

Solution Concept

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

Existing Technical Environment

HipLocal's environment is a mix of on-premises hardware and infrastructure running in Google Cloud Platform. The HipLocal team understands their application well, but has limited experience in global scale applications. Their existing technical environment is as follows:

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

Business Requirements

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

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

Technical Requirements

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

HipLocal 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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

**QUESTION 46****Case Study****Company Overview**

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

Executive Statement

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

Solution Concept

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

Existing Technical Environment

HipLocal's environment is a mix of on-premises hardware and infrastructure running in Google Cloud Platform. The HipLocal team understands their application well, but has limited experience in global scale applications. Their existing technical environment is as follows:

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

Business Requirements

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

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

Technical Requirements

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

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

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

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

**QUESTION 47****Case Study****Company Overview**

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

Executive Statement

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

Solution Concept

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

Existing Technical Environment

HipLocal's environment is a mix of on-premises hardware and infrastructure running in Google Cloud Platform. The HipLocal team understands their application well, but has limited experience in global scale applications. Their existing technical environment is as follows:

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

Business Requirements

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

- Expand availability of the application to new regions.
- Increase the number of concurrent users that can be supported.

- Ensure a consistent experience for users when they travel to different regions.
- Obtain user activity metrics to better understand how to monetize their product.
- Ensure compliance with regulations in the new regions (for example, GDPR).
- Reduce infrastructure management time and cost.
- Adopt the Google-recommended practices for cloud computing.

Technical Requirements

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

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

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

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

Correct Answer: BC

Section: (none)

Explanation

Explanation/Reference:

**QUESTION 48****Case Study****Company Overview**

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

Executive Statement

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

Solution Concept

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

Existing Technical Environment

HipLocal's environment is a mix of on-premises hardware and infrastructure running in Google Cloud Platform. The HipLocal team understands their application well, but has limited experience in global scale applications. Their existing technical environment is as follows:

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

Business Requirements

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

- Expand availability of the application to new regions.
- Increase the number of concurrent users that can be supported.

- Ensure a consistent experience for users when they travel to different regions.
- Obtain user activity metrics to better understand how to monetize their product.
- Ensure compliance with regulations in the new regions (for example, GDPR).
- Reduce infrastructure management time and cost.
- Adopt the Google-recommended practices for cloud computing.

Technical Requirements

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

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

- A. 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.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 49**Case Study****Company Overview**

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

Executive Statement

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

Solution Concept

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

Existing Technical Environment

HipLocal's environment is a mix of on-premises hardware and infrastructure running in Google Cloud Platform. The HipLocal team understands their application well, but has limited experience in global scale applications. Their existing technical environment is as follows:

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

Business Requirements

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

- Expand availability of the application to new regions.
- Increase the number of concurrent users that can be supported.
- Ensure a consistent experience for users when they travel to different regions.
- Obtain user activity metrics to better understand how to monetize their product.

- Ensure compliance with regulations in the new regions (for example, GDPR).
- Reduce infrastructure management time and cost.
- Adopt the Google-recommended practices for cloud computing.

Technical Requirements

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

Which service should HipLocal use for their public APIs?

- A. Cloud Armor
- B. Cloud Functions
- C. Cloud Endpoints
- D. Shielded Virtual Machines

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 50**Case Study****Company Overview**

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

Executive Statement

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

Solution Concept

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

Existing Technical Environment

HipLocal's environment is a mix of on-premises hardware and infrastructure running in Google Cloud Platform. The HipLocal team understands their application well, but has limited experience in global scale applications. Their existing technical environment is as follows:

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

Business Requirements

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

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

- Adopt the Google-recommended practices for cloud computing.

Technical Requirements

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

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

Which configuration should they choose?

- A. 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.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

