

PROFESSIONAL-CLOUD- DEVELOPER^{Q&As}

Professional Cloud Developer

**Pass Google PROFESSIONAL-CLOUD-DEVELOPER
Exam with 100% Guarantee**

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.leads4pass.com/professional-cloud-developer.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Google
Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



QUESTION 1

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.

Correct Answer: B

<https://cloud.google.com/build/docs/build-config-file-schema>

QUESTION 2

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.updateStrategy.type to RollingUpdate. Use Cloud Monitoring to look for performance issues, and run the kubectl rollback command if there are any issues.

Correct Answer: C

https://cloud.google.com/architecture/implementing-deployment-and-testing-strategies-on-gke#perform_a_bluegreen_deployment

QUESTION 3

You are responsible for deploying a new API. That API will have three different URL paths:

- 1.

<https://yourcompany.com/students>

2.

<https://yourcompany.com/teachers>

3.

<https://yourcompany.com/classes>

You need to configure each API URL path to invoke a different function in your code. What should you do?

- A. Create one Cloud Function as a backend service exposed using an HTTPS load balancer.
- B. Create three Cloud Functions exposed directly.
- C. Create one Cloud Function exposed directly.
- D. Create three Cloud Functions as three backend services exposed using an HTTPS load balancer.

Correct Answer: D

<https://cloud.google.com/load-balancing/docs/https/setup-global-ext-https-serverless>

QUESTION 4

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.

Correct Answer: D

QUESTION 5

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

QUESTION 6

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.

Correct Answer: D

Reference: <https://cloud.google.com/blog/products/operations/troubleshoot-gke-faster-with-monitoring-data-in-your-logs>

QUESTION 7

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.

Correct Answer: B

QUESTION 8

You plan to deploy a new application revision with a Deployment resource to Google Kubernetes Engine (GKE) in production. The container might not work correctly. You want to minimize risk in case there are issues after deploying the revision. You want to follow Google-recommended best practices. What should you do?

- A. Perform a rolling update with a PodDisruptionBudget of 80%.

- B. Perform a rolling update with a HorizontalPodAutoscaler scale-down policy value of 0.
- C. Convert the Deployment to a StatefulSet, and perform a rolling update with a PodDisruptionBudget of 80%.
- D. Convert the Deployment to a StatefulSet, and perform a rolling update with a HorizontalPodAutoscaler scale-down policy value of 0.

Correct Answer: A

<https://cloud.google.com/blog/products/containers-kubernetes/ensuring-reliability-and-uptime-for-your-gke-cluster>

Setting PodDisruptionBudget ensures that your workloads have a sufficient number of replicas, even during maintenance. Using the PDB, you can define a number (or percentage) of pods that can be terminated, even if terminating them brings the current replica count below the desired value. With PDB configured, Kubernetes will drain a node following the configured disruption schedule. New pods will be deployed on other available nodes. This approach ensures Kubernetes schedules workloads in an optimal way while controlling the disruption based on the PDB configuration.

<https://blog.knoldus.com/how-to-avoid-outages-in-your-kubernetes-cluster-using-pdb/>

QUESTION 9

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.

Correct Answer: B

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

QUESTION 10

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.

Correct Answer: B

QUESTION 11

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.

Correct Answer: A

<https://cloud.google.com/build/docs/automating-builds/create-manage-triggers#:~:text=A%20Cloud%20Build%20trigger%20automatically,changes%20that%20match%20certain%20criteria.>

QUESTION 12

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.

Correct Answer: C

QUESTION 13

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.

Correct Answer: C

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

QUESTION 14

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.

Correct Answer: C

Reference: <https://cloud.google.com/sql/docs/mysql/sql-proxy>

QUESTION 15

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.

Correct Answer: A

<https://cloud.google.com/pubsub/docs/subscriber> "Multiple subscribers can make pull calls to the same "shared" subscription. Each subscriber will receive a subset of the messages."

[PROFESSIONAL-CLOUD-DEVELOPER Practice Test](#)

[PROFESSIONAL-CLOUD-DEVELOPER Exam Questions](#)

[PROFESSIONAL-CLOUD-DEVELOPER Braindumps](#)