



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.pass4itsure.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 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

Correct Answer: C

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

QUESTION 2

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>

Correct Answer: D

QUESTION 3

For this question, refer to the HipLocal case study.

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

- A. Create a service account and download its key. Use the key to authenticate to Cloud Key Management Service



(KMS) to obtain the database credentials.

B. Create a service account and download its key. Use the key to authenticate to Cloud Key Management Service (KMS) to obtain a key used to decrypt the database credentials.

C. Create a service account and grant it the roles/iam.serviceAccountUser role. Impersonate as this account and authenticate using the Cloud SQL Proxy.

D. Grant the roles/secretmanager.secretAccessor role to the Compute Engine service account. Store and access the database credentials with the Secret Manager API.

Correct Answer: D

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

QUESTION 4

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

The Dockerfile is as follows: FROM python:3.7-alpine COPY ./app WORKDIR /app RUN pip install -r requirements.txt CMD ["unicorn", "-w 4", "main:app"] You notice that Cloud Build runs are taking longer than expected to complete. You want to decrease the build time. What should you do? (Choose two.)

A. Select a virtual machine (VM) size with higher CPU for Cloud Build runs.

B. Deploy a Container Registry on a Compute Engine VM in a VPC, and use it to store the final images.

C. Cache the Docker image for subsequent builds using the -- cache-from argument in your build config file.

D. Change the base image in the Dockerfile to ubuntu:latest, and install Python 3.7 using a package manager utility.

E. Store application source code on Cloud Storage, and configure the pipeline to use gsutil to download the source code.

Correct Answer: AC

<https://cloud.google.com/build/docs/optimize-builds/increase-vcpu-for-builds>

By default, Cloud Build runs your builds on a standard virtual machine (VM). In addition to the standard VM, Cloud Build provides several high-CPU VM types to run builds. To increase the speed of your build, select a machine with a higher

vCPU to run builds. Keep in mind that although selecting a high vCPU machine increases your build speed, it may also increase the startup time of your build as Cloud Build only starts non-standard machines on demand.

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

The easiest way to increase the speed of your Docker image build is by specifying a cached image that can be used for subsequent builds. You can specify the cached image by adding the --cache-from argument in your build config file,

which will instruct Docker to build using that image as a cache source.



QUESTION 5

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.

Correct Answer: A

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

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

QUESTION 6

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.

Correct Answer: B

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

QUESTION 7

You have deployed a Java application to Cloud Run. Your application requires access to a database hosted on Cloud SQL. Due to regulatory requirements: your connection to the Cloud SQL instance must use its internal IP address. How should you configure the connectivity while following Google-recommended best practices?

- A. Configure your Cloud Run service with a Cloud SQL connection.



- B. Configure your Cloud Run service to use a Serverless VPC Access connector
- C. Configure your application to use the Cloud SQL Java connector
- D. Configure your application to connect to an instance of the Cloud SQL Auth proxy

Correct Answer: A

QUESTION 8

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

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

QUESTION 9

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

QUESTION 10

You work for a financial services company that has a container-first approach. Your team develops microservices applications. You have a Cloud Build pipeline that creates a container image, runs regression tests, and publishes the



image to Artifact Registry You need to ensure that only containers that have passed the regression tests are deployed to Google Kubernetes Engine (GKE) clusters You have already enabled Binary Authorization on the GKE clusters What should you do next?

- A. Deploy Voucher Server and Voucher Client Components. After a container image has passed the regression tests, run Voucher Client as a step in the Cloud Build pipeline.
- B. Set the Pod Security Standard level to Restricted for the relevant namespaces Digitally sign the container images that have passed the regression tests as a step in the Cloud Build pipeline.
- C. Create an attestor and a policy. Create an attestation for the container images that have passed the regression tests as a step in the Cloud Build pipeline.
- D. Create an attestor and a policy Run a vulnerability scan to create an attestation for the container image as a step in the Cloud Build pipeline.

Correct Answer: C

QUESTION 11

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 12

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

QUESTION 13

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.

Correct Answer: D

Reference: <https://cloud.google.com/blog/products/application-development/least-privilege-for-cloud-functions-using-cloud-iam>

QUESTION 14

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

Explanation: <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."

QUESTION 15

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.

Correct Answer: C

[Latest PROFESSIONAL-CLOUD-DEVELOPER Dumps](#)

[PROFESSIONAL-CLOUD-DEVELOPER VCE Dumps](#)

[PROFESSIONAL-CLOUD-DEVELOPER Study Guide](#)