

Google Professional-Cloud-Developer Real Exam Questions and Answers FREE [Q110-Q128]



Google Professional-Cloud-Developer Real Exam Questions and Answers FREE
Exam Dumps Professional-Cloud-Developer Practice Free Latest Google Practice Tests

Google Professional-Cloud-Developer, also known as Google Certified Professional - Cloud Developer, is a certification exam offered by Google. Professional-Cloud-Developer exam is designed to test the candidate's knowledge and skills in developing, designing, and managing applications using Google Cloud Platform. Professional-Cloud-Developer exam is intended for developers who have experience in developing and deploying applications on Google Cloud Platform.


Q110. 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?

- * Use HTTP signed URLs to securely provide access to the required resources.
- * Use the instance's service account Application Default Credentials to authenticate to the required resources.
- * 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.


* 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.

Q111. You are designing a chat room application that will host multiple rooms and retain the message history for each room. You have selected Firestore as your database. How should you represent the data in Firestore?


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




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



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



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



Explanation

<https://firebase.google.com/docs/firestore/data-model#hierarchical-data>

Q112. 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 [ "python -m gunicorn -w 4 -k uvicorn: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.)

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

Explanation

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

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

Q113. You are developing a single-player mobile game backend that has unpredictable traffic patterns as users interact with the game throughout the day and night. You want to optimize costs by ensuring that you have enough resources to handle requests, but minimize over-provisioning. You also want the system to handle traffic spikes efficiently. Which compute platform should you use?

- * Cloud Run
- * Compute Engine with managed instance groups
- * Compute Engine with unmanaged instance groups
- * Google Kubernetes Engine using cluster autoscaling

Q114. 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?

- * Deployment
- * StatefulSet

- * ReplicaSet
- * ReplicaController

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

Q115. You want to migrate an on-premises container running in Knative to Google Cloud. You need to make sure that the migration doesn't affect your application's deployment strategy, and you want to use a fully managed service. Which Google Cloud service should you use to deploy your container?

- * Cloud Run
- * Compute Engine
- * Google Kubernetes Engine
- * App Engine flexible environment

Explanation

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

Q116. 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?

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

Q117. You are designing a deployment technique for your new applications on Google Cloud. As part of your deployment planning, you want to use live traffic to gather performance metrics for both new and existing applications. You need to test against the full production load prior to launch. What should you do?

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

Q118. You are planning to deploy hundreds of microservices in your Google Kubernetes Engine (GKE) cluster. How should you secure communication between the microservices on GKE using a managed service?

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

Explanation

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

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

Q119. HipLocal's data science team wants to analyze user reviews.

How should they prepare the data?

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

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

Explanation

<https://cloud.google.com/dlp/docs/deidentify-sensitive-data>

Q120. 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?

- * Deployment
- * StatefulSet
- * ReplicaSet
- * ReplicaController

Reference:

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

Q121. You are evaluating developer tools to help drive Google Kubernetes Engine adoption and integration with your development environment, which includes VS Code and IntelliJ. What should you do?

- * Use Cloud Code to develop applications.
- * Use the Cloud Shell integrated Code Editor to edit code and configuration files.
- * Use a Cloud Notebook instance to ingest and process data and deploy models.
- * Use Cloud Shell to manage your infrastructure and applications from the command line.

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

Q122. You need to migrate an internal file upload API with an enforced 500-MB file size limit to App Engine.

What should you do?

- * Use FTP to upload files.
- * Use CPanel to upload files.
- * Use signed URLs to upload files.
- * Change the API to be a multipart file upload API.

Q123. You have recently instrumented a new application with OpenTelemetry, and you want to check the latency of your application requests in Trace. You want to ensure that a specific request is always traced. What should you do?

- * Wait 10 minutes, then verify that Trace captures those types of requests automatically.
- * Write a custom script that sends this type of request repeatedly from your dev project.
- * Use the Trace API to apply custom attributes to the trace.
- * Add the X-Cloud-Trace-Context header to the request with the appropriate parameters.

<https://cloud.google.com/trace/docs/setup#force-trace>

Cloud Trace doesn't sample every request. To force a specific request to be traced, add an X-Cloud-Trace-Context header to the request.

Q124. 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?

- * Enable private IP for the Cloud SQL instance.

- * Whitelist a project to access Cloud SQL, and add Compute Engine instances in the whitelisted project.
- * Create a role in Cloud SQL that allows access to the database from external instances, and assign the Compute Engine instances to that role.
- * 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.

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

Q125. Your team recently deployed an application on Google Kubernetes Engine (GKE). You are monitoring your application and want to be alerted when the average memory consumption of your containers is under 20% or above 80% How should you configure the alerts?

- * Create a Cloud Function that consumes the Monitoring API. Create a schedule to trigger the Cloud Function hourly and alert you if the average memory consumption is outside the defined range
- * In Cloud Monitoring, create an alerting policy to notify you if the average memory consumption is outside the defined range
- * Create a Cloud Function that runs on a schedule, executes kubectl top on all the workloads on the cluster, and sends an email alert if the average memory consumption is outside the defined range
- * Write a script that pulls the memory consumption of the instance at the OS level and sends an email alert if the average memory consumption is outside the defined range

Q126. You are a developer at a social media company The company runs their social media website on-premises and uses MySQL as a backend to store user profiles and user posts. Your company plans to migrate to Google Cloud, and your team will migrate user profile information to Firestore. You are tasked with designing the Firestore collections. What should you do?

- * Create one root collection for user profiles, and store each user's post as a nested list in the user profile document.
- * Create one root collection for user profiles, and create one root collection for user posts.
- * Create one root collection for user profiles, and create one subcollection for each user's posts.
- * Create one root collection for user posts, and create one subcollection for each user's profile

Q127. You are developing an application hosted on Google Cloud that uses a MySQL relational database schema.

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

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

<https://cloud.google.com/spanner/docs/migrating-mysql-to-spanner#migration-process> Cloud SQL: Cloud SQL is a web service that allows you to create, configure, and use relational databases that live in Google's cloud. It is a fully-managed service that maintains, manages, and administers your databases, allowing you to focus on your applications and services.

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

Q128. You are running an application on App Engine that you inherited. You want to find out whether the application is using insecure binaries or is vulnerable to XSS attacks. Which service should you use?

- * Cloud Armor
- * Stackdriver Debugger
- * Cloud Security Scanner
- * Stackdriver Error Reporting

Verified Professional-Cloud-Developer Exam Dumps Q&As - Provide Professional-Cloud-Developer with Correct Answers:
https://www.test4engine.com/Professional-Cloud-Developer_exam-latest-braindumps.html