

Q&As

Professional Cloud Architect on Google Cloud Platform

Pass Google PROFESSIONAL-CLOUD-ARCHITECT Exam with 100% Guarantee

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

<https://www.pass2lead.com/professional-cloud-architect.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 creating an App Engine application that uses Cloud Datastore as its persistence layer. You need to retrieve several root entities for which you have the identifiers. You want to minimize the overhead in operations performed by Cloud Datastore. What should you do?

- A. Create the Key object for each Entity and run a batch get operation
- B. Create the Key object for each Entity and run multiple get operations, one operation for each entity
- C. Use the identifiers to create a query filter and run a batch query operation
- D. Use the identifiers to create a query filter and run multiple query operations, one operation for each entity

Correct Answer: A

QUESTION 2

Mountkirk Games\' gaming servers are not automatically scaling properly. Last month, they rolled out a new feature, which suddenly became very popular. A record number of users are trying to use the service, but many of them are getting 503 errors and very slow response times. What should they investigate first?

- A. Verify that the database is online
- B. Verify that the project quota hasn\'t been exceeded
- C. Verify that the new feature code did not introduce any performance bugs
- D. Verify that the load-testing team is not running their tool against production

Correct Answer: B

503 is service unavailable error. If the database was online everyone would get the 503 error.

QUESTION 3

As part of their new application experience, Dress4Wm allows customers to upload images of themselves.

The customer has exclusive control over who may view these images.

Customers should be able to upload images with minimal latency and also be shown their images quickly on the main application page when they log in.

Which configuration should Dress4Win use?

- A. Store image files in a Google Cloud Storage bucket. Use Google Cloud Datastore to maintain metadata that maps each customer\'s ID and their image files.
- B. Store image files in a Google Cloud Storage bucket. Add custom metadata to the uploaded images in Cloud Storage that contains the customer\'s unique ID.

C. Use a distributed file system to store customers' images. As storage needs increase, add more persistent disks and/or nodes. Assign each customer a unique ID, which sets each file's owner attribute, ensuring privacy of images.

D. Use a distributed file system to store customers' images. As storage needs increase, add more persistent disks and/or nodes. Use a Google Cloud SQL database to maintain metadata that maps each customer's ID to their image files.

Correct Answer: A

QUESTION 4

Your architecture calls for the centralized collection of all admin activity and VM system logs within your project. How should you collect these logs from both VMs and services?

A. All admin and VM system logs are automatically collected by Stackdriver.

B. Stackdriver automatically collects admin activity logs for most services. The Stackdriver Logging agent must be installed on each instance to collect system logs.

C. Launch a custom syslogd compute instance and configure your GCP project and VMs to forward all logs to it.

D. Install the Stackdriver Logging agent on a single compute instance and let it collect all audit and access logs for your environment.

Correct Answer: B

Reference <https://cloud.google.com/logging/docs/agent/>

QUESTION 5

The development team has provided you with a Kubernetes Deployment file. You have no infrastructure yet and need to deploy the application. What should you do?

A. Use gcloud to create a Kubernetes cluster. Use Deployment Manager to create the deployment.

B. Use gcloud to create a Kubernetes cluster. Use kubectl to create the deployment.

C. Use kubectl to create a Kubernetes cluster. Use Deployment Manager to create the deployment.

D. Use kubectl to create a Kubernetes cluster. Use kubectl to create the deployment.

Correct Answer: B

<https://cloud.google.com/kubernetes-engine/docs/how-to/creating-a-cluster>

QUESTION 6

You need to develop procedures to test a disaster plan for a mission-critical application. You want to use Google-recommended practices and native capabilities within GCP. What should you do?

A. Use Deployment Manager to automate service provisioning. Use Activity Logs to monitor and debug your tests.

- B. Use Deployment Manager to automate provisioning. Use Stackdriver to monitor and debug your tests.
- C. Use gcloud scripts to automate service provisioning. Use Activity Logs monitor and debug your tests.
- D. Use automated scripts to automate service provisioning. Use Activity Logs monitor and debug your tests.

Correct Answer: B

<https://cloud.google.com/solutions/dr-scenarios-planning-guide>

QUESTION 7

Your team will start developing a new application using microservices architecture on Kubernetes Engine. As part of the development lifecycle, any code change that has been pushed to the remote develop branch on your GitHub repository should be built and tested automatically. When the build and test are successful, the relevant microservice will be deployed automatically in the development environment. You want to ensure that all code deployed in the development environment follows this process. What should you do?

- A. Have each developer install a pre-commit hook on their workstation that tests the code and builds the container when committing on the development branch. After a successful commit, have the developer deploy the newly built container image on the development cluster.
- B. Install a post-commit hook on the remote git repository that tests the code and builds the container when code is pushed to the development branch. After a successful commit, have the developer deploy the newly built container image on the development cluster.
- C. Create a Cloud Build trigger based on the development branch that tests the code, builds the container, and stores it in Container Registry. Create a deployment pipeline that watches for new images and deploys the new image on the development cluster. Ensure only the deployment tool has access to deploy new versions.
- D. Create a Cloud Build trigger based on the development branch to build a new container image and store it in Container Registry. Rely on Vulnerability Scanning to ensure the code tests succeed. As the final step of the Cloud Build process, deploy the new container image on the development cluster. Ensure only Cloud Build has access to deploy new versions.

Correct Answer: C

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

Create a Cloud Build trigger based on the development branch that tests the code, builds the container, and stores it in Container Registry. Create a deployment pipeline that watches for new images and deploys the new image on the development cluster. Ensure only the deployment tool has access to deploy new versions.

QUESTION 8

You want to allow your operations team to store logs from all the production projects in your Organization, without including logs from other projects. All of the production projects are contained in a folder. You want to ensure that all logs for existing and new production projects are captured automatically. What should you do?

- A. Create an aggregated export on the Production folder. Set the log sink to be a Cloud Storage bucket in an operations project.
- B. Create an aggregated export on the Organization resource. Set the log sink to be a Cloud Storage bucket in an

operations project.

C. Create log exports in the production projects. Set the log sinks to be a Cloud Storage bucket in an operations project.

D. Create log exports in the production projects. Set the log sinks to be BigQuery datasets in the production projects, and grant IAM access to the operations team to run queries on the datasets.

Correct Answer: A

QUESTION 9

You need to deploy a stateful workload on Google Cloud. The workload can scale horizontally, but each instance needs to read and write to the same POSIX filesystem. At high load, the stateful workload needs to support up to 100 MB/s of writes. What should you do?

A. Use a persistent disk for each instance.

B. Use a regional persistent disk for each instance.

C. Create a Cloud Filestore instance and mount it in each instance.

D. Create a Cloud Storage bucket and mount it in each instance using gcsfuse.

Correct Answer: C

Cloud Filestore: Cloud Filestore is a scalable and highly available shared file service fully managed by Google. Cloud Filestore provides persistent storage ideal for shared workloads. It is best suited for enterprise applications requiring persistent, durable, shared storage which is accessed by NFS or requires a POSIX compliant file system.

<https://cloud.google.com/storage/docs/gcs-fuse#notes> Reference: <https://cloud.google.com/storage/docs/gcs-fuse>

QUESTION 10

The TerramEarth development team wants to create an API to meet the company's business requirements. You want the development team to focus their development effort on business value versus creating a custom framework.

Which method should they use?

A. Use Google App Engine with Google Cloud Endpoints. Focus on an API for dealers and partners

B. Use Google App Engine with a JAX-RS Jersey Java-based framework. Focus on an API for the public

C. Use Google App Engine with the Swagger (Open API Specification) framework. Focus on an API for the public

D. Use Google Container Engine with a Django Python container. Focus on an API for the public

E. Use Google Container Engine with a Tomcat container with the Swagger (Open API Specification) framework. Focus on an API for dealers and partners

Correct Answer: A

Develop, deploy, protect and monitor your APIs with Google Cloud Endpoints. Using an Open API Specification or one

of our API frameworks, Cloud Endpoints gives you the tools you need for every phase of API development.

From scenario: Business Requirements Decrease unplanned vehicle downtime to less than 1 week, without increasing the cost of carrying surplus inventory Support the dealer network with more data on how their customers use their equipment to better position new products and services Have the ability to partner with different companies ?especially with seed and fertilizer suppliers in the fast-growing agricultural business ?to create compelling joint offerings for their customers.

Reference: <https://cloud.google.com/certification/guides/cloud-architect/casestudy-terramearth>

QUESTION 11

JencoMart wants to move their User Profiles database to Google Cloud Platform. Which Google Database should they use?

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

Correct Answer: D

Common workloads for Google Cloud Datastore:

1.
User profiles
2.
Product catalogs
3.
Game state

References:

<https://cloud.google.com/storage-options/>

<https://cloud.google.com/datastore/docs/concepts/overview>

QUESTION 12

Your customer is moving an existing corporate application to Google Cloud Platform from an on-premises data center. The business owners require minimal user disruption. There are strict security team requirements for storing passwords. What authentication strategy should they use?

- A. Use G Suite Password Sync to replicate passwords into Google

- B. Federate authentication via SAML 2.0 to the existing Identity Provider
- C. Provision users in Google using the Google Cloud Directory Sync tool
- D. Ask users to set their Google password to match their corporate password

Correct Answer: B

<https://cloud.google.com/solutions/authenticating-corporate-users-in-a-hybrid-environment>

QUESTION 13

Mountkirk Games has deployed their new backend on Google Cloud Platform (GCP). You want to create a through testing process for new versions of the backend before they are released to the public. You want the testing environment to scale in an economical way. How should you design the process?

- A. Create a scalable environment in GCP for simulating production load
- B. Use the existing infrastructure to test the GCP-based backend at scale
- C. Build stress tests into each component of your application using resources internal to GCP to simulate load
- D. Create a set of static environments in GCP to test different levels of load ?for example, high, medium, and low

Correct Answer: A

From scenario: Requirements for Game Backend Platform

1.
Dynamically scale up or down based on game activity
 2.
Connect to a managed NoSQL database service
 3.
Run customize Linux distro
-

QUESTION 14

JencoMart has built a version of their application on Google Cloud Platform that serves traffic to Asia. You want to measure success against their business and technical goals.

Which metrics should you track?

- A. Error rates for requests from Asia
- B. Latency difference between US and Asia
- C. Total visits, error rates, and latency from Asia

- D. Total visits and average latency for users from Asia
- E. The number of character sets present in the database

Correct Answer: C

From scenario: Business Requirements include: Expand services into Asia Technical Requirements include: Decrease latency in Asia

QUESTION 15

You want to optimize the performance of an accurate, real-time, weather-charting application. The data comes from 50,000 sensors sending 10 readings a second, in the format of a timestamp and sensor reading. Where should you store the data?

- A. Google BigQuery
- B. Google Cloud SQL
- C. Google Cloud Bigtable
- D. Google Cloud Storage

Correct Answer: C

Google Cloud Bigtable is a scalable, fully-managed NoSQL wide-column database that is suitable for both real-time access and analytics workloads.

Good for: Low-latency read/write access High-throughput analytics Native time series support

Common workloads: IoT, finance, adtech Personalization, recommendations Monitoring Geospatial datasets Graphs

References: <https://cloud.google.com/storage-options/>

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

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

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