

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.certbus.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 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/>

QUESTION 2

You are running a cluster on Kubernetes Engine to serve a web application. Users are reporting that a specific part of the application is not responding anymore. You notice that all pods of your deployment keep restarting after 2 seconds. The application writes logs to standard output. You want to inspect the logs to find the cause of the issue. Which approach can you take?

- A. Review the Stackdriver logs for each Compute Engine instance that is serving as a node in the cluster.
- B. Review the Stackdriver logs for the specific Kubernetes Engine container that is serving the unresponsive part of the application.
- C. Connect to the cluster using gcloud credentials and connect to a container in one of the pods to read the logs.
- D. Review the Serial Port logs for each Compute Engine instance that is serving as a node in the cluster.

Correct Answer: C

QUESTION 3

Your company has just recently activated Cloud Identity to manage users. The Google Cloud Organization has been configured as well. The security team needs to secure projects that will be part of the Organization. They want to prohibit IAM users outside the domain from gaining permissions from now on. What should they do?

- A. Configure an organization policy to restrict identities by domain.
- B. Configure an organization policy to block creation of service accounts.

C. Configure Cloud Scheduler to trigger a Cloud Function every hour that removes all users that don't belong to the Cloud Identity domain from all projects.

D. Create a technical user (e.g., crawler@yourdomain.com), and give it the project owner role at root organization level. Write a bash script that: Lists all the IAM rules of all projects within the organization. Deletes all users that do not belong to the company domain. Create a Compute Engine instance in a project within the Organization and configure gcloud to be executed with technical user credentials. Configure a cron job that executes the bash script every hour.

Correct Answer: D

QUESTION 4

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 5

You are migrating your on-premises solution to Google Cloud in several phases. You will use Cloud VPN to maintain a connection between your on-premises systems and Google Cloud until the migration is completed. You want to make sure all your on-premises systems remain reachable during this period. How should you organize your networking in Google Cloud?

- A. Use the same IP range on Google Cloud as you use on-premises
- B. Use the same IP range on Google Cloud as you use on-premises for your primary IP range and use a secondary range that does not overlap with the range you use on-premises
- C. Use an IP range on Google Cloud that does not overlap with the range you use on-premises
- D. Use an IP range on Google Cloud that does not overlap with the range you use on-premises for your primary IP range and use a secondary range with the same IP range as you use on-premises

Correct Answer: D

QUESTION 6

For this question, refer to the Mountkirk Games case study. Mountkirk Games wants to design their solution for the future in order to take advantage of cloud and technology improvements as they become available. Which two steps

should they take? (Choose two.)

- A. Store as much analytics and game activity data as financially feasible today so it can be used to train machine learning models to predict user behavior in the future.
- B. Begin packaging their game backend artifacts in container images and running them on Kubernetes Engine to improve the availability to scale up or down based on game activity.
- C. Set up a CI/CD pipeline using Jenkins and Spinnaker to automate canary deployments and improve development velocity.
- D. Adopt a schema versioning tool to reduce downtime when adding new game features that require storing additional player data in the database.
- E. Implement a weekly rolling maintenance process for the Linux virtual machines so they can apply critical kernel patches and package updates and reduce the risk of 0-day vulnerabilities.

Correct Answer: CE

QUESTION 7

Your customer is moving their corporate applications to Google Cloud Platform. The security team wants detailed visibility of all projects in the organization. You provision the Google Cloud Resource Manager and set up yourself as the org admin.

What Google Cloud Identity and Access Management (Cloud IAM) roles should you give to the security team?

- A. Org viewer, project owner
- B. Org viewer, project viewer
- C. Org admin, project browser
- D. Project owner, network admin

Correct Answer: B

QUESTION 8

The JencoMart security team requires that all Google Cloud Platform infrastructure is deployed using a least privilege model with separation of duties for administration between production and development resources.

What Google domain and project structure should you recommend?

- A. Create two G Suite accounts to manage users: one for development/test/staging and one for production. Each account should contain one project for every application
- B. Create two G Suite accounts to manage users: one with a single project for all development applications and one with a single project for all production applications
- C. Create a single G Suite account to manage users with each stage of each application in its own project
- D. Create a single G Suite account to manage users with one project for the development/test/staging environment and

one project for the production environment

Correct Answer: D

Note: The principle of least privilege and separation of duties are concepts that, although semantically different, are intrinsically related from the standpoint of security. The intent behind both is to prevent people from having higher privilege levels than they actually need Principle of Least Privilege: Users should only have the least amount of privileges required to perform their job and no more. This reduces authorization exploitation by limiting access to resources such as targets, jobs, or monitoring templates for which they are not authorized. Separation of Duties: Beyond limiting user privilege level, you also limit user duties, or the specific jobs they can perform. No user should be given responsibility for more than one related function. This limits the ability of a user to perform a malicious action and then cover up that action.

References: <https://cloud.google.com/kms/docs/separation-of-duties>

QUESTION 9

Your organization has decided to restrict the use of external IP addresses on instances to only approved instances. You want to enforce this requirement across all of your Virtual Private Clouds (VPCs). What should you do?

- A. Remove the default route on all VPCs. Move all approved instances into a new subnet that has a default route to an internet gateway.
- B. Create a new VPC in custom mode. Create a new subnet for the approved instances, and set a default route to the internet gateway on this new subnet.
- C. Implement a Cloud NAT solution to remove the need for external IP addresses entirely.
- D. Set an Organization Policy with a constraint on constraints/compute.vmExternallpAccess. List the approved instances in the allowedValues list.

Correct Answer: D

Reference: <https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address>

QUESTION 10

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 11

Your company is designing its data lake on Google Cloud and wants to develop different ingestion pipelines to collect unstructured data from different sources. After the data is stored in Google Cloud, it will be processed in several data pipelines to build a recommendation engine for end users on the website. The structure of the data retrieved from the source systems can change at any time. The data must be stored exactly as it was retrieved for reprocessing purposes in case the data structure is incompatible with the current processing pipelines. You need to design an architecture to support the use case after you retrieve the data. What should you do?

A. Send the data through the processing pipeline, and then store the processed data in a BigQuery table for reprocessing.

B. Store the data in a BigQuery table. Design the processing pipelines to retrieve the data from the table.

C. Send the data through the processing pipeline, and then store the processed data in a Cloud Storage bucket for reprocessing.

D. Store the data in a Cloud Storage bucket. Design the processing pipelines to retrieve the data from the bucket.

Correct Answer: D

QUESTION 12

Your company has successfully migrated to the cloud and wants to analyze their data stream to optimize operations. They do not have any existing code for this analysis, so they are exploring all their options. These options include a mix of batch and stream processing, as they are running some hourly jobs and live-processing some data as it comes in.

Which technology should they use for this?

A. Google Cloud Dataproc

B. Google Cloud Dataflow

C. Google Container Engine with Bigtable

D. Google Compute Engine with Google BigQuery

Correct Answer: B

Cloud Dataflow is a fully-managed service for transforming and enriching data in stream (real time) and batch (historical) modes with equal reliability and expressiveness -- no more complex workarounds or compromises needed. References: <https://cloud.google.com/dataflow/>

QUESTION 13

During a high traffic portion of the day, one of your relational databases crashes, but the replica is never promoted to a master. You want to avoid this in the future. What should you do?

- A. Use a different database
- B. Choose larger instances for your database
- C. Create snapshots of your database more regularly
- D. Implement routinely scheduled failovers of your databases

Correct Answer: C

Take regular snapshots of your database system.

If your database system lives on a Compute Engine persistent disk, you can take snapshots of your system each time you upgrade. If your database system goes down or you need to roll back to a previous version, you can simply create a

new persistent disk from your desired snapshot and make that disk the boot disk for a new Compute Engine instance. Note that, to avoid data corruption, this approach requires you to freeze the database system's disk while taking a snapshot.

Reference: <https://cloud.google.com/solutions/disaster-recovery-cookbook>

QUESTION 14

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: B

QUESTION 15

A few days after JencoMart migrates the user credentials database to Google Cloud Platform and shuts down the old

server, the new database server stops responding to SSH connections. It is still serving database requests to the application servers correctly.

What three steps should you take to diagnose the problem? Choose 3 answers.

- A. Delete the virtual machine (VM) and disks and create a new one
- B. Delete the instance, attach the disk to a new VM, and investigate
- C. Take a snapshot of the disk and connect to a new machine to investigate
- D. Check inbound firewall rules for the network the machine is connected to
- E. Connect the machine to another network with very simple firewall rules and investigate
- F. Print the Serial Console output for the instance for troubleshooting, activate the interactive console, and investigate

Correct Answer: CDF

D: Handling "Unable to connect on port 22" error message Possible causes include:

There is no firewall rule allowing SSH access on the port. SSH access on port 22 is enabled on all Compute Engine instances by default. If you have disabled access, SSH from the Browser will not work. If you run sshd on a port other than

22, you need to enable the access to that port with a custom firewall rule.

The firewall rule allowing SSH access is enabled, but is not configured to allow connections from GCP Console services. Source IP addresses for browser-based SSH sessions are dynamically allocated by GCP Console and can vary from

session to session.

F: Handling "Could not connect, retrying..." error You can verify that the daemon is running by navigating to the serial console output page and looking for output lines prefixed with the accounts-from-metadata: string. If you are using a standard image but you do not see these output prefixes in the serial console output, the daemon might be stopped. Reboot the instance to restart the daemon.

References: <https://cloud.google.com/compute/docs/ssh-in-browser> <https://cloud.google.com/compute/docs/ssh-in-browser>

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

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

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

To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Try our product !

100% Guaranteed Success

100% Money Back Guarantee

365 Days Free Update

Instant Download After Purchase

24x7 Customer Support

Average 99.9% Success Rate

More than 800,000 Satisfied Customers Worldwide

Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications.
You can view Vendor list of All Certification Exams offered:

<https://www.certbus.com/allproducts>

Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 <p>One Year Free Update Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 <p>Money Back Guarantee To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 <p>Security & Privacy We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p>
---	---	--

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.

Copyright © certbus, All Rights Reserved.