



VCE & PDF

GeekCert.com

<https://www.geekcert.com/professional-cloud-architect.html>

2024 Latest geekcert PROFESSIONAL-CLOUD-ARCHITECT PDF and VCE
dumps Download

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.geekcert.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



VCE & PDF

GeekCert.com

<https://www.geekcert.com/professional-cloud-architect.html>

2024 Latest geekcert PROFESSIONAL-CLOUD-ARCHITECT PDF and VCE dumps Download

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





QUESTION 1

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 2

Auditors visit your teams every 12 months and ask to review all the Google Cloud Identity and Access Management (Cloud IAM) policy changes in the previous 12 months. You want to streamline and expedite the analysis and audit process. What should you do?

- A. Create custom Google Stackdriver alerts and send them to the auditor.
- B. Enable Logging export to Google BigQuery and use ACLs and views to scope the data shared with the auditor.
- C. Use cloud functions to transfer log entries to Google Cloud SQL and use ACLS and views to limit an auditor's view.
- D. Enable Google Cloud Storage (GCS) log export to audit logs into a GCS bucket and delegate access to the bucket.

Correct Answer: D

Export the logs to Google Cloud Storage bucket-Archive Storage, as it will not be used for 1 year, price for which is \$0.004 per GB per Month. The price for long term storage in BigQuery is \$0.01 per GB per Month (250% more). Also for analysis purpose, whenever Auditors are there (once per year), you can use BigQuery and use GCS bucket as external data source. BigQuery supports querying Cloud Storage data from these storage classes: Standard Nearline Coldline Archive

QUESTION 3

For this question, refer to the Dress4Win case study. To be legally compliant during an audit, Dress4Win must be able to give insights in all administrative actions that modify the configuration or metadata of resources on Google Cloud. What should you do?

- A. Use Stackdriver Trace to create a trace list analysis.



- B. Use Stackdriver Monitoring to create a dashboard on the project's activity.
- C. Enable Cloud Identity-Aware Proxy in all projects, and add the group of Administrators as a member.
- D. Use the Activity page in the GCP Console and Stackdriver Logging to provide the required insight.

Correct Answer: D

QUESTION 4

You are moving an application that uses MySQL from on-premises to Google Cloud. The application will run on Compute Engine and will use Cloud SQL. You want to cut over to the Compute Engine deployment of the application with minimal downtime and no data loss to your customers. You want to migrate the application with minimal modification. You also need to determine the cutover strategy. What should you do?

- A. 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL server. 2. Stop the on-premises application. 3. Create a mysqldump of the on-premises MySQL server. 4. Upload the dump to a Cloud Storage bucket. 5. Import the dump into Cloud SQL. 6. Modify the source code of the application to write queries to both databases and read from its local database. 7. Start the Compute Engine application. 8. Stop the on-premises application.
- B. 1. Set up Cloud SQL proxy and MySQL proxy. 2. Create a mysqldump of the on-premises MySQL server. 3. Upload the dump to a Cloud Storage bucket. 4. Import the dump into Cloud SQL. 5. Stop the on-premises application. 6. Start the Compute Engine application.
- C. 1. Set up Cloud VPN to provide private network connectivity between the Compute Engine application and the on-premises MySQL server. 2. Stop the on-premises application. 3. Start the Compute Engine application, configured to read and write to the on-premises MySQL server. 4. Create the replication configuration in Cloud SQL. 5. Configure the source database server to accept connections from the Cloud SQL replica. 6. Finalize the Cloud SQL replica configuration. 7. When replication has been completed, stop the Compute Engine application. 8. Promote the Cloud SQL replica to a standalone instance. 9. Restart the Compute Engine application, configured to read and write to the Cloud SQL standalone instance.
- D. 1. Stop the on-premises application. 2. Create a mysqldump of the on-premises MySQL server. 3. Upload the dump to a Cloud Storage bucket. 4. Import the dump into Cloud SQL. 5. Start the application on Compute Engine.

Correct Answer: C

External replica promotion migration In the migration strategy of external replica promotion, you create an external database replica and synchronize the existing data to that replica. This can happen with minimal downtime to the existing database. When you have a replica database, the two databases have different roles that are referred to in this document as primary and replica. After the data is synchronized, you promote the replica to be the primary in order to move the management layer with minimal impact to database uptime. In Cloud SQL, an easy way to accomplish the external replica promotion is to use the automated migration workflow. This process automates many of the steps that are needed for this type of migration.

<https://cloud.google.com/architecture/migrating-mysql-to-cloudsql-concept>

The best option for migrating your MySQL database is to use an external replica promotion. In this strategy, you create a replica database and set your existing database as the primary. You wait until the two databases are in sync, and you then promote your MySQL replica database to be the primary. This process minimizes database downtime related to the database migration.



https://cloud.google.com/architecture/migrating-mysql-to-cloudsql-concept#external_replica_promotion_migration

QUESTION 5

Your company is running its application workloads on Compute Engine. The applications have been deployed in production, acceptance, and development environments. The production environment is business-critical and is used 24/7, while the acceptance and development environments are only critical during office hours. Your CFO has asked you to optimize these environments to achieve cost savings during idle times. What should you do?

- A. Create a shell script that uses the gcloud command to change the machine type of the development and acceptance instances to a smaller machine type outside of office hours. Schedule the shell script on one of the production instances to automate the task.
- B. Use Cloud Scheduler to trigger a Cloud Function that will stop the development and acceptance environments after office hours and start them just before office hours.
- C. Deploy the development and acceptance applications on a managed instance group and enable autoscaling.
- D. Use regular Compute Engine instances for the production environment, and use preemptible VMs for the acceptance and development environments.

Correct Answer: B

Reference: <https://cloud.google.com/blog/products/it-ops/best-practices-for-optimizing-your-cloud-costs>

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

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

[PROFESSIONAL-CLOUD-ARCHITECT Braindumps](#)