

# Google Professional Cloud Database Engineer

1. You are developing a new application on a VM that is on your corporate network. The application will use Java Database Connectivity (JDBC) to connect to Cloud SQL for PostgreSQL. Your Cloud SQL instance is configured with IP address 192.168.3.48, and SSL is disabled. You want to ensure that your application can access your database instance without requiring configuration changes to your database.

What should you do?

A. Define a connection string using your Google username and password to point to the external (public) IP address of your Cloud SQL instance.

B. Define a connection string using a database username and password to point to the internal (private) IP address of your Cloud SQL instance.

C. Define a connection string using Cloud SQL Auth proxy configured with a service account to point to the internal (private) IP address of your Cloud SQL instance.

D. Define a connection string using Cloud SQL Auth proxy configured with a service account to point to the external (public) IP address of your Cloud SQL instance.

**Answer(s): C**

---

2. Your digital-native business runs its database workloads on Cloud SQL. Your website must be globally accessible 24/7. You need to prepare your Cloud SQL instance for high availability (HA). You want to follow Google-recommended practices.

What should you do? (Choose two.)

A. Set up manual backups.

B. Create a PostgreSQL database on-premises as the HA option.

C. Configure single zone availability for automated backups.

D. Enable point-in-time recovery.

E. Schedule automated backups.

**Answer(s):** D E

---

3. Your company wants to move to Google Cloud. Your current data center is closing in six months. You are running a large, highly transactional Oracle application footprint on VMWare. You need to design a solution with minimal disruption to the current architecture and provide ease of migration to Google Cloud.

What should you do?

A. Migrate applications and Oracle databases to Google Cloud VMware Engine (VMware Engine).

B. Migrate applications and Oracle databases to Compute Engine.

C. Migrate applications to Cloud SQL.

D. Migrate applications and Oracle databases to Google Kubernetes Engine (GKE).

**Answer(s):** A

---

4. Your customer has a global chat application that uses a multi-regional Cloud Spanner instance. The application has recently experienced degraded performance after a new version of the application was launched. Your customer asked you for assistance. During initial troubleshooting, you observed high read latency.

What should you do?

A. Use query parameters to speed up frequently executed queries.

B. Change the Cloud Spanner configuration from multi-region to single region.

C. Use SQL statements to analyze SPANNER\_SYS.READ\_STATS\* tables.

D. Use SQL statements to analyze SPANNER\_SYS.QUERY\_STATS\* tables.

**Answer(s):** C

---

5. Your company has PostgreSQL databases on-premises and on Amazon Web Services (AWS). You are planning multiple database migrations to Cloud SQL in an effort to reduce costs and downtime. You want to follow Google-recommended practices and use Google native data migration tools. You also want to closely monitor the migrations as part of the cutover strategy. What should you do?

A. Use Database Migration Service to migrate all databases to Cloud SQL.

B. Use Database Migration Service for one-time migrations, and use third-party or partner tools for change data capture (CDC) style migrations.

C. Use data replication tools and CDC tools to enable migration.

D. Use a combination of Database Migration Service and partner tools to support the data migration strategy.

**Answer(s): A**

---

6. You are setting up a Bare Metal Solution environment. You need to update the operating system to the latest version. You need to connect the Bare Metal Solution environment to the internet so you can receive software updates.

What should you do?

A. Setup a static external IP address in your VPC network.

B. Set up bring your own IP (BYOIP) in your VPC.

C. Set up a Cloud NAT gateway on the Compute Engine VM.

D. Set up Cloud NAT service.

**Answer(s): C**

---

7. Your organization is running a MySQL workload in Cloud SQL. Suddenly you see a degradation in database performance. You need to identify the root cause of the performance degradation.

What should you do?

A. Use Logs Explorer to analyze log data.

B. Use Cloud Monitoring to monitor CPU, memory, and storage utilization metrics.

C. Use Error Reporting to count, analyze, and aggregate the data.

D. Use Cloud Debugger to inspect the state of an application.

**Answer(s): B**

---

**8.** You work for a large retail and ecommerce company that is starting to extend their business globally.

Your company plans to migrate to Google Cloud. You want to use platforms that will scale easily, handle transactions with the least amount of latency, and provide a reliable customer experience. You need a storage layer for sales transactions and current inventory levels. You want to retain the same relational schema that your existing platform uses.

What should you do?

A. Store your data in Firestore in a multi-region location, and place your compute resources in one of the constituent regions.

B. Deploy Cloud Spanner using a multi-region instance, and place your compute resources close to the default leader region.

C. Build an in-memory cache in Memorystore, and deploy to the specific geographic regions where your application resides.

D. Deploy a Bigtable instance with a cluster in one region and a replica cluster in another geographic region.

**Answer(s): B**

---

**9.** You host an application in Google Cloud. The application is located in a single region and uses Cloud SQL for transactional data.

Most of your users are located in the same time zone and expect the application to be available 7 days a week, from 6 AM to 10 PM. You want to ensure regular maintenance updates to your Cloud SQL instance without creating downtime for your users.

What should you do?

A. Configure a maintenance window during a period when no users will be on the system. Control the order of update by setting non-production instances to earlier and production instances to later.

B. Create your database with one primary node and one read replica in the region.

C. Enable maintenance notifications for users, and reschedule maintenance activities to a specific time after notifications have been sent.

D. Configure your Cloud SQL instance with high availability enabled.

**Answer(s): A**

---

**10.** Your team recently released a new version of a highly consumed application to accommodate additional user traffic. Shortly after the release, you received an alert from your production monitoring team that there is consistently high replication lag between your primary instance and the read replicas of your Cloud SQL for MySQL instances. You need to resolve the replication lag. What should you do?

A. Identify and optimize slow running queries, or set parallel replication flags.

B. Stop all running queries, and re-create the replicas.

C. Edit the primary instance to upgrade to a larger disk, and increase vCPU count.

D. Edit the primary instance to add additional memory.

**Answer(s): A**

---

**11.** Your organization operates in a highly regulated industry. Separation of concerns (SoC) and security principle of least privilege (PoLP) are critical. The operations team consists of:

Person A is a database administrator.

Person B is an analyst who generates metric reports.

Application C is responsible for automatic backups.

You need to assign roles to team members for Cloud Spanner.

Which roles should you assign?

A. roles/spanner.databaseAdmin for Person A  
roles/spanner.databaseReader for Person B  
roles/spanner.backupWriter for Application C

B. roles/spanner.databaseAdmin for Person Aroles/spanner.databaseReader for Person Broles/spanner.backupAdmin for Application C

C. roles/spanner.databaseAdmin for Person Aroles/spanner.databaseUser for Person Broles/spanner.databaseReader for Application C

D. roles/spanner.databaseAdmin for Person Aroles/spanner.databaseUser for Person Broles/spanner.backupWriter for Application C

**Answer(s): A**

---

**12.** You are designing an augmented reality game for iOS and Android devices. You plan to use Cloud Spanner as the primary backend database for game state storage and player authentication. You want to track in-game rewards that players unlock at every stage of the game. During the testing phase, you discovered that costs are much higher than anticipated, but the query response times are within the SL

A. You want to follow Google-recommended practices. You need the database to be performant and highly available while you keep costs low. What should you do?

B. Manually scale down the number of nodes after the peak period has passed.

C. Use interleaving to co-locate parent and child rows.

D. Use the Cloud Spanner query optimizer to determine the most efficient way to execute the SQL query.

E. Use granular instance sizing in Cloud Spanner and Autoscaler.

**Answer(s): D**

---

**13.** You recently launched a new product to the US market. You currently have two Bigtable clusters in one US region to serve all the traffic. Your marketing team is planning an immediate expansion to APAC. You need to roll out the regional expansion while implementing high availability according to Google-recommended practices. What should you do?

A. Maintain a target of 23% CPU utilization by locating:cluster-a in zone us-central1-a cluster-b in zone europe-west1-d cluster-c in zone asia-east1-b

B. Maintain a target of 23% CPU utilization by locating:cluster-a in zone us-central1-a cluster-b in zone us-central1-b cluster-c in zone us-east1-a

C. Maintain a target of 35% CPU utilization by locating:cluster-a in zone us-central1-a cluster-b in zone australia-southeast1-a cluster-c in zone europe-west1-d cluster-d in zone asia-east1-b

D. Maintain a target of 35% CPU utilization by locating:cluster-a in zone us-central1-a cluster-b in zone us-central2-a cluster-c in zone asia-northeast1-b cluster-d in zone asia-east1-b

**Answer(s): D**

---

**14.** Your ecommerce website captures user clickstream data to analyze customer traffic patterns in real time and support personalization features on your website. You plan to analyze this data using big data tools. You need a low-latency solution that can store 8 TB of data and can scale to millions of read and write requests per second.

What should you do?

A. Write your data into Bigtable and use Dataproc and the Apache Hbase libraries for analysis.

B. Deploy a Cloud SQL environment with read replicas for improved performance. Use Datastream to export data to Cloud Storage and analyze with Dataproc and the Cloud Storage connector.

C. Use Memorystore to handle your low-latency requirements and for real-time analytics.

D. Stream your data into BigQuery and use Dataproc and the BigQuery Storage API to analyze large volumes of data.

**Answer(s): A**

---

**15.** Your company uses Cloud Spanner for a mission-critical inventory management system that is globally available. You recently loaded stock keeping unit (SKU) and product catalog data from a company acquisition and observed hot-spots in the Cloud Spanner database. You want to follow Google-recommended schema design practices to avoid performance degradation.

What should you do? (Choose two.)

- A. Use an auto-incrementing value as the primary key.
- B. Normalize the data model.
- C. Promote low-cardinality attributes in multi-attribute primary keys.
- D. Promote high-cardinality attributes in multi-attribute primary keys.
- E. Use bit-reverse sequential value as the primary key.

**Answer(s):** D E

---

**16.** You are managing multiple applications connecting to a database on Cloud SQL for PostgreSQL. You need to be able to monitor database performance to easily identify applications with long-running and resource-intensive queries.

What should you do?

- A. Use log messages produced by Cloud SQL.
- B. Use Query Insights for Cloud SQL.
- C. Use the Cloud Monitoring dashboard with available metrics from Cloud SQL.
- D. Use Cloud SQL instance monitoring in the Google Cloud Console.

**Answer(s):** B

---

**17.** You are building an application that allows users to customize their website and mobile experiences. The application will capture user information and preferences. User profiles have a dynamic schema, and users can add or delete information from their profile. You need to ensure that user changes automatically trigger updates to your downstream BigQuery data warehouse.

What should you do?

- A. Store your data in Bigtable, and use the user identifier as the key. Use one column family to store user profile data, and use another column family to store user preferences.



B. Use Cloud SQL, and create different tables for user profile data and user preferences from your recommendations model. Use SQL to join the user profile data and preferences

C. Use Firestore in Native mode, and store user profile data as a document. Update the user profile with preferences specific to that user and use the user identifier to query.

D. Use Firestore in Datastore mode, and store user profile data as a document. Update the user profile with preferences specific to that user and use the user identifier to query.

**Answer(s): C**

---

**18.** Your application uses Cloud SQL for MySQL. Your users run reports on data that relies on near-real time; however, the additional analytics caused excessive load on the primary database. You created a read replica for the analytics workloads, but now your users are complaining about the lag in data changes and that their reports are still slow. You need to improve the report performance and shorten the lag in data replication without making changes to the current reports.

Which two approaches should you implement? (Choose two.)

A. Create secondary indexes on the replica.

B. Create additional read replicas, and partition your analytics users to use different read replicas.

C. Disable replication on the read replica, and set the flag for parallel replication on the read replica. Re-enable replication and optimize performance by setting flags on the primary instance.

D. Disable replication on the primary instance, and set the flag for parallel replication on the primary instance. Re-enable replication and optimize performance by setting flags on the read replica.

E. Move your analytics workloads to BigQuery, and set up a streaming pipeline to move data and update BigQuery.

**Answer(s): B C**

---

**19.** You are evaluating Cloud SQL for PostgreSQL as a possible destination for your on-premises PostgreSQL instances. Geography is becoming increasingly relevant to customer privacy worldwide. Your solution must support data residency requirements and include a strategy to: configure where data is stored control where the encryption keys are stored govern the access to

data

What should you do?

A. Replicate Cloud SQL databases across different zones.

B. Create a Cloud SQL for PostgreSQL instance on Google Cloud for the data that does not need to adhere to data residency requirements. Keep the data that must adhere to data residency requirements on-premises. Make application changes to support both databases.

C. Allow application access to data only if the users are in the same region as the Google Cloud region for the Cloud SQL for PostgreSQL database.

D. Use features like customer-managed encryption keys (CMEK), VPC Service Controls, and Identity and Access Management (IAM) policies.

**Answer(s): D**

---

**20.** Your customer is running a MySQL database on-premises with read replicas. The nightly incremental backups are expensive and add maintenance overhead. You want to follow Google-recommended practices to migrate the database to Google Cloud, and you need to ensure minimal downtime.

What should you do?

A. Create a Google Kubernetes Engine (GKE) cluster, install MySQL on the cluster, and then import the dump file.

B. Use the mysqldump utility to take a backup of the existing on-premises database, and then import it into Cloud SQL.

C. Create a Compute Engine VM, install MySQL on the VM, and then import the dump file.

D. Create an external replica, and use Cloud SQL to synchronize the data to the replica.

**Answer(s): D**

---