

Professional Cloud Database Engineer

Certification exam guide

A Professional Cloud Database Engineer is a database professional with two years of Google Cloud experience and five years of overall database and IT experience. The Professional Cloud Database Engineer designs, creates, manages, and troubleshoots Google Cloud databases used by applications to store and retrieve data. The Professional Cloud Database Engineer should be comfortable translating business and technical requirements into scalable and cost-effective database solutions.

Section 1: Design scalable and highly available cloud database solutions (~42% of the exam)

1.1 Analyze relevant variables to perform database capacity and usage planning. Activities include:

- Given a scenario, perform solution sizing based on current environment workload metrics and future requirements
- Evaluate performance and cost tradeoffs of different database configurations (machine types, HDD versus SSD, etc.)
- Size database compute and storage based on performance requirements

1.2 Evaluate database high availability and disaster recovery options given the requirements. Activities include:

- Evaluate tradeoffs between multi-region, region, and zonal database deployment strategies
- Given a scenario, define maintenance windows and notifications based on application availability requirements
- Plan database upgrades for Google Cloud-managed databases

1.3 Determine how applications will connect to the database. Activities include:

- Design scalable, highly available, and secure databases
- Configure network and security (Cloud SQL Auth Proxy, CMEK, SSL certificates)
- Justify the use of session pooler services
- Assess auditing policies for managed services

1.4 Evaluate appropriate database solutions on Google Cloud. Activities include:

- Differentiate between managed and unmanaged database services (self-managed, bare metal, Google-managed databases and partner database offerings)
- Distinguish between SQL and NoSQL business requirements (structured, semi-structured, unstructured)
- Analyze the cost of running database solutions in Google Cloud (comparative analysis)
- Assess application and database dependencies

Section 2: Manage a solution that can span multiple database solutions (~34% of the exam)

2.1 Determine database connectivity and access management considerations. Activities include:

- Determine Identity and Access Management (IAM) policies for database connectivity and access control
- Manage database users, including authentication and access

2.2 Configure database monitoring and troubleshooting options. Activities include:

- Assess slow running queries and database locking and identify missing indexes
- Monitor and investigate database vitals: RAM, CPU storage, I/O, Cloud Logging
- Monitor and update quotas
- Investigate database resource contention
- Set up alerts for errors and performance metrics

2.3 Design database backup and recovery solutions. Activities include:

- Given SLAs and SLOs, recommend backup and recovery options (automatic scheduled backups)
- Configure export and import data for databases
- Design for recovery time objective (RTO) and recovery point objective (RPO)

2.4 Optimize database cost and performance in Google Cloud. Activities include:

- Assess options for scaling up and scaling out.
- Scale database instances based on current and upcoming workload

Google Cloud

- Define replication strategies
- Continuously assess and optimize the cost of running a database solution

2.5 Determine solutions to automate database tasks. Activities include:

- Perform database maintenance
- Assess table fragmentation
- Schedule database exports

Section 3: Migrate data solutions (~14% of the exam)

3.1 Design and implement data migration and replication. Activities include:

- Develop and execute migration strategies and plans, including zero downtime, near-zero downtime, extended outage, and fallback plans
- Reverse replication from Google Cloud to source
- Plan and perform database migration, including fallback plans and schema conversion
- Determine the correct database migration tools for a given scenario

Section 4: Deploy scalable and highly available databases in Google Cloud (~10% of the exam)

4.1 Apply concepts to implement highly scalable and available databases in Google Cloud. Activities include:

- Provision high availability database solutions in Google Cloud
- Test high availability and disaster recovery strategies periodically
- Set up multi-regional replication for databases
- Assess requirements for read replicas
- Automate database instance provisioning