Google Cloud

Professional Cloud Database Engineer

Certification exam guide

A Professional Cloud Database Engineer is a database professional with experience in designing, creating, managing, and troubleshooting 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, resilient, and cost-effective database solutions.

Section 1: Design innovative, scalable, and highly available cloud database solutions (~32% of the exam)

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

- Perform solution sizing based on current environment workload metrics and future requirements
- Evaluate performance and cost tradeoffs of different database configurations (e.g., machine types, storage types)
- Size database compute and storage based on performance requirements

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

- Evaluate tradeoffs between multi-regional, regional, and zonal database deployment strategies
- Define maintenance windows and notifications based on application availability requirements

1.3 Determine how applications will connect to the database
Considerations include:

- Configure networking, key management, encryption, and security
- Justify the use of session pooler services
- Assess auditing policies for managed services

1.4 Evaluate appropriate database solutions on Google Cloud

Considerations include:

- Differentiate between managed and unmanaged database services (e.g., self-managed, bare metal, Google-managed, Google Cloud native and partner database offerings)
- Distinguish between SQL and NoSQL business requirements (e.g., structured, semi-structured, unstructured, vector)
- Analyze the cost of running database solutions in Google Cloud (comparative analysis)
- Assess application and database dependencies
- Identify solutions to support regulatory and compliance requirements
- Understand implications of organizational policies on database strategy
- Consider solutions that span multiple database technologies (e.g. federation,
- exports, hybrid deployments)
- Leverage database technologies to support generative AI and LLM use cases

Section 2: Manage a solution that can span multiple database technologies (~25% of the exam)

2.1 Determine database connectivity and access management considerations
Considerations include:

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

2.2 Configure database monitoring and troubleshooting options
Considerations include:

- Assess slow running queries, database locking - identify missing indexes
- Monitor and investigate database vitals - RAM, CPU storage, I/O, and audit 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
Considerations include:

- Given requirements, recommend backup and recovery options (automatic scheduled backups)
- Configure export and import data for databases
- Design for RTO, RPO, and PITR

- Manage data retention

2.4 Optimize database cost and performance in Google Cloud
Considerations include:

- Assess scaling up and scaling out options
- Scale database instances based on current and upcoming workload
- Define replication strategies
- Continuously assess and optimize the cost of running a database solution
- Optimize queries for cost and performance

2.5 Automate common database tasks
Considerations include:

- Perform database maintenance (e.g., rebuilding indexes, data exports)
- Schedule database exports
- Manage upgrades for Google Cloud-managed databases
- Monitor database SLA/SLOs

Section 3: Migrate data solutions (~23% of the exam)

3.1 Design and implement data migration and replication
Considerations include:

- Develop and execute migration strategies and plans, including zero/near-zero downtime, extended outage, and fallback
- Reverse replication from Google Cloud to source
- Plan and perform database migration, including fallback plans and DDL/DML conversion
- Determine the correct database migration tools for a given scenario (e.g., databases hosted outside of Google Cloud)

Section 4: Deploy scalable and highly available databases in Google Cloud (~20% of the exam)

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

- Provision highly available database solutions in Google Cloud
- Test high availability and disaster recovery strategies

- Set up multi-regional replication for databases
- Deploy and scale read replicas
- Automate database instance provisioning
- Configure monitoring for highly available databases