



Relational Databases

GCP Professional Database Engineer | Cloud SQL, Cloud Spanner, AlloyDB & Schema Design

Cloud SQL

Cloud SQL is Google Cloud’s fully managed relational database service supporting MySQL, PostgreSQL, and SQL Server. It handles replication, backups, encryption, patching, and capacity management while you focus on application logic.

Engines	MySQL 5.7/8.0, PostgreSQL 12-16, SQL Server 2017/2019/2022 (Web, Standard, Enterprise).
Machine types	Shared-core (db-f1-micro, db-g1-small) for dev. Custom (db-custom-CPU-RAM) for production. Up to 96 vCPU, 624 GB RAM.
High Availability	Regional HA with synchronous replication to standby in different zone. Automatic failover (60s). Heartbeat-based detection.
Read replicas	Same-region or cross-region async replication. Scale reads. Cross-region for DR and lower latency. Up to 10 replicas.
Automated backups	Daily automated backups. Configurable retention (1-365 days). Point-in-time recovery (PITR) using WAL/binlog replay.
Maintenance windows	Configurable day/hour for patches. Deny maintenance period for blackout windows. Updates are rolling (HA instances).

Connectivity Options


Method	Security	Best For
Public IP + Auth Proxy	IAM + encrypted	Simple setup, Cloud Run/Functions
Private IP	VPC-internal only	GKE, Compute Engine, low latency
Auth Proxy sidecar	IAM + encrypted	Cloud Run, GKE pods
Cloud SQL connectors	IAM + encrypted	Application code (Python, Java, Go, Node)
PSC (Private Service Connect)	Private endpoint	Cross-VPC, no peering needed

5 nines

Spanner multi-region SLA

99.999% availability

Cloud Spanner



TrueTime	GPS + atomic clock time synchronization. Enables external consistency (globally ordered transactions). Unique to Spanner.
Splits	Data automatically sharded across nodes. Split by primary key prefix. Redistributed as data grows. No manual sharding.
Configurations	Regional (3 zones in 1 region, 99.999% SLA). Multi-region (5+ zones across regions, 99.999% SLA). Dual-region also available.
Processing units	100 PU = ~1 node. Scale by 100 PU increments (small) or 1000 PU (large). Autoscaler available.
Schema design	Avoid monotonically increasing PKs (hotspots). Use UUIDs, bit-reversed sequences, or composite keys. Interleaved tables for parent-child co-location.
Interleaved tables	Child rows physically stored with parent rows. Avoid cross-split joins. Define with INTERLEAVE IN PARENT. ON DELETE CASCADE or NO ACTION.

Spanner Best Practices

- Primary key design is the most critical performance decision
- Avoid sequential PKs (timestamps, auto-increment) — causes write hotspots
- Use stale reads (bounded or exact staleness) for non-critical reads — better performance
- Batch mutations for bulk writes — reduces RPC overhead
- Use read-only transactions when possible — no locks, can run on any replica
- Change streams: stream data changes to BigQuery, GCS, Pub/Sub for CDC pipelines

99.999%	64 TB	100 PU
Spanner Multi-Region SLA	Cloud SQL Max Storage	Spanner Min Scale Unit

Cloud SQL	Cloud Spanner
MySQL, PostgreSQL, SQL Server	Proprietary (GoogleSQL)
Regional (HA with failover)	Global (multi-region)
Up to 96 vCPUs, 624 GB RAM	Horizontally scalable
Automated backups + PITR	External consistency

AlloyDB for PostgreSQL

Architecture	Disaggregated compute + storage. Log-based storage layer. Primary instance + read pool instances. 100% PostgreSQL compatible.
Columnar engine	Automatic columnar cache for analytical queries. 100x faster than standard PostgreSQL for analytics. No configuration needed.
Adaptive autovacuum	Intelligent vacuum scheduling. Prevents table bloat. Better than standard PostgreSQL autovacuum for large tables.
HA	Regional HA with synchronous replication. <60s failover. No maintenance windows for HA failovers.
Read pool	Read-only instances for scaling reads. Horizontally scalable. Cross-zone for HA. Node-level or pool-level read endpoints.
AlloyDB AI	Built-in pgvector for vector similarity search. Vertex AI integration for calling ML models from SQL.

Feature	Cloud SQL PostgreSQL	AlloyDB
Performance	Standard PG	4x OLTP, 100x analytics
Storage	SSD/HDD attached	Disaggregated, log-based
Columnar	No	Yes (automatic)
Vacuum	Standard	Adaptive (ML-based)
Vector search	pgvector extension	Native pgvector + Vertex AI
Price	Lower	Higher (enterprise tier)

Database Migration

Database Migration Service (DMS)	Managed service for migrating databases to Cloud SQL, AlloyDB, or Spanner. Continuous replication during migration. Minimal downtime.
Homogeneous migration	Same engine (MySQL to Cloud SQL MySQL). DMS handles replication. Cutover when ready.
Heterogeneous migration	Different engines (Oracle to PostgreSQL, SQL Server to Spanner). Requires schema conversion. Use Harbourbridge (Spanner) or pgLoader.
Harbourbridge	Open-source tool for migrating to Spanner. Schema conversion + data migration from MySQL, PostgreSQL, Oracle, SQL Server.

Cloud SQL vs AlloyDB	
✓ Pros	✗ Cons
+ AlloyDB 4x faster OLTP	- AlloyDB higher cost
+ AlloyDB 100x faster analytics	- AlloyDB PostgreSQL only
+ Adaptive autovacuum	- Cloud SQL supports MySQL/SQL Server
+ Built-in pgvector	- Cloud SQL simpler setup

Official resources

 **Check the current exam guide**

Google Cloud is updating its exams to the products announced at Google Cloud Next '26 (Gemini Enterprise Agent Platform, the data and analytics stack), per the certification hub checked 2026-10-05: read the current exam guide before these sheets.



Professional Cloud
Database Engineer exam...
[services.google...english.pdf](#)
domains and topics



Professional Cloud
Database Engineer...
[cloud.google...e-engineer](#)
*format, price, sample
questions link, renewal*



Professional Cloud
Database Engineer...
[skills.google/paths/22](#)
*Google Skills (ex Cloud
Skills Boost)*



gcloud CLI cheat sheet
[cloud.google...Cheatsheet](#)
the commands on one page



Cloud SQL
documentation
[cloud.google.com/sql/docs](#)



AlloyDB documentation
[cloud.google...loydb/docs](#)



Spanner documentation
[cloud.google...anner/docs](#)



Database Migration
Service
[cloud.google...ation/docs](#)