

Google Cloud Spanner is a fully managed, **scalable, relational database service** offered by Google Cloud Platform (GCP). It's designed for **global-scale, high availability**, and **strong consistency**, and is often used in scenarios that need both SQL capabilities and NoSQL-like scalability.

Key Features

- **Fully Managed:** No need to worry about hardware, replication, sharding, or failover.
- **Global Distribution:** You can create multi-region instances for high availability and low latency.
- **Horizontal Scalability:** Scales across thousands of nodes, ideal for large workloads.
- **SQL Support:** Supports ANSI SQL with extensions, and offers schemas, tables, indexes, etc.
- **Strong Consistency:** Even across global deployments.
- **High Availability:** 99.999% SLA for multi-region instances.

Use Cases

- **Global financial applications**
- **Gaming backends**
- **Retail and eCommerce**
- **Telecommunications systems**
- **Any system that needs both relational data and global scale**

```
gcloud spanner instances create my-instance \
```

```
--config=regional-us-central1 \
```

```
--description="Test Instance" \
```

```
--nodes=1
```

```
gcloud spanner databases create my-database --instance=my-instance
```

```
gcloud spanner databases execute-sql my-database \  
--instance=my-instance \  
--sql='SELECT * FROM MyTable'
```

```
gcloud spanner databases execute-sql my-database \  
--instance=my-instance \  
--sql='SELECT * FROM MyTable'
```

- ❑ Small or simple apps that don't need horizontal scaling.
- ❑ Apps with no need for high availability or global consistency.
- ❑ When you prefer open-source DBs or require tight control over infrastructure.

Real-World Example: Global Ride-Sharing App (like Uber or Lyft)

Imagine you're building a **ride-sharing platform** that operates in **multiple countries**, and your app needs to:

1. Handle **millions of concurrent users** across the globe.
2. Provide **real-time** driver-passenger matching.
3. Ensure **strong consistency** in bookings and payments (you don't want double bookings or lost transactions).
4. **Scale** up during peak times and **scale down** during off-peak hours.
5. Offer **99.999% uptime** — no downtime is acceptable.

Why Not Use MySQL or PostgreSQL?

- **Vertical scaling limits:** Traditional relational databases can only scale so far vertically (CPU, RAM).
- **Complex sharding:** To handle large volumes, you'd need to manually shard your database, which increases complexity.
- **Global latency:** A single-region database can't provide good performance to users around the world.
- **No built-in HA:** You'd need to set up and manage replication, failover, and backups yourself.

Is there only *one* database in a multi-region Spanner instance?

Yes — in Google Cloud Spanner, even in a **multi-region setup**, there is still **one single logical database**.

However, the **data is automatically replicated** across multiple regions **under the hood**. You don't have to manage or see the replication — Spanner handles it for you.

What does multi-region really mean?

When you create a **multi-region instance**, Spanner:

- **Replicates your data** across multiple data centers (typically 3+ regions).
- Uses **synchronous replication** with a **TrueTime API** to ensure **global consistency**.
- Performs **leader election** for each data split, so reads/writes can happen from the nearest low-latency region, depending on the configuration.

Analogy

Imagine your Spanner database is like **a single brain** that's **thinking in multiple places at once**. It doesn't matter where your user is — they interact with the same database, and Spanner makes sure every part of the world sees **one consistent version** of your data.

Example: Multi-region Setup

Say you choose the **nam3** configuration (U.S. multi-region):

- Primary Region: Iowa (us-central1)
- Replica Region 1: South Carolina (us-east1)
- Replica Region 2: Oregon (us-west1)

Spanner will:

- **Write** to the primary region but **synchronously replicate** the data to the others.
- **Read** from the closest replica (if it's a stale-tolerant read) or route to the leader for strong reads.

Benefits of Multi-Region

Feature	Explanation
High availability	Automatic failover if a region goes down
Low latency reads	Local replicas can serve fast, consistent reads
Global consistency	TrueTime guarantees consistent transactions

It's more like adding more power (CPU, memory, and storage capacity) to your existing single logical database — which stays unified and consistent.

What is Sharding?

Sharding is the process of **splitting a large database into smaller, faster, and more manageable pieces** called **shards**.

Each **shard**:

- Is a **subset of your total data**.
 - Lives on a **separate database or server**.
 - Handles only a portion of the load (queries, reads/writes).
-

Why Use Sharding?

When a single database server can't:

- Store all your data,
- Handle all your queries (CPU/RAM limits),

...you **shard** it across multiple machines.

Real-Life Analogy: Library System

Imagine you run a **giant library** with 10 million books.

Without Sharding:

- All books are in one giant room.
- One librarian handles everything (slow, crowded, hard to search).

With Sharding:

- You split books into sections (shards) by category or author.
- Each section has its own room and librarian.
- Readers go to the correct section directly (faster, less load per room).

Example in Tech

Imagine a user database with 100 million users.

Before Sharding:

- All users are stored in one table on one machine:

sql

CopyEdit

```
SELECT * FROM users WHERE user_id = 723948;
```

After Sharding by User ID:

- You split users across 4 databases:
 - Shard 1: user_id 1–25M
 - Shard 2: user_id 25M–50M
 - Shard 3: user_id 50M–75M
 - Shard 4: user_id 75M–100M
- Your app routes queries to the correct shard.

Spanner does sharding

Historically, Spanner scaled in **full nodes** (each node = 1000 PUs).

```
CREATE TABLE Users (
  UserId STRING(36) NOT NULL,
  Name STRING(100),
  Email STRING(100),
  CreatedAt TIMESTAMP OPTIONS (allow_commit_timestamp=true),
) PRIMARY KEY(UserId);
```

```
import random
```

```
import uuid
```

```
from google.cloud import spanner
```

```
import threading
```

```
import time
```

```
# Initialize Spanner client
```

```
spanner_client = spanner.Client()
```

```
instance = spanner_client.instance('mytest')
```

```
database = instance.database('test')
```

```
QUERY = """
```

```
SELECT COUNT(*)
```

```
FROM Users,
```

```
UNNEST(GENERATE_ARRAY(1, 5000)) AS x
WHERE LENGTH(Name) > 2
```

```
"""
```

```
def generate_user():
    """Create fake user data."""
    return (
        str(uuid.uuid4()),
        f"User{random.randint(1000, 9999)}",
        f"user{random.randint(1000, 9999)}@example.com",
        spanner.COMMIT_TIMESTAMP
    )
```

```
def insert_users(batch_size=10):
    """Insert users in batch."""
    try:
        with database.batch() as batch:
            rows = [generate_user() for _ in range(batch_size)]
            batch.insert(
                table='Users',
                columns=('UserId', 'Name', 'Email', 'CreatedAt'),
                values=rows
            )
        print(f"Inserted {batch_size} users.")
    except Exception as e:
        print("Insert failed:", e)
```

```
def write_worker(batch_size=10, delay=0.5):
    """Worker that writes users in a loop."""
    while True:
        insert_users(batch_size=batch_size)
        time.sleep(delay) # simulate continuous but controlled traffic
```

```
def run_continuous_load(thread_count=20, batch_size=10, delay=0.5):
    """Start threads that continuously generate insert load."""
    print(f"Starting {thread_count} threads...")
    threads = []
    for _ in range(thread_count):
        t = threading.Thread(target=write_worker, args=(batch_size, delay), daemon=True)
        t.start()
        threads.append(t)
```

```
# Keep main thread alive forever
try:
    while True:
        time.sleep(60)
except KeyboardInterrupt:
    print("Stopped by user.")
```

```

if __name__ == "__main__":
    run_continuous_load(thread_count=20, batch_size=10, delay=0.2)

def generate_users(n):
    first_names = ["Alice", "Bob", "Charlie", "Diana", "Eve", "Frank", "Grace", "Hank", "Ivy", "Jack"]
    last_names = ["Smith", "Johnson", "Lee", "Brown", "Williams", "Miller", "Jones", "Davis", "Garcia",
                  "Martinez"]

    users = []
    for _ in range(n):
        uid = str(uuid.uuid4())
        name = f"{random.choice(first_names)} {random.choice(last_names)}"
        email = f"{name.replace(' ', '.').lower()}@example.com"
        created_at = datetime.now(UTC) - timedelta(days=random.randint(0, 365))
        users.append((uid, name, email, created_at.isoformat()))
    return users

# Simulate CPU stress by generating multiple large insert statements
def generate_insert_batches(batch_size, batch_count):
    batches = []
    for _ in range(batch_count):
        users = generate_users(batch_size)
        insert_values = ",\n".join(
            f'("{uid}", "{name}", "{email}", "{created_at}Z")'
            for uid, name, email, created_at in users
        )
        sql = (
            "INSERT INTO Users (UserId, Name, Email, CreatedAt) VALUES\n" +
            insert_values + ";"
        )
        batches.append(sql)
    return batches

# Generate 10 batches of 500 rows each (5,000 rows total)
stress_batches = generate_insert_batches(batch_size=500, batch_count=10)

# Return the first few characters of one example for demonstration
stress_batches[0][:1000] # Show only the beginning for readability

```