



Deepti Srivastava
Product Manager

Cloud Spanner



Google Cloud

What is tricky with many databases?

1

Scale

Vertical scaling
Sharding

2

Operations

HW breakdowns
SW upgrades
End-of-life
Tweaking
Rebuild indexes

3

Replication

Read only replicas
Disaster recovery replicas

4

Transactions

Strong consistency

Cloud Spanner



**Relational
semantics**

Schemas, ACID
transactions, SQL



**Horizontal
scale**

99.999% SLA, fully
managed, and scalable

What workloads fit Cloud Spanner best?

01

Sharded RDBMS

Manually sharding is difficult. People do it to achieve scale. Cloud Spanner gives you relational data and scale.

02

Scalable relational data

Cloud Spanner is a scalable relational database. Instead of moving to NoSQL, move from one relational database to a more scalable relational database.

03

Manageability/HA

Cloud Spanner is highly automated. Schema changes and patching are all online operations. Cloud Spanner does not have any planned downtime and comes with up to a 99.999% availability SLA.

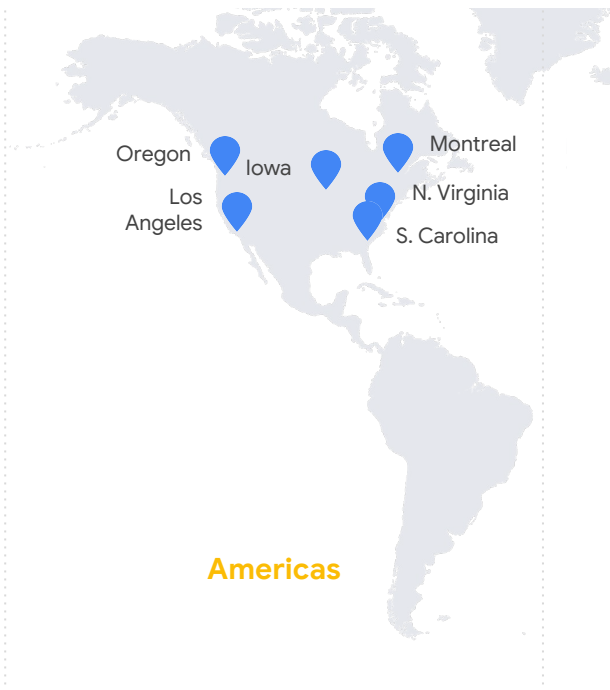
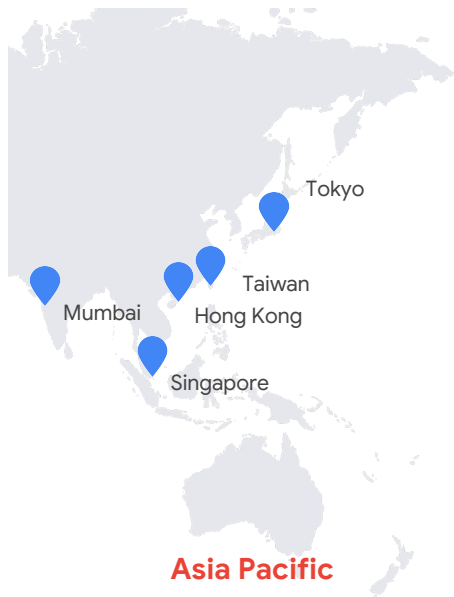
04

Multi-region

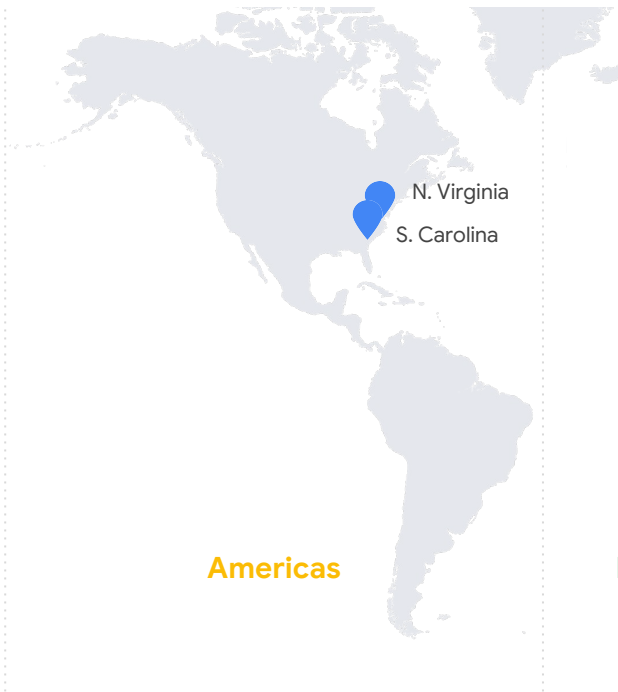
Write once and have Cloud Spanner automatically replicate your data to multiple regions.

Most customers use regional instances, but multi-region is there if you need it.

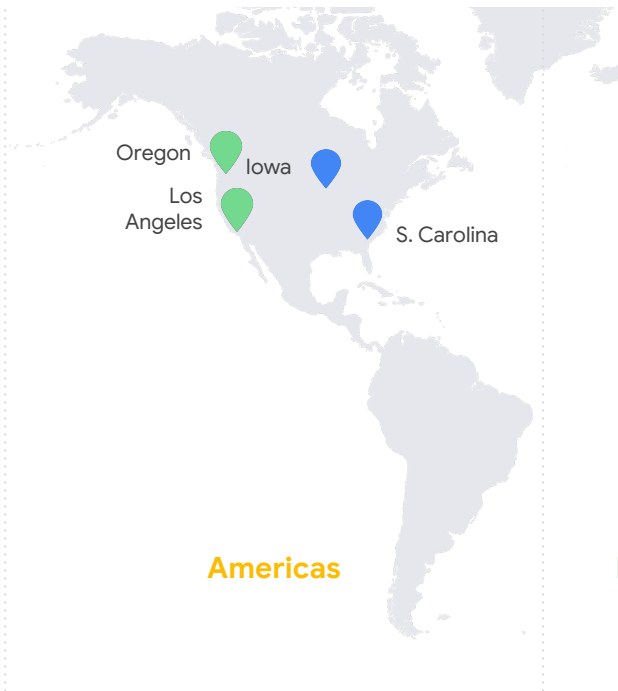
Most commonly used: regional configurations



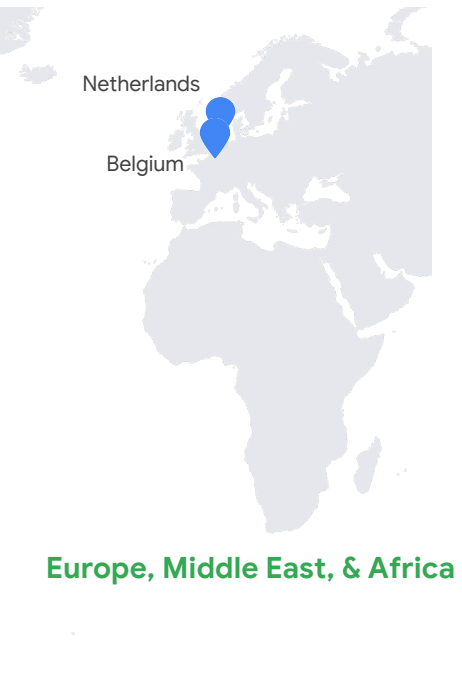
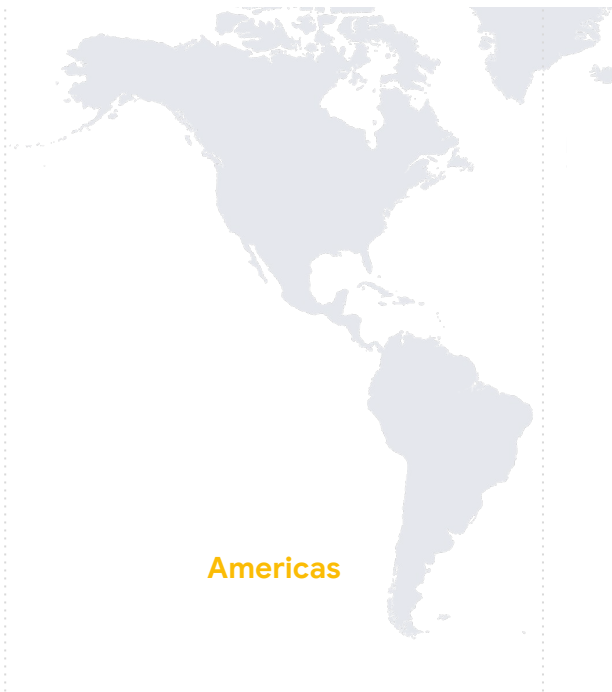
Multi-region configuration (nam3)



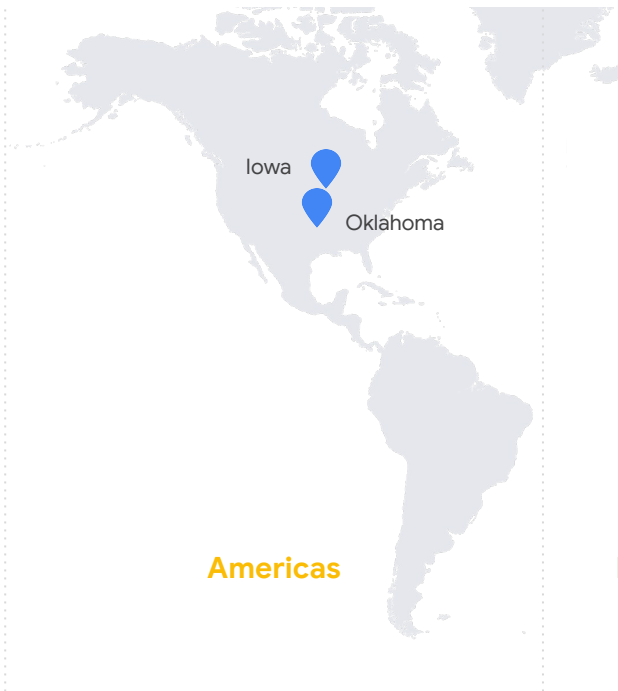
Multi-region configuration (nam6)



Multi-region configuration (eur3)



Multi-region configuration (nam-eur-asia1)



When Cloud Spanner fits less well

1

Lift and shift

2

Lots of
in-database
business logic
(triggers, stored
procedures)

3

Compatibility
needed

4

App is very sensitive
to very low latency
(micro/nano/low
single digit ms)

Common misconceptions

Fully managed – yes, really!

- No messing with tablespaces
- No software patching
- No upgrades
- No downtime
- No index rebuilds
- No resharding
- Dynamically scalable (up and down) within seconds



Common misconceptions

External (strong) consistency –
yes, really!

- Low-latency synchronous replication — data is available on commit
- True ACID database
- How is this possible? Through Google's private high-speed redundant global network and Cloud Spanner's optimized stack





Common issues

Key choice in tables and indexes

- Sequential- or timestamp-based keys can cause hotspots limiting performance
- Random UUIDs or hashes distribute data better

Lack of: Triggers, stored procedures, and check constraints

- Not Null constraint exists, Keys are Unique, Unique column values via Indexes.

Foreign key constraints and on-delete cascades only with parent-child interleaved tables

No sequence generation for keys



Common issues

Indexes are tables too

- Same constraints for index keys
- Querying non-covering index is a table join

Primary key is the ROWID

- Uniquely identifies the row
- Used to determine where the row is in the storage

Backup and DR

- Managed export / import
- One hour of historical data snapshots
- System-managed backup coming soon

How Cloud Spanner works



Architecture overview

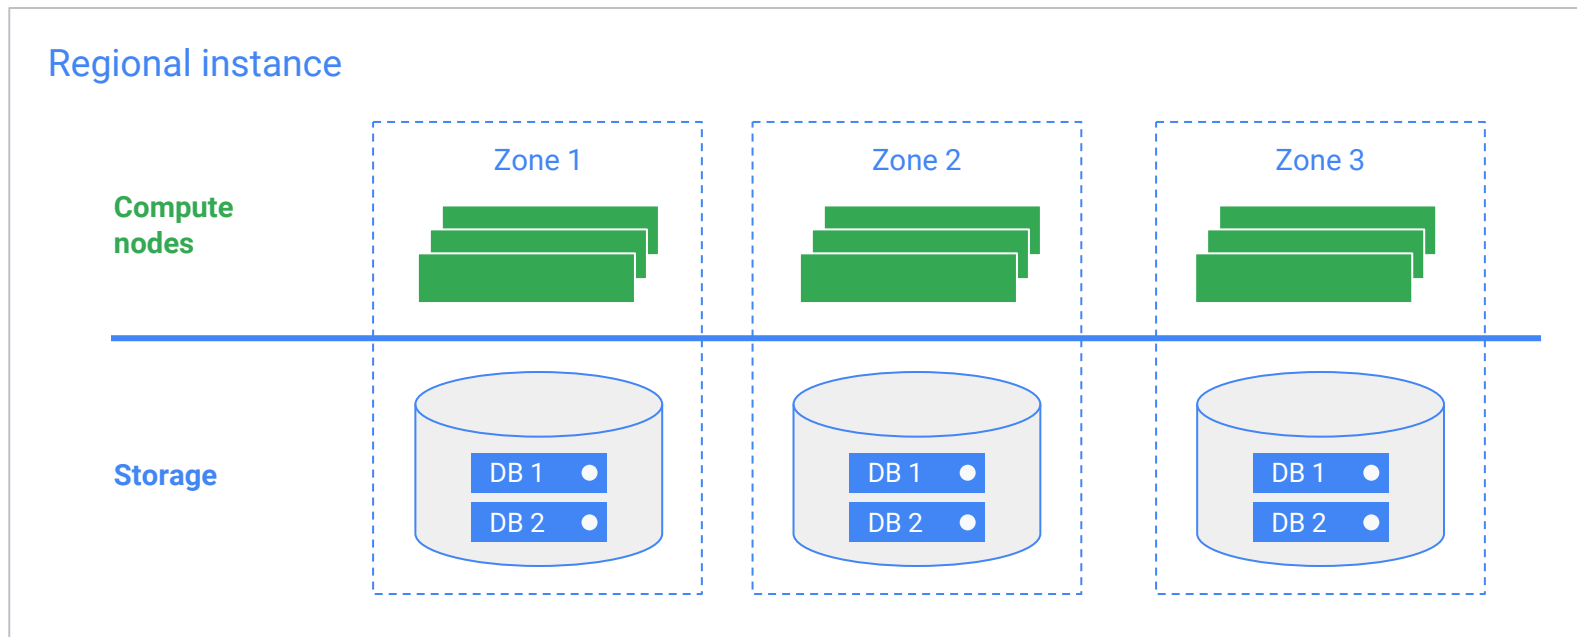


Table splits

KeyRange	Split ID	
$[-\infty, 3]$	0	Node A
[4, 224]	1	
[225, 712]	2	
[713, 717]	3	Node B
[718, 1265]	4	
[1725, 1997]	5	
[1998, 2456]	6	Node C
[1998, 2456]	7	
[2457, ∞]	8	

Table splits

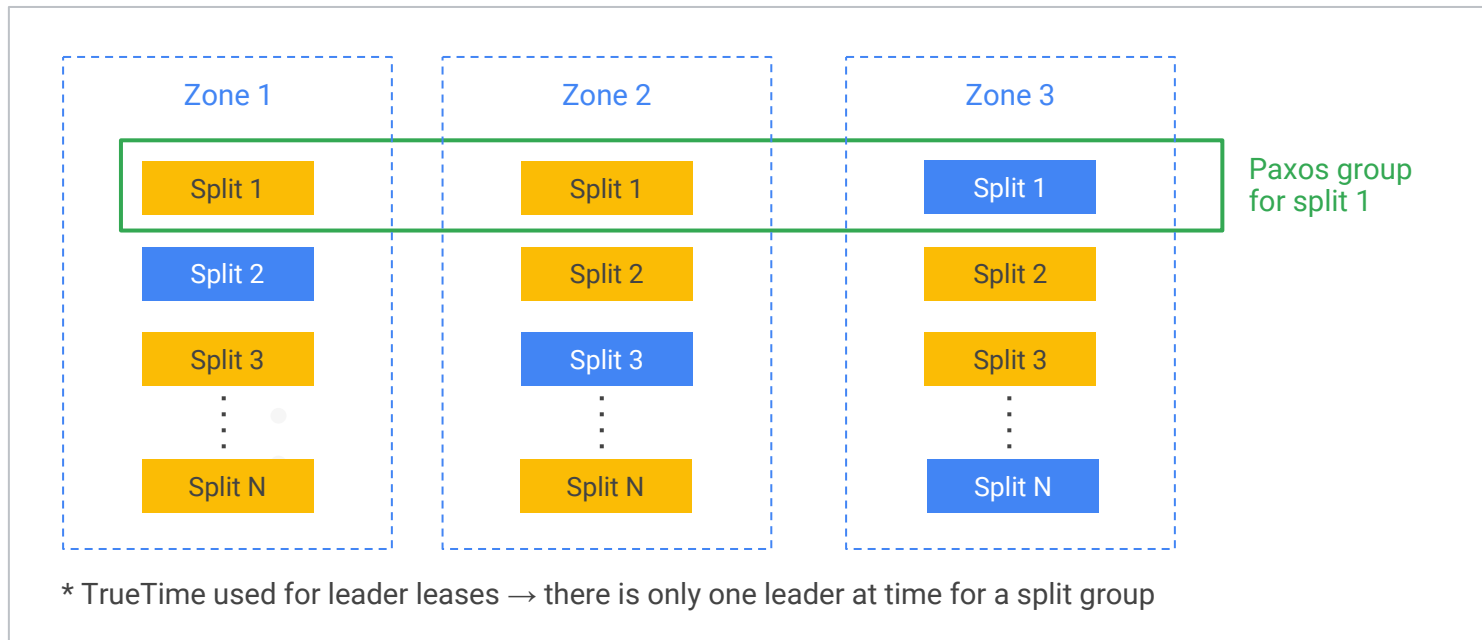
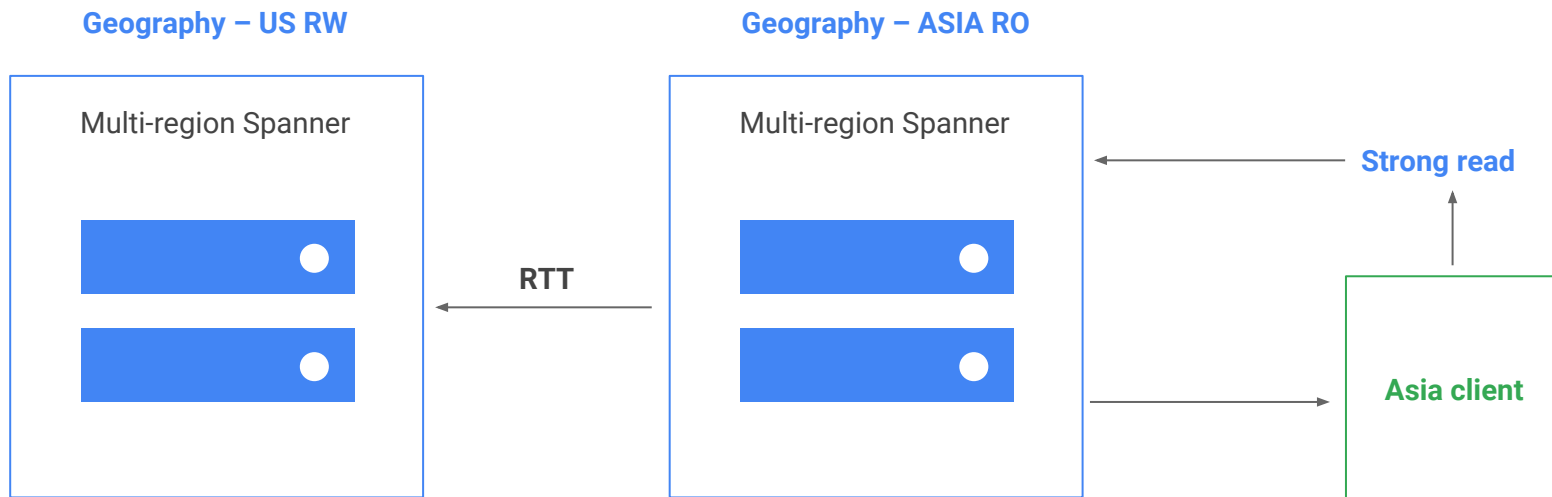


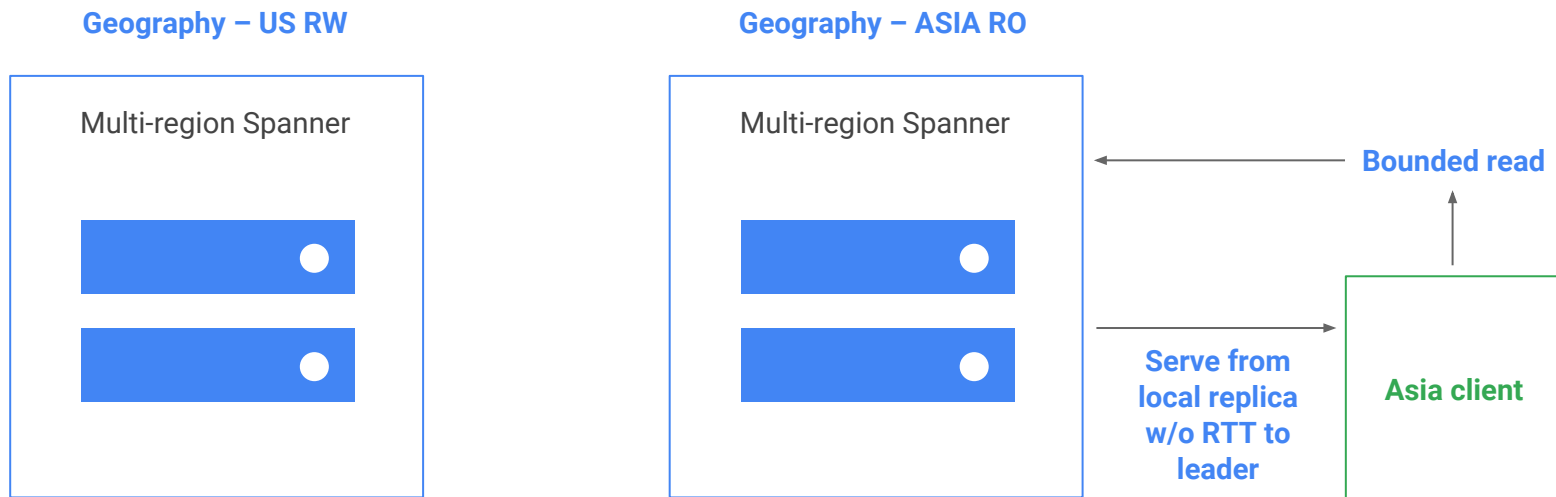
Table splits

KeyRange	Split ID
$[-\infty, 3]$	0
[4, 224]	1
[225, 712]	2
[713, 717]	3
[718, 1265]	4
[1725, 1997]	5
[1998, 2456]	6
[1998, 2456]	7
[2457, ∞]	8

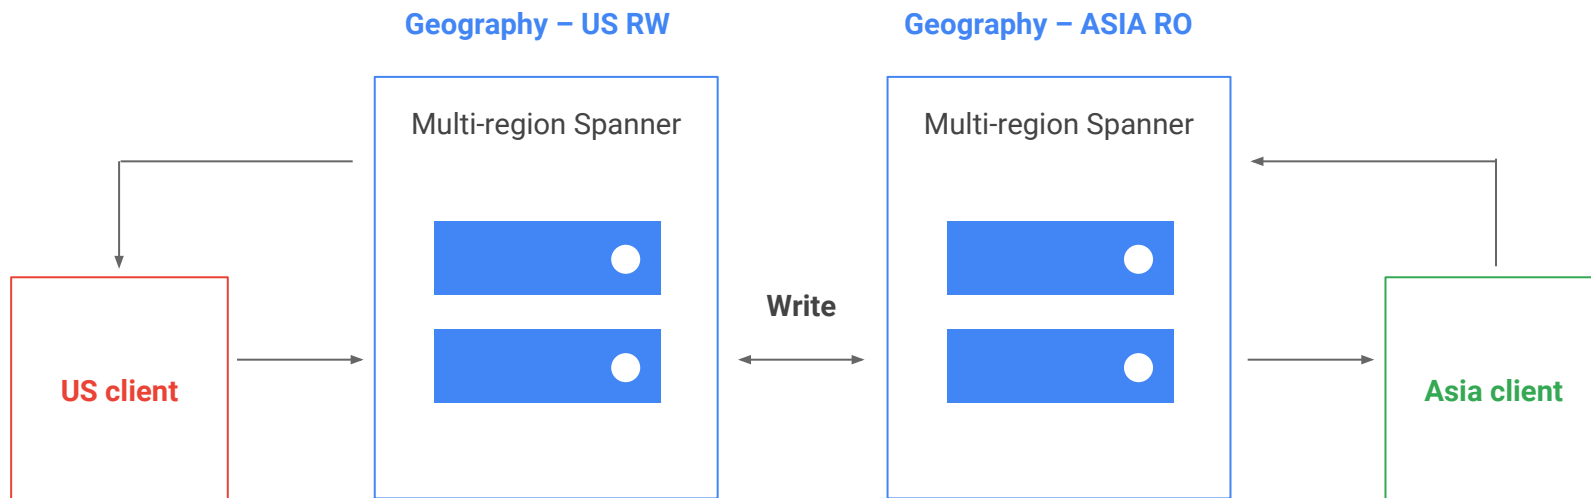
Multi-region strong read request



Multi-region stale/bounded read request



Multi-region write request



Portfolio of fully managed services

Great apps run on great databases.

In-memory



Cloud Memorystore

Managed Redis

Non-relational



Cloud Firestore

Serverless, document database service



Cloud Bigtable

Wide-column database service

Relational



Cloud SQL

Managed MySQL & PostgreSQL



Cloud Spanner

Scalable relational database service

Object



Cloud Storage

Object storage, data lake

Warehouse



BigQuery

Enterprise data warehouse



Q&A

Google Cloud



Thank you

More info

cloud.google.com/spanner

Google Cloud