Google Cloud Next

Cloud Spanner

Rohit Gupta, Solutions Engineer
@rohitforcloud

Today’s goals

- Provide a brief history of Spanner at Google
- Provide an explanation of Cloud Spanner
- Do a demo!

Google Cloud Platform

Built on the same infrastructure that powers Google

- Super-flexible compute
- Superior economics
- Always available storage
- Fastest, most reliable network
- Robust, easy to use Big Data solutions

Google Cloud Platform

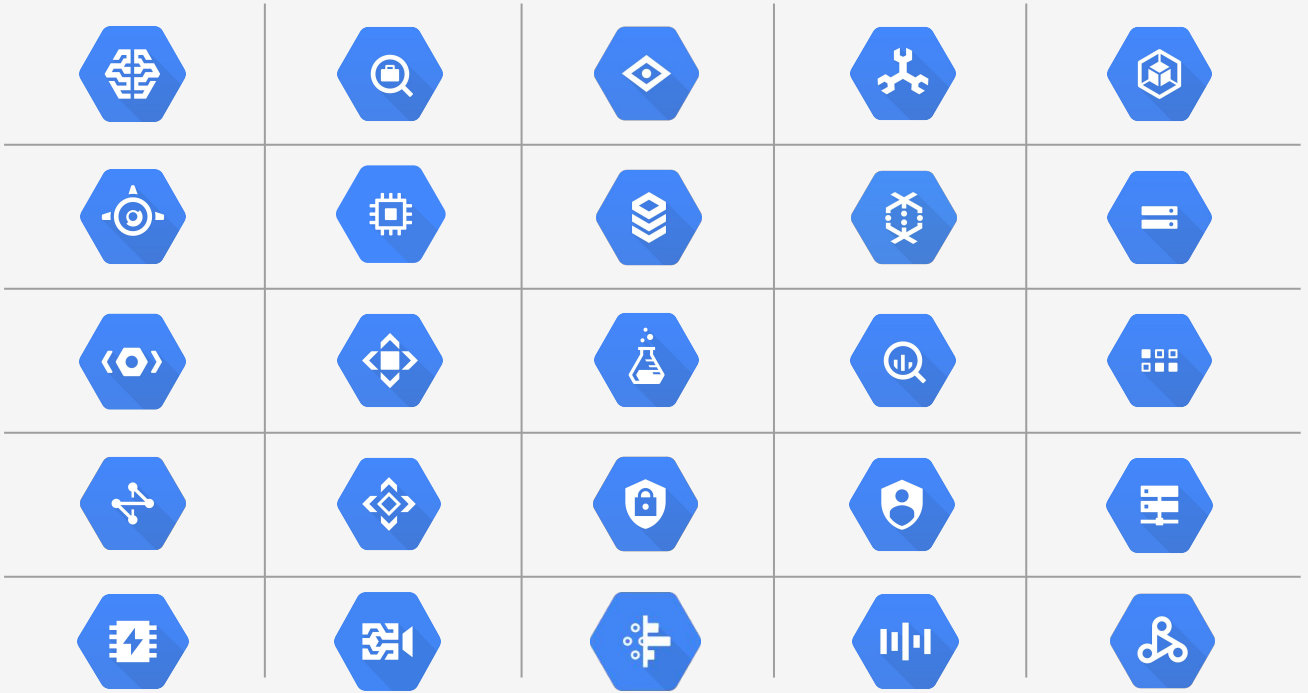

GCP Regions

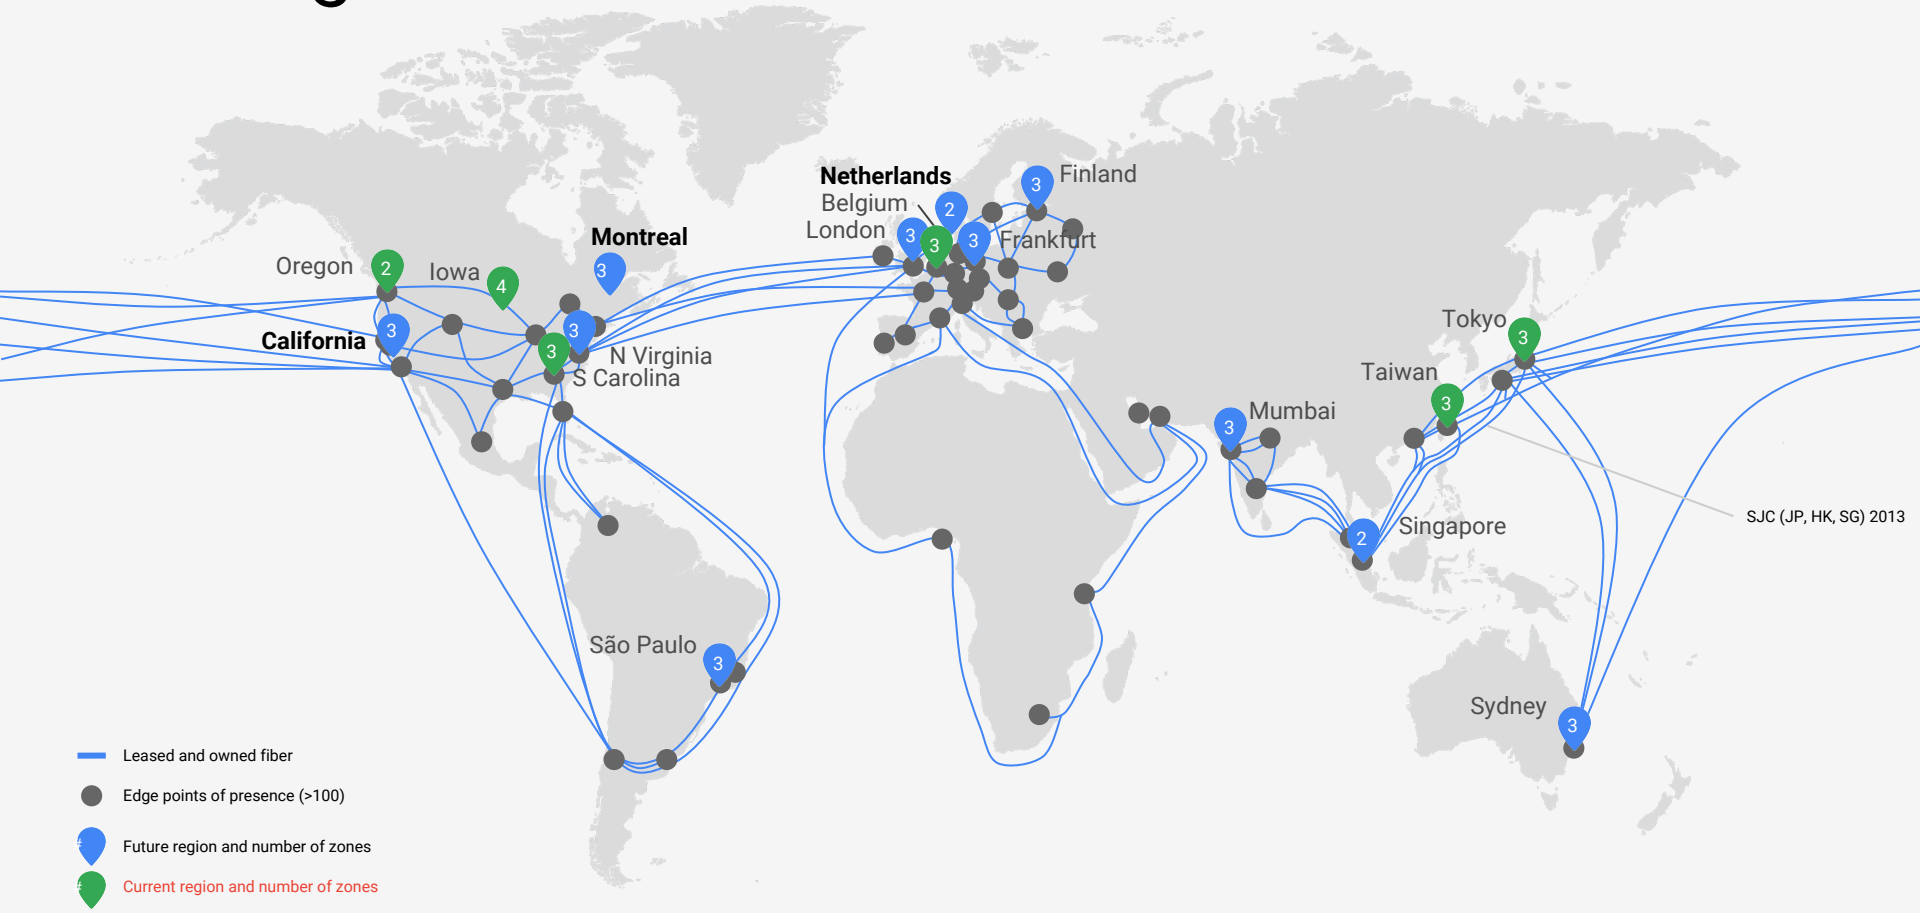
Oregon
Iowa
Montreal
California
N Virginia
S Carolina
Netherlands
Belgium
London
Finland
Frankfurt
Mumbai
Taiwan
Tokyo
Singapore
SJC (JP, HK, SG) 2013
São Paulo
Sydney
Leased and owned fiber
Edge points of presence (>100)
Future region and number of zones
Current region and number of zones

Background

Why build Spanner?

Technical infrastructure at Google

It’s 2005...

Google’s needs

- Horizontally Scaling Database
- ACID Transactions with global consistency
- No downtime!

Overview

What is Cloud Spanner?

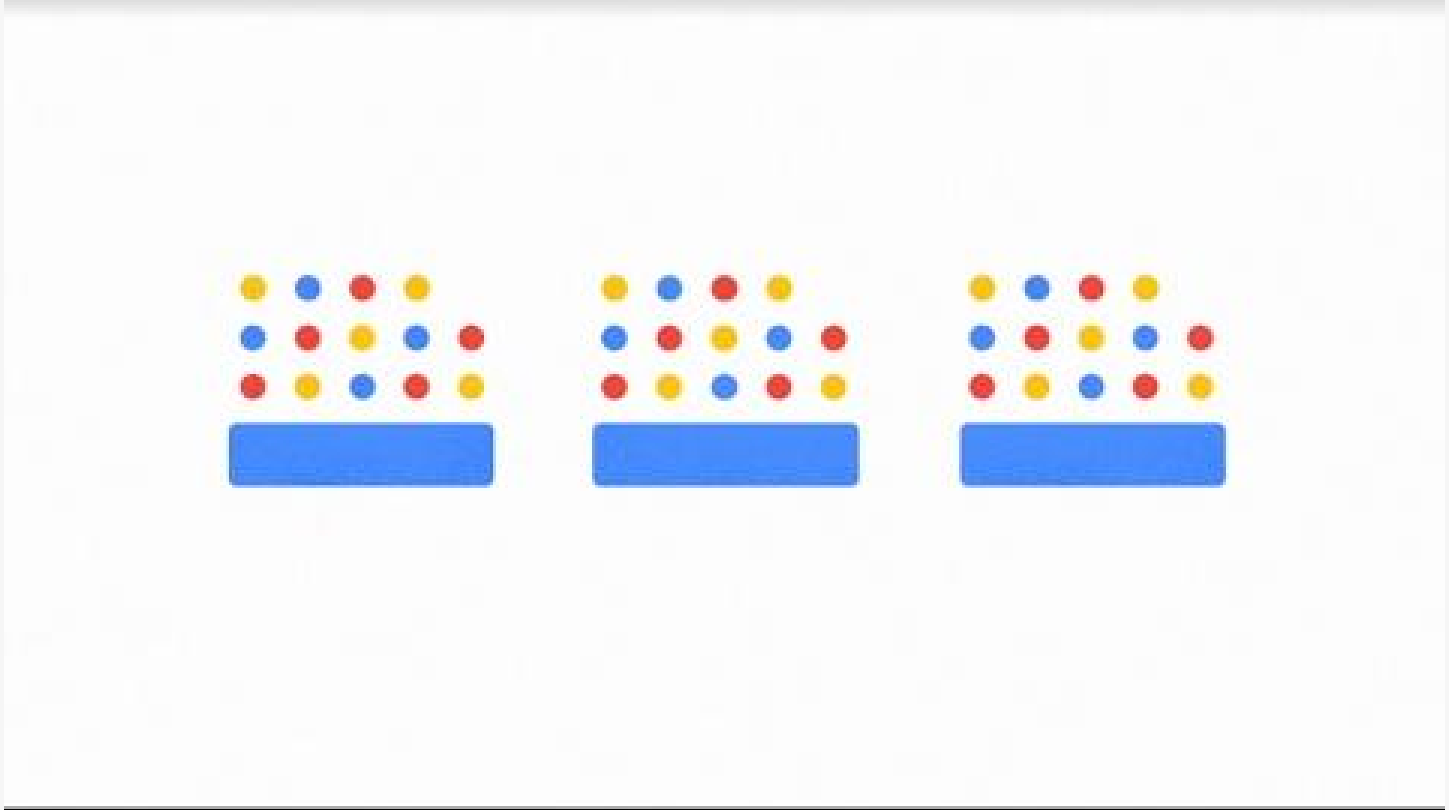

What is Cloud Spanner?

Google's mission-critical scalable relational Database Service

- Fully managed, database service with global scale
- Traditional relational semantics: schemas, ACID transactions, SQL
- Automatic, synchronous replication within and across regions for availability
- Battle-tested within Google for 5+ yrs (AdWords, GooglePlay)

How does it compare?

	CLOUD SPANNER	TRADITIONAL RELATIONAL	TRADITIONAL NON-RELATIONAL
Schema	✓ **Yes**	✓ Yes	✗ No
SQL	✓ **Yes**	✓ Yes	✗ No
Consistency	✓ **Strong**	✓ Strong	✗ Eventual
Availability	✓ **High**	✗ Failover	✓ High
Scalability	✓ **Horizontal**	✗ Vertical	✓ Horizontal
Replication	✓ **Automatic**	↻ Configurable	↻ Configurable

Open standards

- Standard SQL (ANSI 2011)
- Encryption, Audit logging, Identity and Access Management
- Client libraries in popular languages (Java, Python, Go, Node.js, etc.)
- JDBC driver

Types of workloads

Transactional

Companies that have outgrown their single-instance RDBMS and have already moved to NoSQL solution, but need transactional consistency, or they are looking to move to a scalable solution

Scale-out

Companies currently sharding databases because they need more read or write throughput than can be placed on a single node

Global data plane

Companies and/or developers building applications that have global data and need strong consistency

Database consolidation

Companies that store their business data in multiple database products with variable maintenance overheads and capabilities and need consolidation of their data

Architecture overview

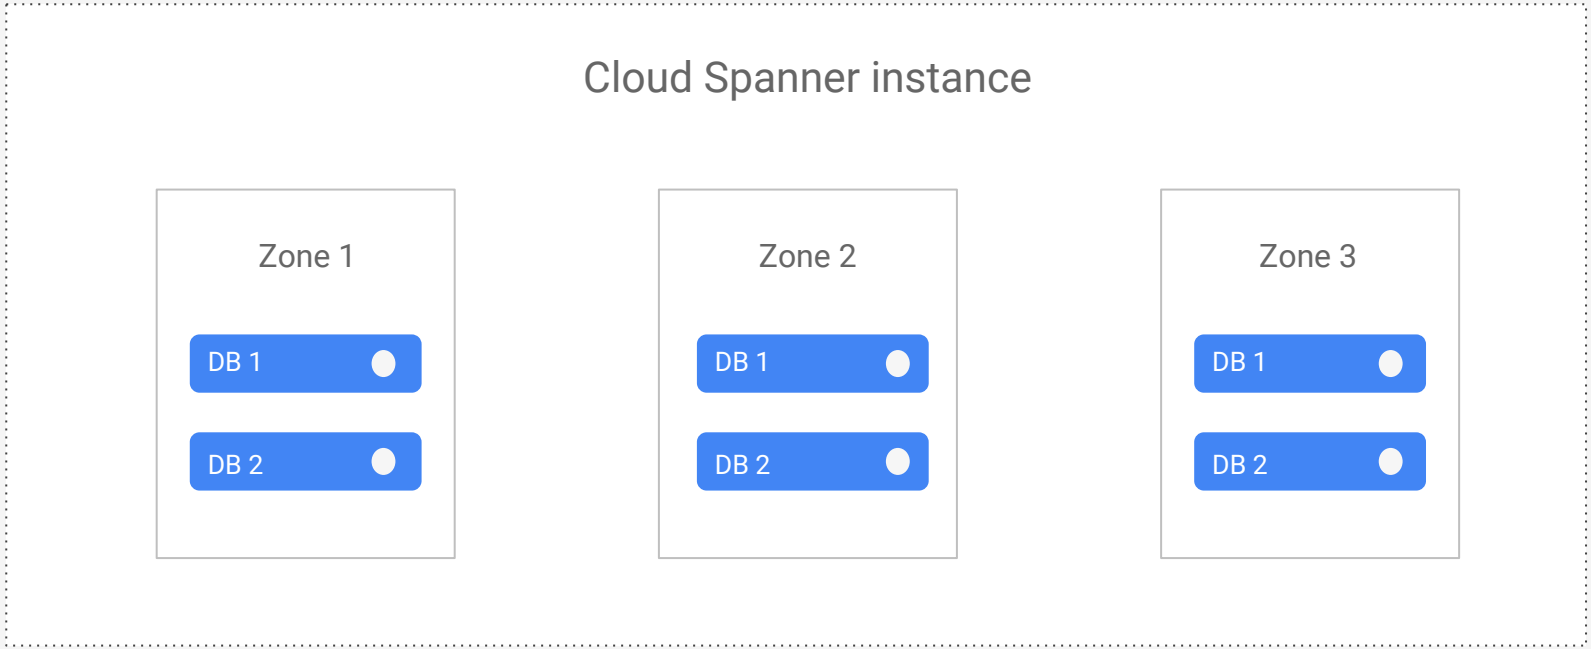

Data replication

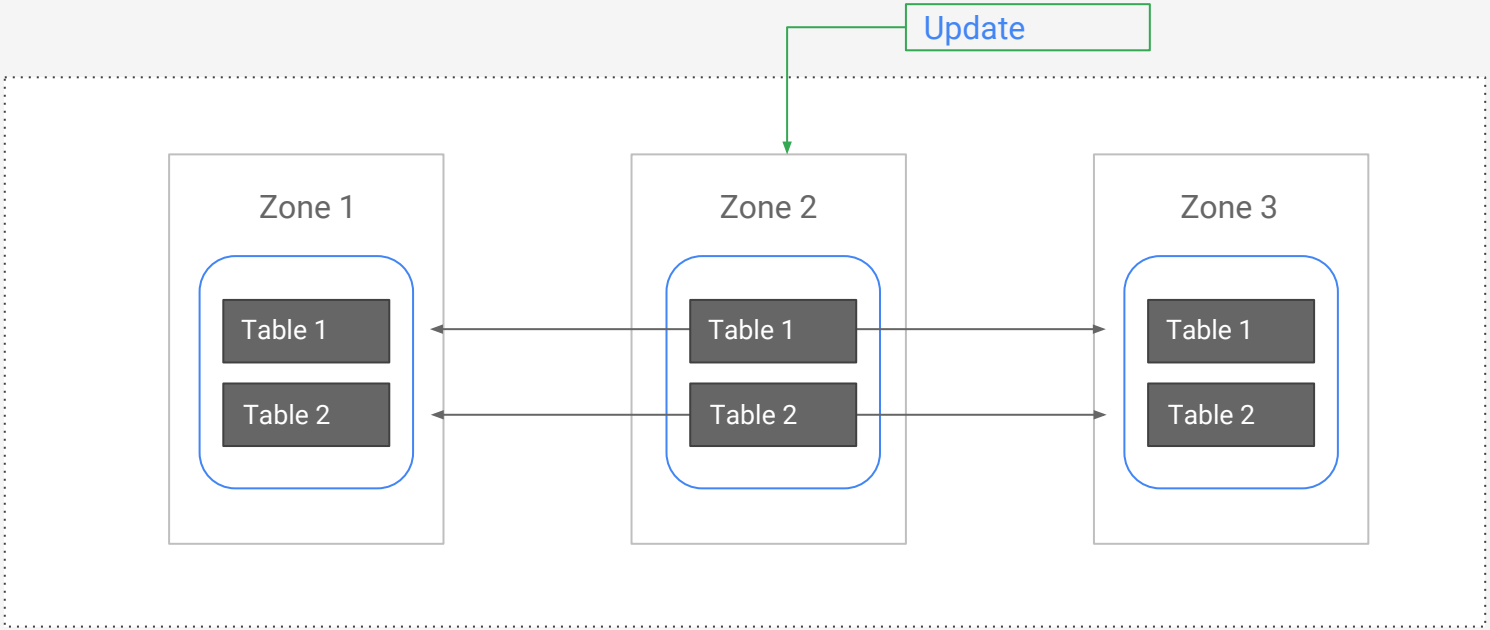

Relational Data Layout

SingerId	SingerName
1	Beatles
2	U2
3	Pink Floyd

SingerId	AlbumId	AlbumName
1	1	Help!
1	2	Abbey Road
3	1	The Wall

Interleave Data Layout

1	Beatles	
1	1	Help!
1	2	Abbey Road
2	U2	
3	Pink Floyd	
3	1	The Wall

Relational data model

```
CREATE TABLE Singers (
  SingerId INT64 NOT NULL,
  SingerName STRING(MAX),
) PRIMARY KEY(SingerId);

CREATE TABLE Albums (
  SingerId INT64 NOT NULL,
  AlbumId INT64 NOT NULL,
  AlbumName STRING(MAX),
) PRIMARY KEY(SingerId, AlbumId)
INTERLEAVE IN PARENT Singers;
```

No Downtime
Schema Migrations

```
ALTER TABLE Singers
ADD COLUMN Age INT64;
```

Getting started in Python

```python
# Imports the Google Cloud Client Library.
from google.cloud import spanner

# Instantiate a client.
spanner_client = spanner.Client()

# Your Cloud Spanner instance ID.
instance_id = 'my-instance-id'

# Get a Cloud Spanner instance by ID.
instance = spanner_client.instance(instance_id)

# Your Cloud Spanner database ID.
database_id = 'my-database-id'

# Get a Cloud Spanner database by ID.
database = instance.database(database_id)

# Execute a simple SQL statement.
results = database.execute_sql('SELECT 1')

for row in results:
    print(row)
```

Storage & Database Portfolio

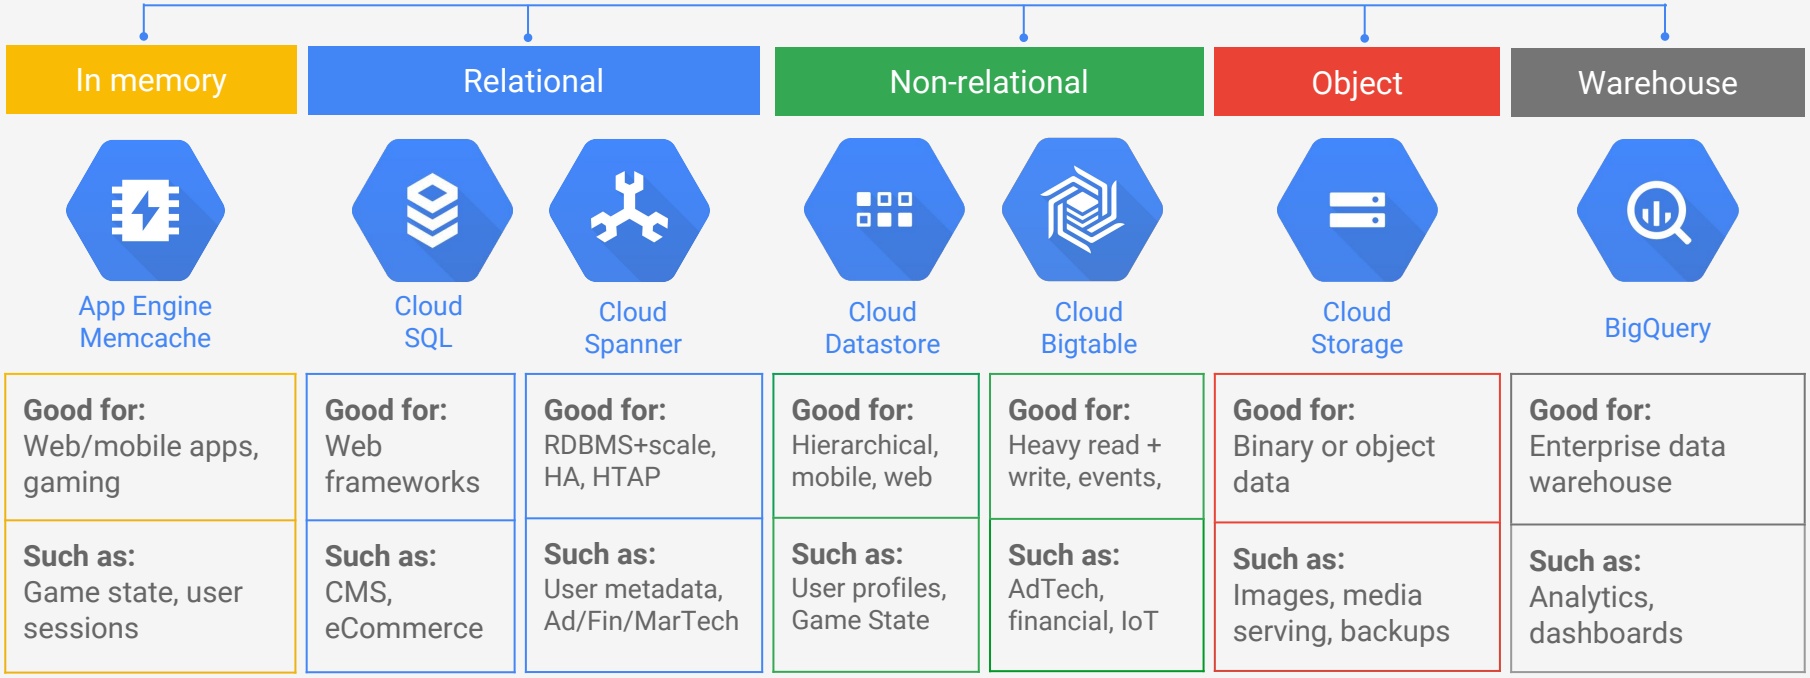

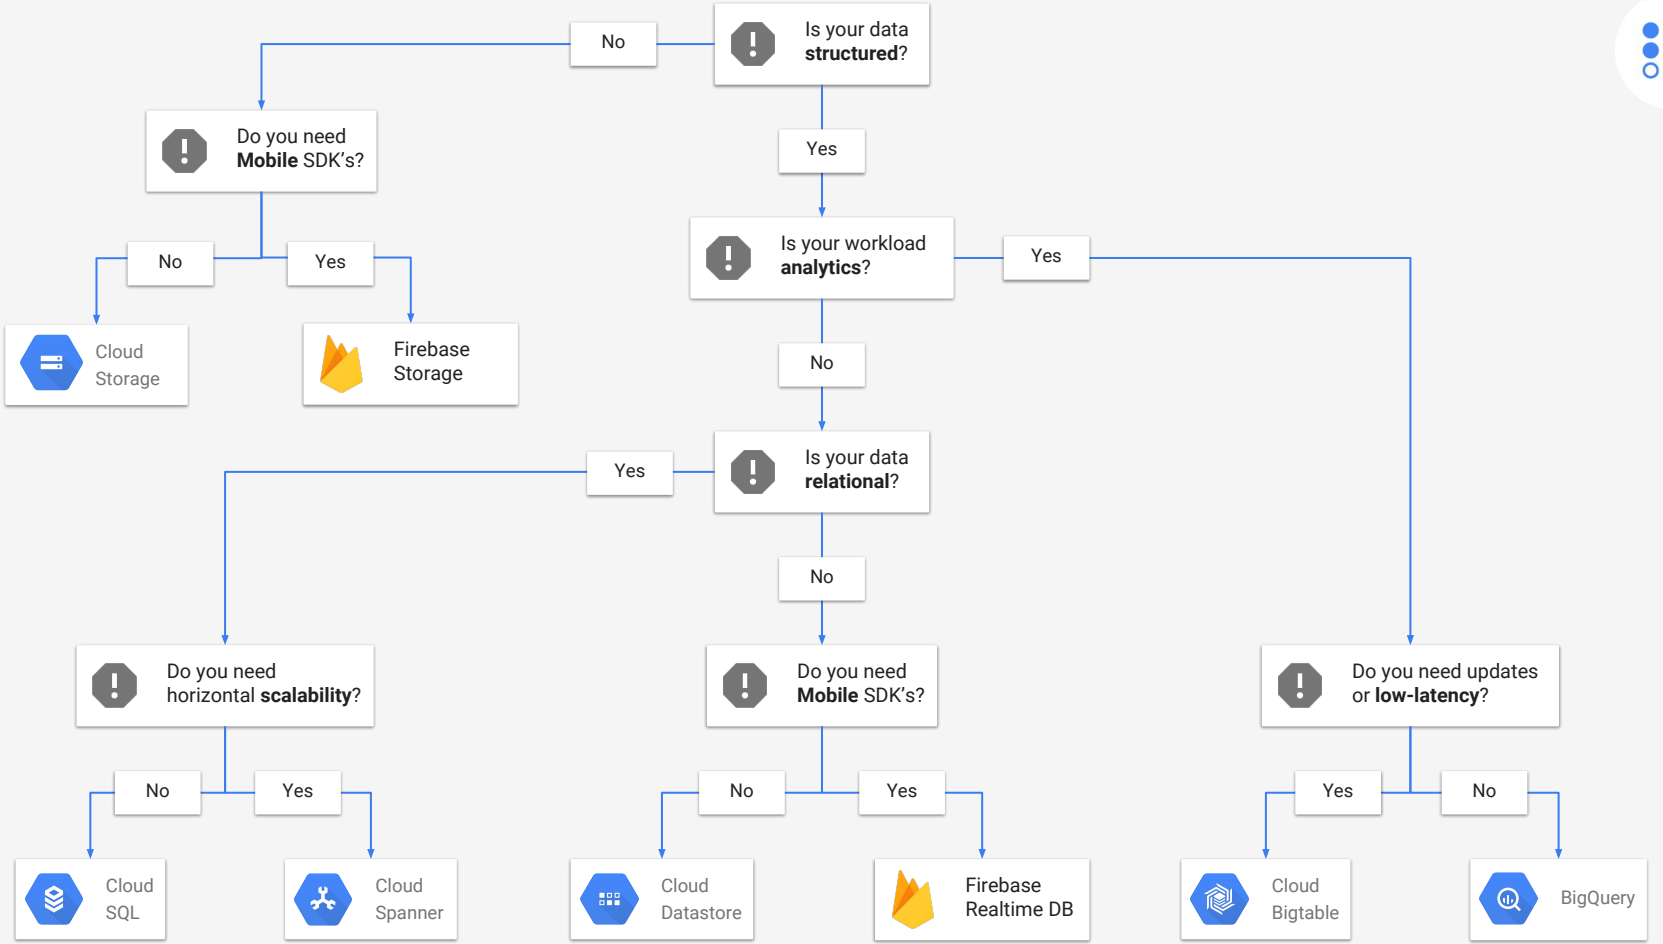
Is your data structured?
No
Yes
Do you need Mobile SDK's?
No
Yes
Cloud Storage
Firebase Storage
Is your workload analytics?
Yes
No
Is your data relational?
Yes
No
Do you need horizontal scalability?
No
Yes
Cloud SQL
Cloud Spanner
Do you need Mobile SDK's?
No
Yes
Cloud Datastore
Firebase Realtime DB
Do you need updates or low-latency?
Yes
No
Cloud Bigtable
BigQuery

Partner integrations

Launch partners

Committed partners

Whats Next

Checkout Google Cloud Spanner at cloud.google.com/spanner

Cloud Spanner 201 - youtu.be/Tzhe7sUNDbg

Cloud Spanner Case Study - Quizlet Case Study

Thank you