Building for Growth: Future of Databases Outlook...Cloudy

Dominic Preuss
Director of Product Management - Storage and Databases
Google Cloud Platform

February 22, 2018

Google Cloud

It’s very early...

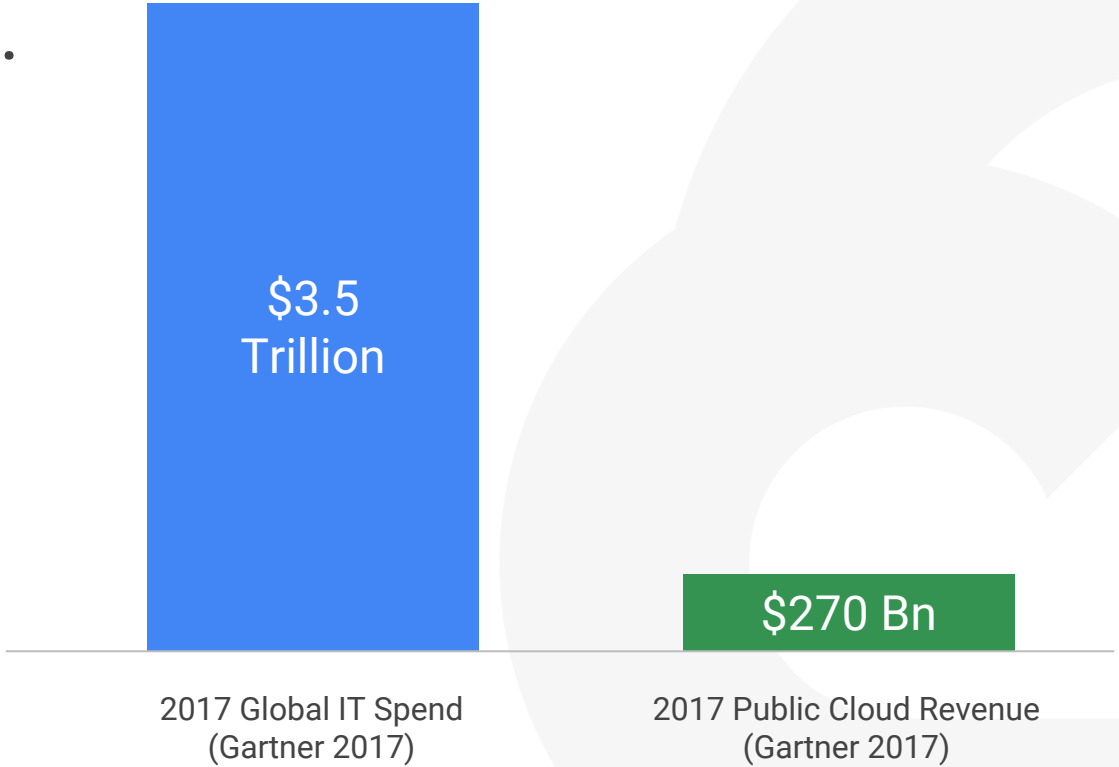

Google Cloud

Source: Gartner Press Release, October 2017 https://www.gartner.com/newsroom/id/3811363

...very, very early

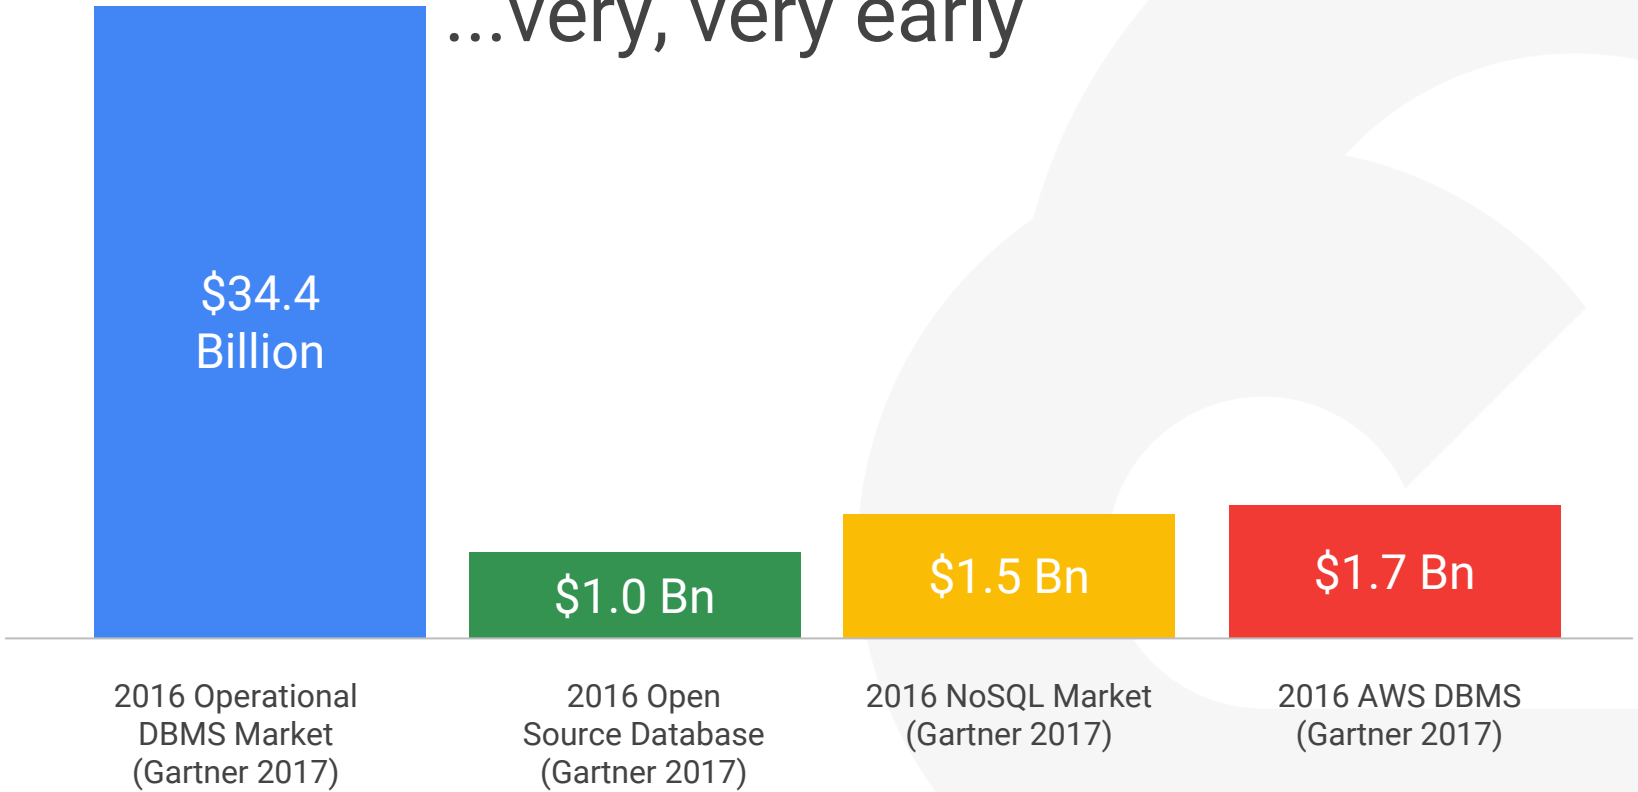

Google Cloud

Gartner: "Magic Quadrant for Operational Database Management Systems" - November 2, 2017, Nick Heudecker, Donald Feinberg, Merv Adrian


What does it mean for an
app to be **mission critical?**
Google Cloud

Current State

Common Growth Challenges

Companies manage their own inflexible database infrastructure on-prem.

Companies run and manage their own database solutions using somewhat less inflexible cloud infrastructure.

Companies pay a managed service provider to run their database solution in the cloud.

Companies run into limitations around RDBMS scale-up/out and sharding solutions.

Companies deal with administration or consistency challenges with scalable NoSQL solutions on-premises or in the cloud.

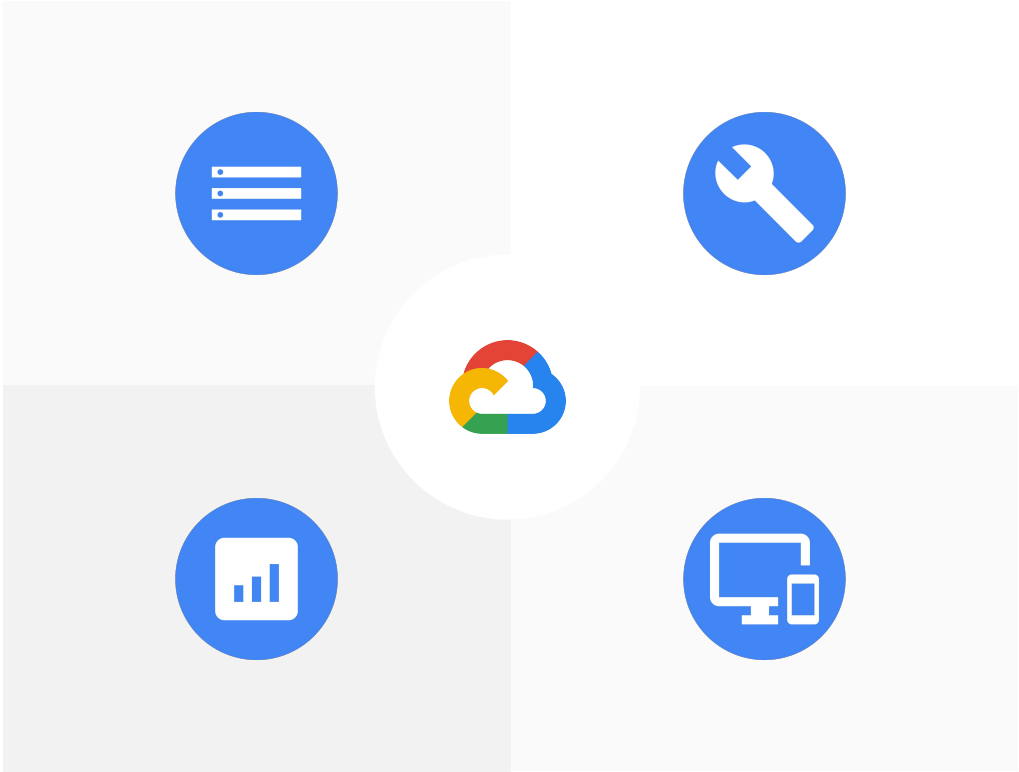

Compromises of Traditional Databases...

		Traditional Relational	Traditional Non-relational
Business Requirements	Schema	✓ Yes	☒ No
	SQL	✓ Yes	☒ No
	Consistency	✓ Strong	☒ Eventual
Technical Requirements	Availability	☒ Failover	✓ High
	Scalability	☒ Vertical	✓ Horizontal
	Replication	↻ Configurable	↻ Configurable

Lead to Complex Architectures

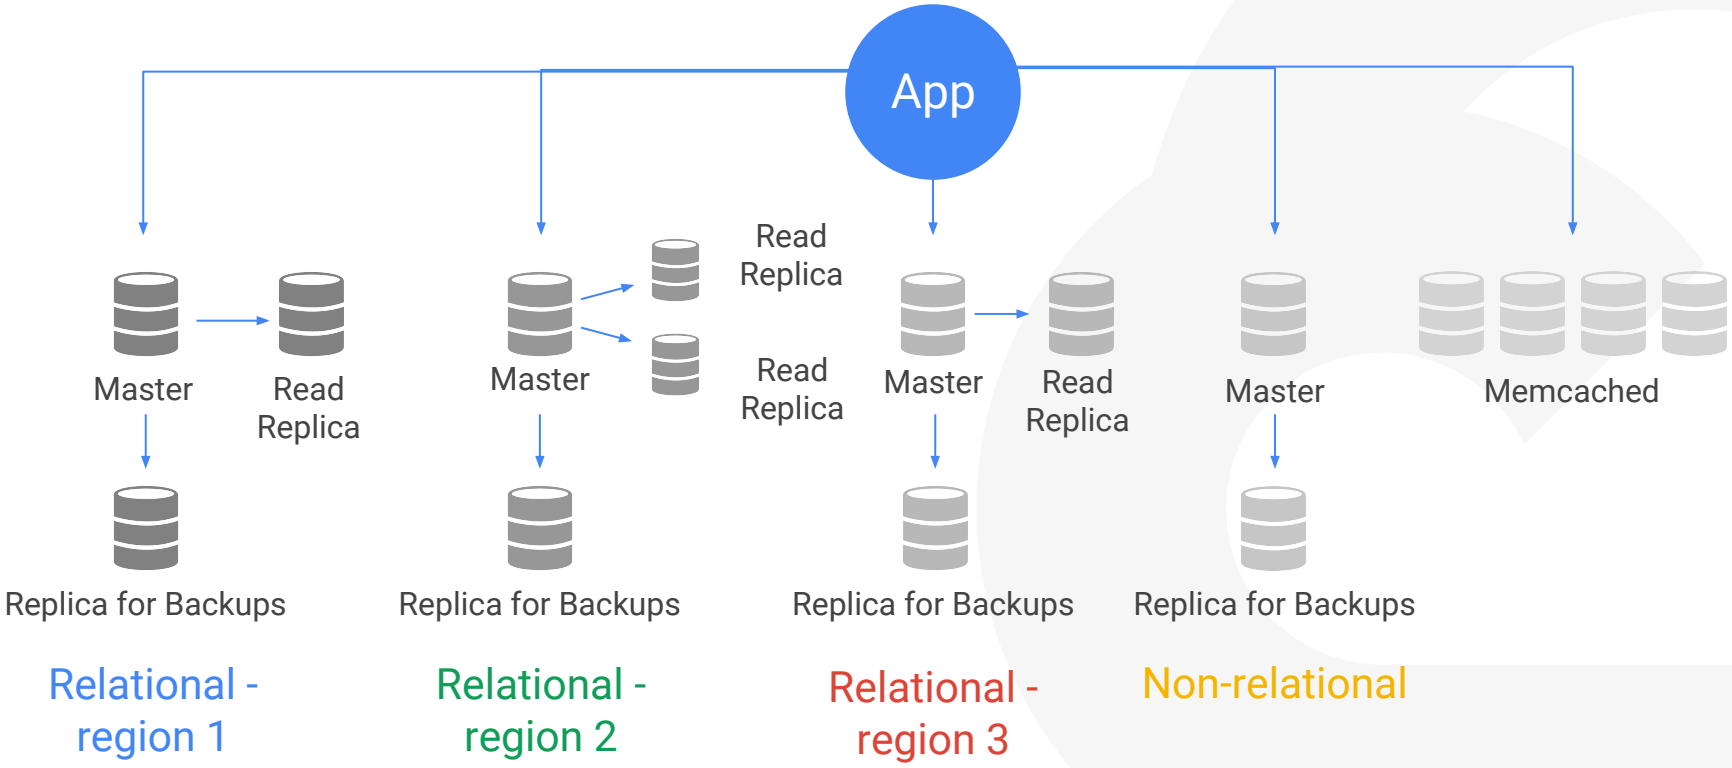

What do customers want?

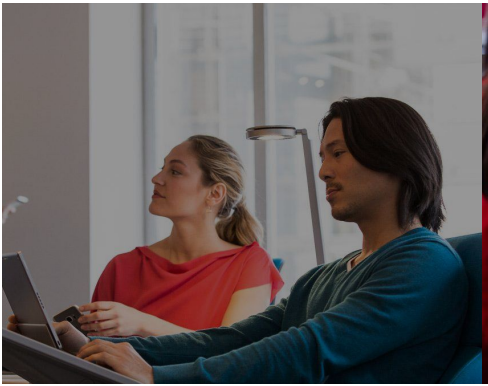


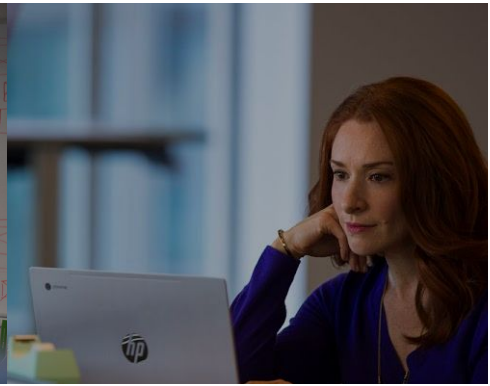

Availability
Scalability
Consistency

Simplified
Operations

Standards
and Ecosystem

Single Point-
of-Contact
for Support

Google Cloud

Technical infrastructure at Google

It’s 2005...

Google’s needs

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

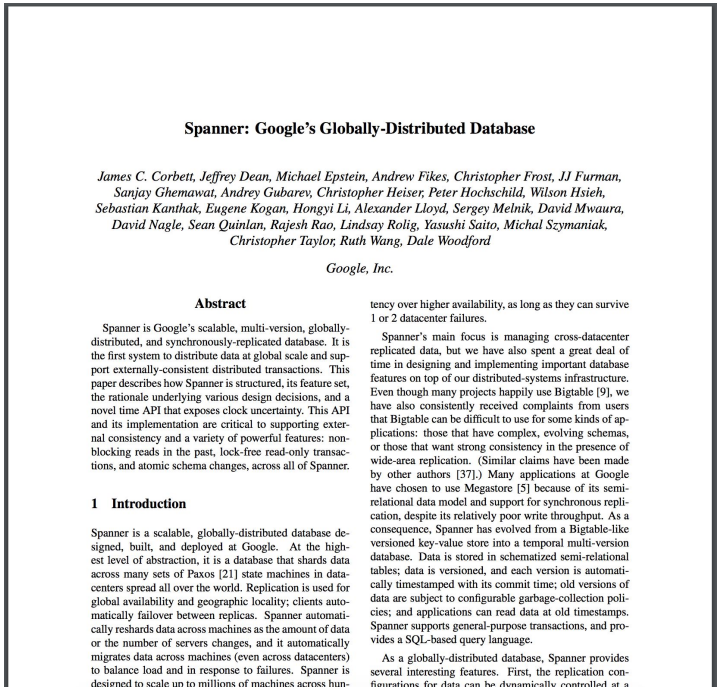

Spanner: Google’s Globally-Distributed Database

James C. Corbett, Jeffrey Dean, Michael Epstein, Andrew Fikes, Christopher Frost, JJ Furman, Sanjay Ghemawat, Andrey Gubarev, Christopher Heiser, Peter Hochschild, Wilson Hsieh, Sebastian Kanthak, Eugene Kogan, Hongyi Li, Alexander Lloyd, Sergey Melnik, David Mwaura, David Nagle, Sean Quinlan, Rajesh Rao, Lindsay Rolig, Yasushi Saito, Michal Szymaniak, Christopher Taylor, Ruth Wang, Dale Woodford

Google, Inc.

Abstract

Spanner is Google’s scalable, multi-version, globally-distributed, and synchronously-replicated database. It is the first system to distribute data at global scale and support externally-consistent distributed transactions. This paper describes how Spanner is structured, its feature set, the rationale underlying various design decisions, and a novel time API that exposes clock uncertainty. This API and its implementation are critical to supporting external consistency and a variety of powerful features: non-blocking reads in the past, lock-free read-only transactions, and atomic schema changes, across all of Spanner.

1 Introduction

Spanner is a scalable, globally-distributed database designed, built, and deployed at Google. At the highest level of abstraction, it is a database that shards data across many sets of Paxos [21] state machines in datacenters spread all over the world. Replication is used for global availability and geographic locality; clients automatically failover between replicas. Spanner automatically reshards data across machines as the amount of data or the number of servers changes, and it automatically migrates data across machines (even across datacenters) to balance load and in response to failures. Spanner is designed to scale up to millions of machines across hun-

tency over higher availability, as long as they can survive 1 or 2 datacenter failures.

Spanner’s main focus is managing cross-datacenter replicated data, but we have also spent a great deal of time in designing and implementing important database features on top of our distributed-systems infrastructure. Even though many projects happily use Bigtable [9], we have also consistently received complaints from users that Bigtable can be difficult to use for some kinds of applications: those that have complex, evolving schemas, or those that want strong consistency in the presence of wide-area replication. (Similar claims have been made by other authors [37].) Many applications at Google have chosen to use Megastore [5] because of its semi-relational data model and support for synchronous replication, despite its relatively poor write throughput. As a consequence, Spanner has evolved from a Bigtable-like versioned key-value store into a temporal multi-version database. Data is stored in schematized semi-relational tables; data is versioned, and each version is automatically timestamped with its commit time; old versions of data are subject to configurable garbage-collection policies; and applications can read data at old timestamps. Spanner supports general-purpose transactions, and provides a SQL-based query language.

As a globally-distributed database, Spanner provides several interesting features. First, the replication con-

Google Cloud

Cloud Spanner

An enterprise-grade, globally-distributed, and strongly-consistent database service built for the cloud specifically to combine the benefits of relational database structure with non-relational horizontal scale.

Google Cloud

Relational semantics

Schemas, ACID transactions, SQL

And

Horizontal-scale

Up to 99.999% Availability SLA,
Fully managed, and Easily scalable

Cloud Spanner makes it simple

	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

Indexes

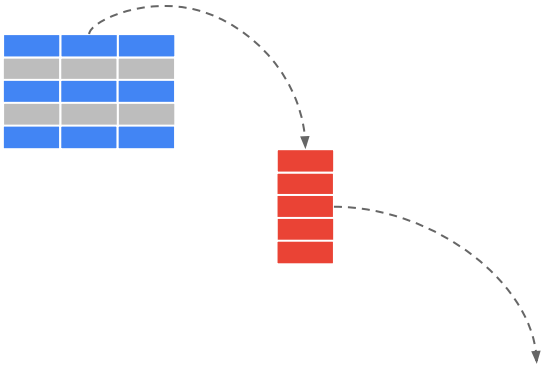

Standard indexes are tables

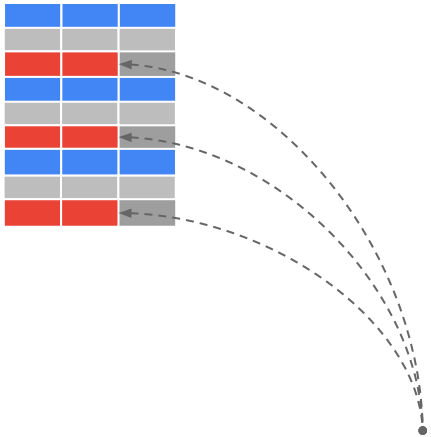

Interleaved indexes

Standard Relational Logical Data Layout

SingerId	SingerName
12ae6da78...	Beatles
b562cc13a...	U2
b7e47ab8d...	Pink Floyd

SingerId	AlbumId	AlbumName
12ae6da78...	011dc128c...	Help!
12ae6da78...	02be0ee5f...	Abbey Road
b7e47ab8d...	3647f2240...	The Wall

```
CREATE TABLE Singers (
  SingerId STRING(32) NOT NULL,
  SingerName STRING(MAX),
) PRIMARY KEY(SingerId);
```

```
CREATE TABLE Albums (
  SingerId STRING(32) NOT NULL,
  AlbumId STRING(32) NOT NULL,
  AlbumName STRING(MAX),
) PRIMARY KEY(SingerId, AlbumId);
```

Interleaving Logical Data Layout

12ae6da78...	Beatles	
12ae6da78...	011dc128c...	Help!
12ae6da78...	02be0ee5f...	Abbey Road
b562cc13a...	U2	
b7e47ab8d...	Pink Floyd	
b7e47ab8d...	3647f2240...	The Wall

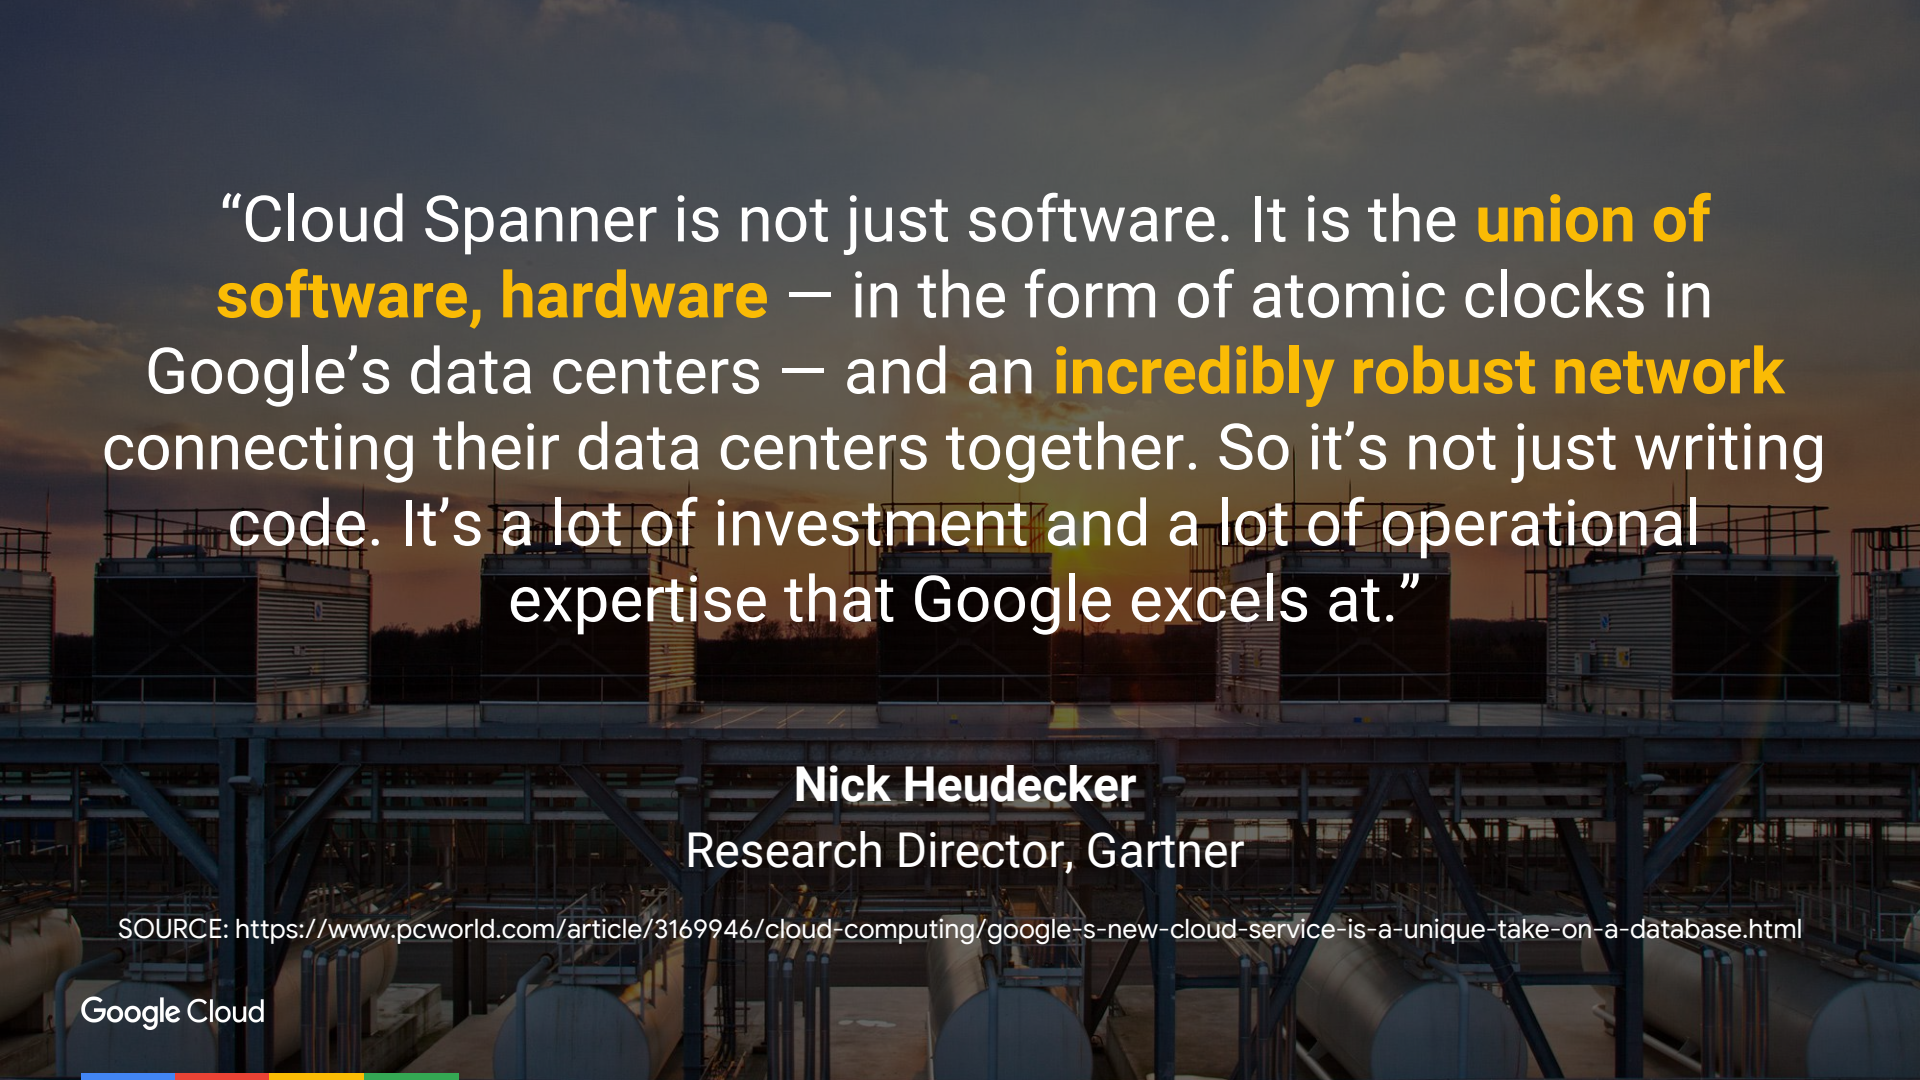
“Cloud Spanner is not just software. It is the union of software, hardware — in the form of atomic clocks in Google’s data centers — and an incredibly robust network connecting their data centers together. So it’s not just writing code. It’s a lot of investment and a lot of operational expertise that Google excels at.”
Nick Heudecker
Research Director, Gartner
SOURCE: https://www.pcworld.com/article/3169946/cloud-computing/google-s-new-cloud-service-is-a-unique-take-on-a-database.html
Google Cloud

“

We want to move from our current on-prem per-customer deployment model to the cloud to **improve performance and reliability**, which is extremely important to us and our customers. With Cloud Spanner, we can process **ten times more transactions per second** (using a current benchmark of 55k transactions per second), allowing us to better serve customers, with a dramatically **reduced total cost of ownership**.

– Danielle Royston, CEO, Optiva

Optiva™

It was **essential for us to have order sequence in our app logic**, and with Cloud Spanner, it's built in. When we started looking at GCP, we quickly identified Cloud Spanner as the solution, as it provided relational semantics and incredible scalability within a managed service. **We hadn't found a Cloud Spanner-like product in other clouds**. We ran a successful POC and plan to move several massive services to Cloud Spanner. We look forward to Multi-Region configurations, as they give us the **ability to expand globally and reduce latencies for customers on the other side of the world**.

– Manoj Goyal, Chief Product Officer, Marketo

“

At our size, problems such as **scalability and reliability don't have a simple answer**, Cloud Spanner is a **transformational technology** choice for us. It will give us a regionally distributed database storage layer for our customers data that can **scale as we continue to grow**. Our whole technology team is excited to bring this into production in the coming months.

– Ben McCormack, Vice President of Operations, Evernote

Google Cloud Database Portfolio

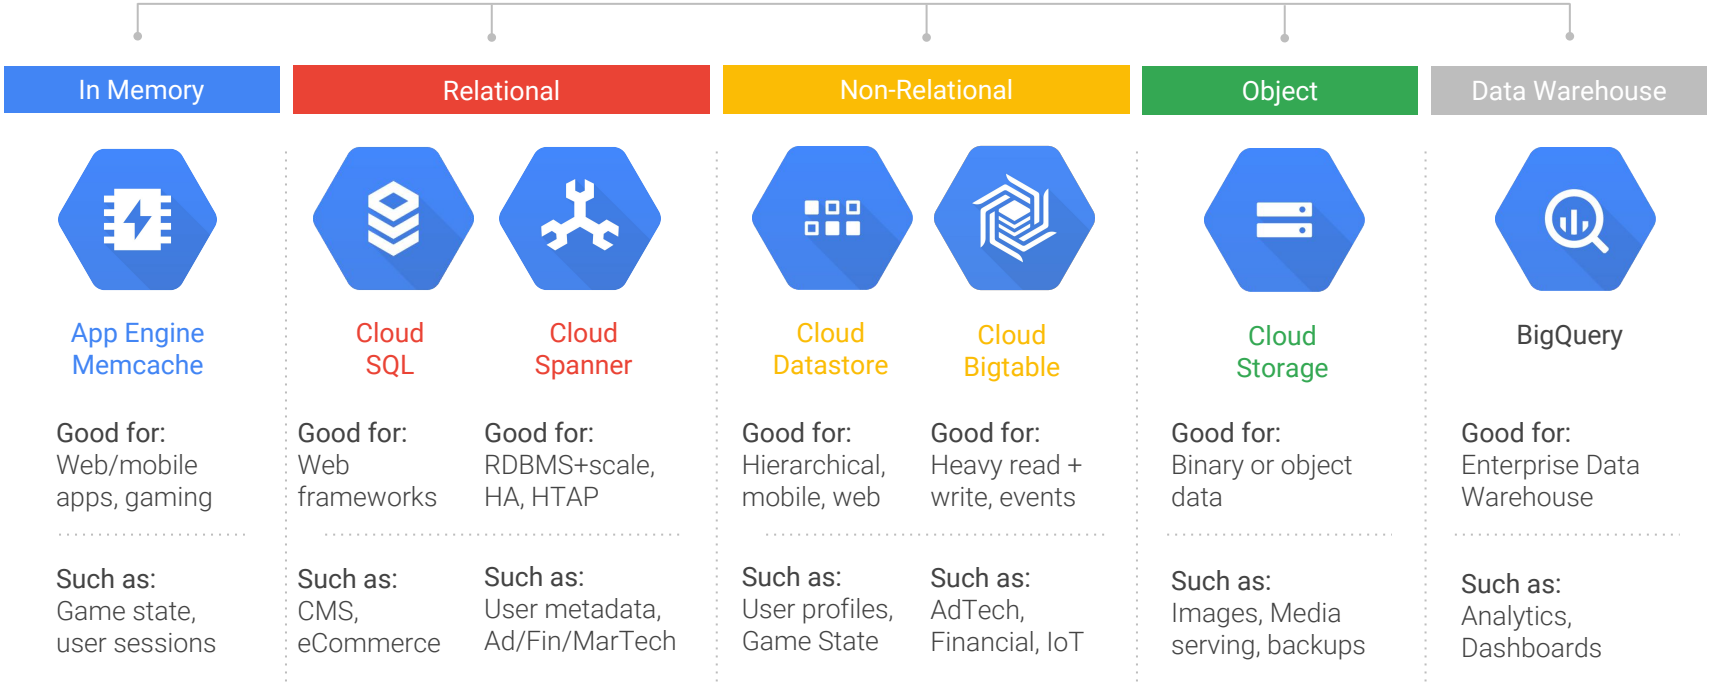

Google Cloud

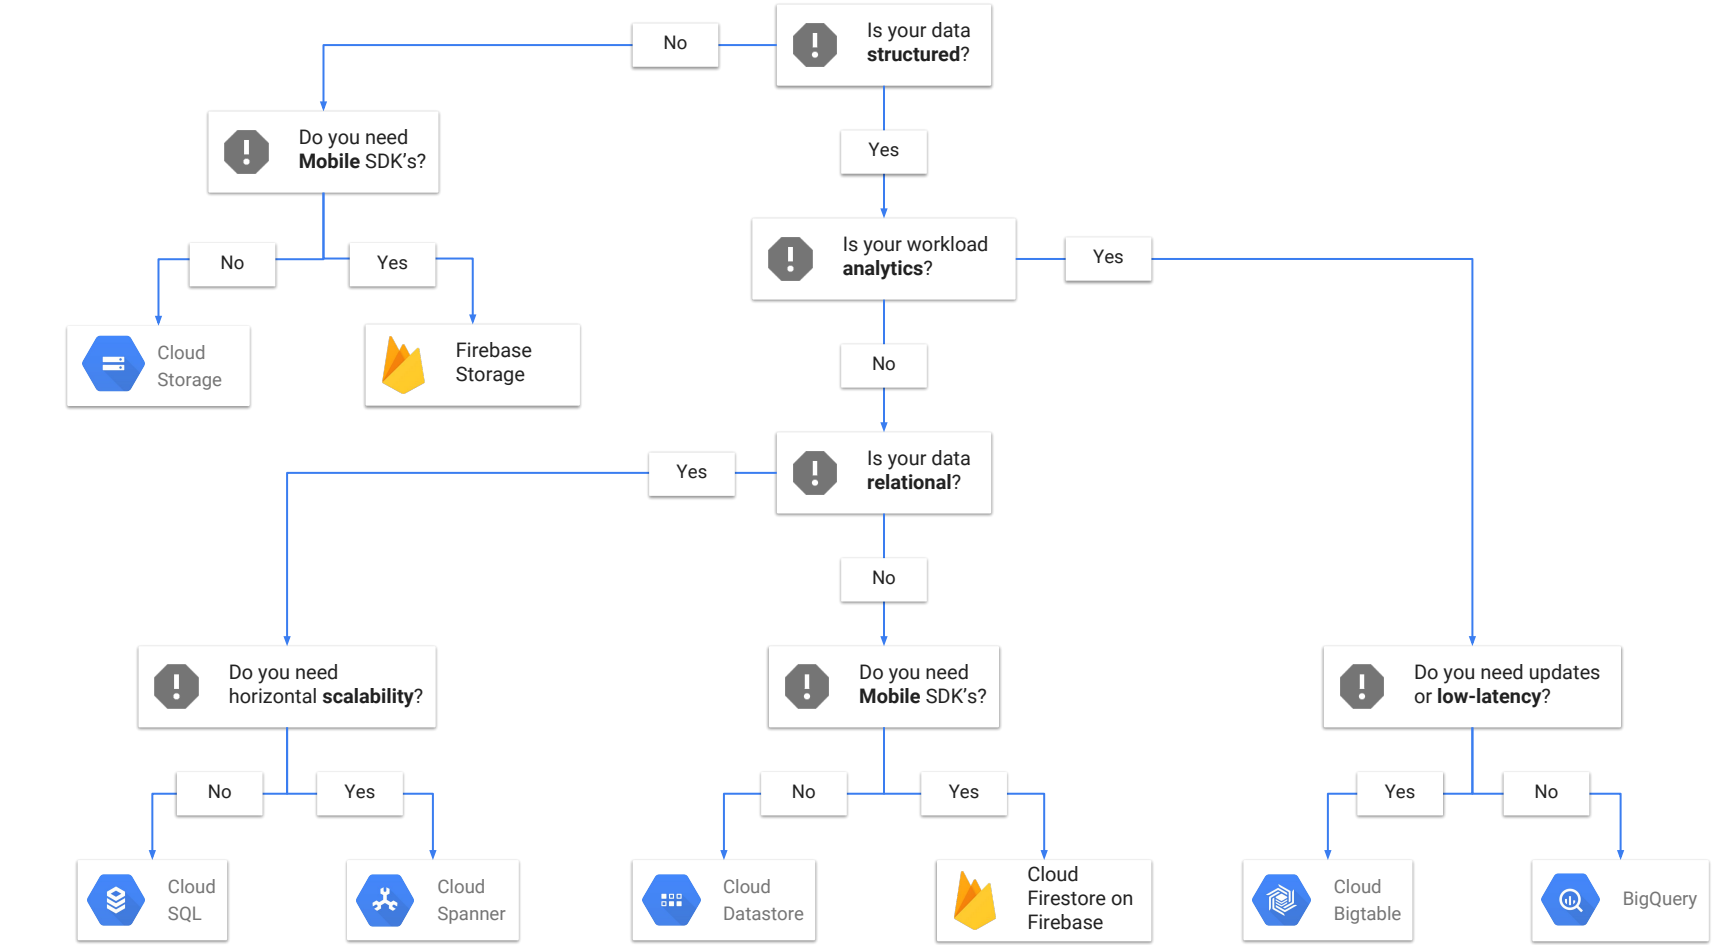

Google Cloud

Thank you

Google Cloud