Distributed Databases

an exploration of approaches and best practices

Julia Ferraioli
Developer Advocate

Brian Dorsey
Developer Programs Engineer

Your Hosts

Julia Ferraioli
Developer Advocate
@juliaferraioli

Brian Dorsey
Developer Programs Engineer
@briandorsey

Why Distributed Databases?

Image courtesy of Allie Brosh of *Hyperbole and a Half*

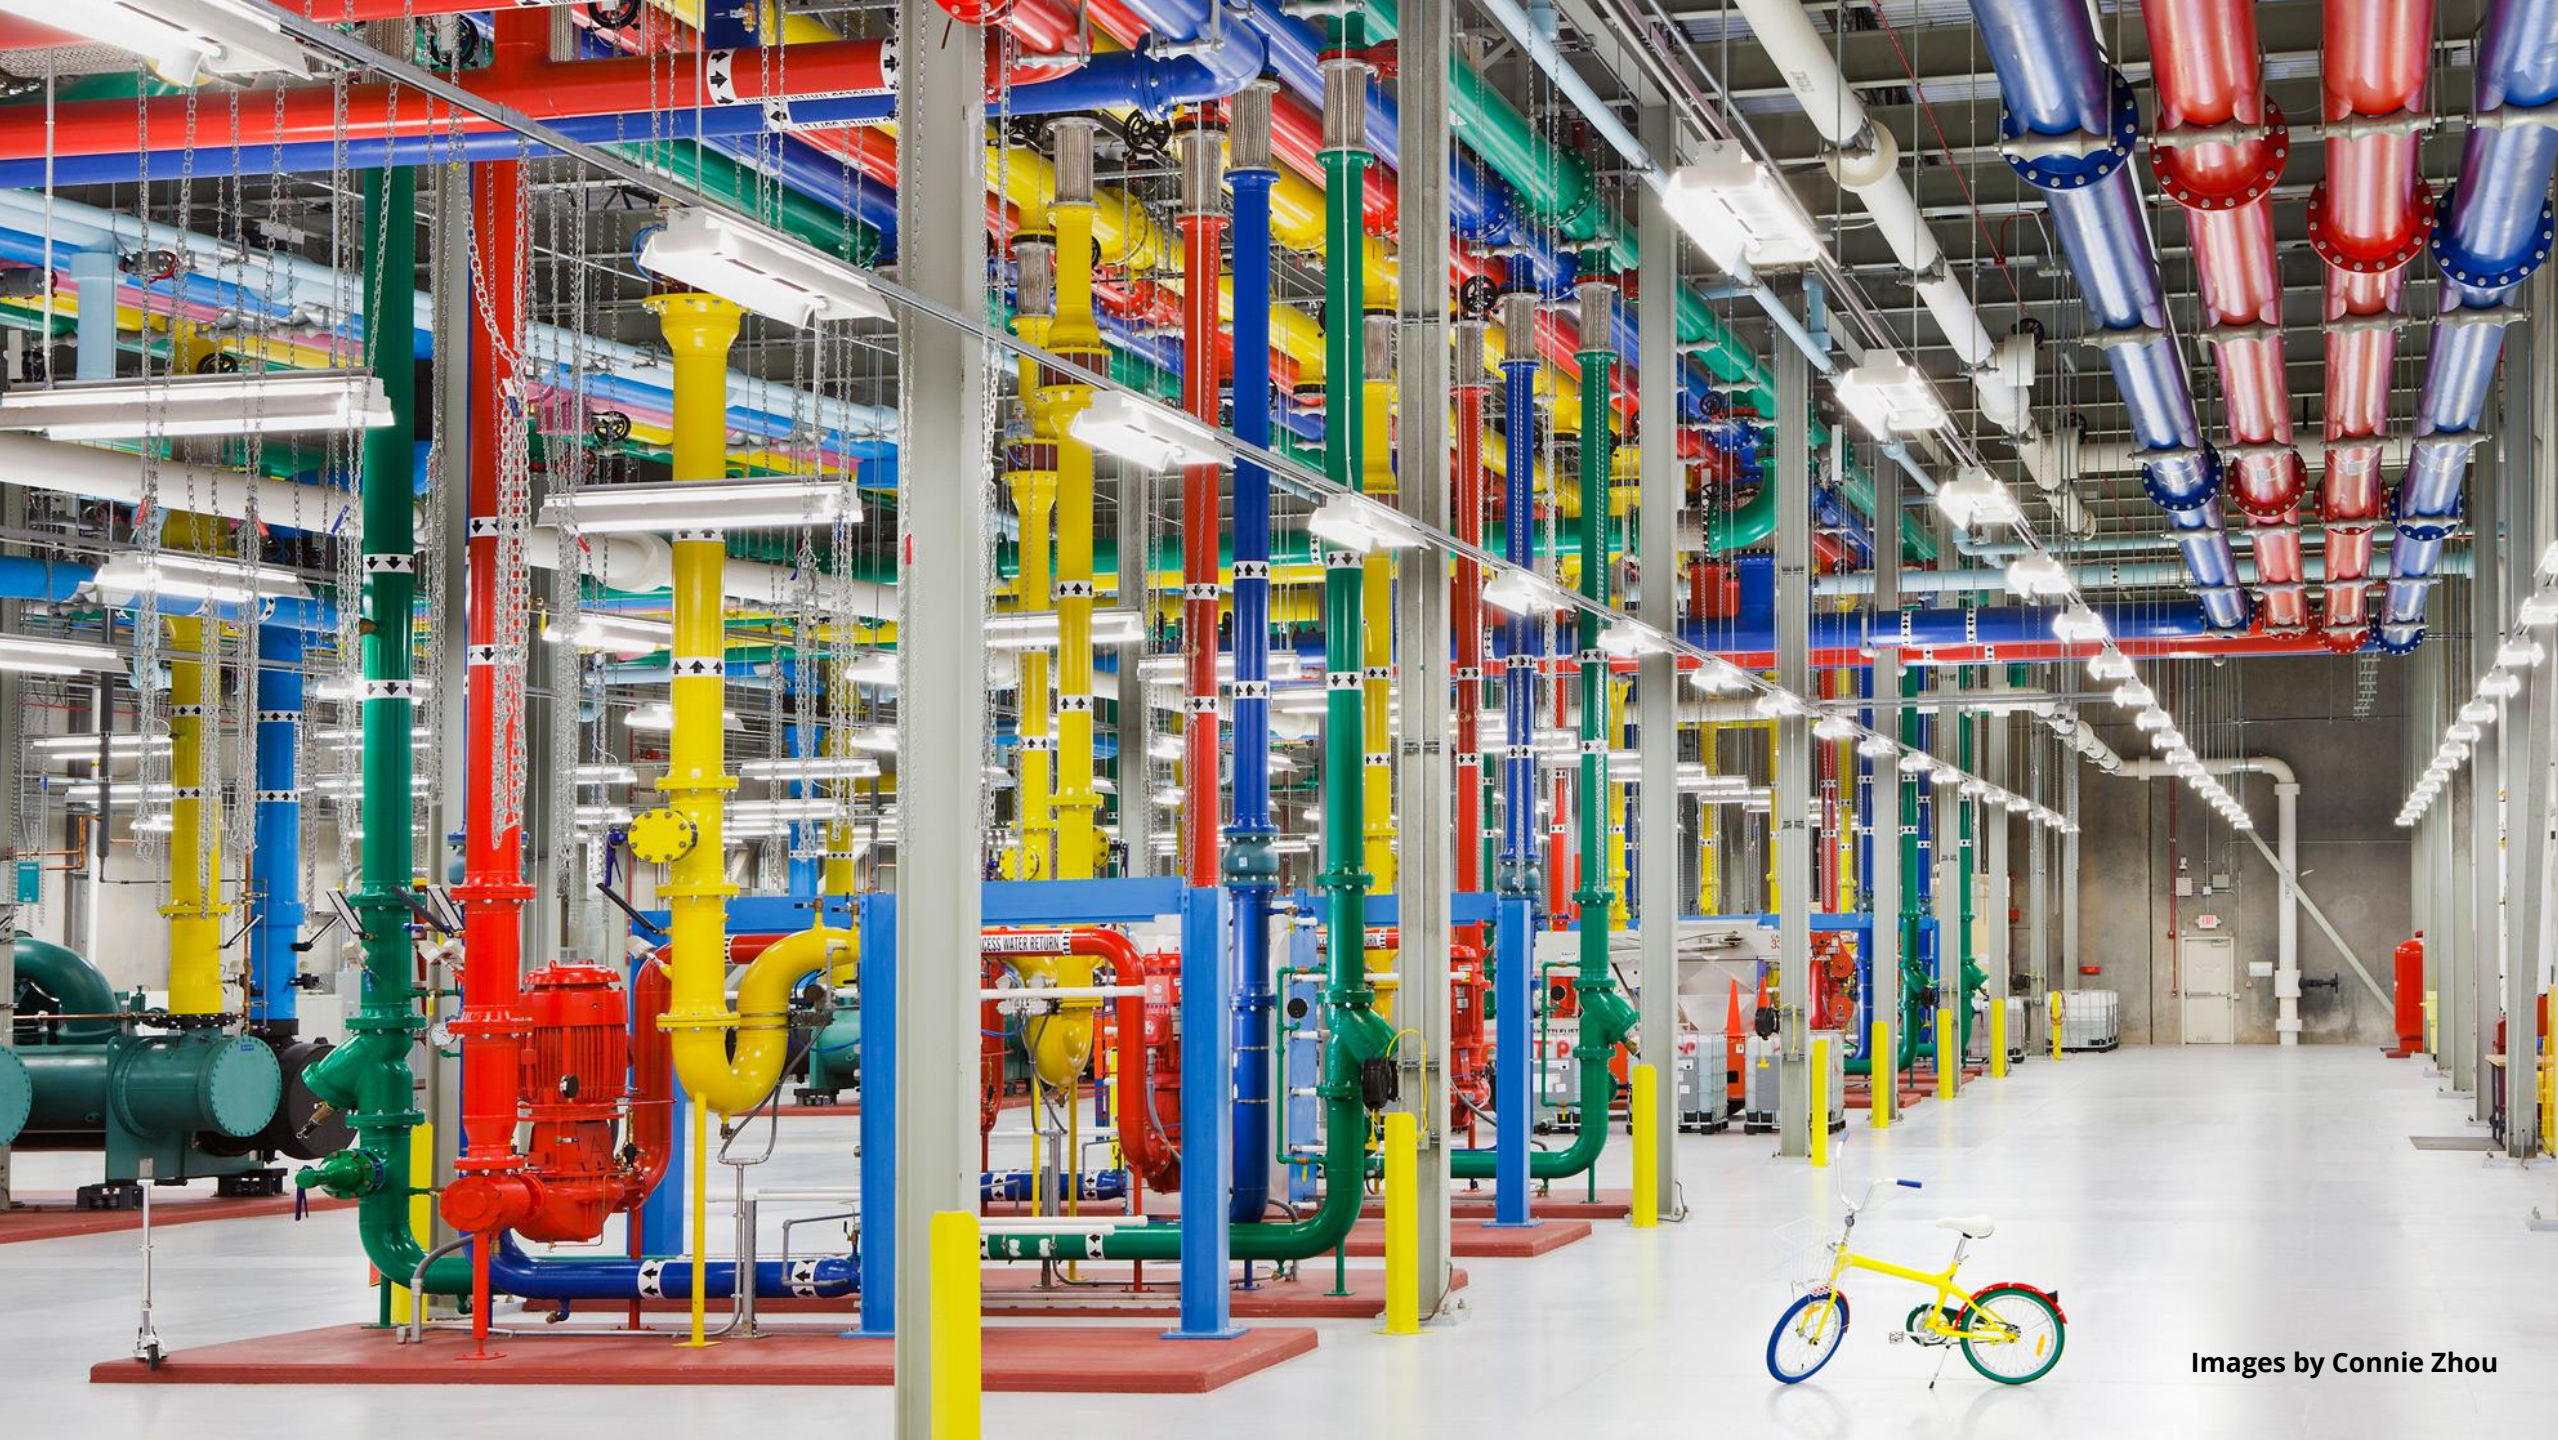

Images by Connie Zhou

Images by Connie Zhou

Your Panelists

Tyler Hannan

@tylerhannan

Mike Miller

@mlmilleratmit

Google Cloud Datastore

Chris Ramsdale

@cramsdale

Will Shulman

@willshulman

Riak: An Open Source, Distributed Key/Value Database

Basho Technologies

Tyler Hannan

About Basho Technologies

Who we are, what we do

- Founded January 2008
- ~ 140 employees worldwide
- Headquarters in Cambridge, MA with offices in Reston, San Francisco, London, and Tokyo
- A distributed company building distributed systems
- Basho makes Riak & Riak CS

What Is Riak?

The Benefits of Riak

Riak is an Ops-friendly database that is:

- Fault-tolerant
- Highly-available
- Scalable
- Self healing

How Does That Work?

The Properties of a Distributed Database

Riak is a key/value store that is:

- Open source
- Distributed
- Masterless
- Eventually consistent

Riak is a Key/Value Store

Simple Operations, Opaque Values, Layered with Extras

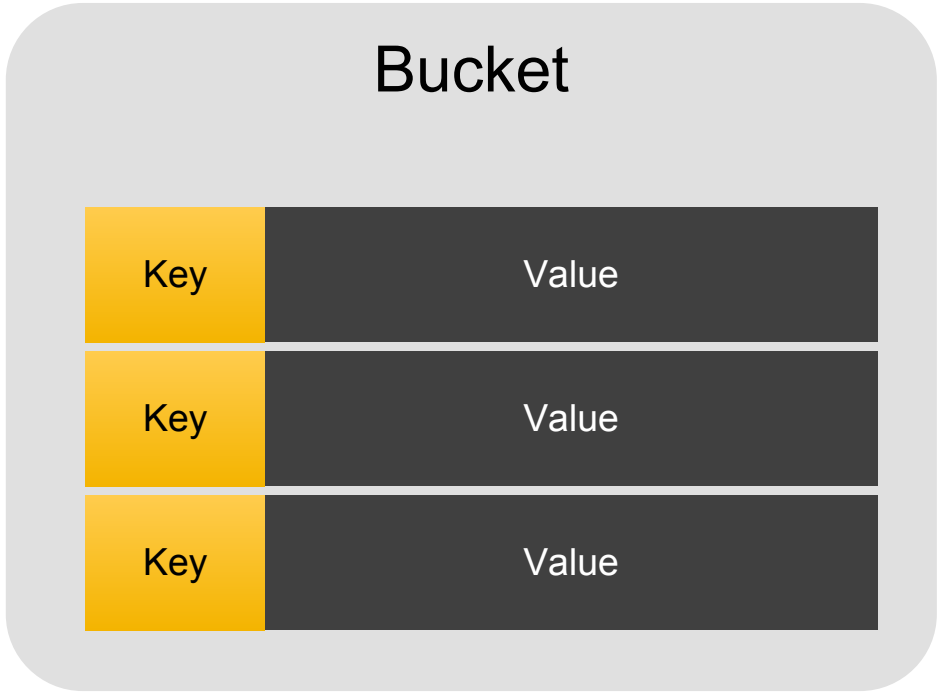

- GET / PUT / DELETE
- Value is mostly opaque
- HTTP & Protobufs API + Client Libraries
- Extras:
  - MapReduce
  - Full-text search
  - Secondary indices
  - Pre/post-commit hooks

Riak is Masterless

Deployed as a Cluster of Nodes

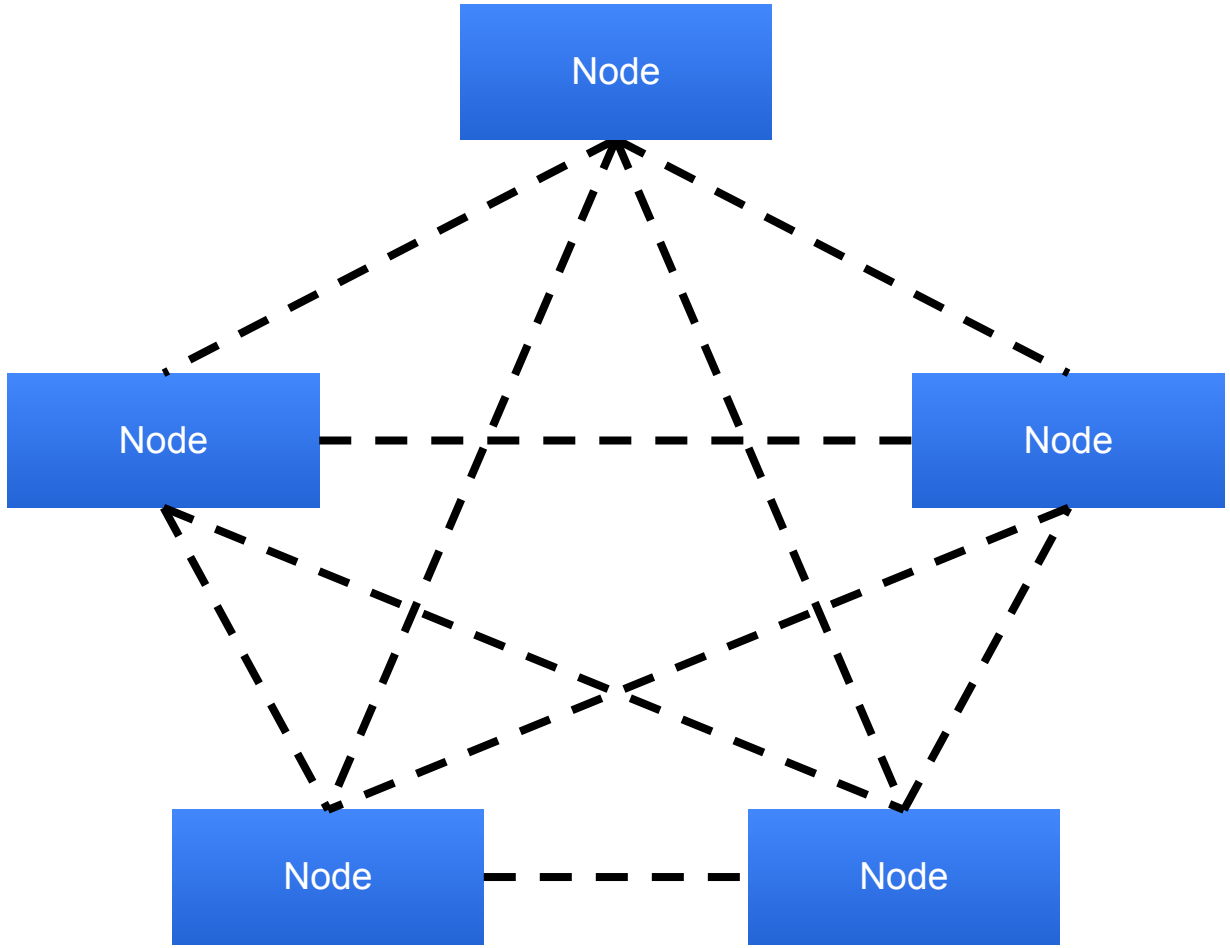

- Based on principles of Dynamo specification
- Any node can serve any request
- Data and load are spread evenly
- Gossip protocol (mesh network)
- Hinted handoff
- Achieve near-linear scale by adding hardware

"Big Data", "Web Scale", "Other Terms"

When Your Data Is Critical, Scalability Is Critical

Shell

```
$ gcutil --project=RiakCluster addinstance \
  riak5 --machine_type=n1-standard-4
$ gcutil --project=RiakCluster ssh riak5

# Install Riak programmatically or via startup
script

$ riak-admin cluster join riak1@192.168.2.2
$ riak-admin cluster plan
$ riak-admin cluster commit
```

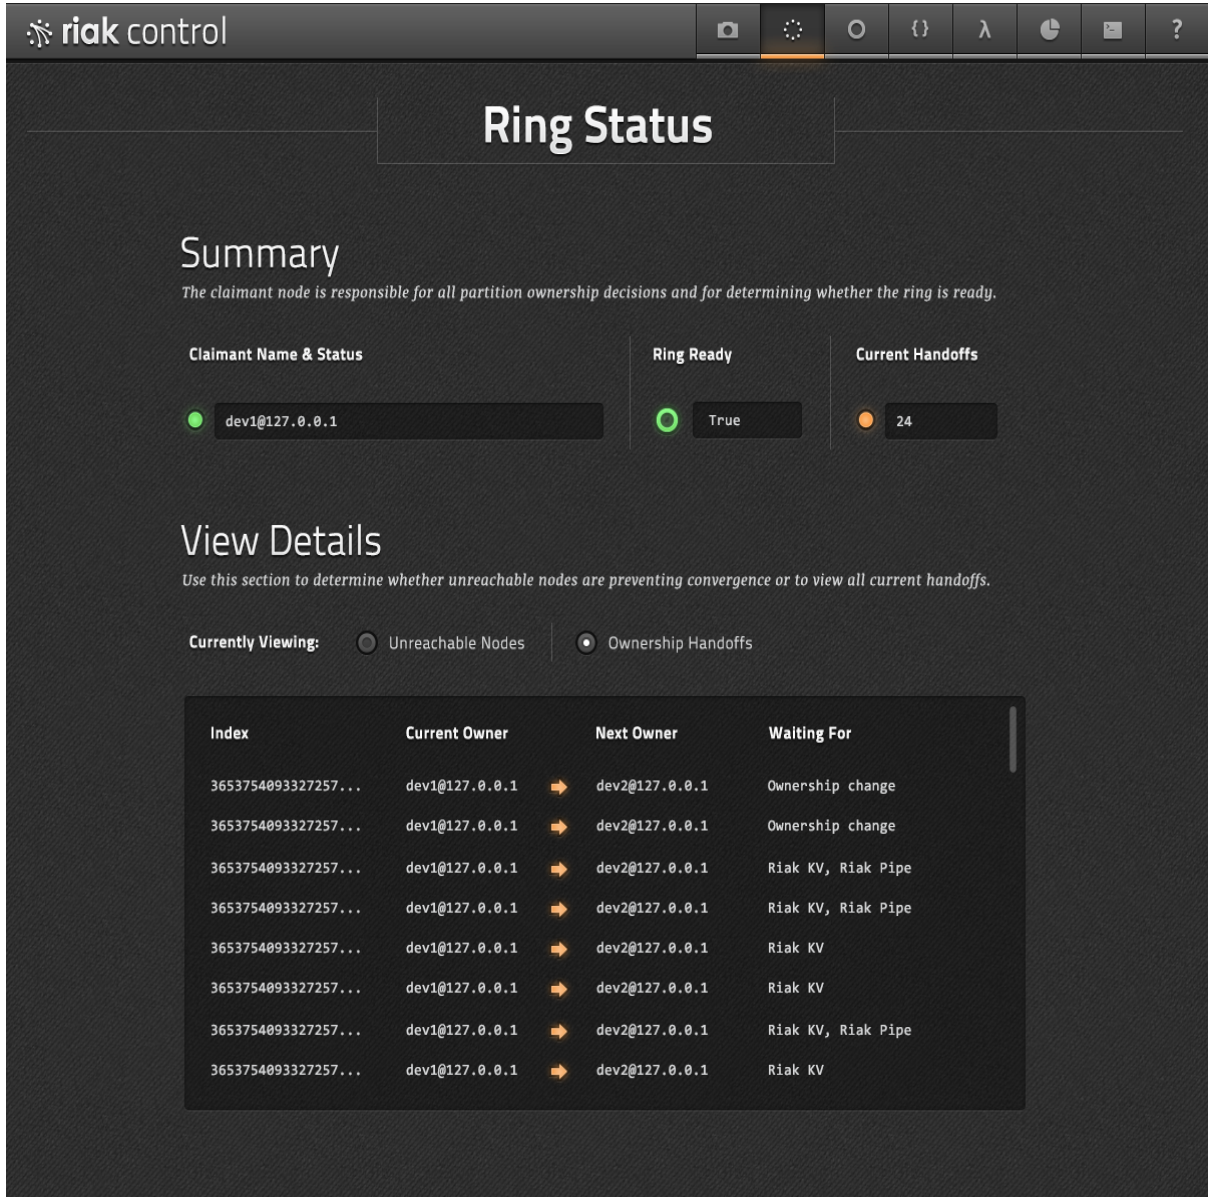

When Would I Use Riak on Google Compute?

The situations & the circumstances

Operationally-friendly database

- combined with -

Operationally-scalable compute platform

for gaming, social, mobile, retail, advertising, etc.

Getting to Know Cloudant

Your Friendly Neighborhood NoSQL Database Service

Mike Miller
Co-Founder, Chief Scientist


CLOUDANT IS THE...
DISTRIBUTED DATABASE as a SERVICE
Ships with a mobile strategy

{Install: ‘Cloudant’}

Step 1	Step 2	Step 3
You do this:	We give you:	Done!
Sign Up	https://<username>.cloudant.com	

{Write: ‘Local’, Sync: ‘Later’}

Embedded,
Edge, Satellites

Desktop,
Browser

Cloud

Google Cloud Datastore

Scale with your users, not your servers

Chris Ramsdale
Product Manager, Google Cloud Platform

Google Cloud Platform Storage

Family of Managed Storage Services

Cloud Storage
blob data

Cloud SQL
relational data

Cloud Datastore
non-relational data

Google Cloud Platform Storage

Family of Managed Storage Services

Cloud Storage
blob data

Cloud SQL
relational data

Cloud Datastore
non-relational data

Announcing the Google Cloud Datastore

Fully Managed Schemaless Storage

App Engine High Replication Datastore (HRD)

Google Cloud Datastore

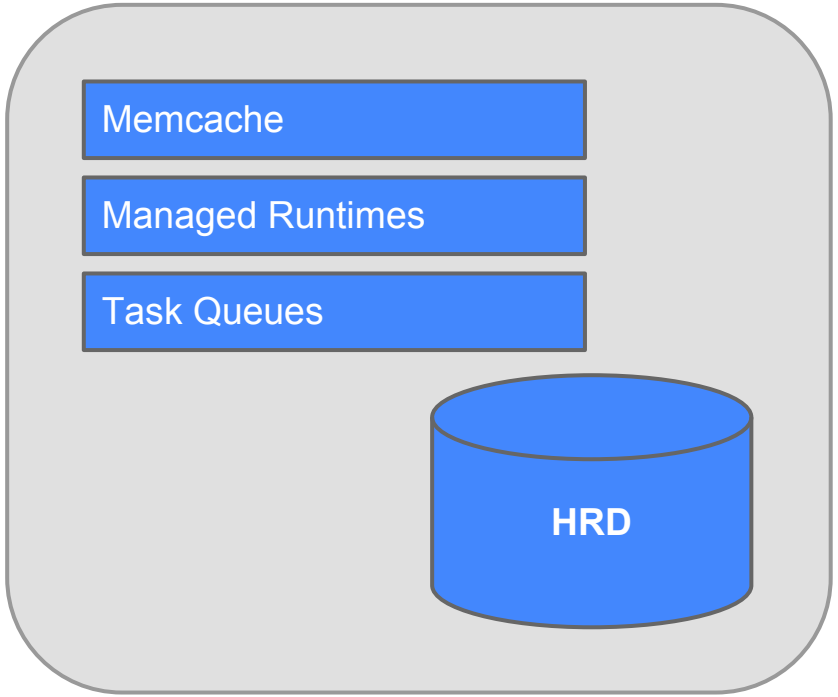

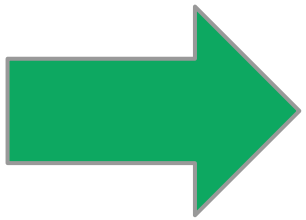

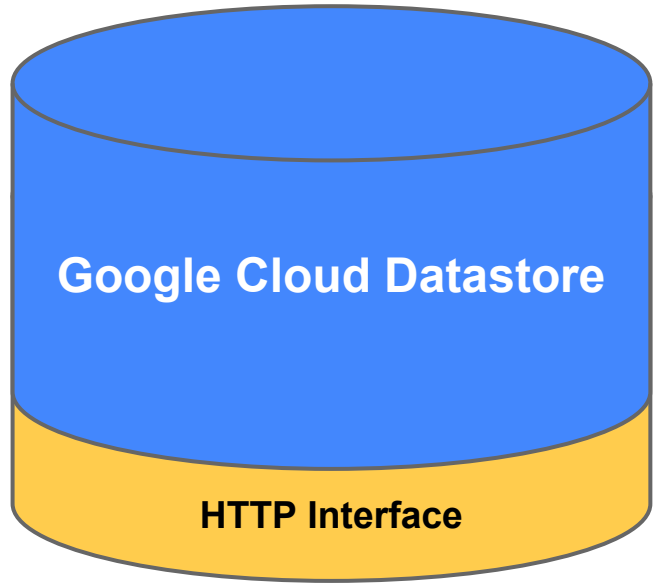

Google Cloud Datastore

Bringing Google Infrastructure to Developers

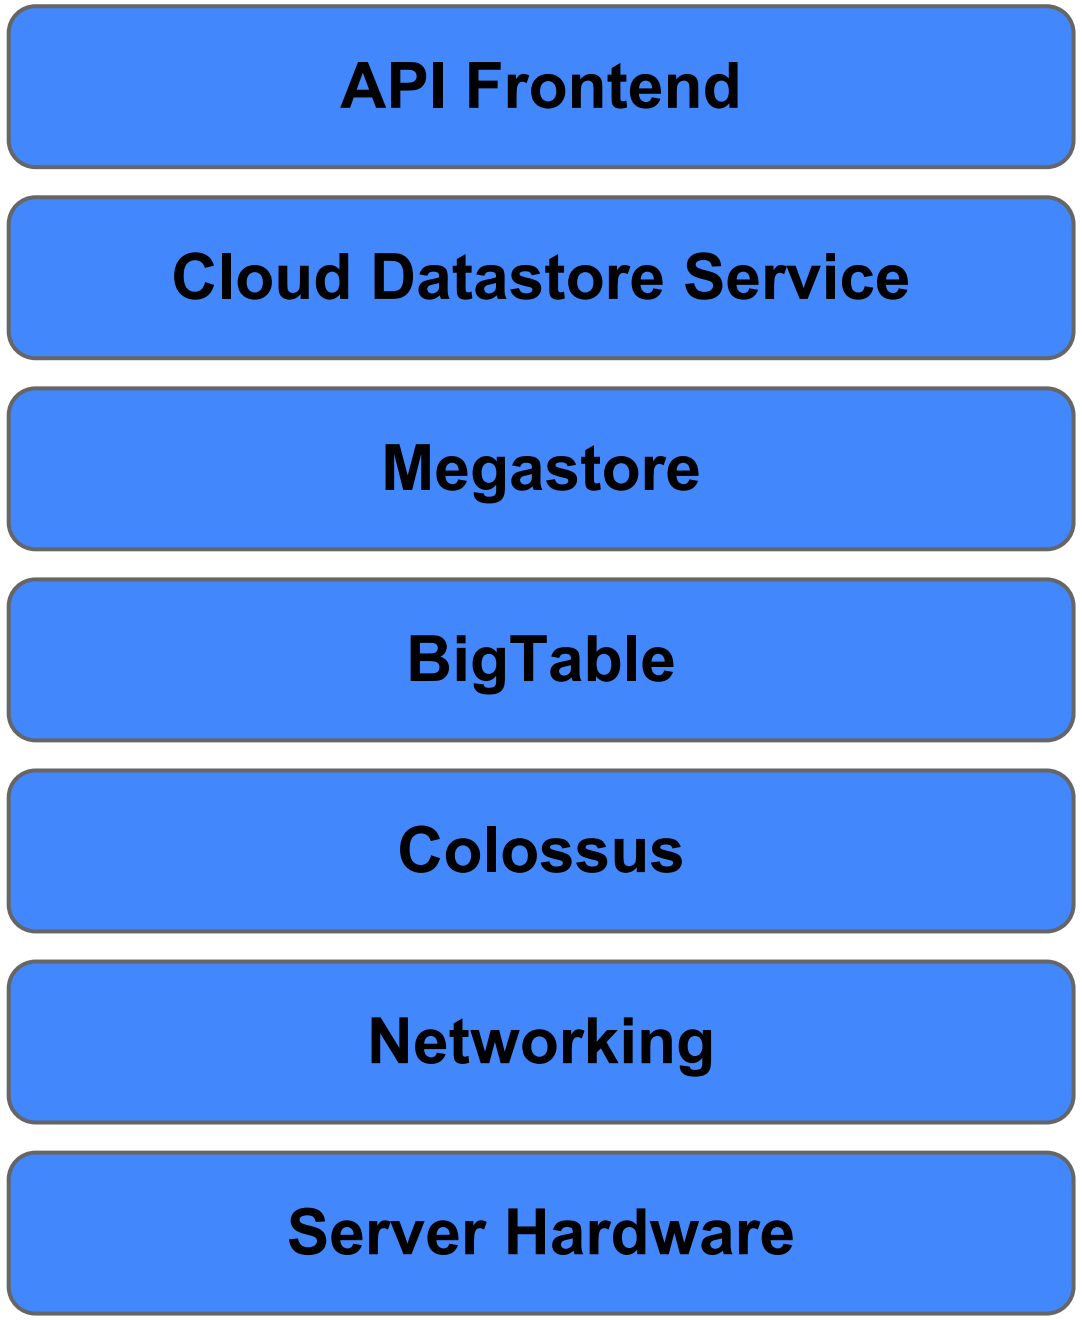

Google Cloud Datastore

Bringing Google Infrastructure to Developers

Google maps

- API Frontend
- Cloud Datastore Service
- Megastore
- BigTable
- Colossus
- Networking
- Server Hardware

Google Cloud Datastore

High Availability

- Auto-replication across multiple datacenters
- Paxos consensus
- Strong and Eventual consistency

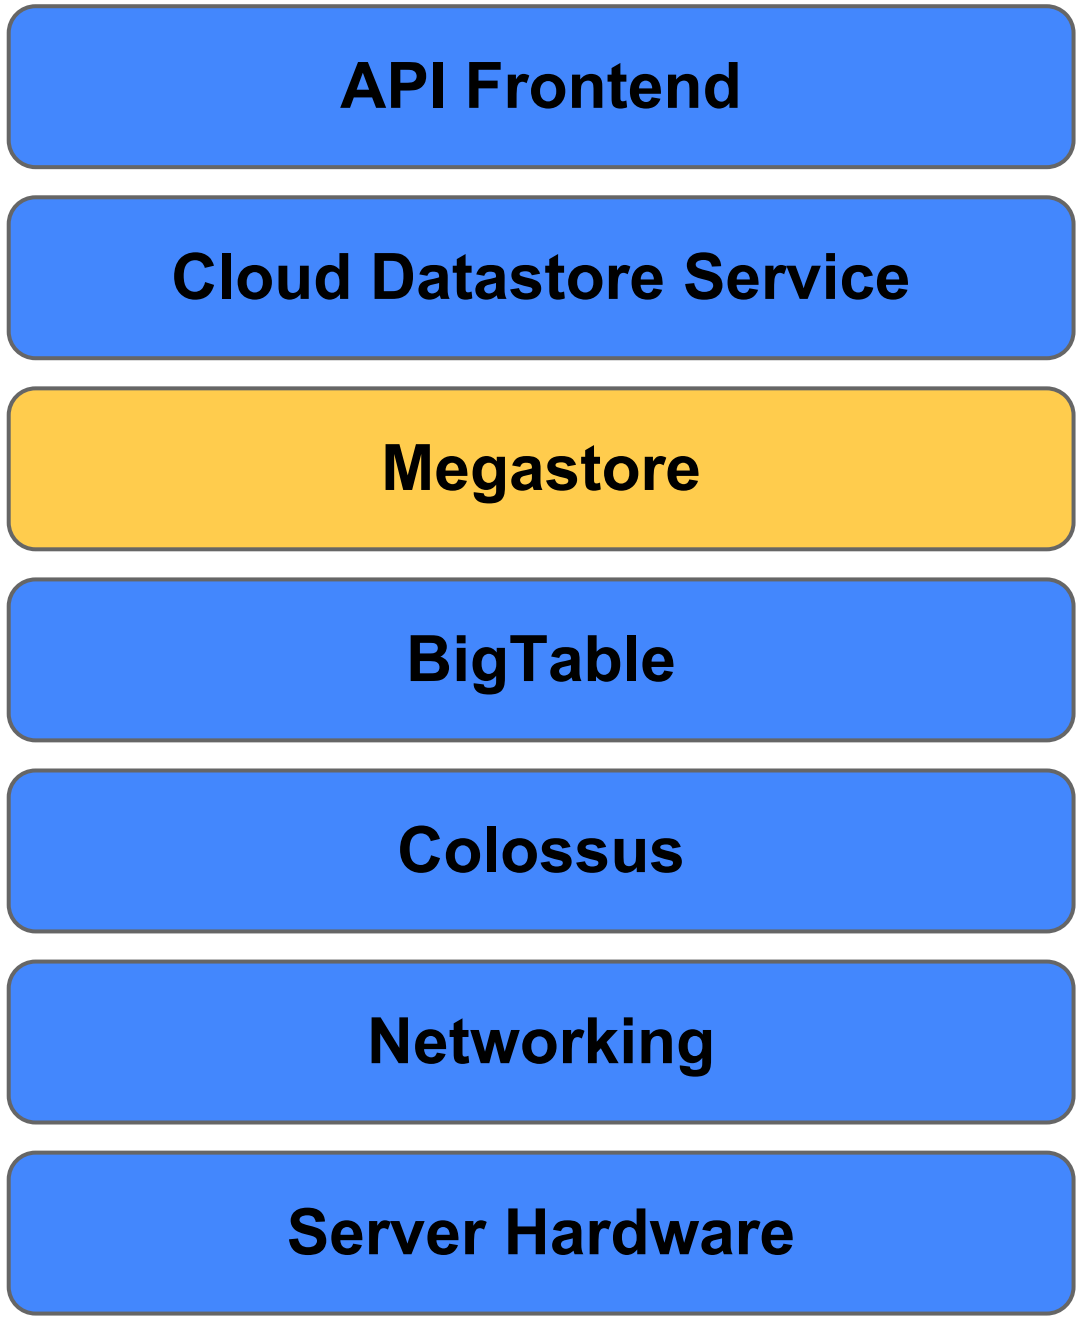

Google Cloud Datastore

High Scalability

- Horizontal auto-scaling
- Huge capacity
- High durability

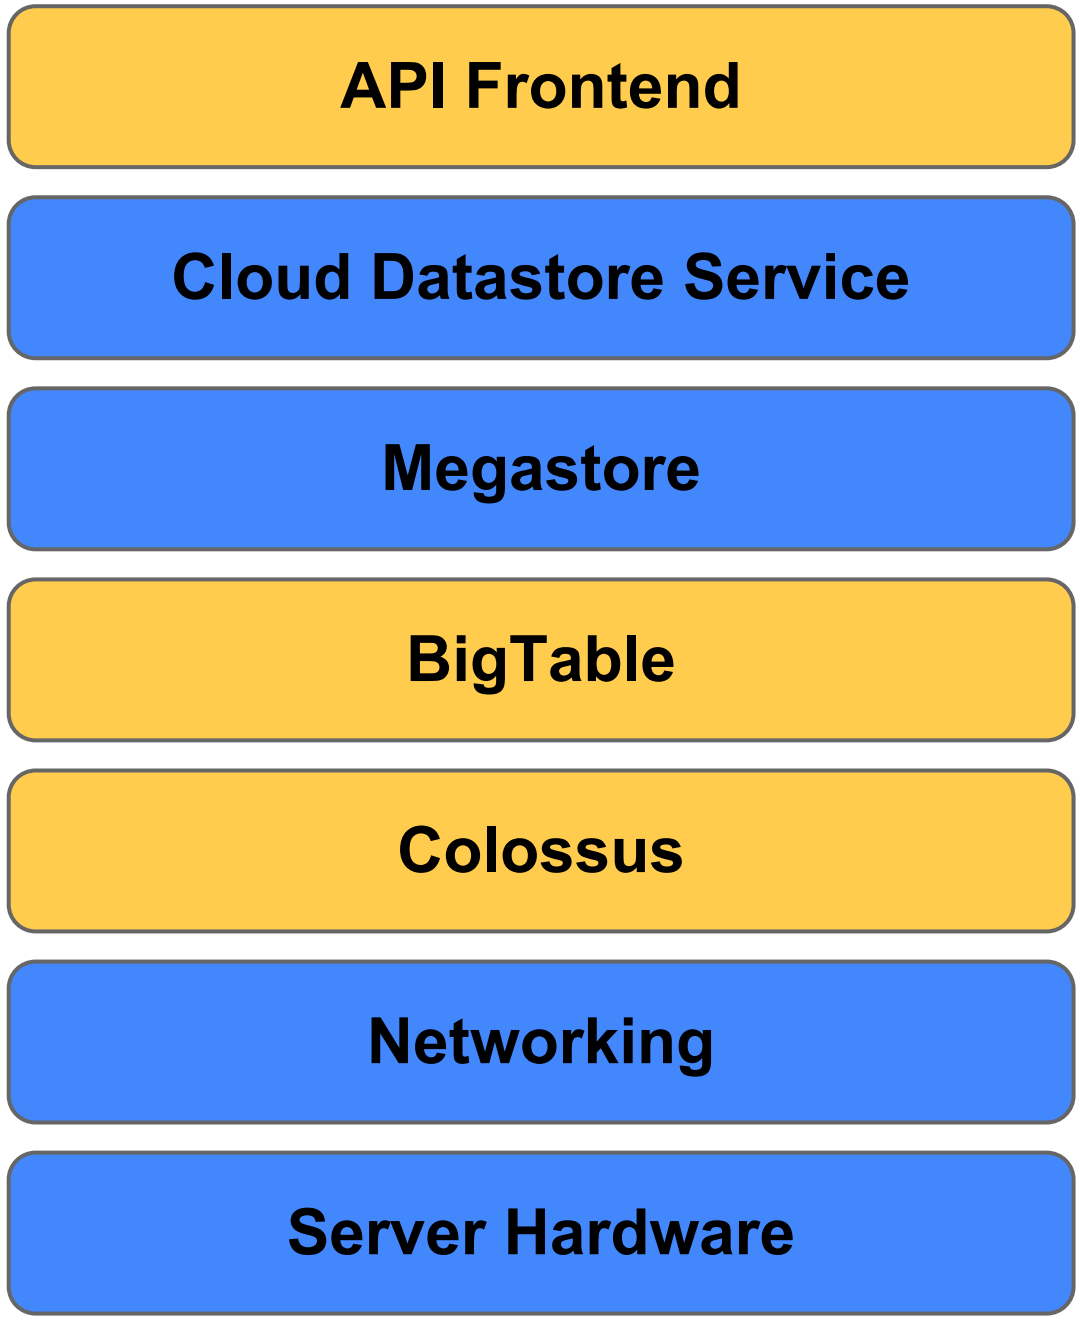

Google Cloud Datastore

Access from Anywhere

Managed Frontend

App Engine SDK

Unmanaged Backend

Cloud Datastore API

API Frontend

Cloud Datastore Service

Megastore

BigTable

Colossus

Networking

Server Hardware

Google Cloud Datastore

Fully Managed

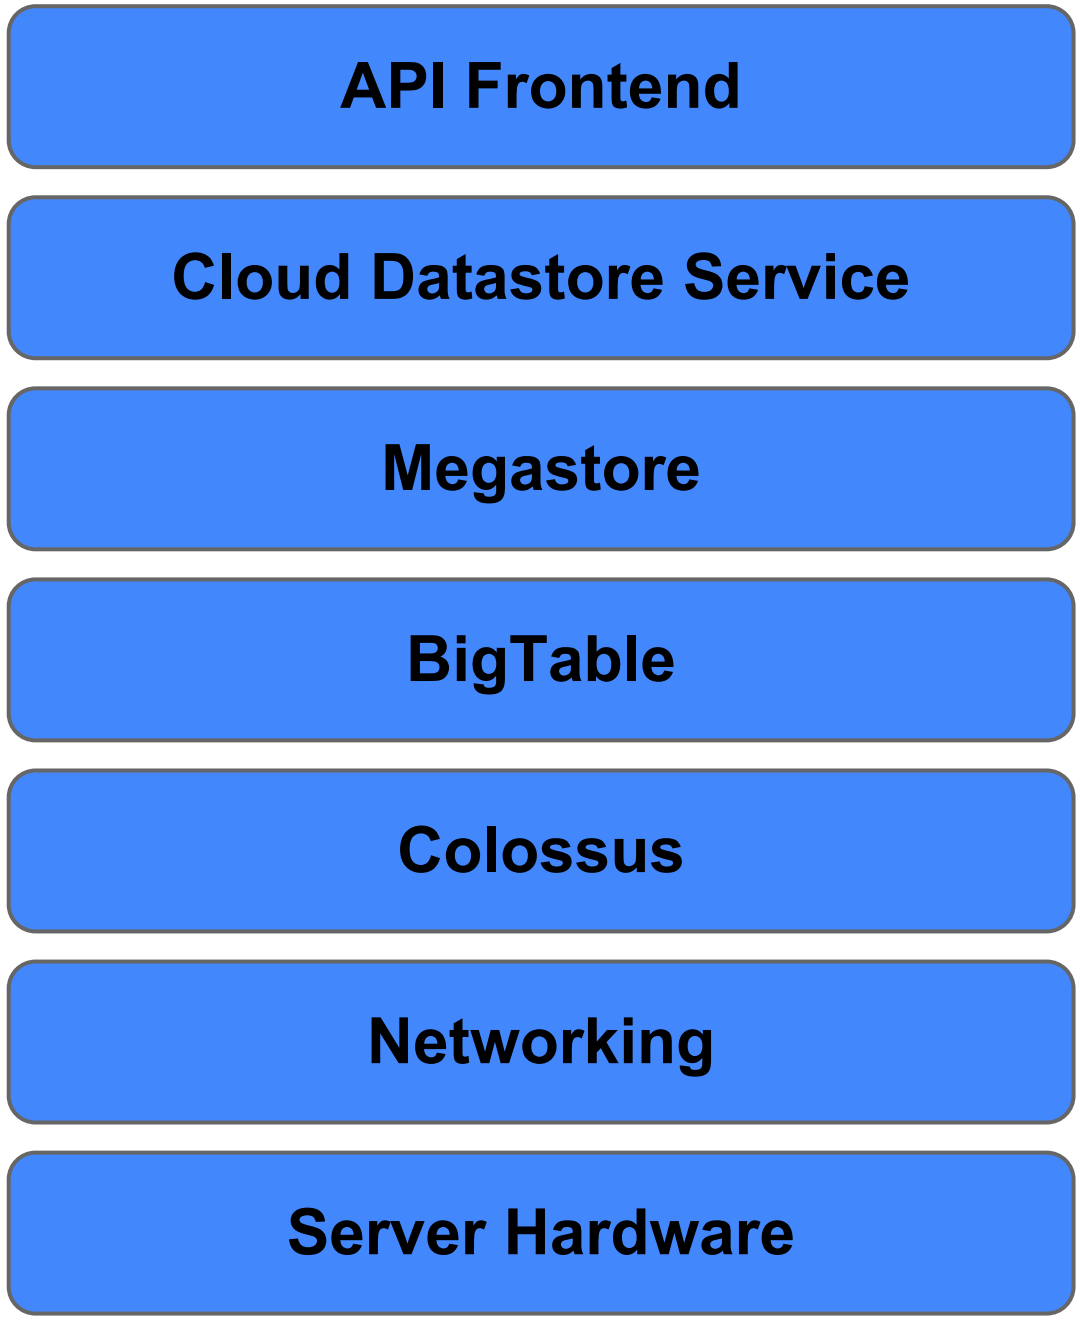

Google Cloud Datastore

Fully Managed

API Frontend

Cloud Datastore Service

Megastore

BigTable

Colossus

Networking

Server Hardware

Google Cloud Datastore

Fully Managed

API Frontend

Cloud Datastore Service

Megastore

BigTable

Colossus

Networking

Server Hardware

An intro to MongoDB and MongoLab

in < 5 minutes

Will Shulman
CEO MongoLab

What is MongoDB?

MongoDB is an open source, high-performance, distributed, and document-oriented database.

MongoDB is document-oriented

a.k.a. object-oriented

```
{
   _id: 1234,
   author: { name: "Bob Davis", email : bob@davis.com },
   post: "In these troubled times I like to …",
   date: { $date: "2010-07-12 13:23UTC" },
   location: [ -121.2322, 42.1223222 ],
   rating: 2.2,
   comments: [
    { user: "jgs32@gmail.com", upVotes: 22, downVotes: 14, text: "Great point" },
    { user: "holly.lu@gmail.com", upVotes: 421, downVotes: 22, text: "You're a moron" }
   ],
   tags: [ "Politics", "Virginia" ]
}
```

MongoDB is great as an operational data store

. . . with a rich query language

```
db.posts.find({ author.name: "mike" })

db.posts.find({ rating: { $gt: 2 }})

db.posts.find({ tags: "Software" })

db.posts.find().sort({date: -1}).limit(10)

db.places.find({loc: {$within : {$center : [[40,40],10]}}})

db.places.aggregate({$group: { _id: "$state", pop: { $sum: "$pop" }}})
```

MongoDB is great as an operational data store

. . . with support for indexes on any field

```
db.posts.ensureIndex({ author.name : 1 })

db.posts.find({ author.name: "mike" })
```

MongoDB is a distributed database

. . . with high availability via Replica Set clusters

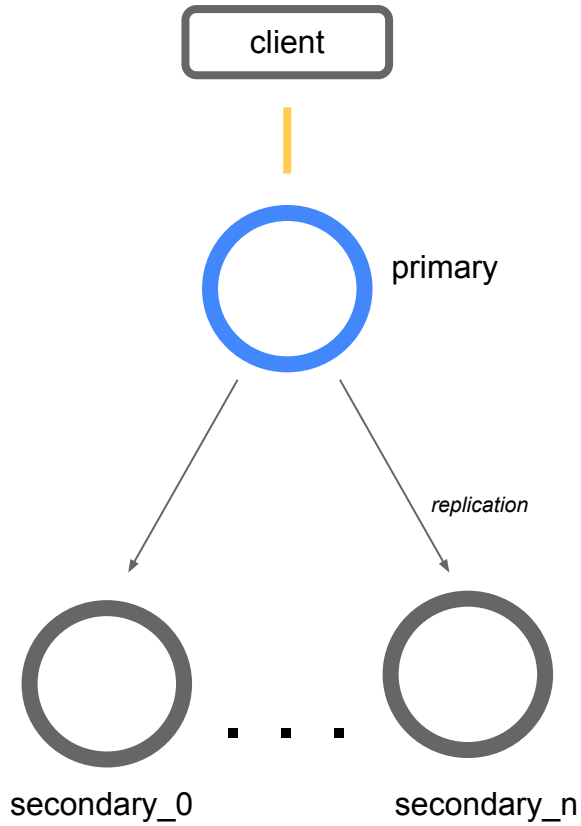

- Single master (read / write)
- Multiple secondaries (read)
- Automatic failover
- Strong consistency or eventual consistency
- Configurable write-concerns
  - w = 1
  - w = 3
  - w = "majority"

MongoDB is a distributed database

. . . with horizontal scalability via Sharded Clusters

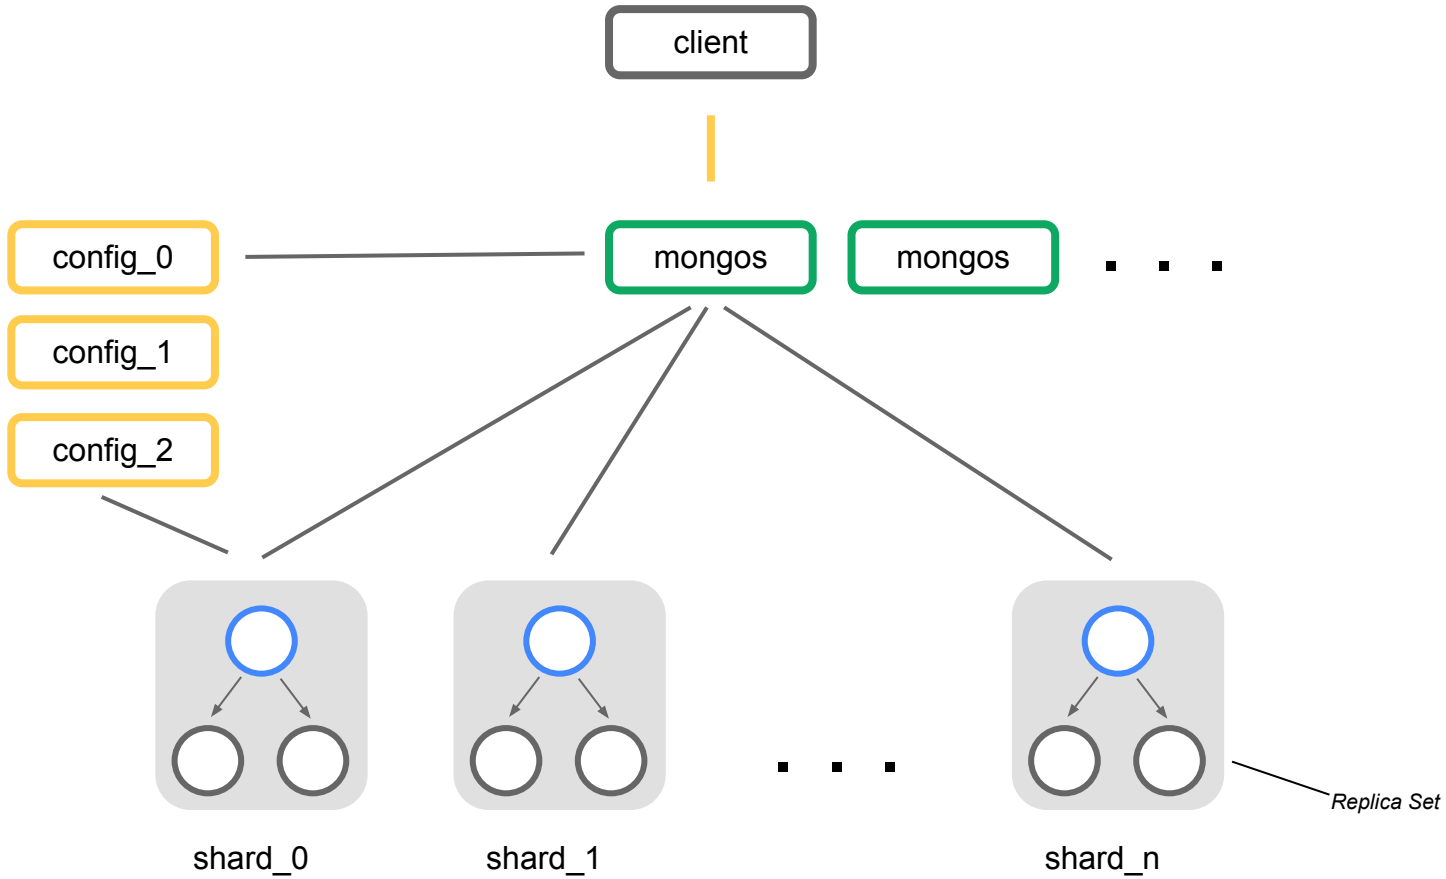

What is MongoLab?

MongoLab is MongoDB-as-a-Service

MongoLab is MongoDB-as-a-Service

We automate the operational aspects of running MongoDB (so you don't have to)

Features/benefits

- provisioning and scaling
- replication and backups
- monitoring and alerting
- rich web UI and tools
- expert support

Product offering

- shared and dedicated VM plans
- SSD plans
- single-node and multi-zone Replica Set clusters
- support for Sharded Clusters in 2014

MongoLab is MongoDB-as-a-Service

We support all the major cloud providers

New as of today!

SELECT questions FROM audience

Tyler Hannan

@tylerhannan

Mike Miller

@mlmilleratmit

Google Cloud Datastore

Chris Ramsdale

@cramsdale

Will Shulman

@willshulman

<Thank You!>

jrf@google.com
google.com/+JuliaFerraioli
@juliaferraioli

briandorsey@google.com
google.com/+BrianDorsey
@briandorsey


Google
Developers