LSST catalogs inside Google

Fabrice Jammes
Scalable Data Systems Expert
IN2P3

Credits: **Ross Thomson,** Solution architect at Google

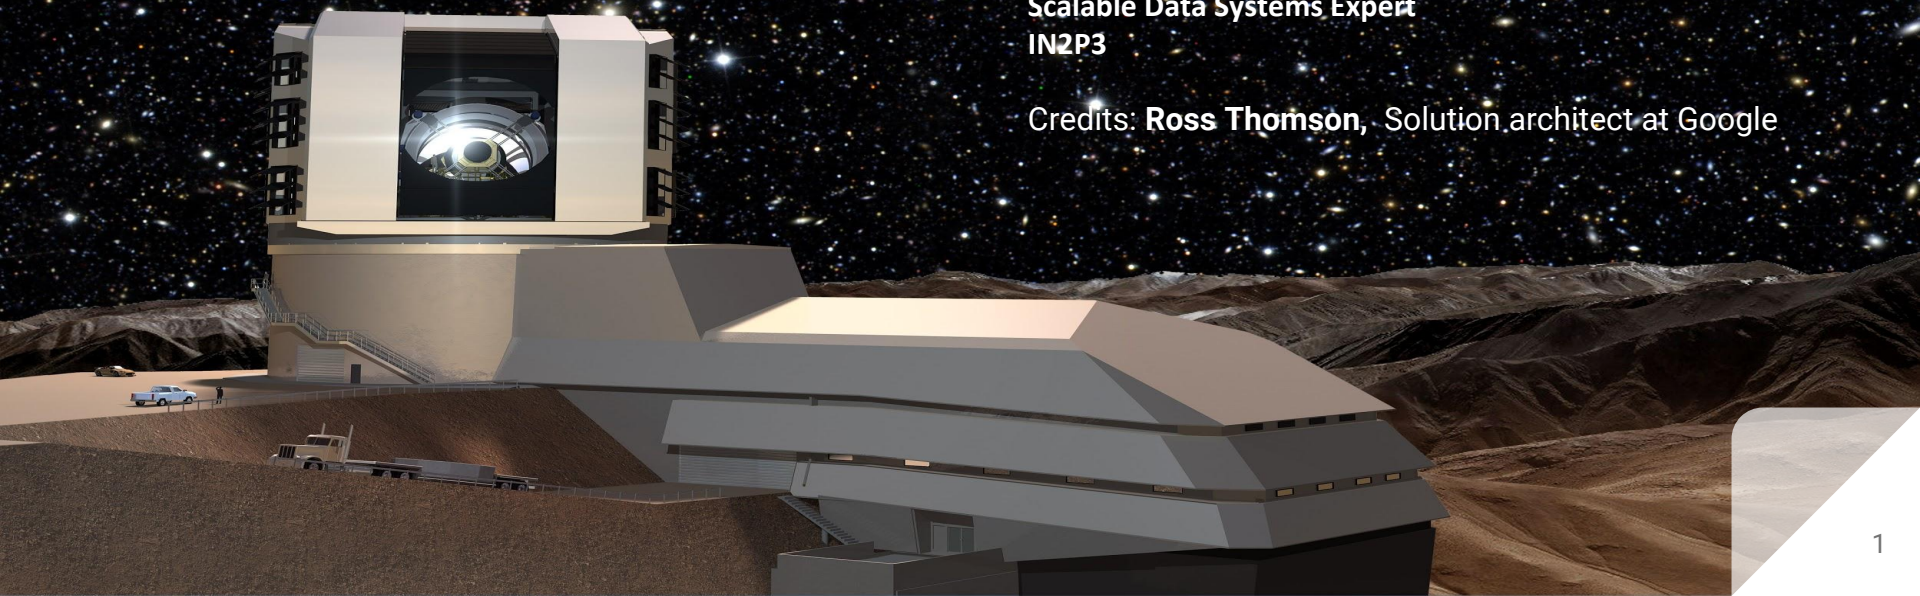

Agenda

1. Qserv
2. Kubernetes
3. Google Kubernetes Engine
4. Google Big Query

+80 PB celestial catalogs

Relational database, 100% open source

Space partionning with overlap

Highly distributed “Map Reduce” model

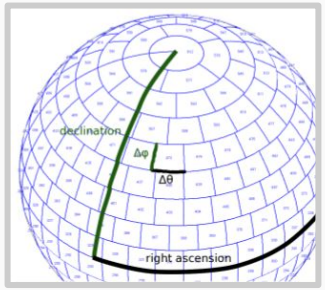

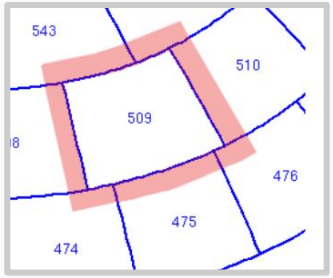

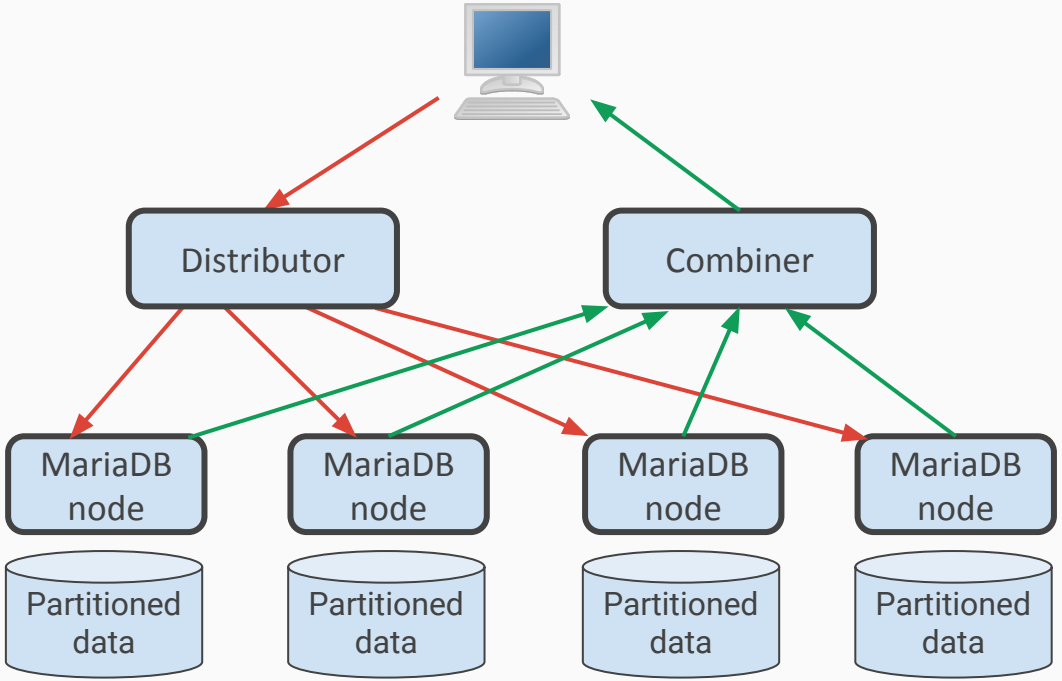

Cloud Native Deployment

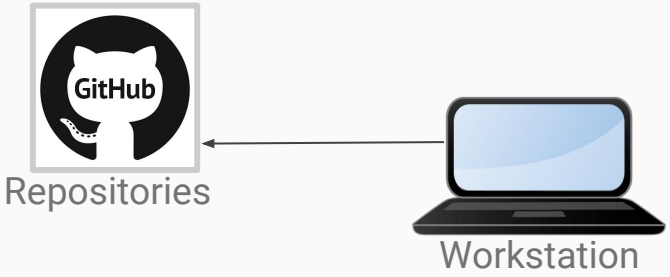

Cloud Native Deployment

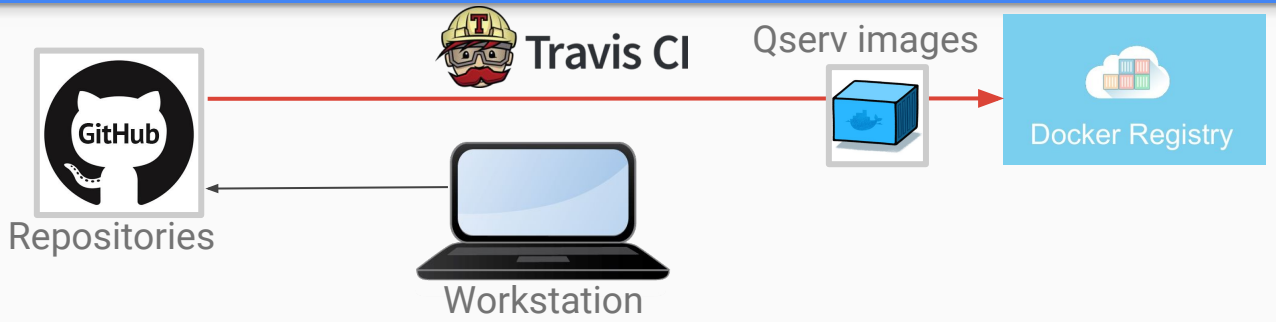

Cloud Native Deployment

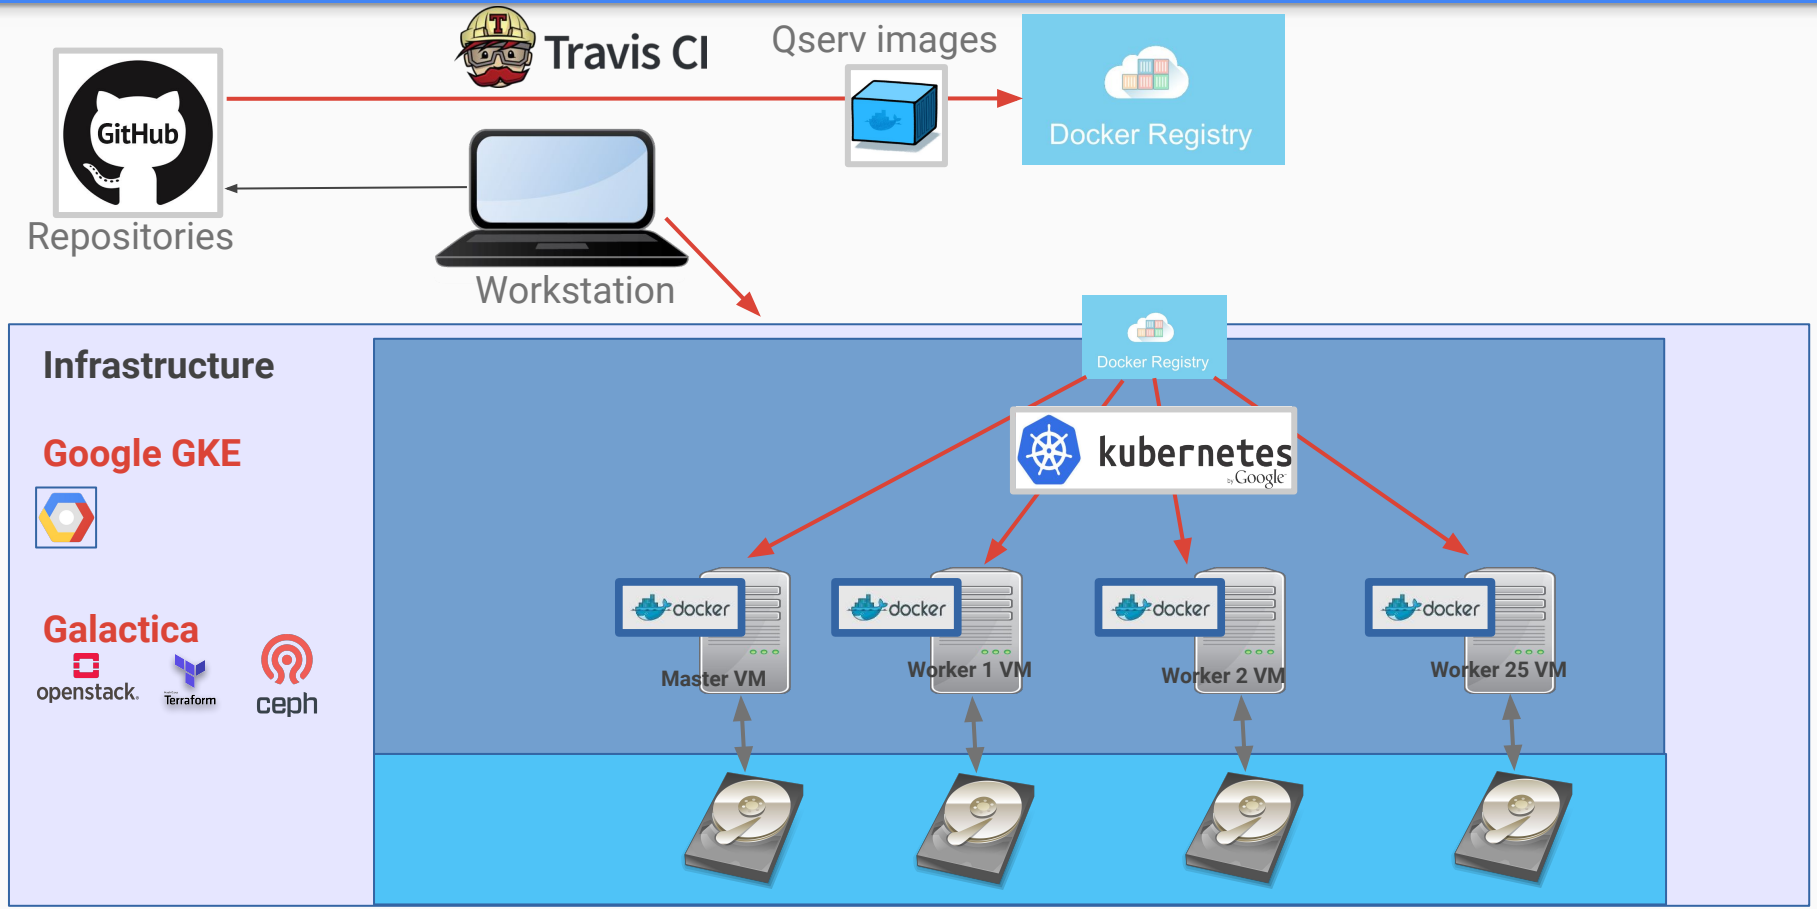

Cloud Native Deployment

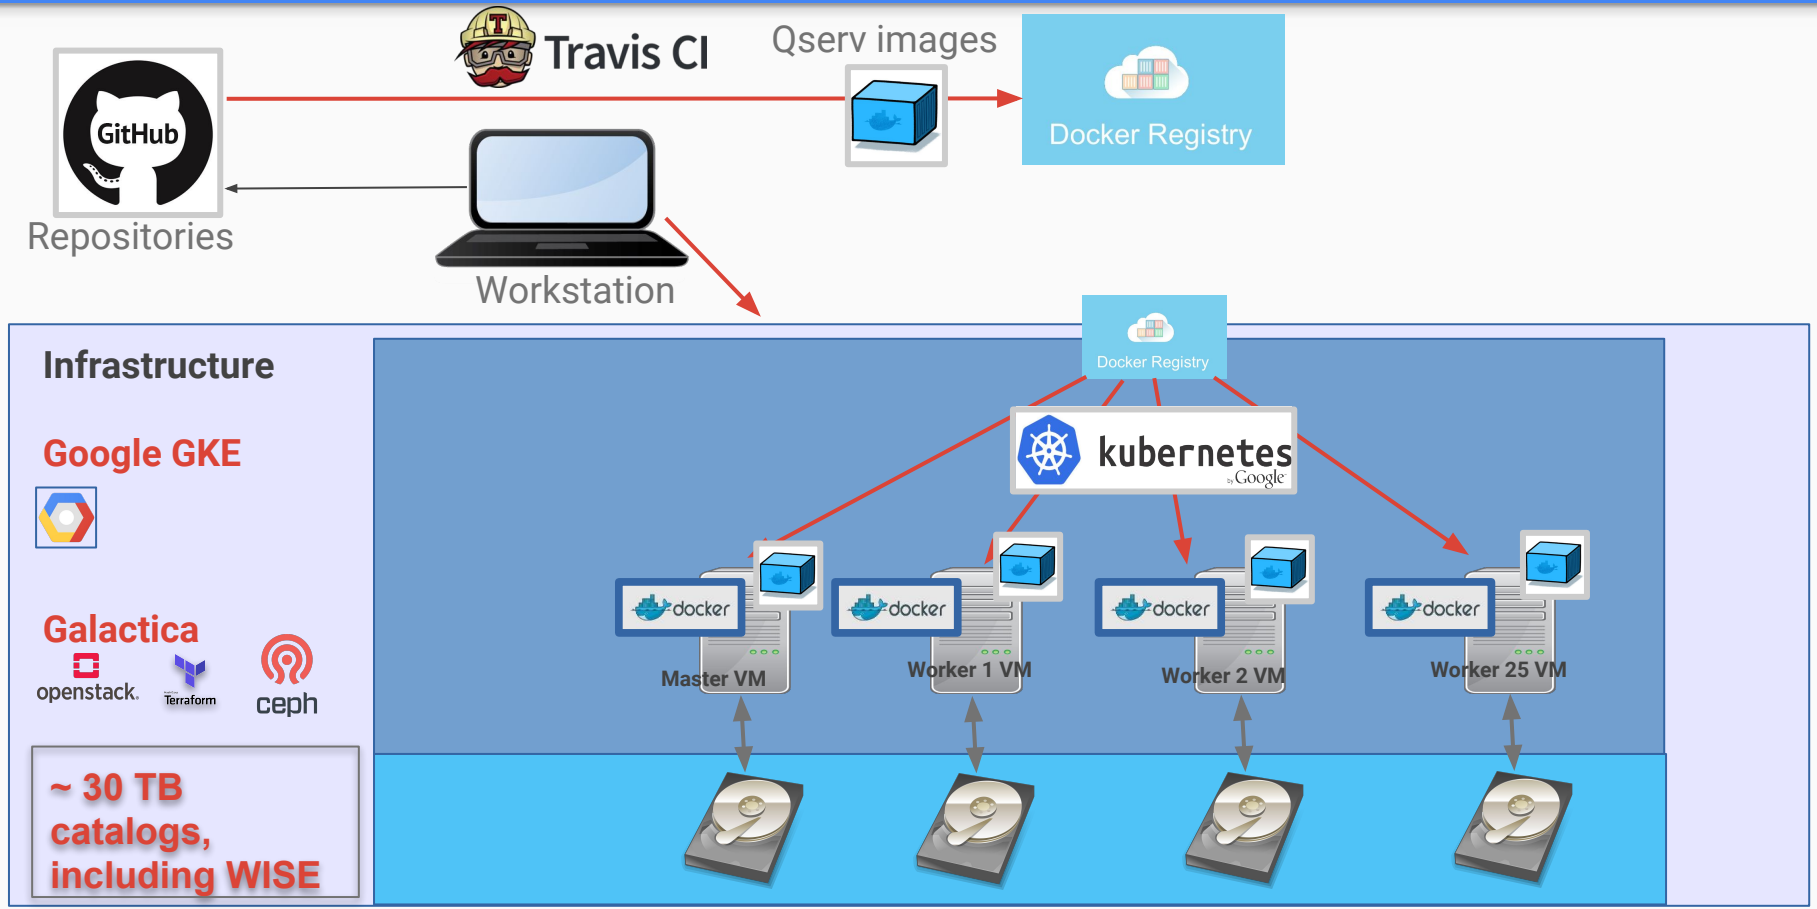

Automated deployment: bare-metal

CC-IN2P3@Lyon (~250TB)
NCSA@Illinois

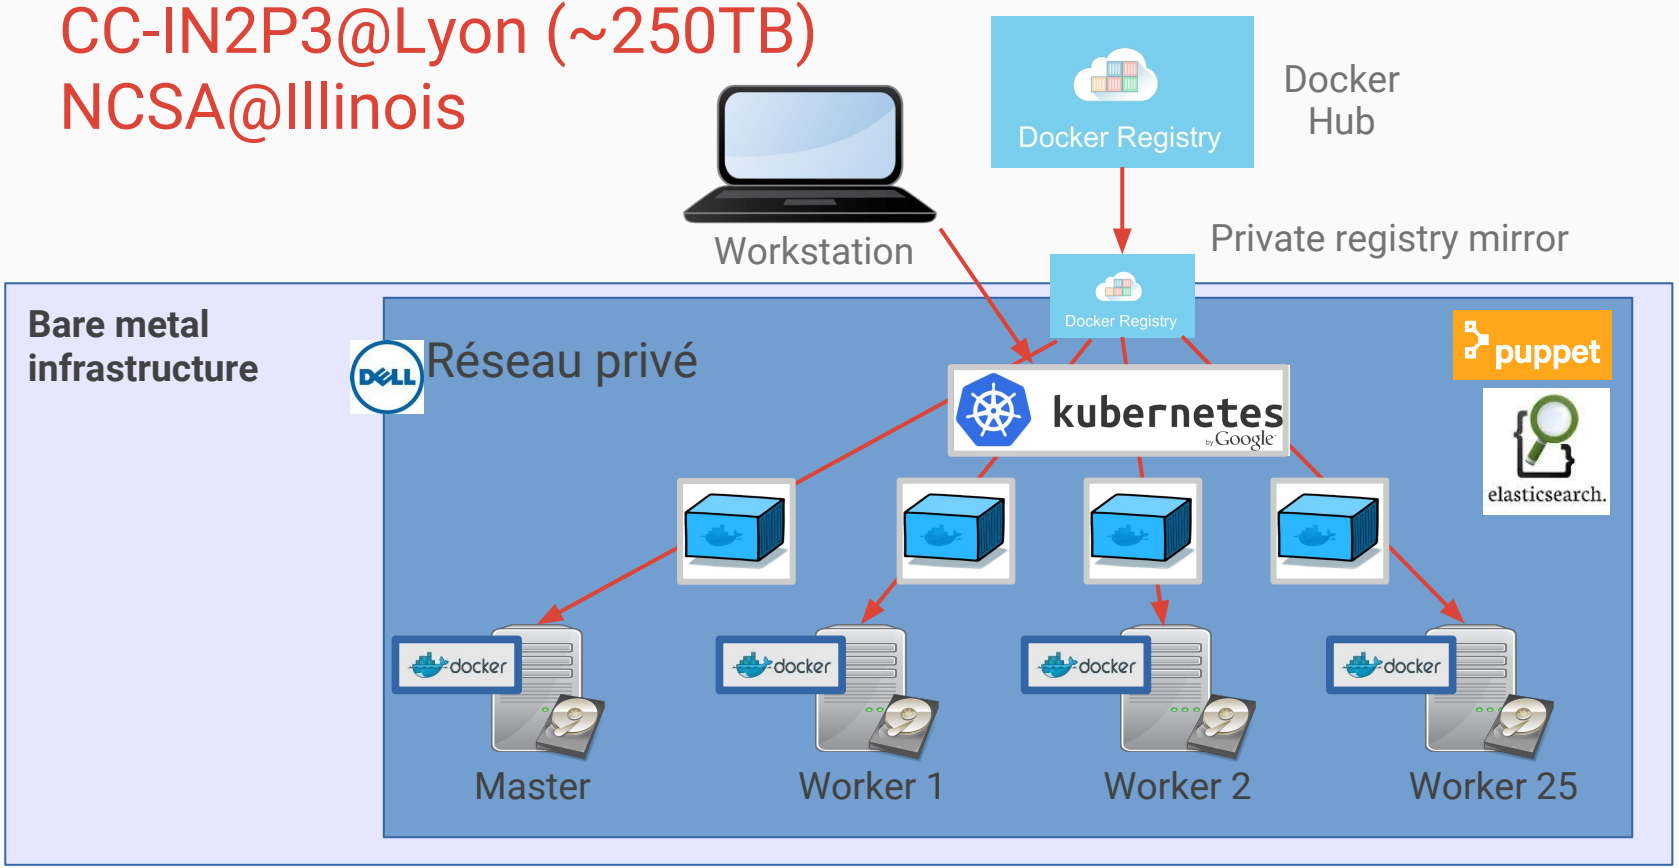

Deployment schemes

Legend:

Managed by Google
Managed by LSST

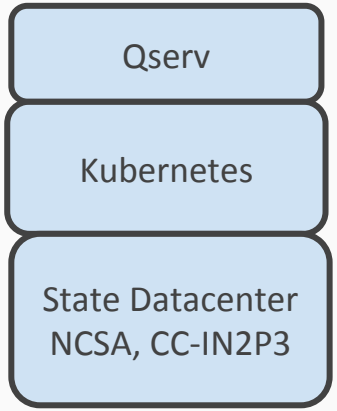

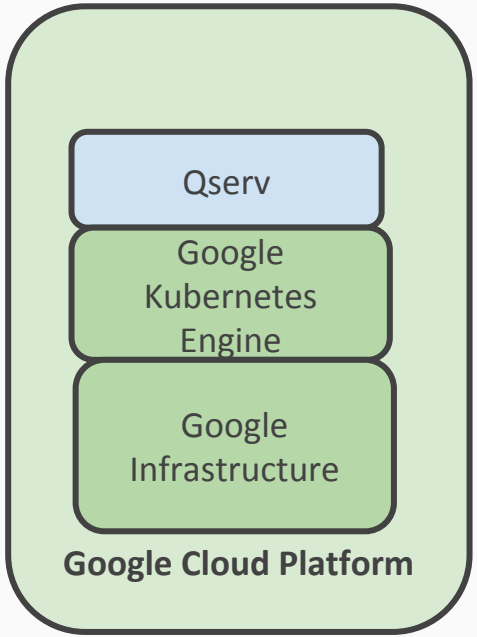

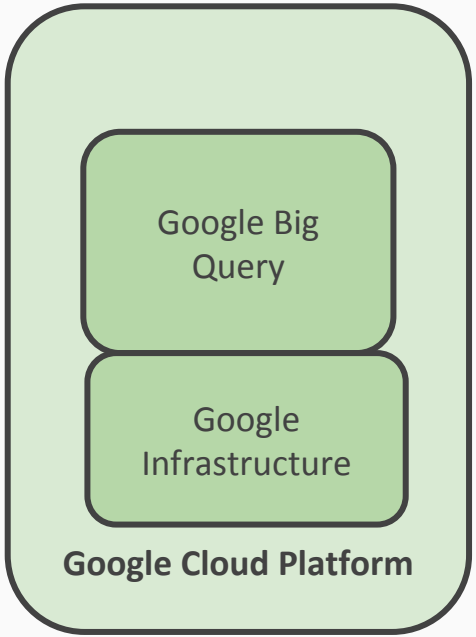

EVERYONE'S EXCITED ABOUT

KUBERNETES

Kubernetes: a modern version of LHC grid

- A portable, open-source, **container-centric** management platform
- Built-in primitives for **deployments, rolling upgrades, scaling, monitoring, and more**
- Inspired by **Google's internal systems**
- Get true **workload portability** and increased **infrastructure efficiency**

Google Cloud

CLOUD NATIVE COMPUTING FOUNDATION

= Open Source Cloud Computing For Applications

All you really care about

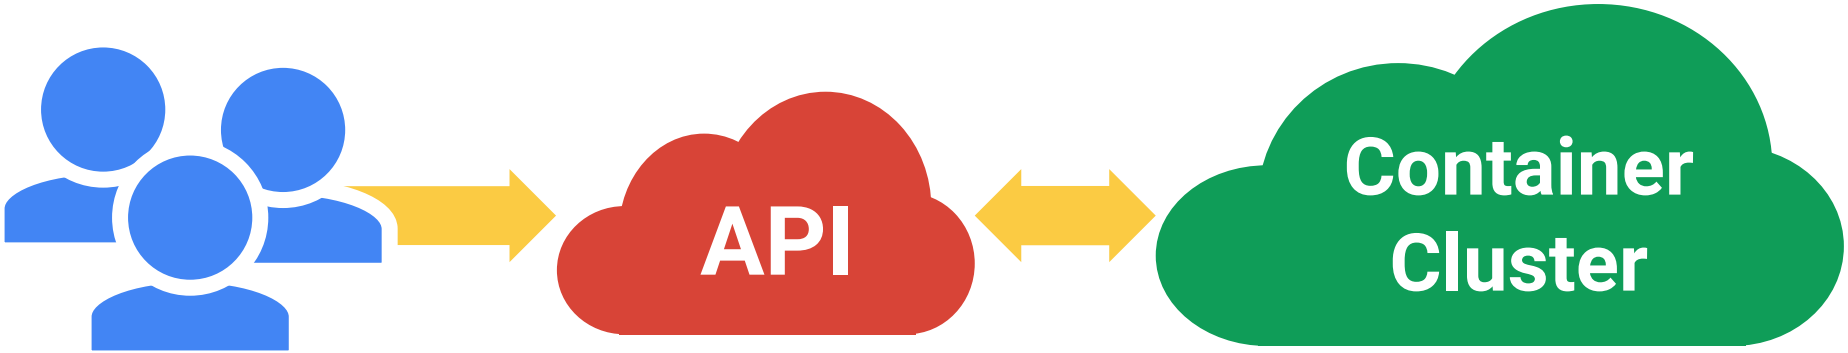

Workload portability

Result: Portability

Put your app on wheels and move it whenever and wherever you need

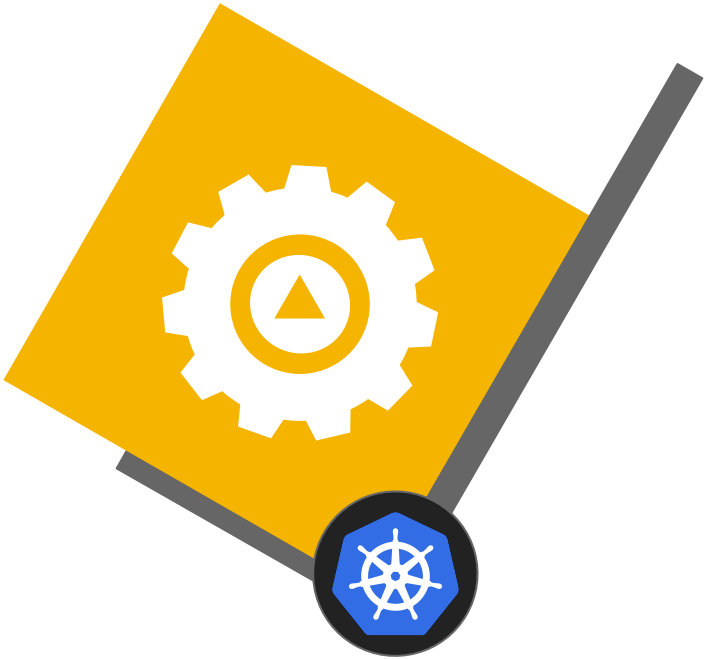

Kubernetes: a highest velocity Open Source Project

- Kubernetes 1.0 launched in **July 2015**
- New minor version every **3 months**. Currently on 1.10.
- Official **CNCF** Project

32 000+
pull requests
the latest year

60,000+
commits
the latest year

~23 PRs
merges/day
in the core repo

Google Cloud

30 Highest Velocity Open Source Projects

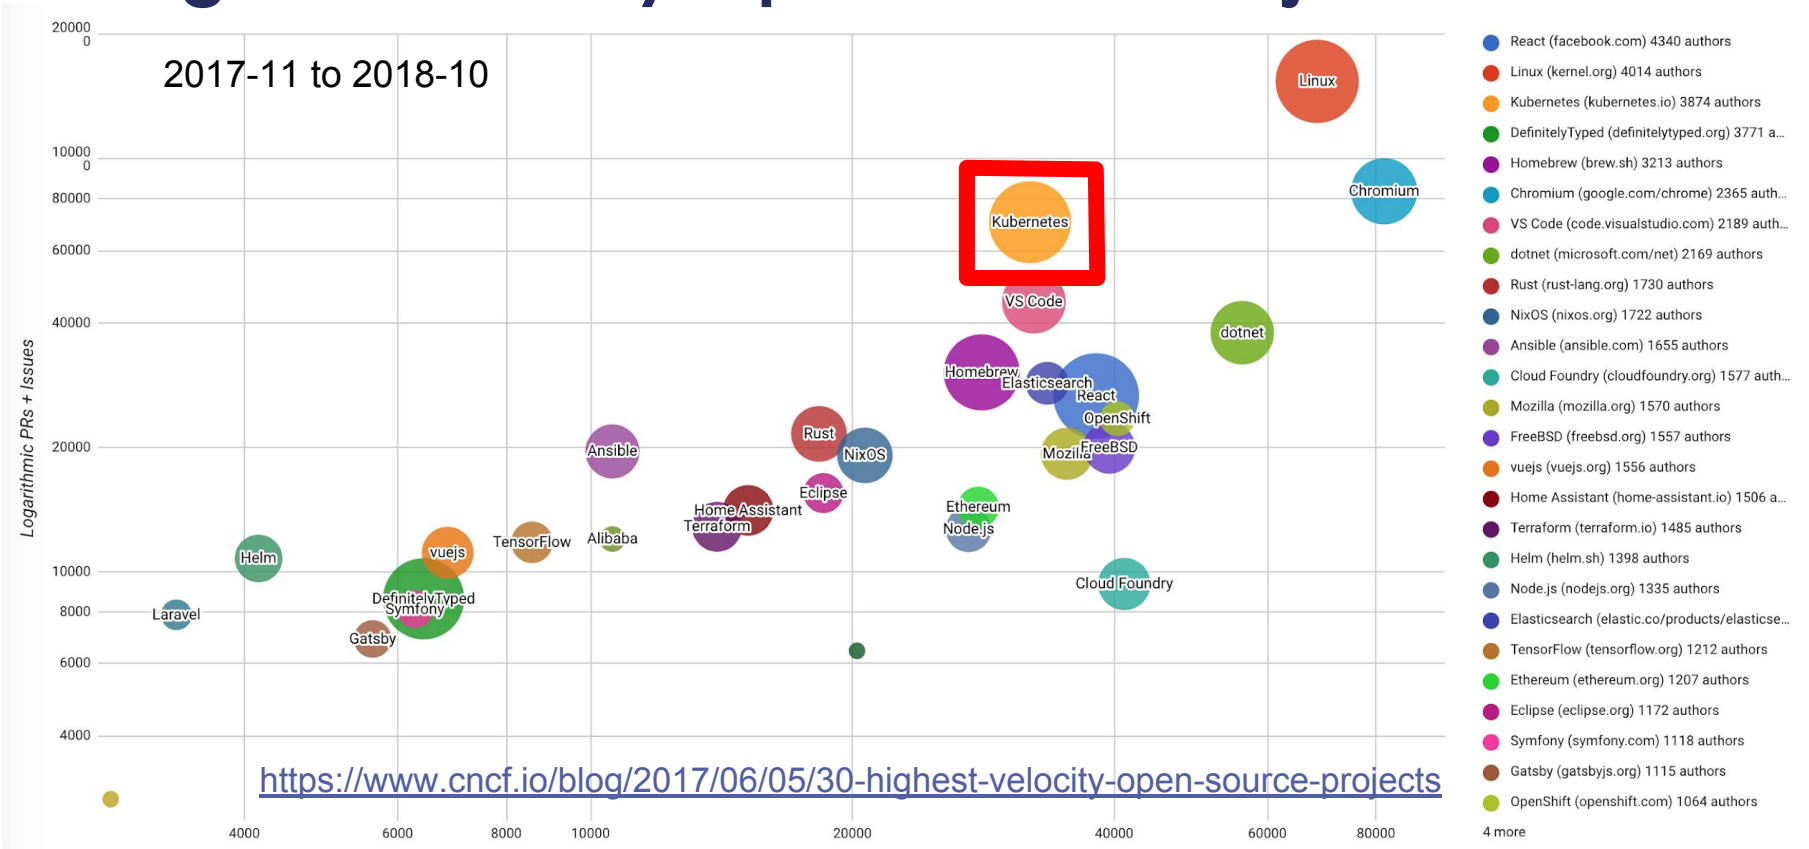

GKE Overview

Containers at Google

Each week, Google launches more than four billion containers across its data centers around the world. These containers house the full range of applications Google runs, including user-facing applications such as Search, Gmail, and YouTube.

Kubernetes was directly inspired by Google's cluster manager, internally known as Borg. Borg allows Google to direct hundreds of thousands of software tasks across vast clusters of machines numbering in the tens of thousands — supporting seven businesses with over one billion users each. Borg and Kubernetes are the culmination of Google's experience deploying resilient applications at scale.

But getting started can be challenging

Kubernetes The Hard Way

This tutorial walks you through setting up Kubernetes the hard way. This guide is not for people looking for a fully automated command to bring up a Kubernetes cluster. If that's you then check out Google Kubernetes Engine, or the Getting Started Guides.

- Prerequisites
- Installing the Client Tools
- Provisioning Compute Resources
- Provisioning the CA and Generating TLS Certificates
- Generating Kubernetes Configuration Files for Authentication
- Generating the Data Encryption Config and Key
- Bootstrapping the etcd Cluster
- Bootstrapping the Kubernetes Control Plane
- Bootstrapping the Kubernetes Worker Nodes
- Configuring kubectl for Remote Access
- Provisioning Pod Network Routes
- Deploying the DNS Cluster Add-on
- Smoke Test
- Cleaning Up

And don't forget "Day 2" ops

- Managing components
- Encrypting and securing etcd
- Configuring HA
- Rolling out security patches
- Backups and disaster recovery
- Bootstrapping TLS
- Managing users and policies

Kubernetes the Easy Way

Start a cluster with one-click

View your clusters and workloads in a single pane of glass

Google keeps your cluster up and running

Google Cloud

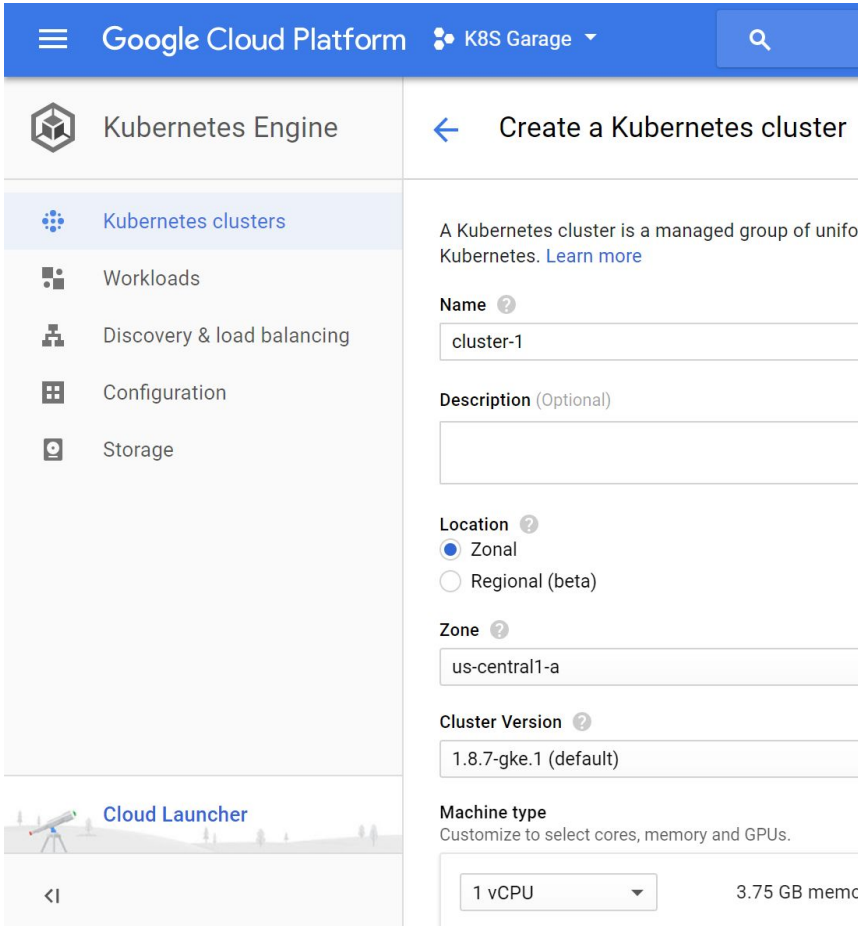

Google
Kubernetes Engine

Google Cloud

Production Ready Kubernetes

Get Started Quickly

Use the Cloud Console to **easily** create a cluster and view your workloads

Reliable and Available

GKE automatically **repairs**, **upgrades,** and **scales** your cluster with a **high-availability** control plane

Integrated with GCP

Take advantage of the **deep integration** with Google Cloud Platform features and services

Kelsey Hightower ✓
@kelseyhightower

Following

Google Container Engine is the cheat code for Kubernetes. Thanks to my fellow Googlers I get to focus on using Kubernetes; not managing it.


lyft


shopify


NIANTIC


ocado


PHILIPS


ebay


Meetup


mixpanel


The New York Times


CANACCORD Genuity


Google Cloud

Kubernetes on Google Cloud Platform

- Use with Preemptible VMs, custom machine types, GPUs
- Live migration and patching
- Best-in-class networking
- Sustained use discounts, committed use discounts, and per-second billing mean you only pay for what you use
- Integration with Google Cloud Platform services such as BigQuery, PubSub, Cloud SQL, and more

Compute and Storage

Can use Preemptible VMs, custom machine size, and GPUs

Provision volumes with GCP Persistent Disk, SSD, and Local SSD

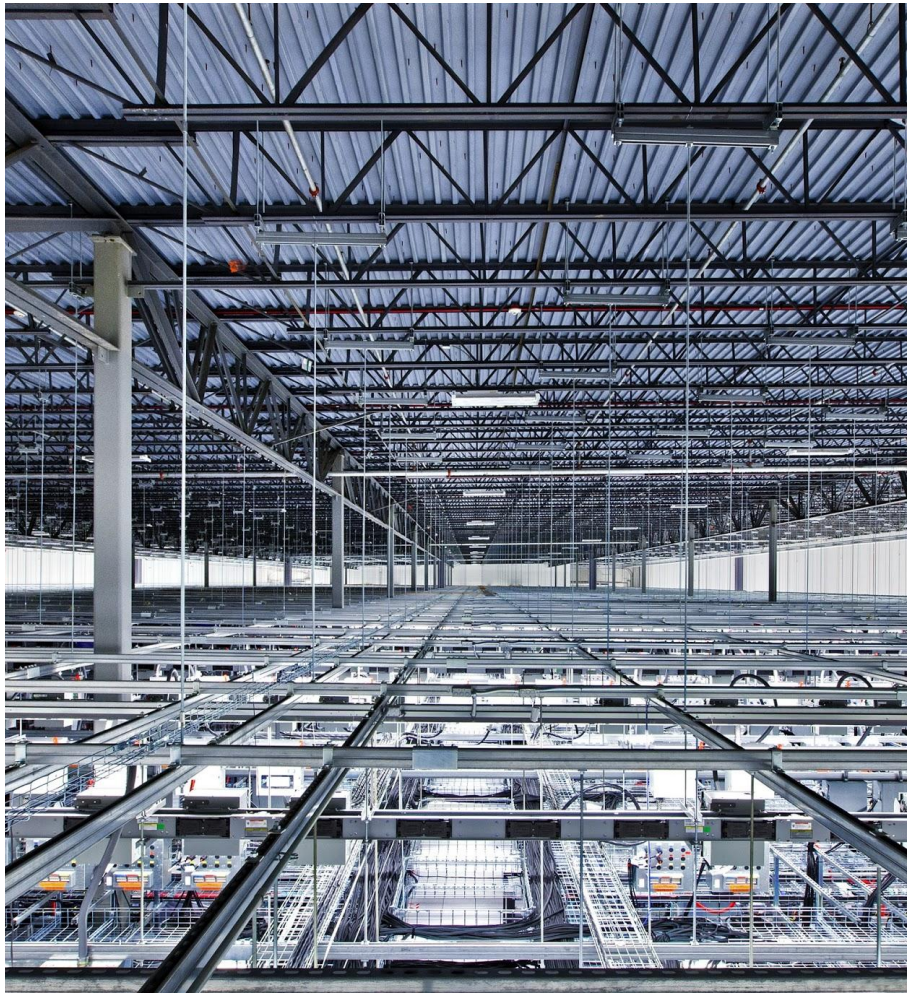

Google Cloud

GKE Cluster Features

Google Cloud

Auto Kubernetes

Auto-repair

Automatically initiate repair process for nodes that fail a health check.

Auto-upgrade

Keep the control plane and nodes in the cluster up-to-date with the latest stable version

Auto-scale

Cluster autoscaling handles increased demand and scales back as needed

Horizontal Pod Autoscaler

Automatically add (or remove) pods when needed.

Based on CPU metrics and custom metrics with Stackdriver

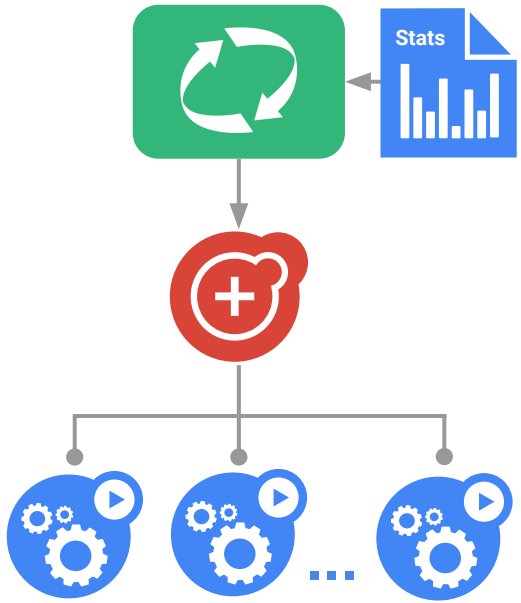

Cluster Autoscaler

Cluster Autoscaler will add nodes when pods are failing to be scheduled due to insufficient resources (IE CPU/Memory).

Kubernetes uses requests to schedule pods, therefore scaling will occur based on pod resource requests not pod resource utilization.

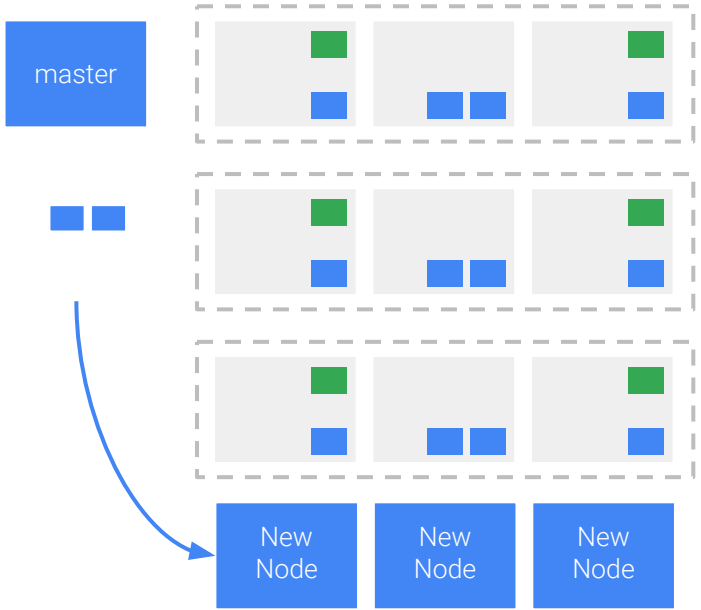

Google Cloud

Cluster Autoscaler Balancing Across Zones

When **scaling up** across different zones, Cluster Autoscaler attempts to **balance the size of each node pool**

This helps prevent an uneven number of nodes when distributing traffic across multiple zones in a region

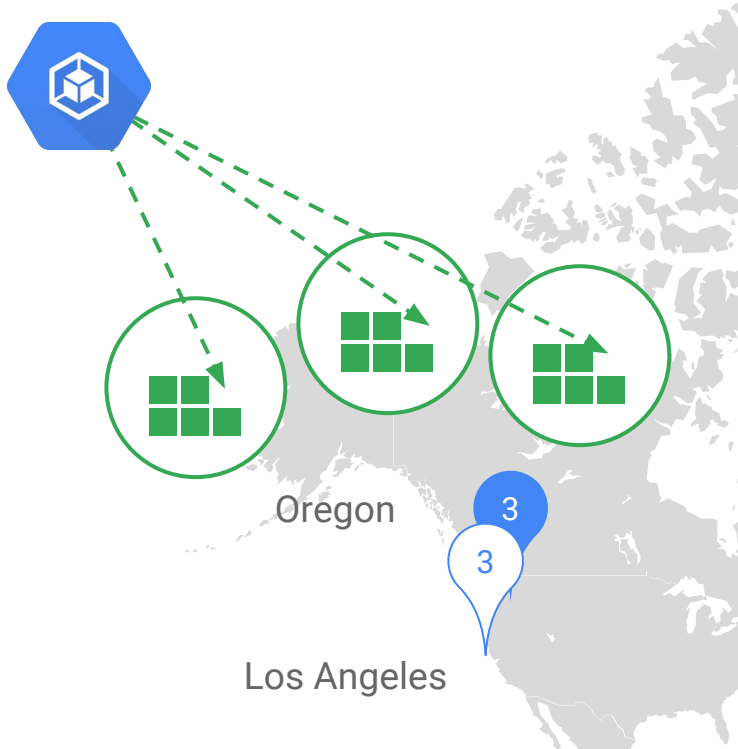

Google Cloud

Security

- All GKE components are **encrypted at rest**. This includes etcd where secrets are stored.
- **TLS** for master-to-master and node-to-master communication
- Integrated with Calico to provide **network policy** to control pod-to-pod communication
- **Private clusters** makes your master inaccessible from the public internet
- **Metadata concealment** isolates workloads from node metadata

Integrated with GCP Networking

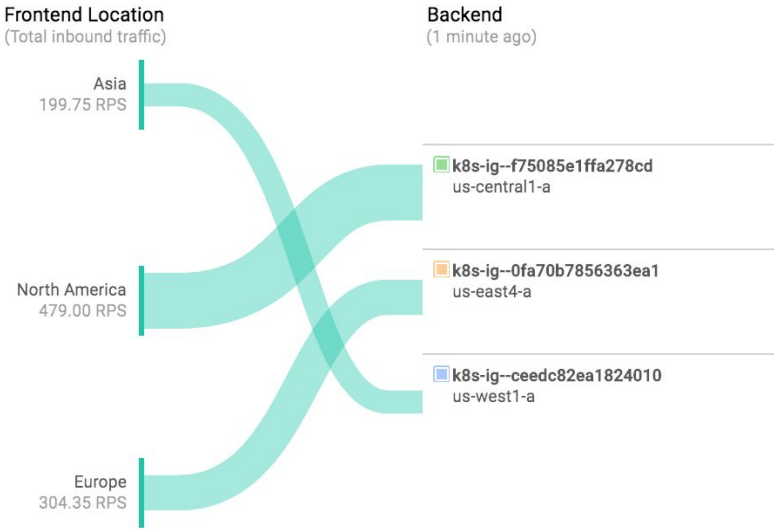

Google Cloud

Logging and Monitoring

Use Stackdriver Logging to automatically collect, process, and store your container and system logs

Integrated with Audit Logging

Stackdriver Monitoring will monitor your cluster's CPU and memory as well as custom metrics for your application

Wrapping Up

1. Kubernetes is a container orchestration platform that allows you to focus on your **application** instead of infrastructure
2. GKE is the **easiest way** to get started with Kubernetes
3. **GKE can be installed on your own infrastructure (early access)**

BigQuery

How can we help LSST

SLAC, October 2018

Google Cloud

Google Innovation in Big Data Analytics @ Scale

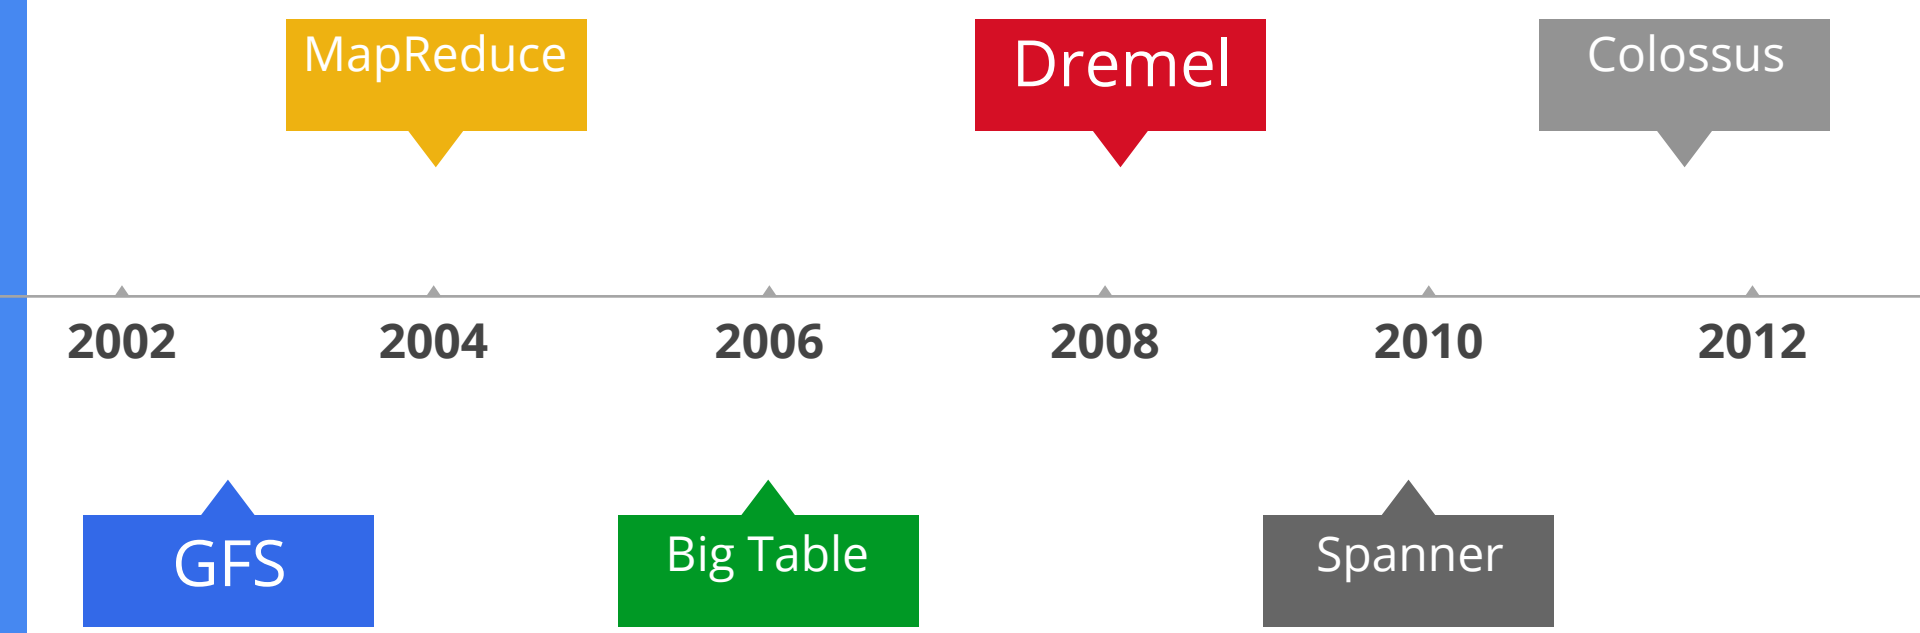

Google Cloud

Good things about BigQuery

- Data encrypted at rest
- No keys or indexes
- **Caching**
- Free batch ingest
- Streaming ingest, real time
- Serverless
- Easy access sharing, control
- **100PB storage**

Google Cloud

BigQuery architecture

Decoupled storage and compute

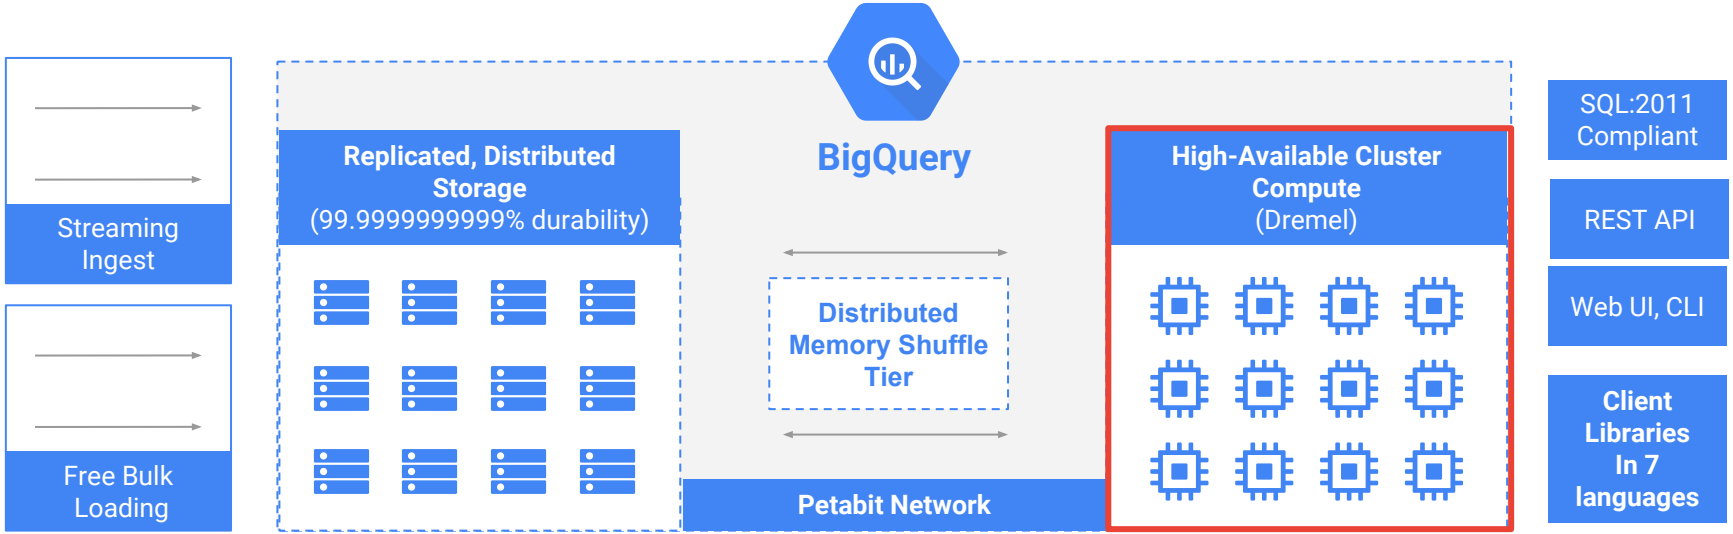

Google Cloud

Simple query execution

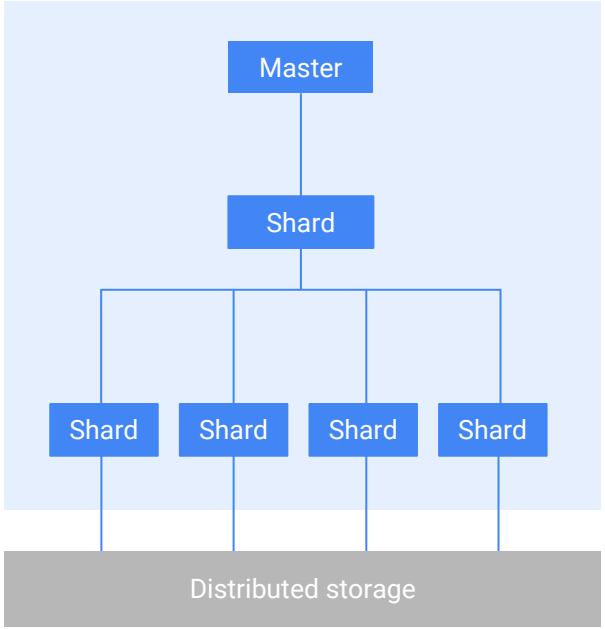


```
SELECT COUNT(*) FROM
wikipedia_benchmark.Wiki1B
WHERE title LIKE "G%o%o"
```

Stage 2: Sum (1 slot)

Stage 1: Filter, Count (220 slots)

Google Cloud

BigQuery by the Numbers

Google Cloud

Big Stuff and Costs

- Active storage
  - $0.02 per GB ($20/mo per TB < 1 VM and 1 TB persistent disk)
  - The first 10 GB (yay!) is free each month. See Storage pricing for details.
- Long-term storage
  - $0.01 per GB
  - The first 10 GB is free each month. See Storage pricing for details.
- Streaming Inserts
  - $0.01 per 200 MB
  - You are charged for rows that are successfully inserted. Individual rows are calculated using a 1 KB minimum size. See Streaming pricing for details.
- Queries (analysis)
  - $5 per TB
  - First 1 TB per month is free, see On-demand pricing for details. Flat-rate pricing is also available for high-volume customers.

Spherical Geometry -- GIS Queries

Google Cloud

GIS Queries - Supports Spherical Geometry

- Right Ascension and Declination on celestial sphere
- Arbitrary geometries can be described with Geo objects
- Variety of operations:
  - ST_DWithin
  - ST_Intersects
  - ST_Contains
  - ST_Within
  - ST_Covers
  - ST_Equals

Google Cloud

UDFs -- User Defined Functions

Google Cloud

Two Flavors of UDF -- SQL is Better

Javascript

```
CREATE TEMP FUNCTION plusOne(x
FLOAT64)
RETURNS FLOAT64
LANGUAGE js
AS "return x+1;";
```

SQL

```
CREATE TEMP FUNCTION
addFourAndDivide(x INT64, y
INT64) AS ((x + 4) / y);
```

Google Cloud

Tracks

1. Google provides high level services
2. GKE: ease a lot engineer work
3. Google Big Query: robust and powerful
4. Total cost is easy to retrieve and should be compared to regular datacenters
5. **What does LSST want to delegate to Google?**

Thanks!

Contact:

Fabrice JAMMES
LPC
Clermont-Ferrand

fabrice.jammes@in2p3.fr