Cluster management at Google

2015-02

john wilkes / johnwilkes@google.com
Principal Software Engineer

Google Cloud Platform

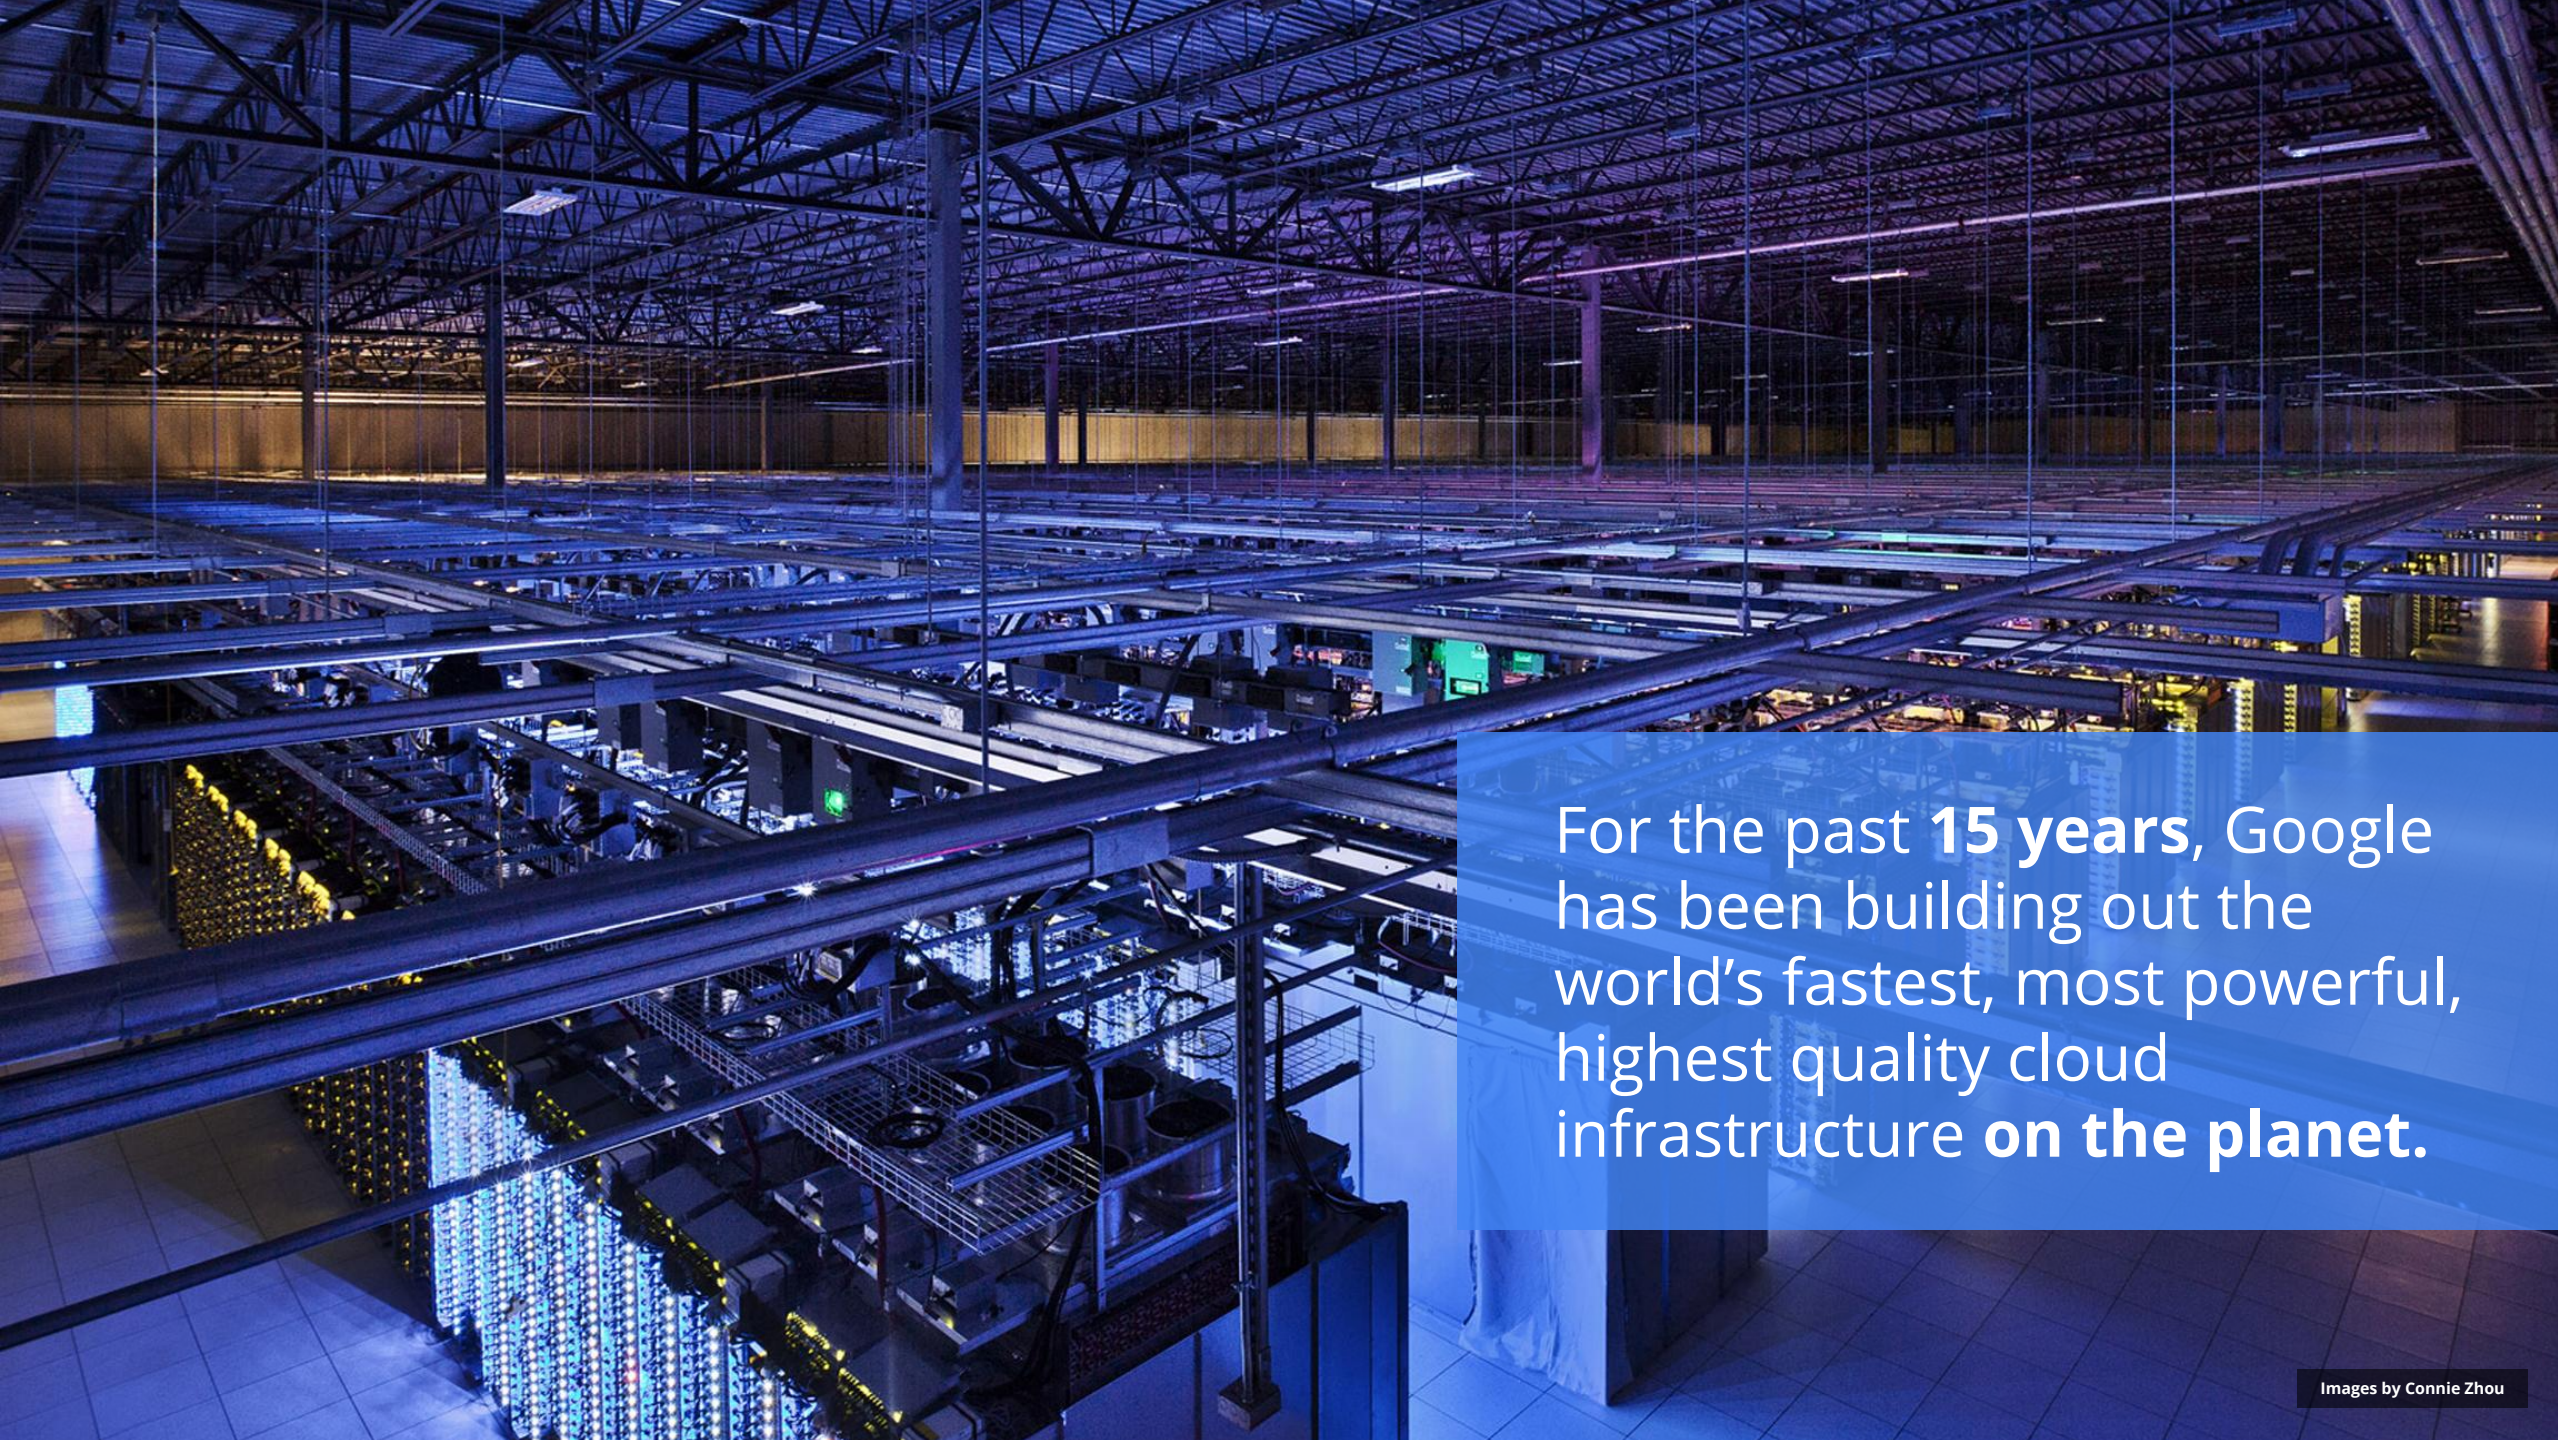

For the past **15 years**, Google has been building out the world's fastest, most powerful, highest quality cloud infrastructure **on the planet.**

Images by Connie Zhou

Hello World

```
job hello_world = {
   runtime = { cell = 'ic' }                // What cluster should we run in?
   binary = '.../hello_world_webserver'     // What program are we to run?
    args = { port = '%port%' }              // Command line parameters
   requirements = {      // Resource requirements
     ram = 100M
     disk = 100M
     cpu = 0.1
   }
   replicas =  10000    // Number of tasks
}
```

Hello World

```
> borgcfg .../hello_world_webserver.borg up
...
About to affect 10000 tasks and 1 packages on cell IC.
Do you wish to continue (yes/no) [no]? yes

==== Staging package hello_world_webserver.63ce1b965155c75e/johnwilkes on ic... SUCCESS
==== Making  package hello_world_webserver.63ce1b965155c75e/johnwilkes on ic... SUCCESS
==== Starting job hello_world on ic... SUCCESS
```

Hello World


Running tasks
10000
7500
5000
2500
0
0:00:30
0:01:00
0:01:30
0:02:00
0:02:30
0:03:00
Elapsed time

Hello World

What just happened?

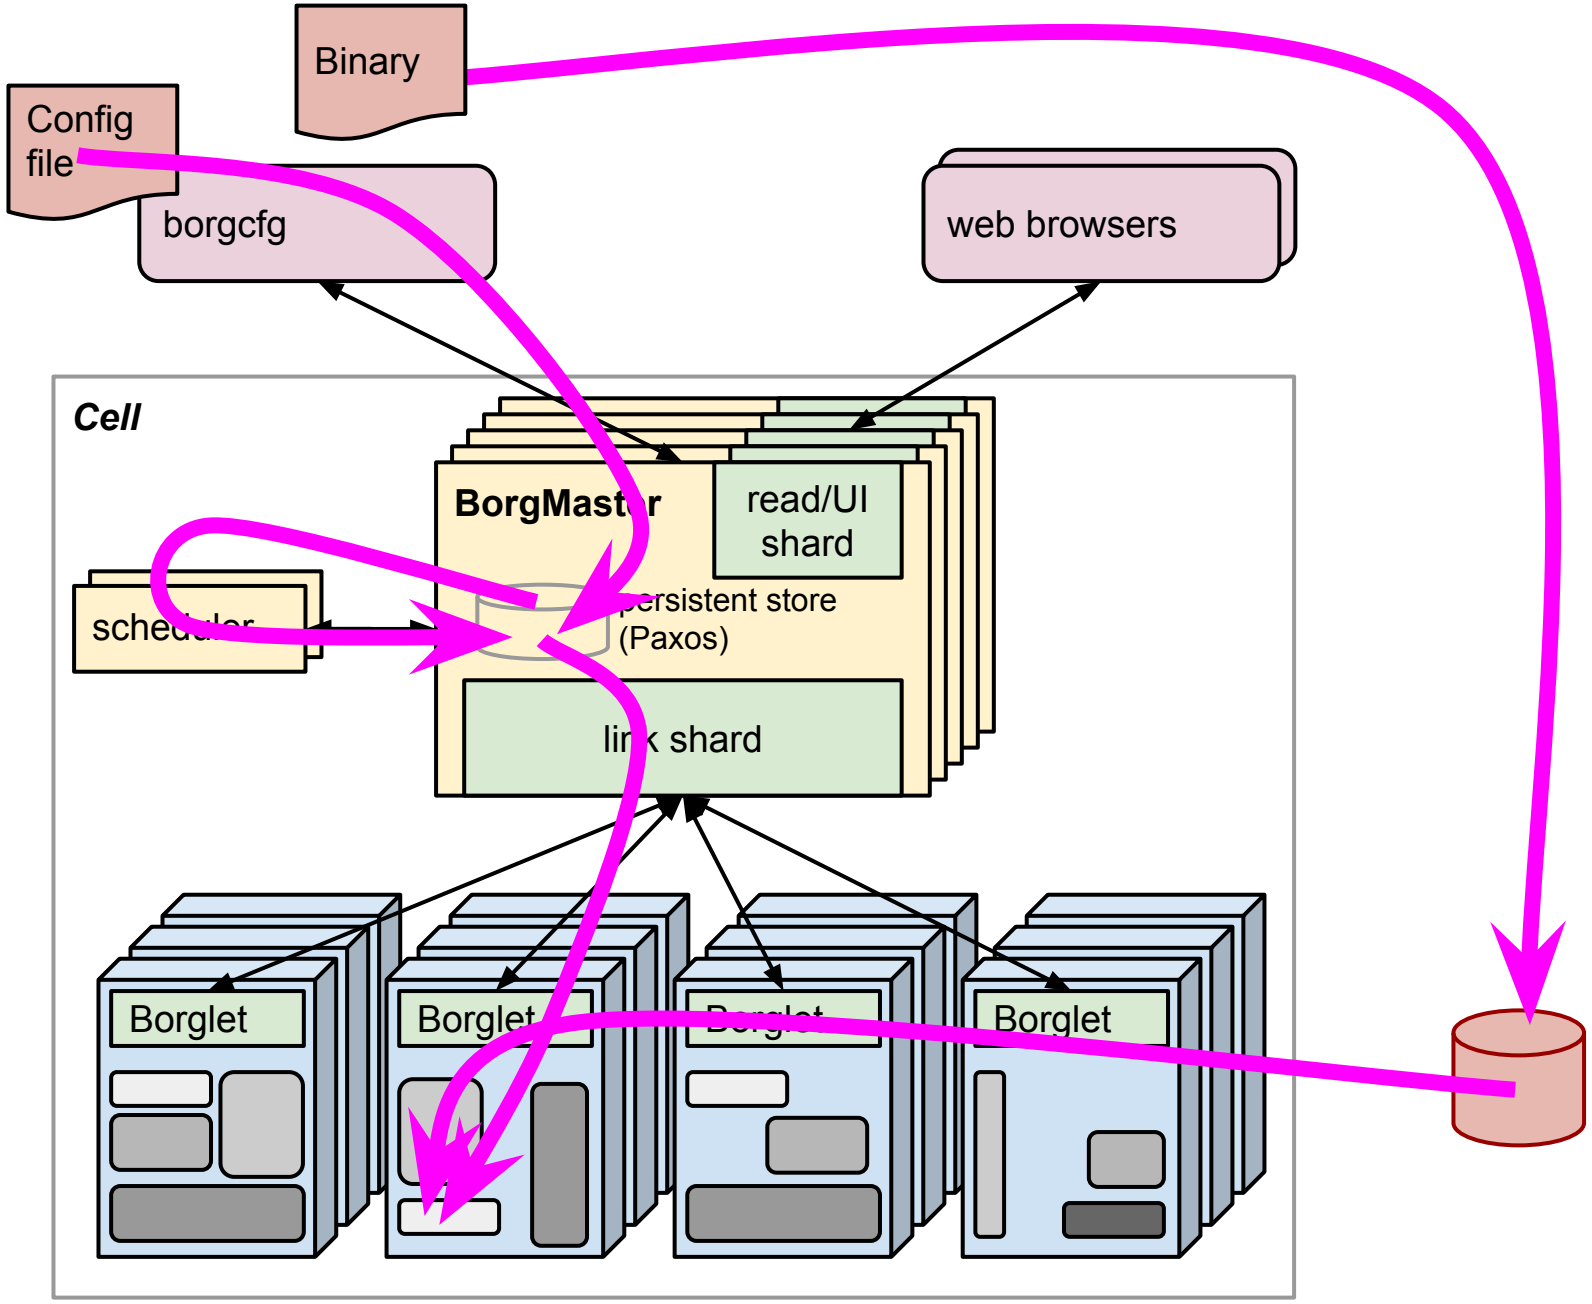

Hello World

Images by Connie Zhou

Hello World

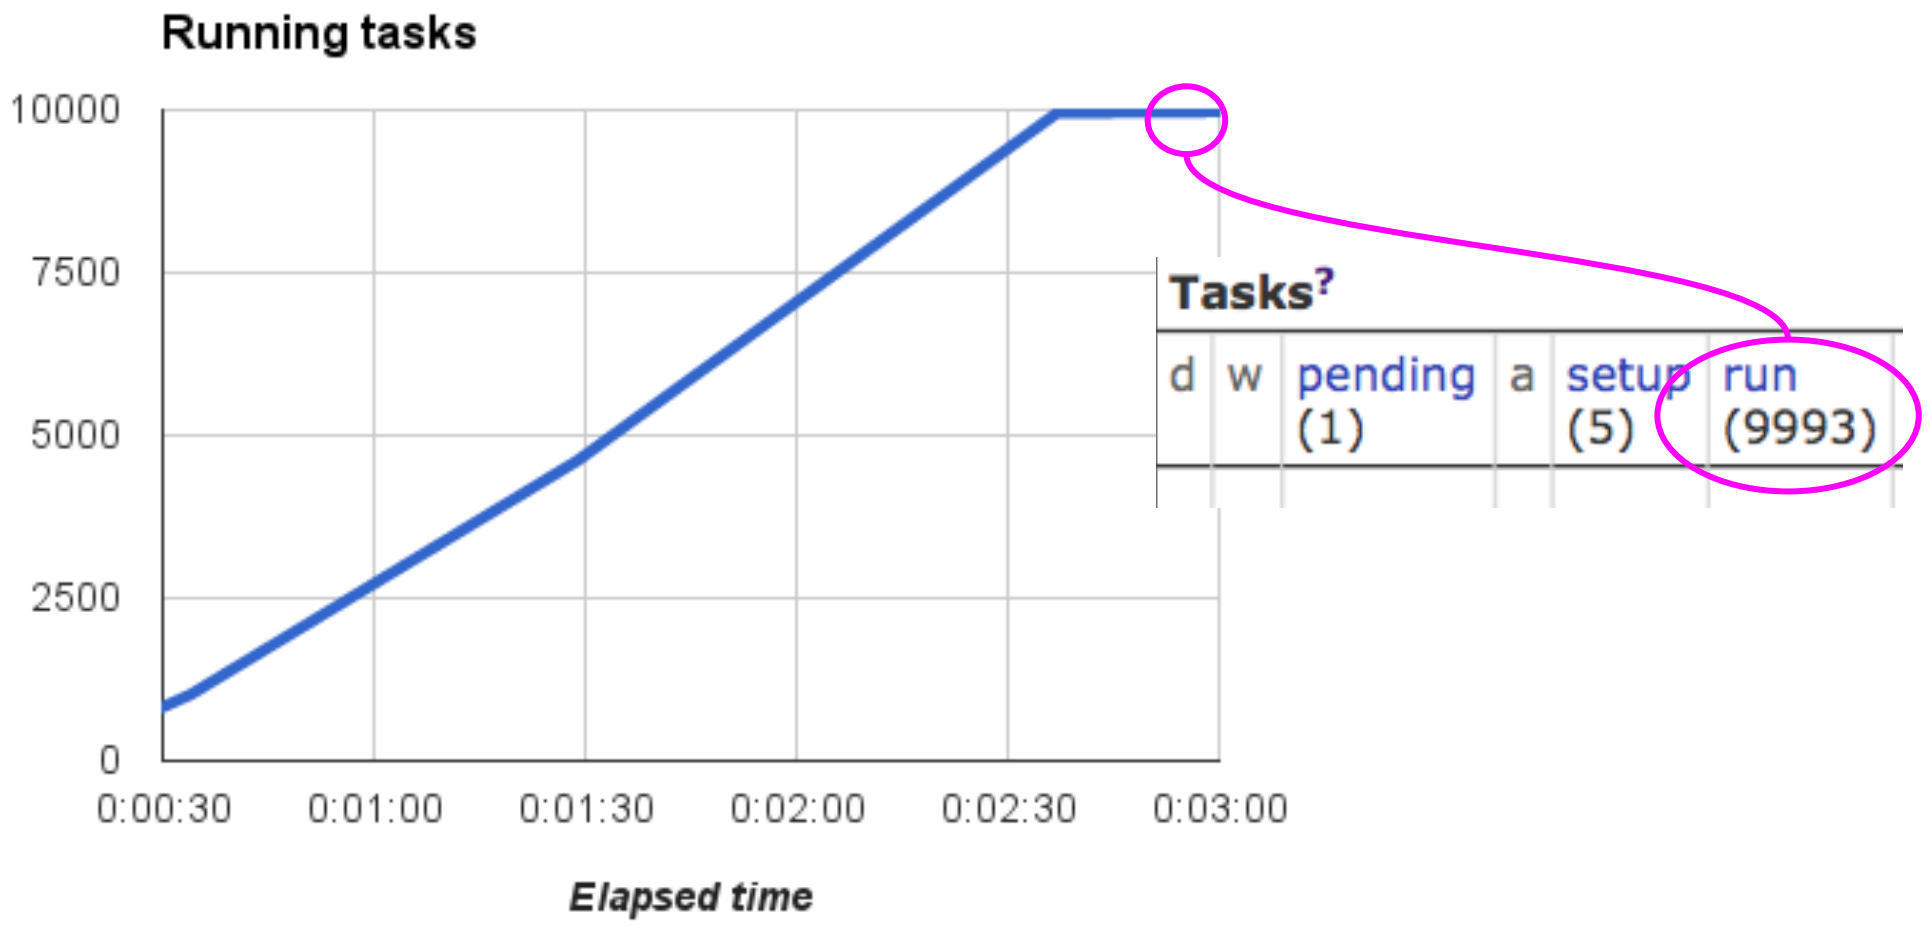
Running tasks
10000
7500
5000
2500
0
0:00:30
0:01:00
0:01:30
0:02:00
0:02:30
0:03:00
Elapsed time
Tasks?
d
w
pending (1)
a
setup (5)
run (9993)

Failures

task-eviction rates
and causes

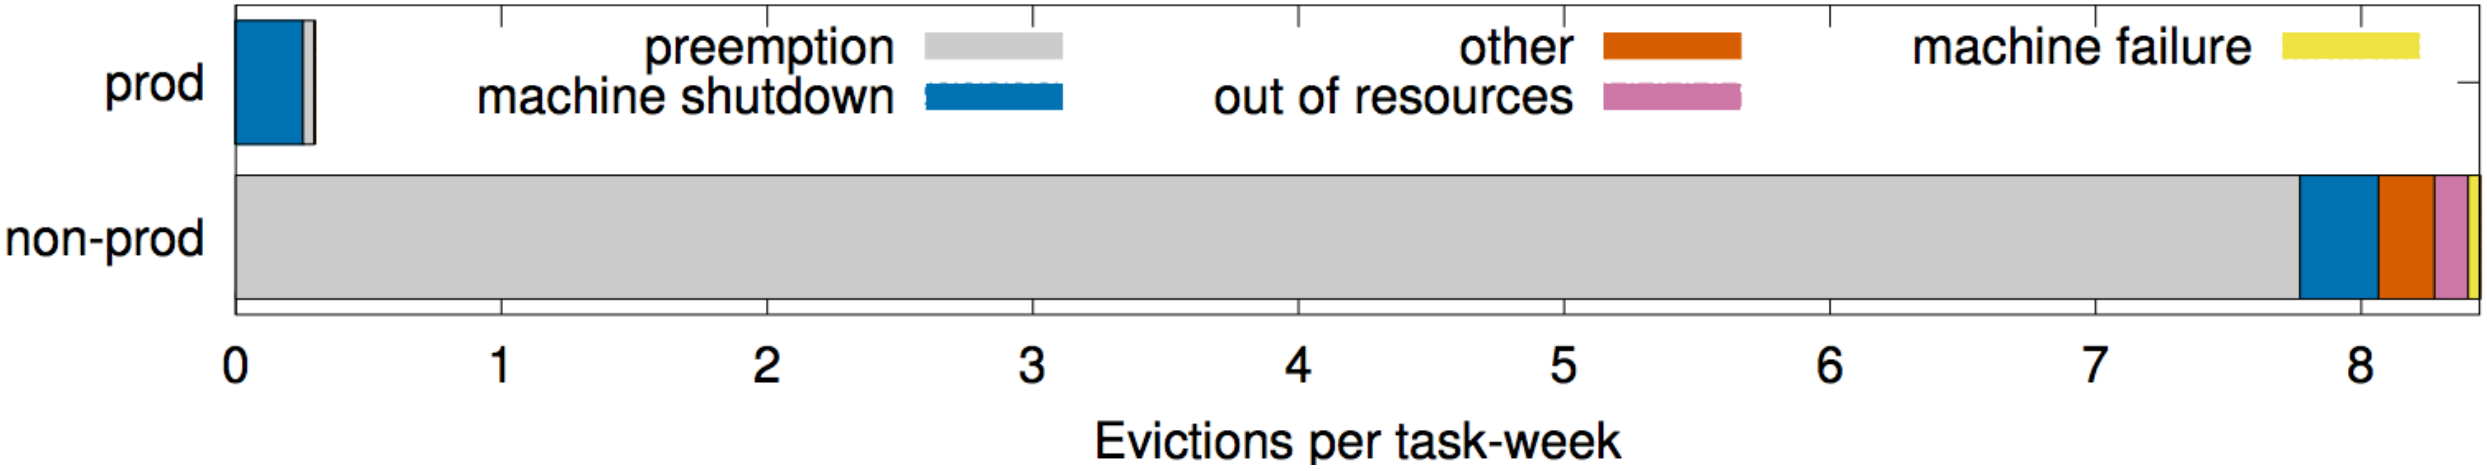

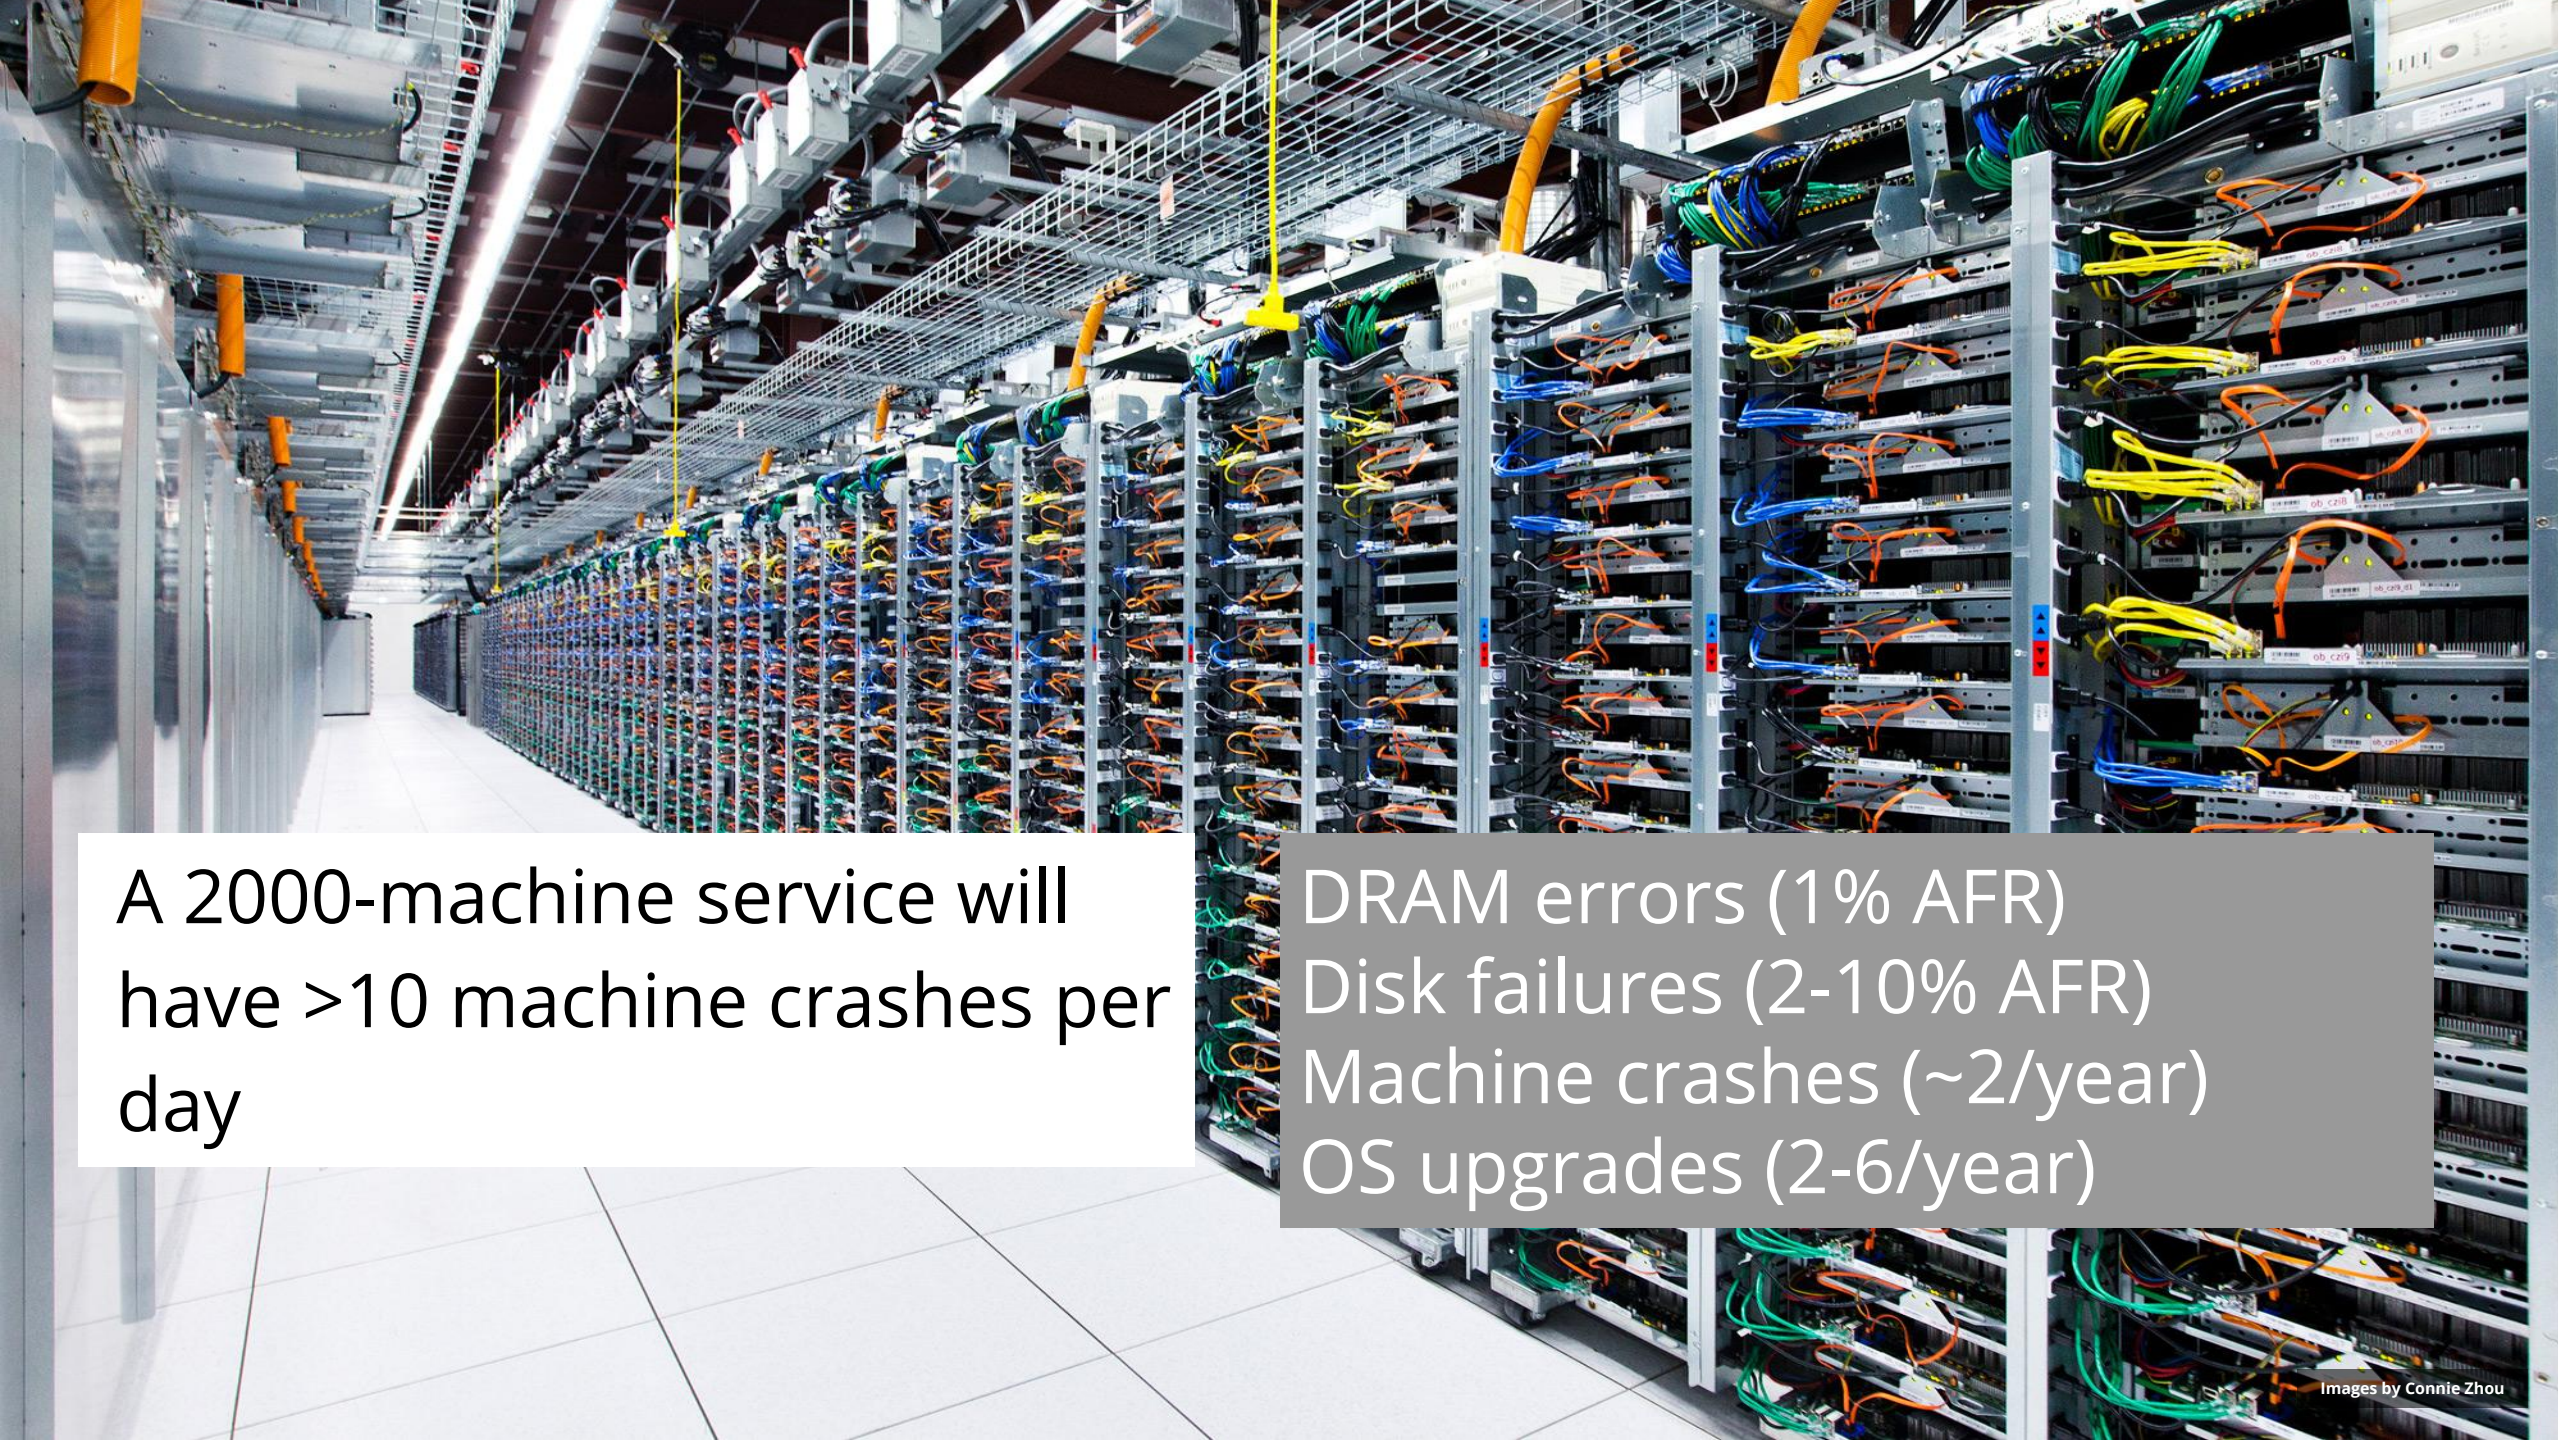

A 2000-machine service will have >10 machine crashes per day

DRAM errors (1% AFR)
Disk failures (2-10% AFR)
Machine crashes (~2/year)
OS upgrades (2-6/year)

Images by Connie Zhou

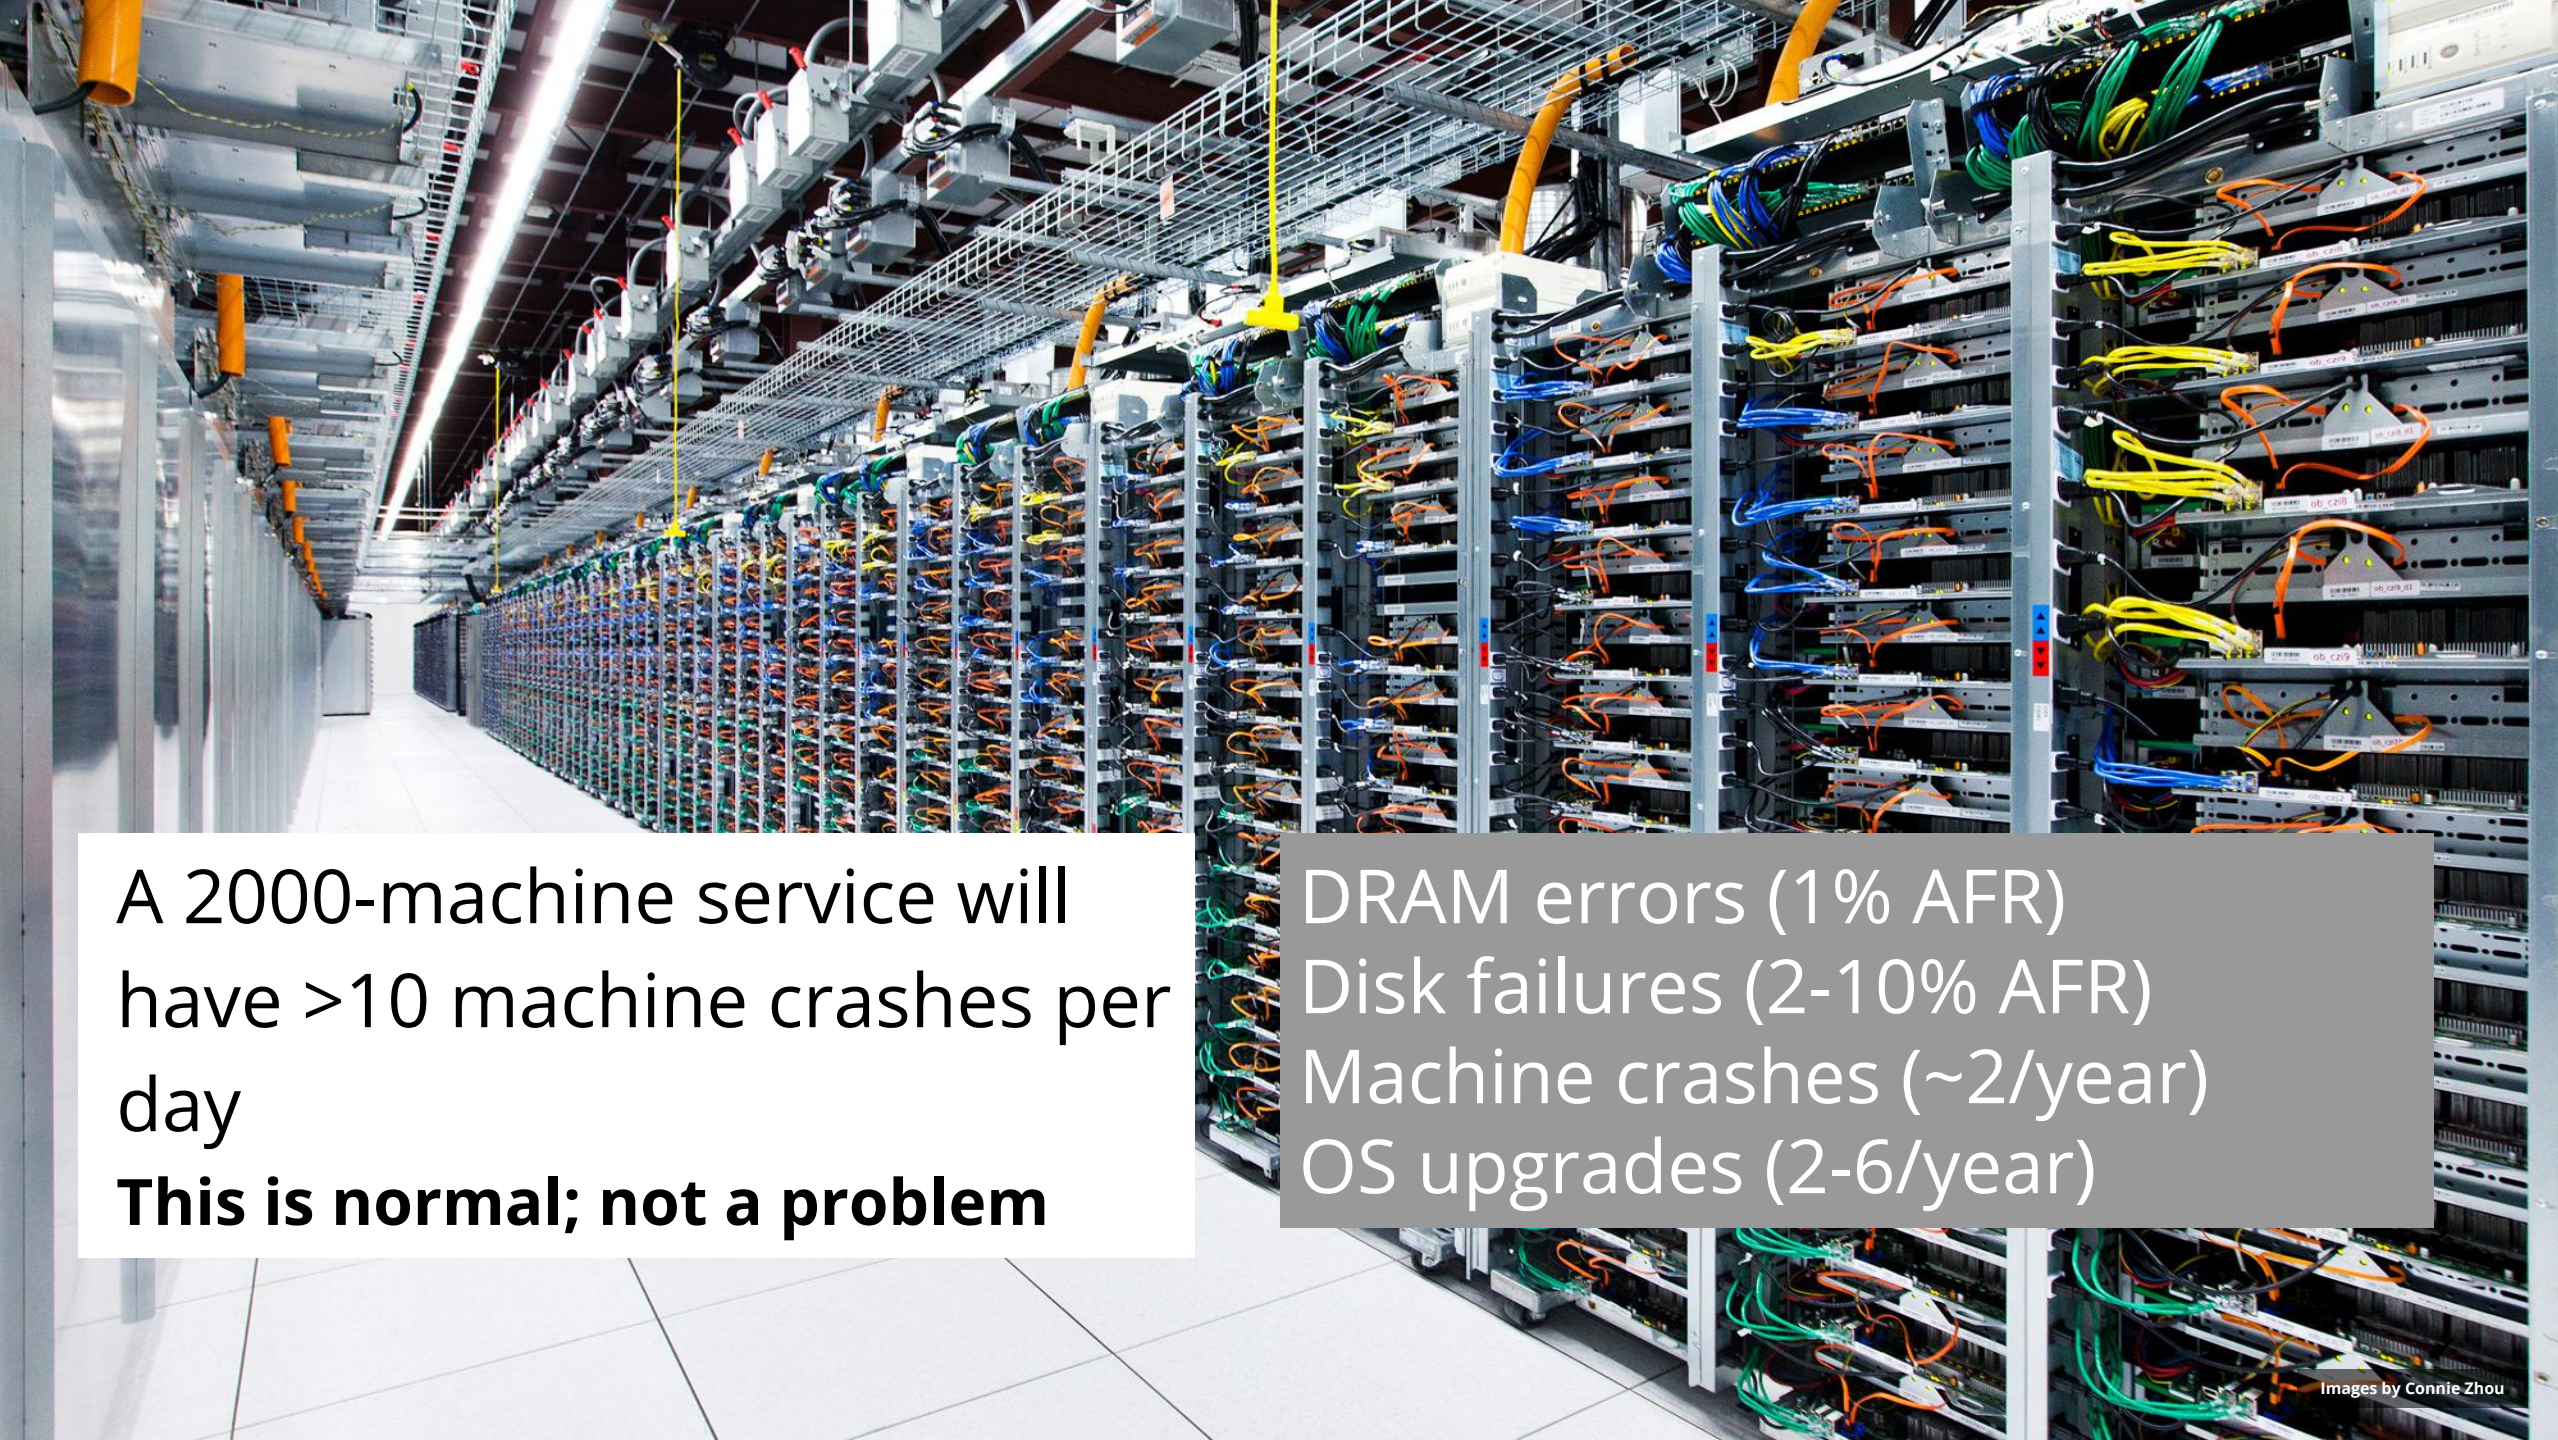

A 2000-machine service will have >10 machine crashes per day

This is normal; not a problem

DRAM errors (1% AFR)
Disk failures (2-10% AFR)
Machine crashes (~2/year)
OS upgrades (2-6/year)

Images by Connie Zhou

Efficiency

Advanced bin-packing algorithms

Experimental placement of production VM workload, July 2014

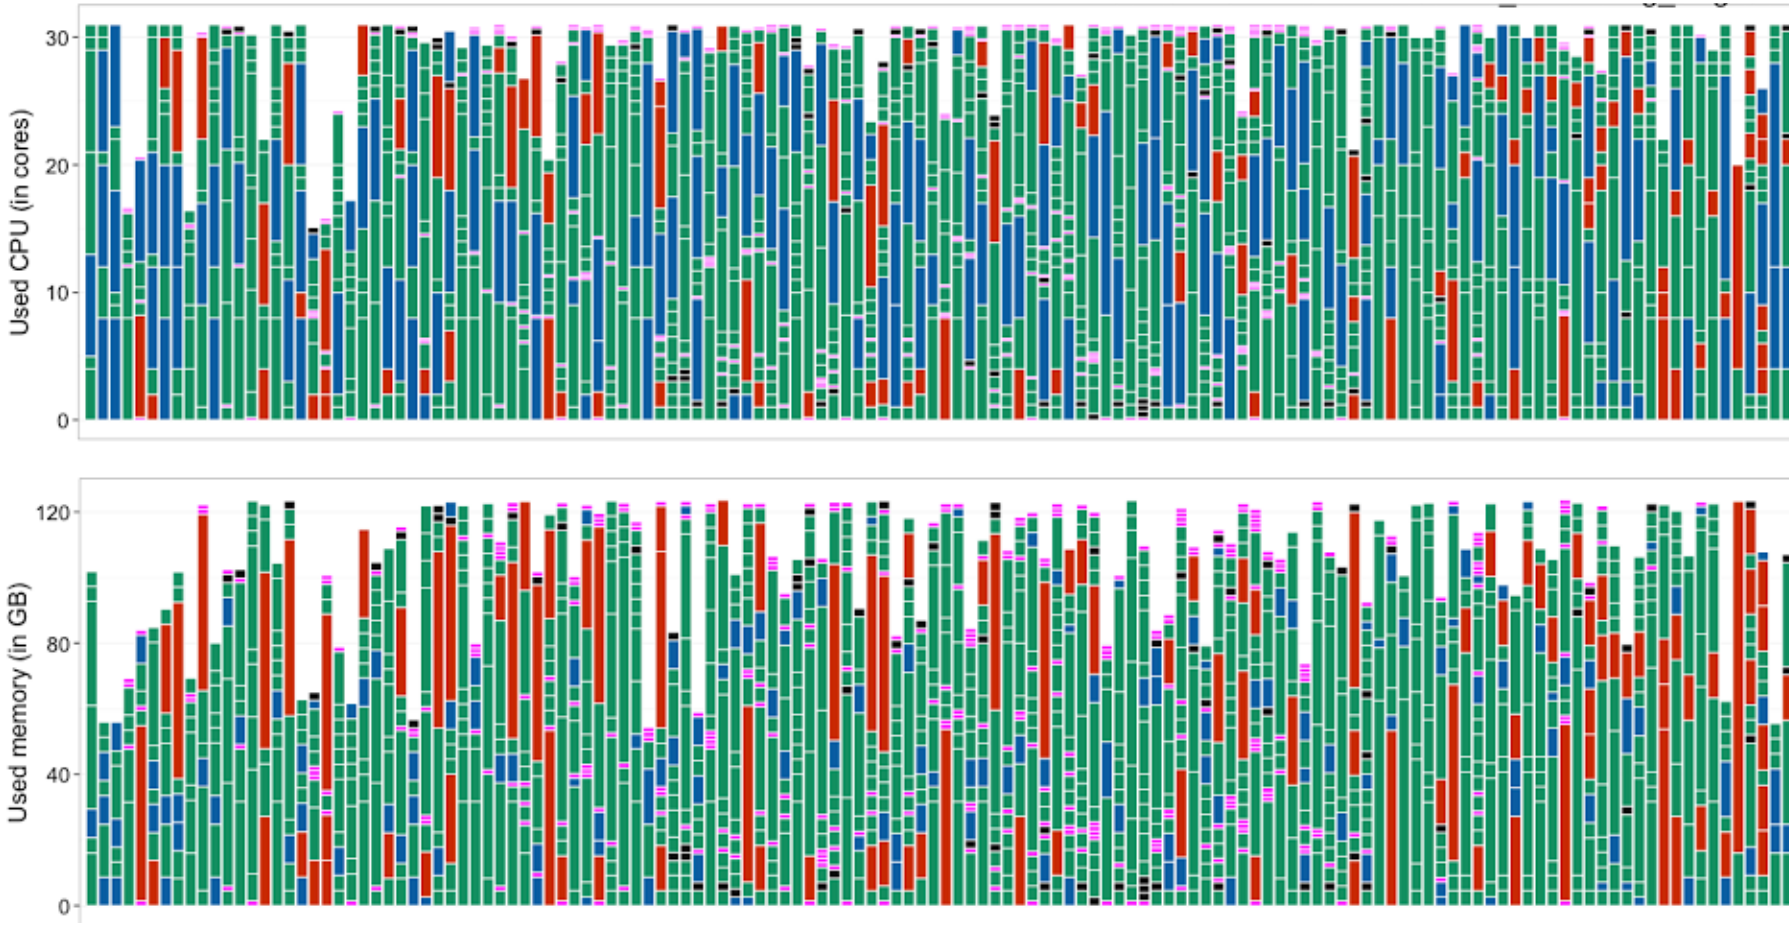

Efficiency

Advanced bin-packing algorithms

There are no obvious bucket sizes (cf. cloud VMs)

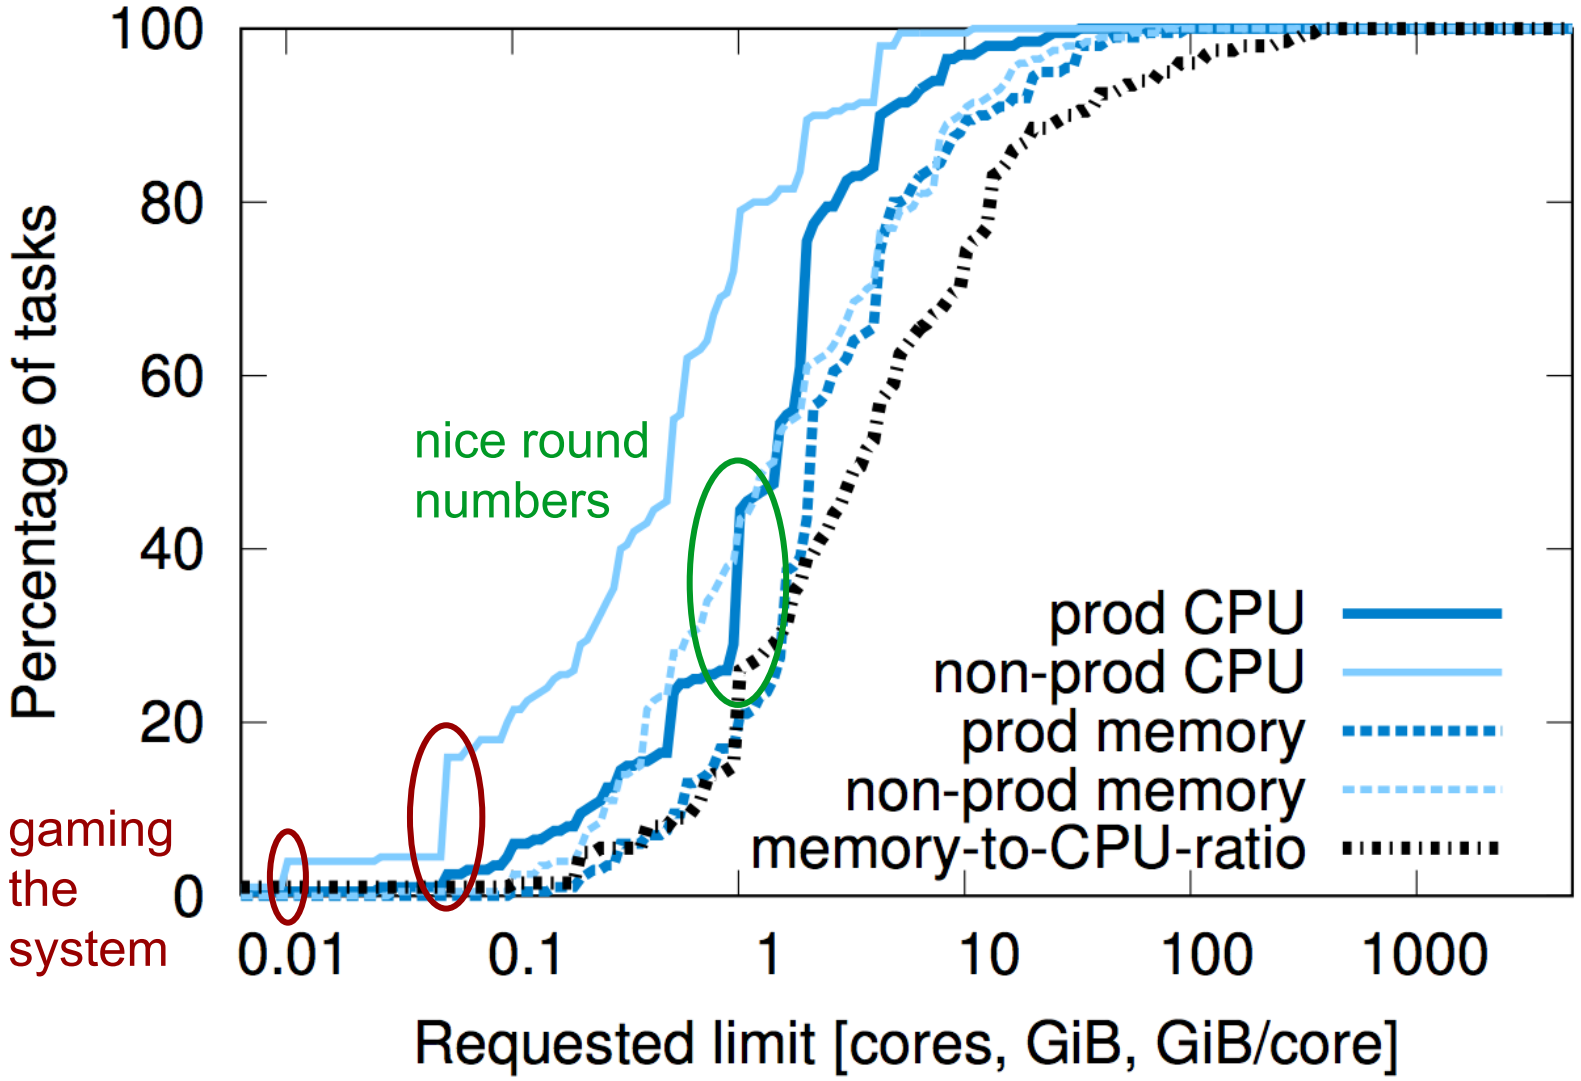

Efficiency

Advanced bin-packing algorithms

Heterogeneous workloads, May 2011

Omega paper, EuroSys 2013

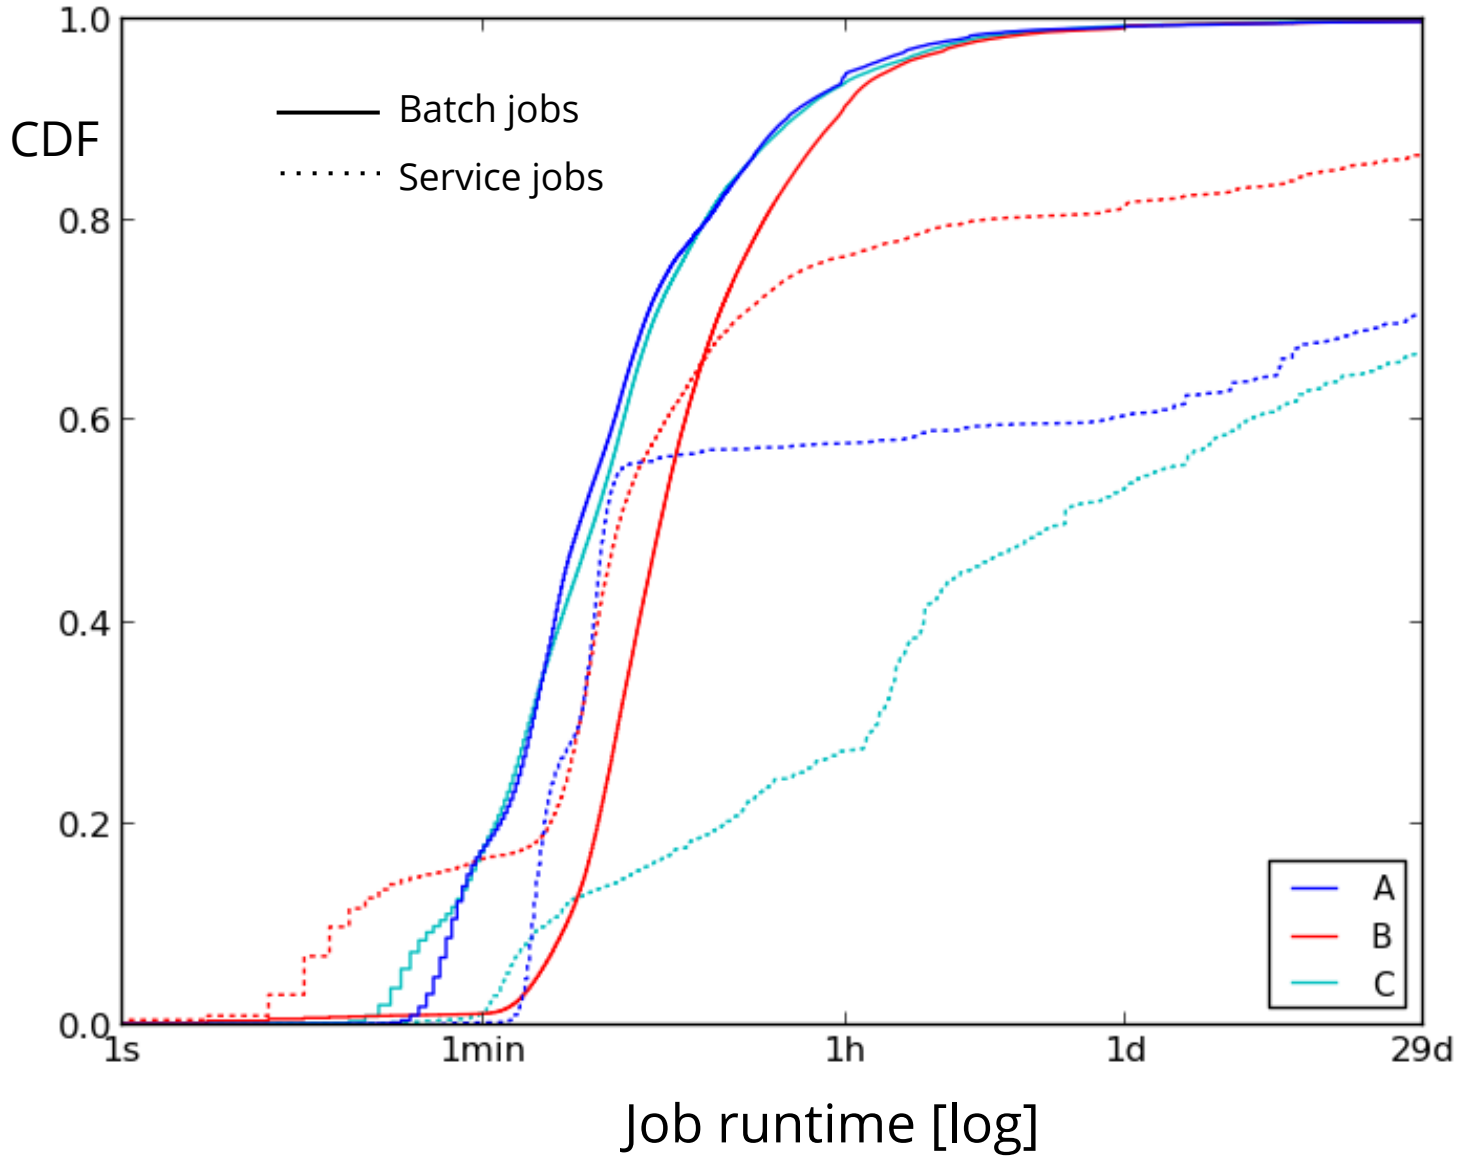

Efficiency

Utilization: sharing clusters between prod/batch helps

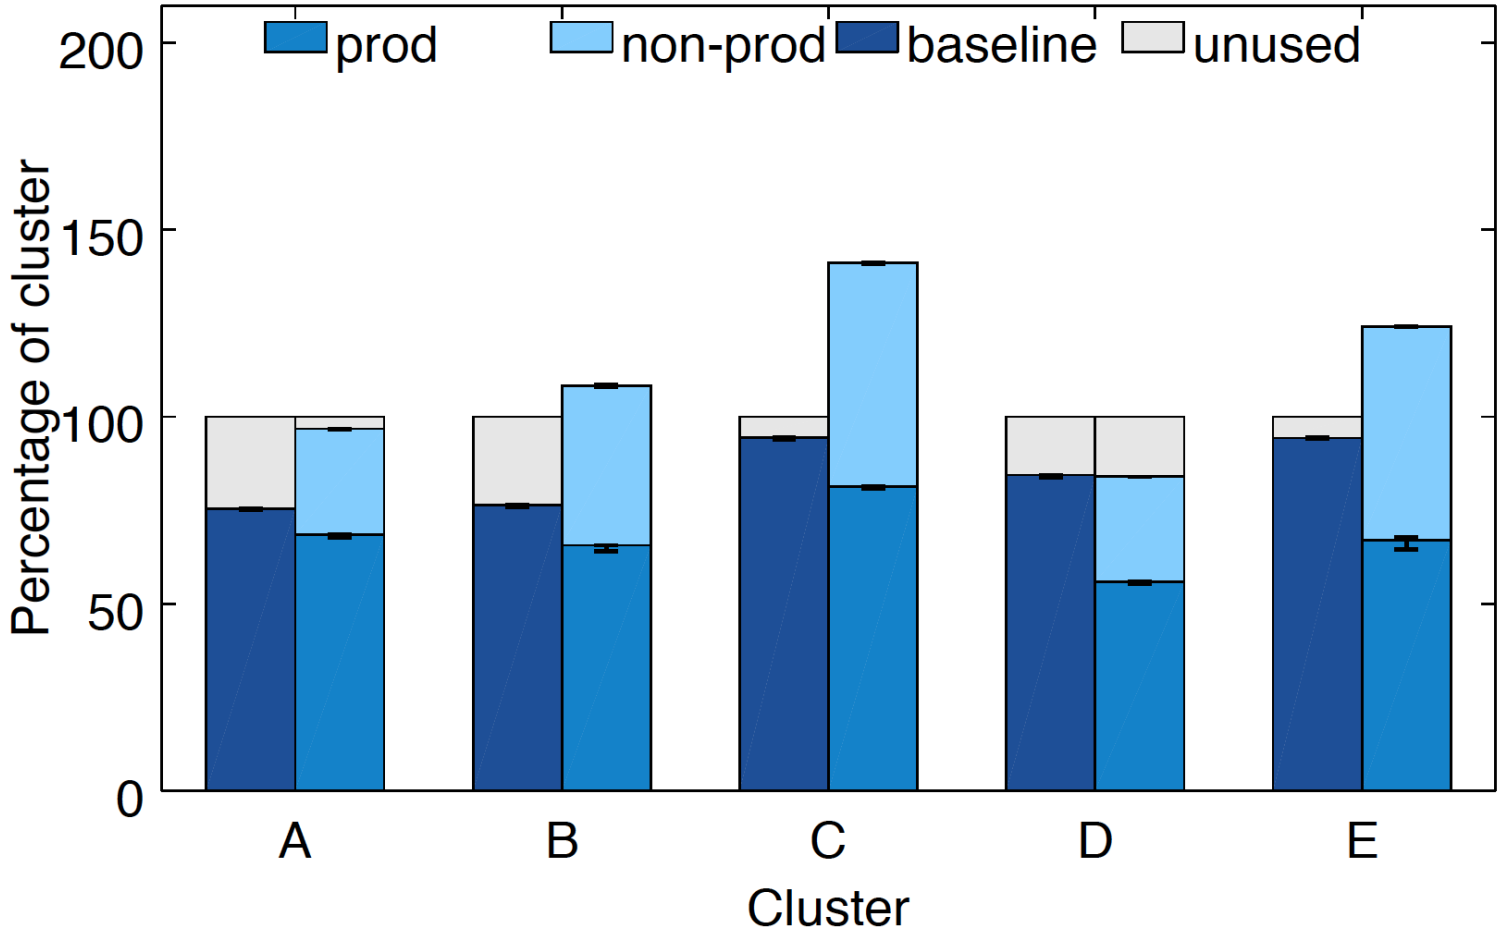

Efficiency

Utilization: sharing clusters between prod/batch helps

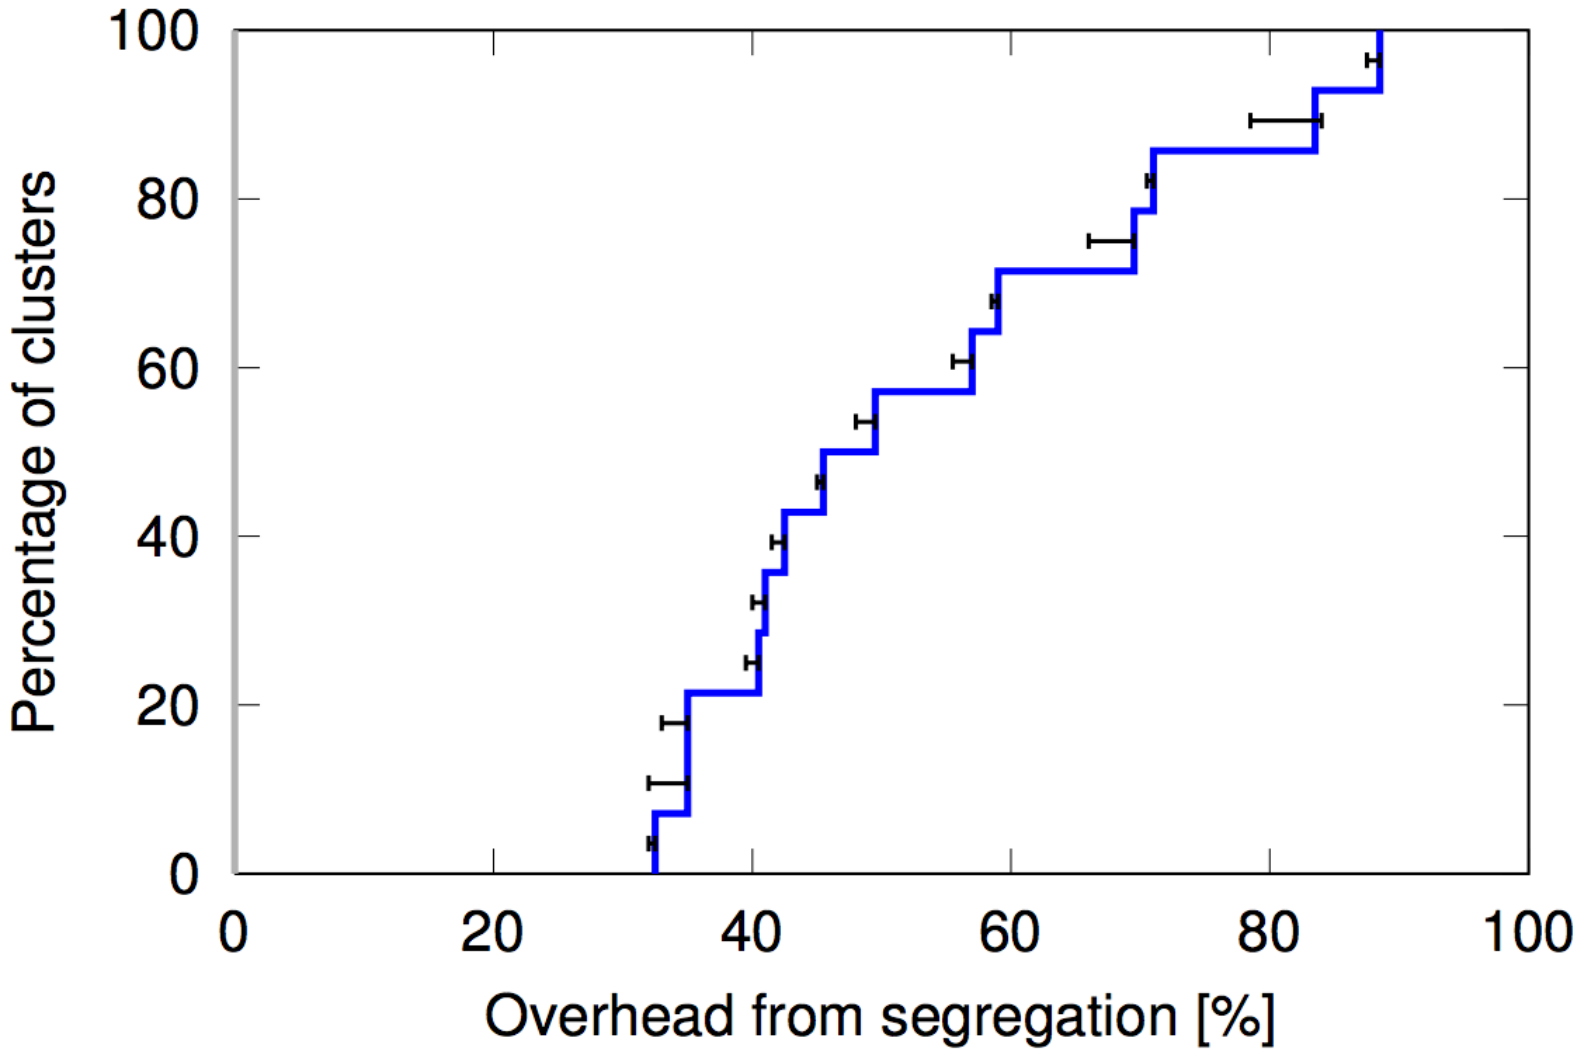

Efficiency

Advanced bin-packing algorithms

Data from a cluster with 12k machines, May 2011

Trace is publicly available

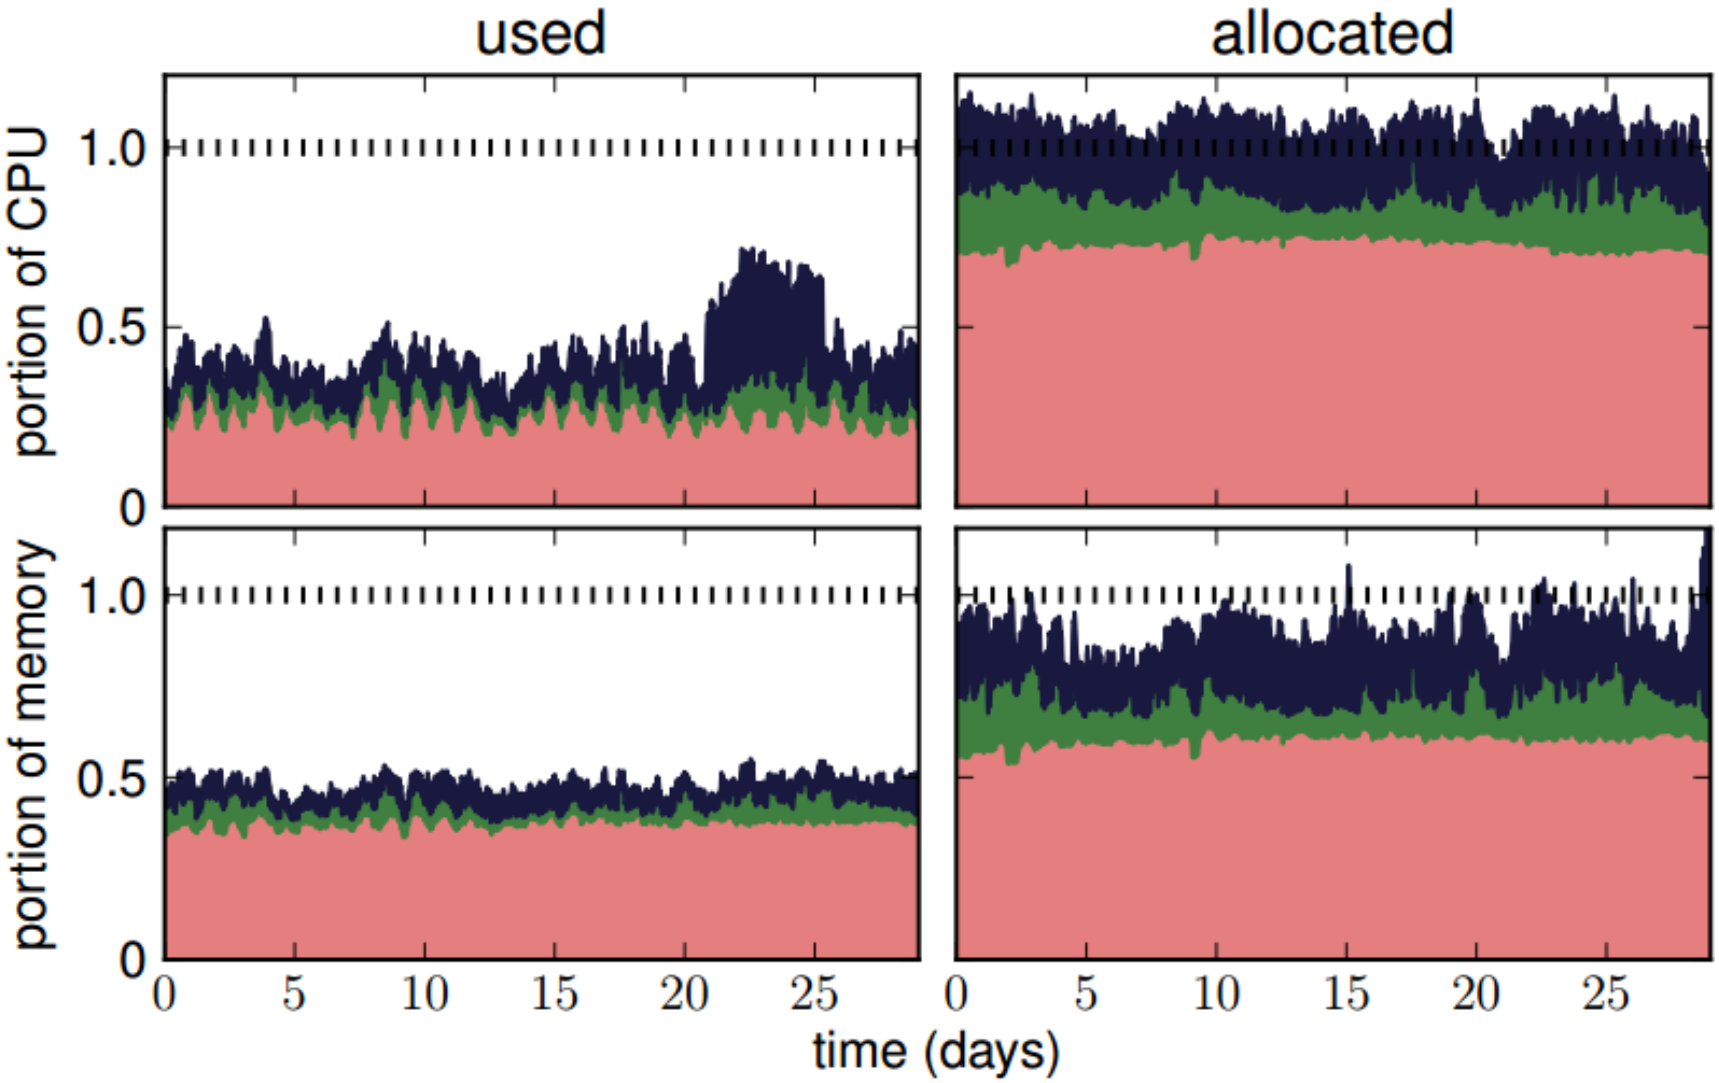

Heterogeneity and dynamicity of clouds at scale: Google trace analysis. SoCC'12

Efficiency

Resource reclamation could be more aggressive

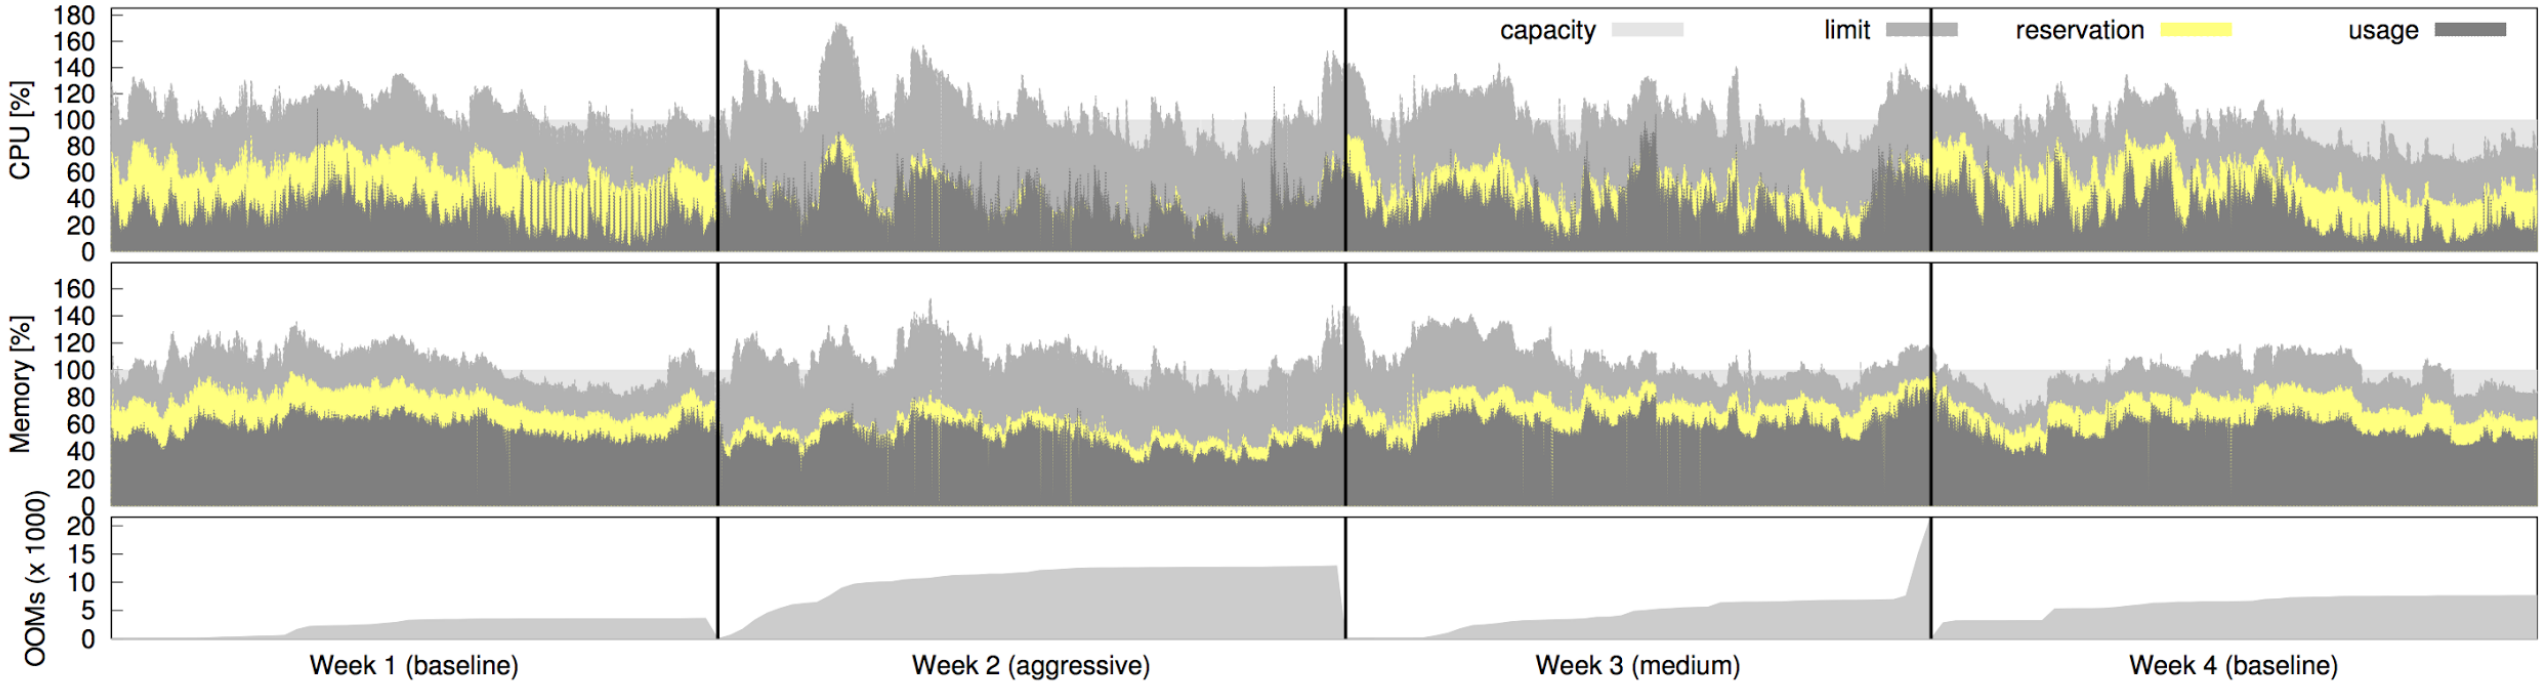

Nov/Dec 2013

Efficiency

Multiple applications per machine

CPI^2 paper, EuroSys 2013

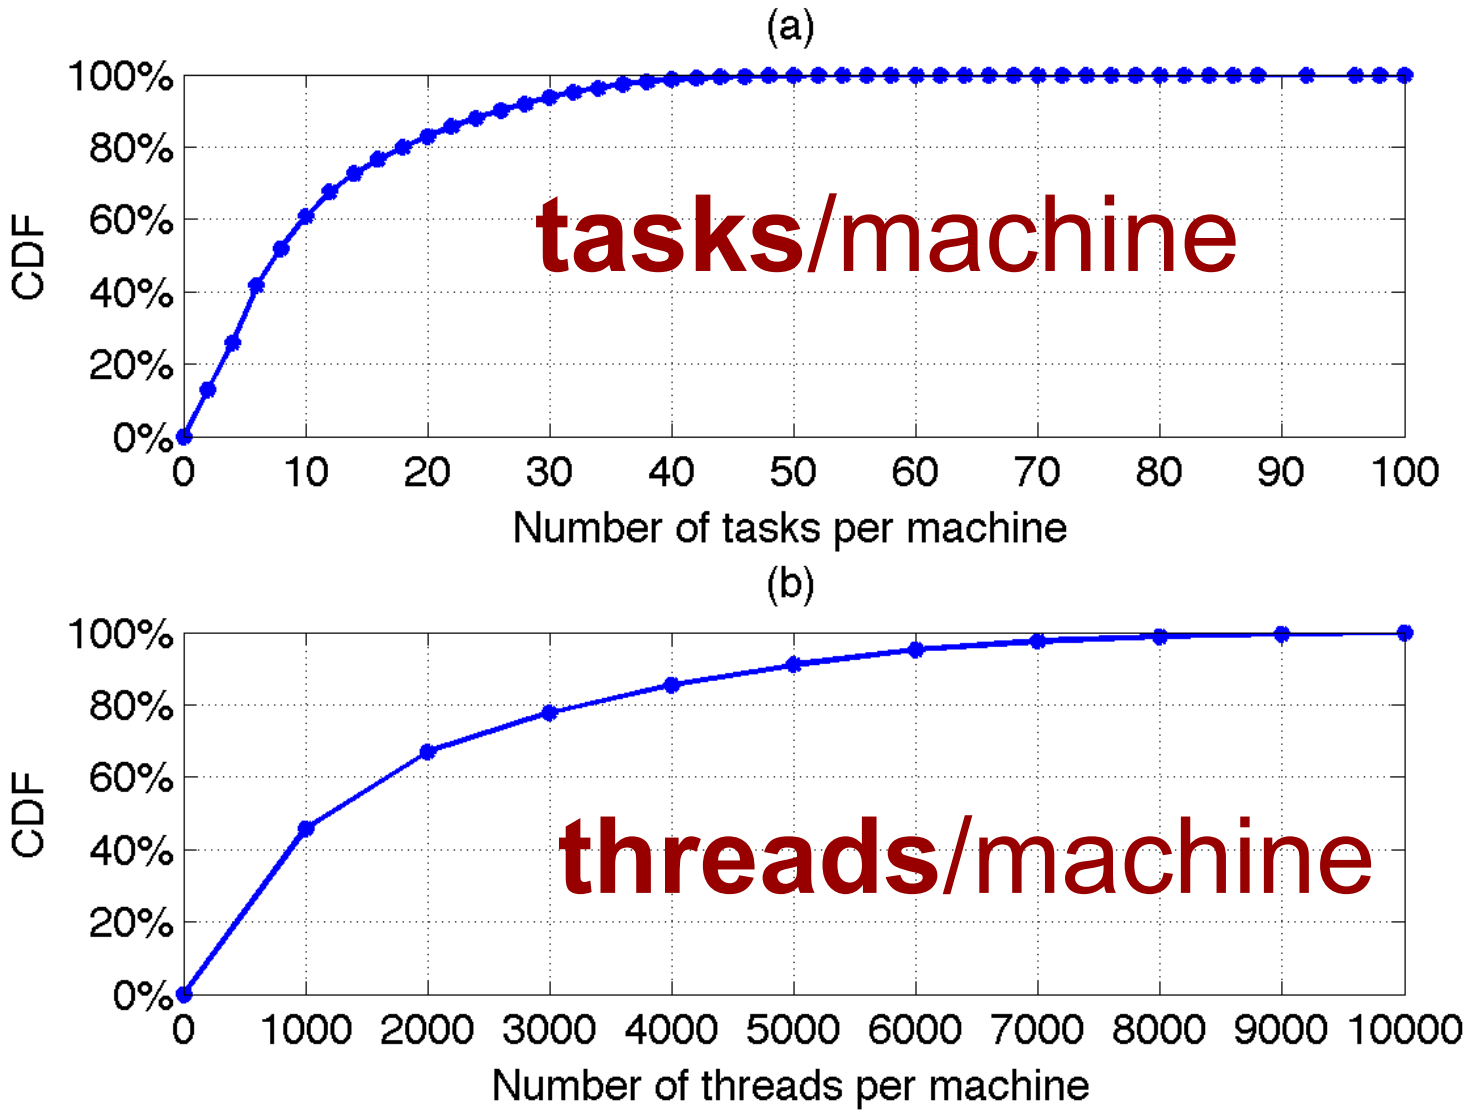

Efficiency

Multiple applications per machine

CPI^2 paper,
EuroSys 2013

1. Gather CPI for all the tasks in a job
2. Find outliers
3. Take action

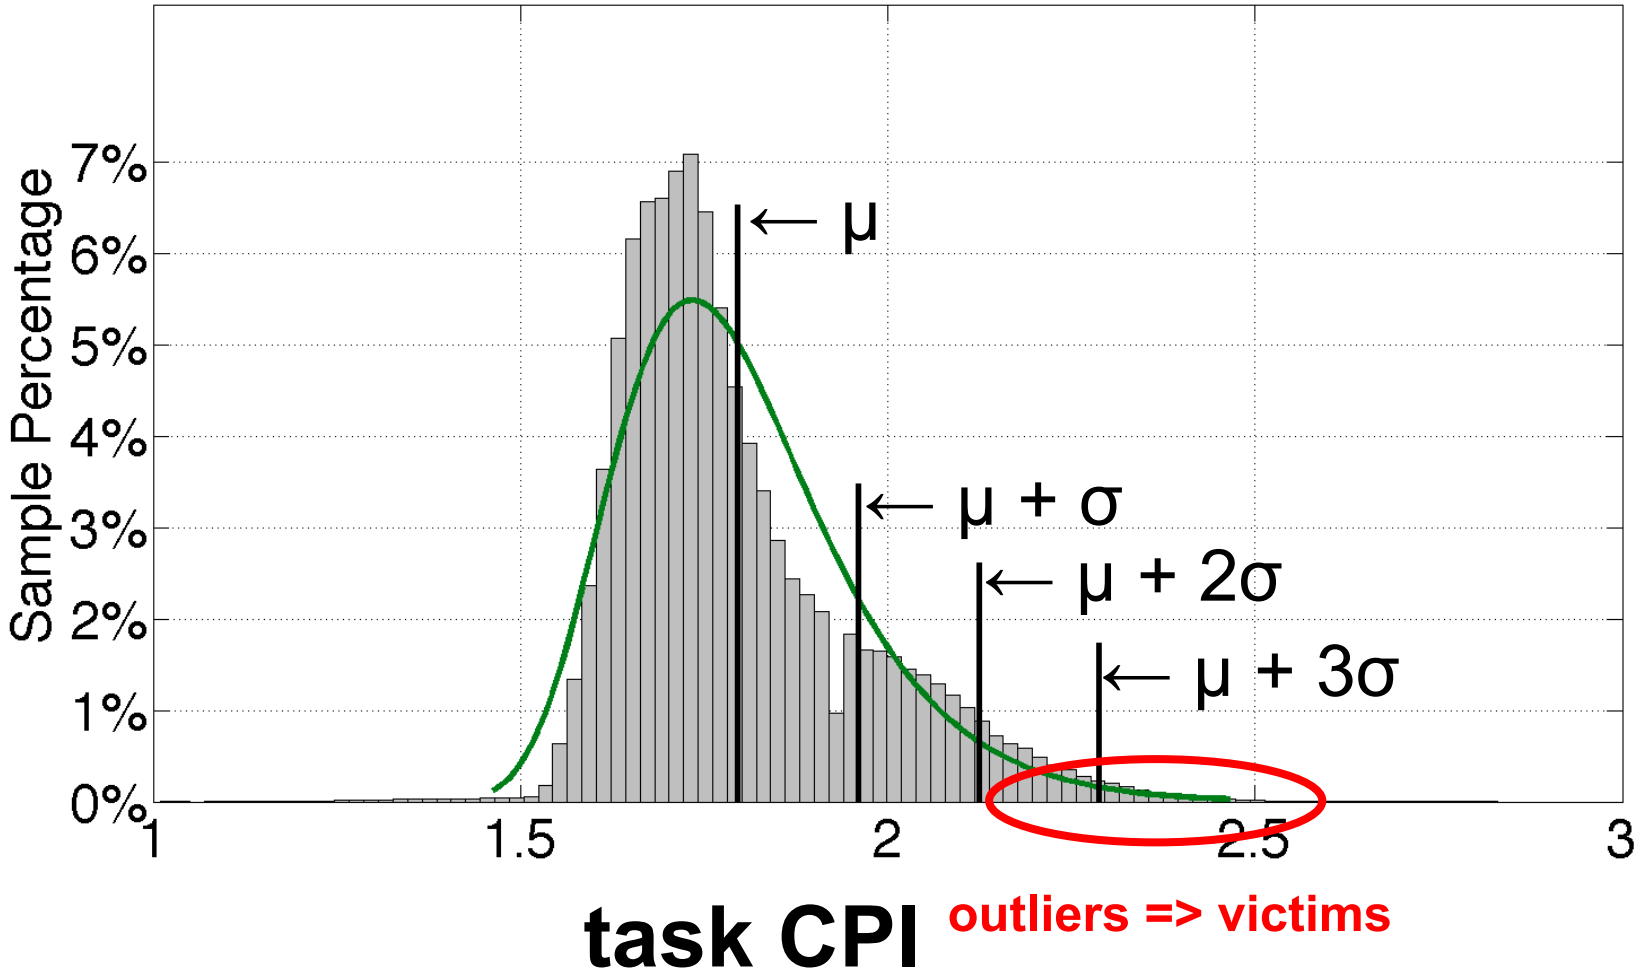

Achieving desired behavior

Exposing mechanisms is fragile

Better: **declarative intents**

Achieving desired behavior

Service level objective (SLO)

Examples:

- availability
- obtainability
- reliability
- velocity

- freshness?
- accuracy?
- security?

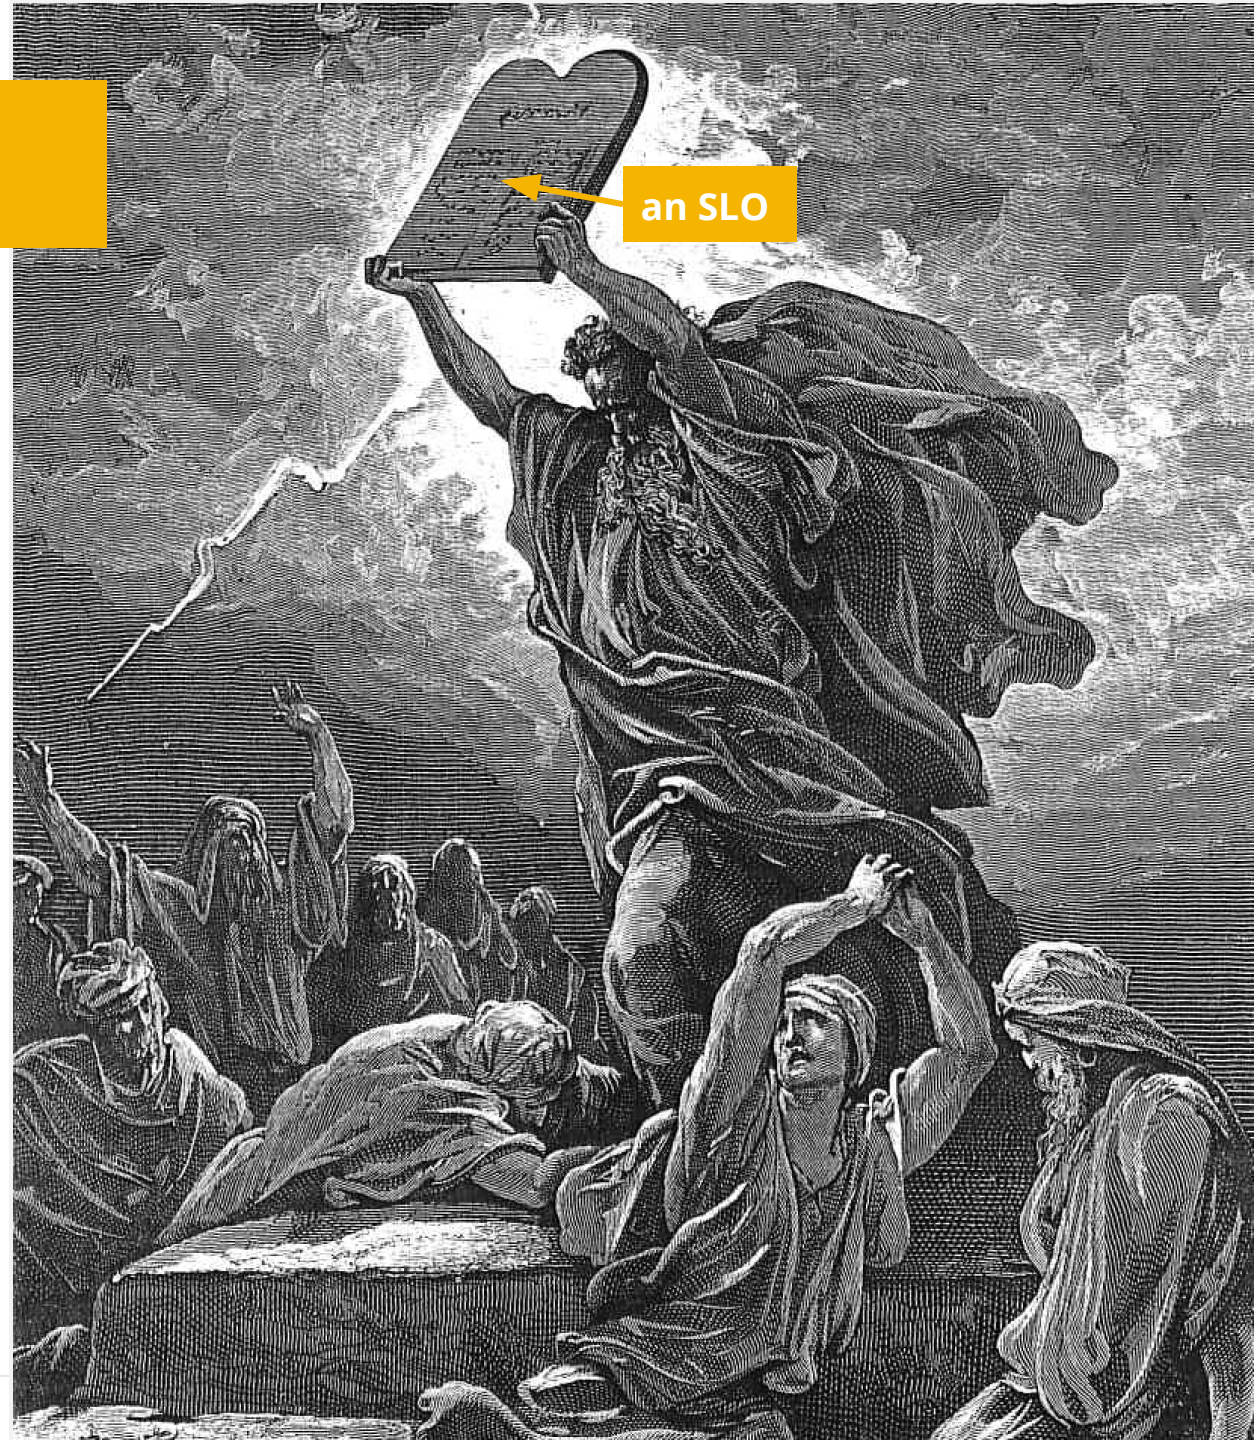

A few other moving parts

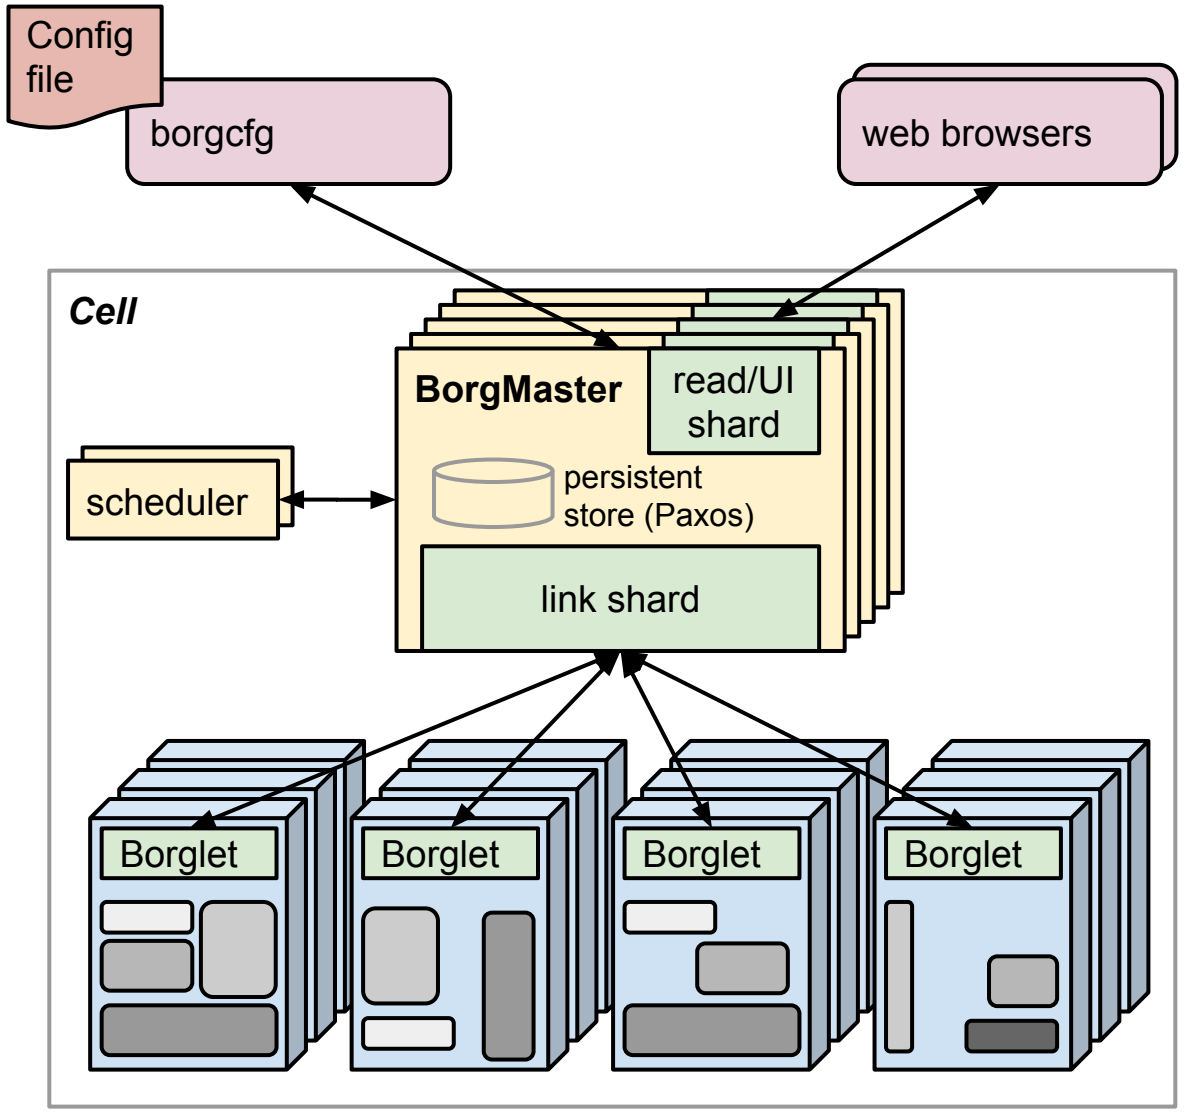

A few other moving parts

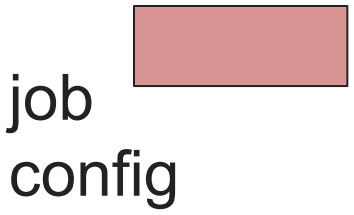

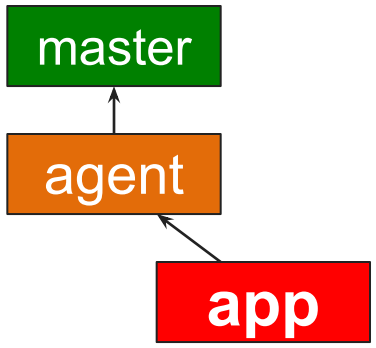

A few other moving parts

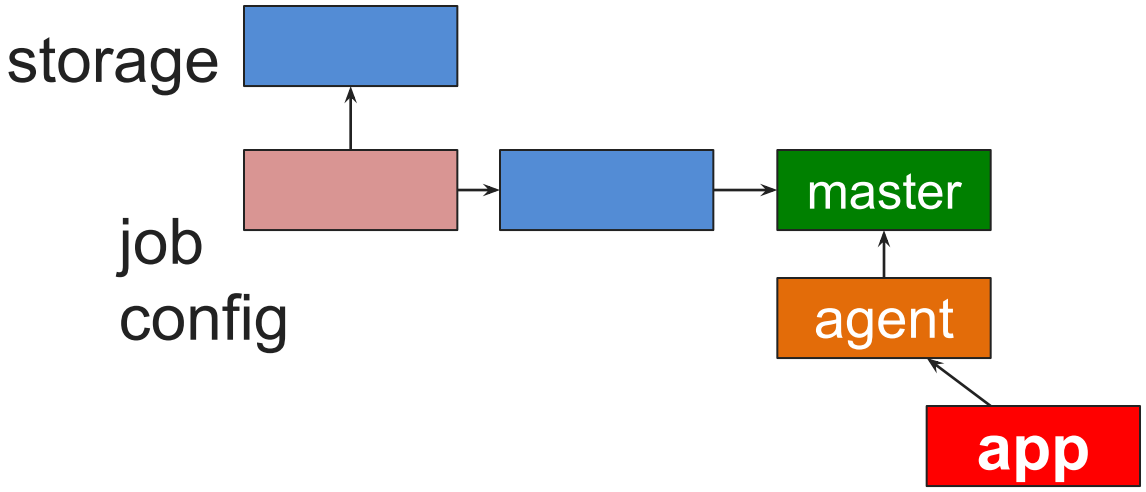

A few other moving parts

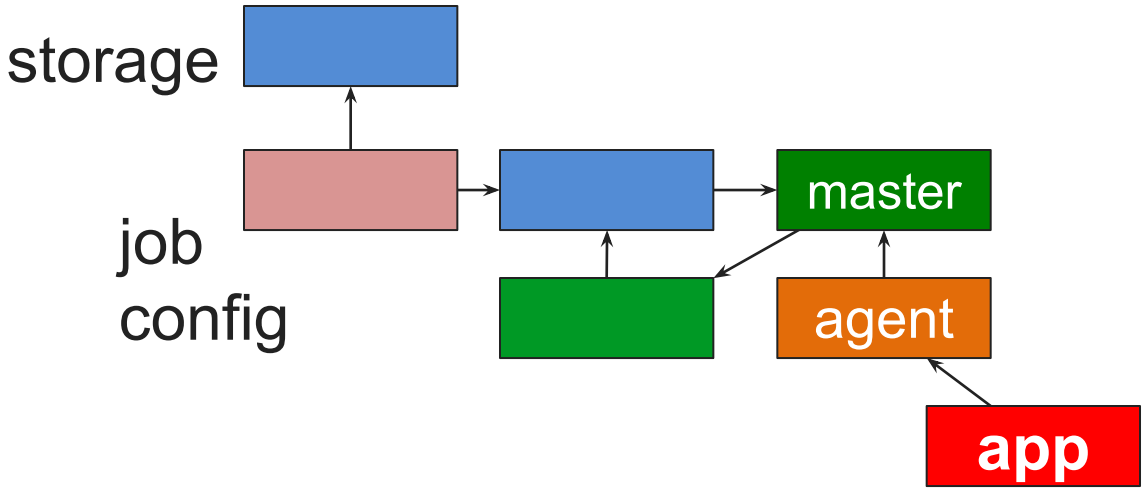

A few other moving parts

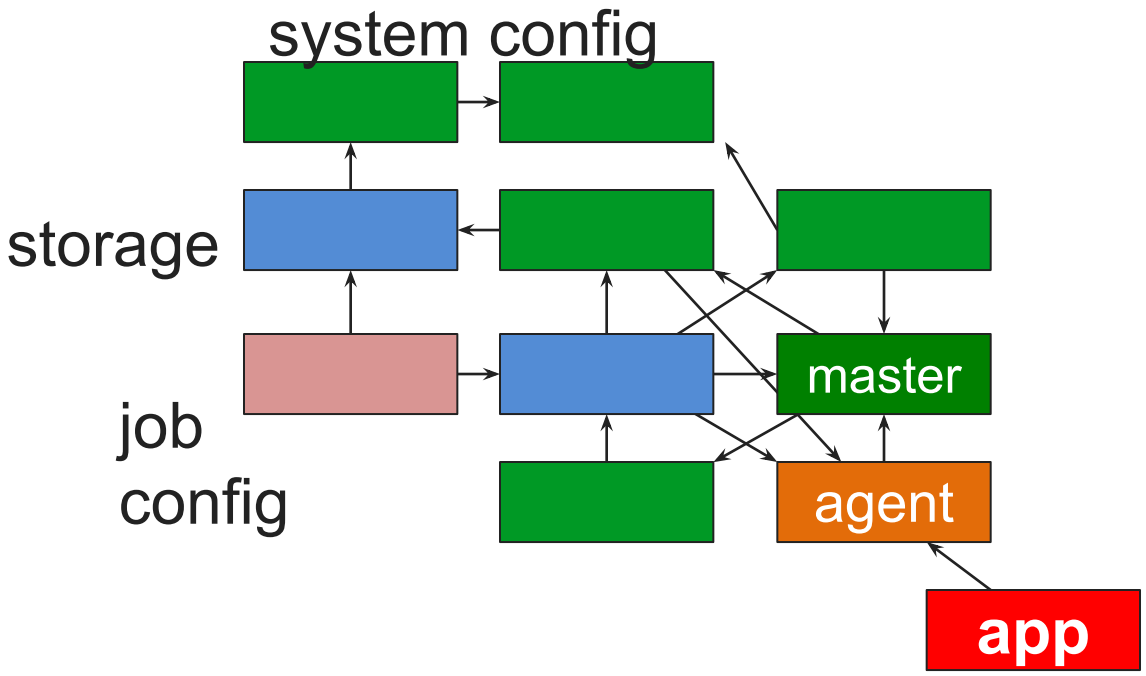

A few other moving parts

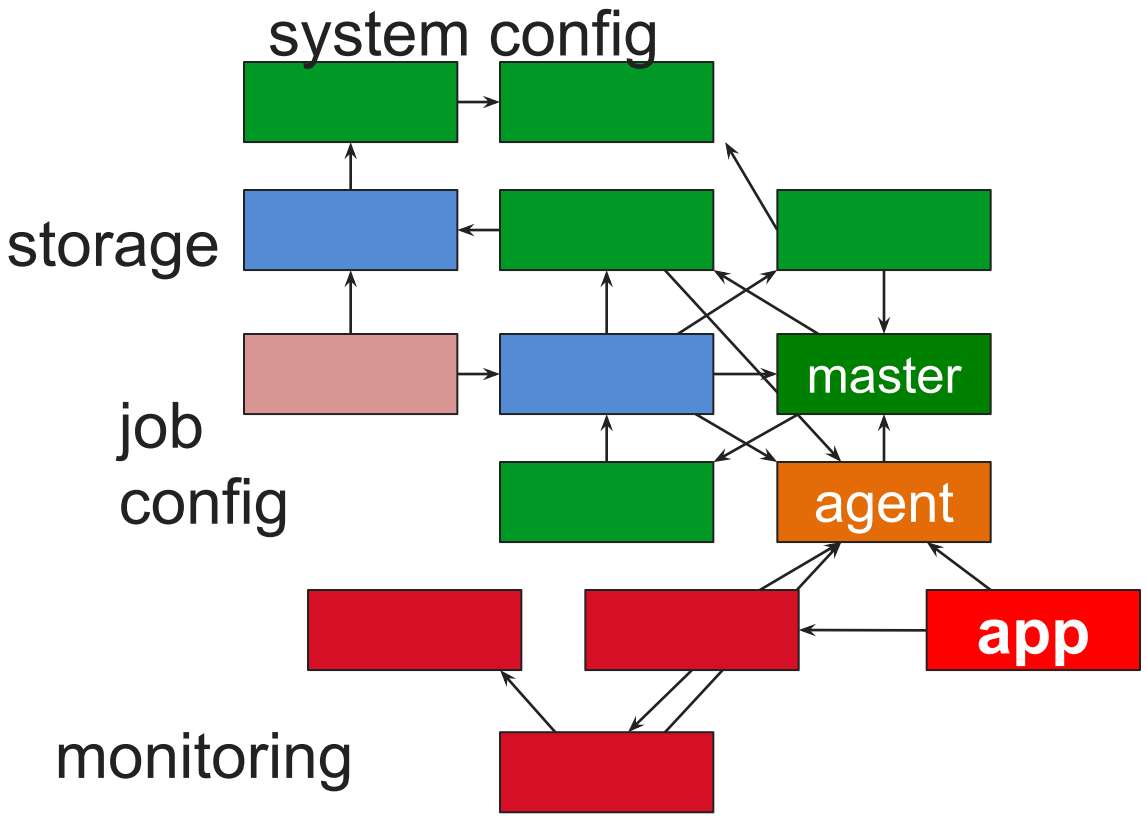

A few other moving parts

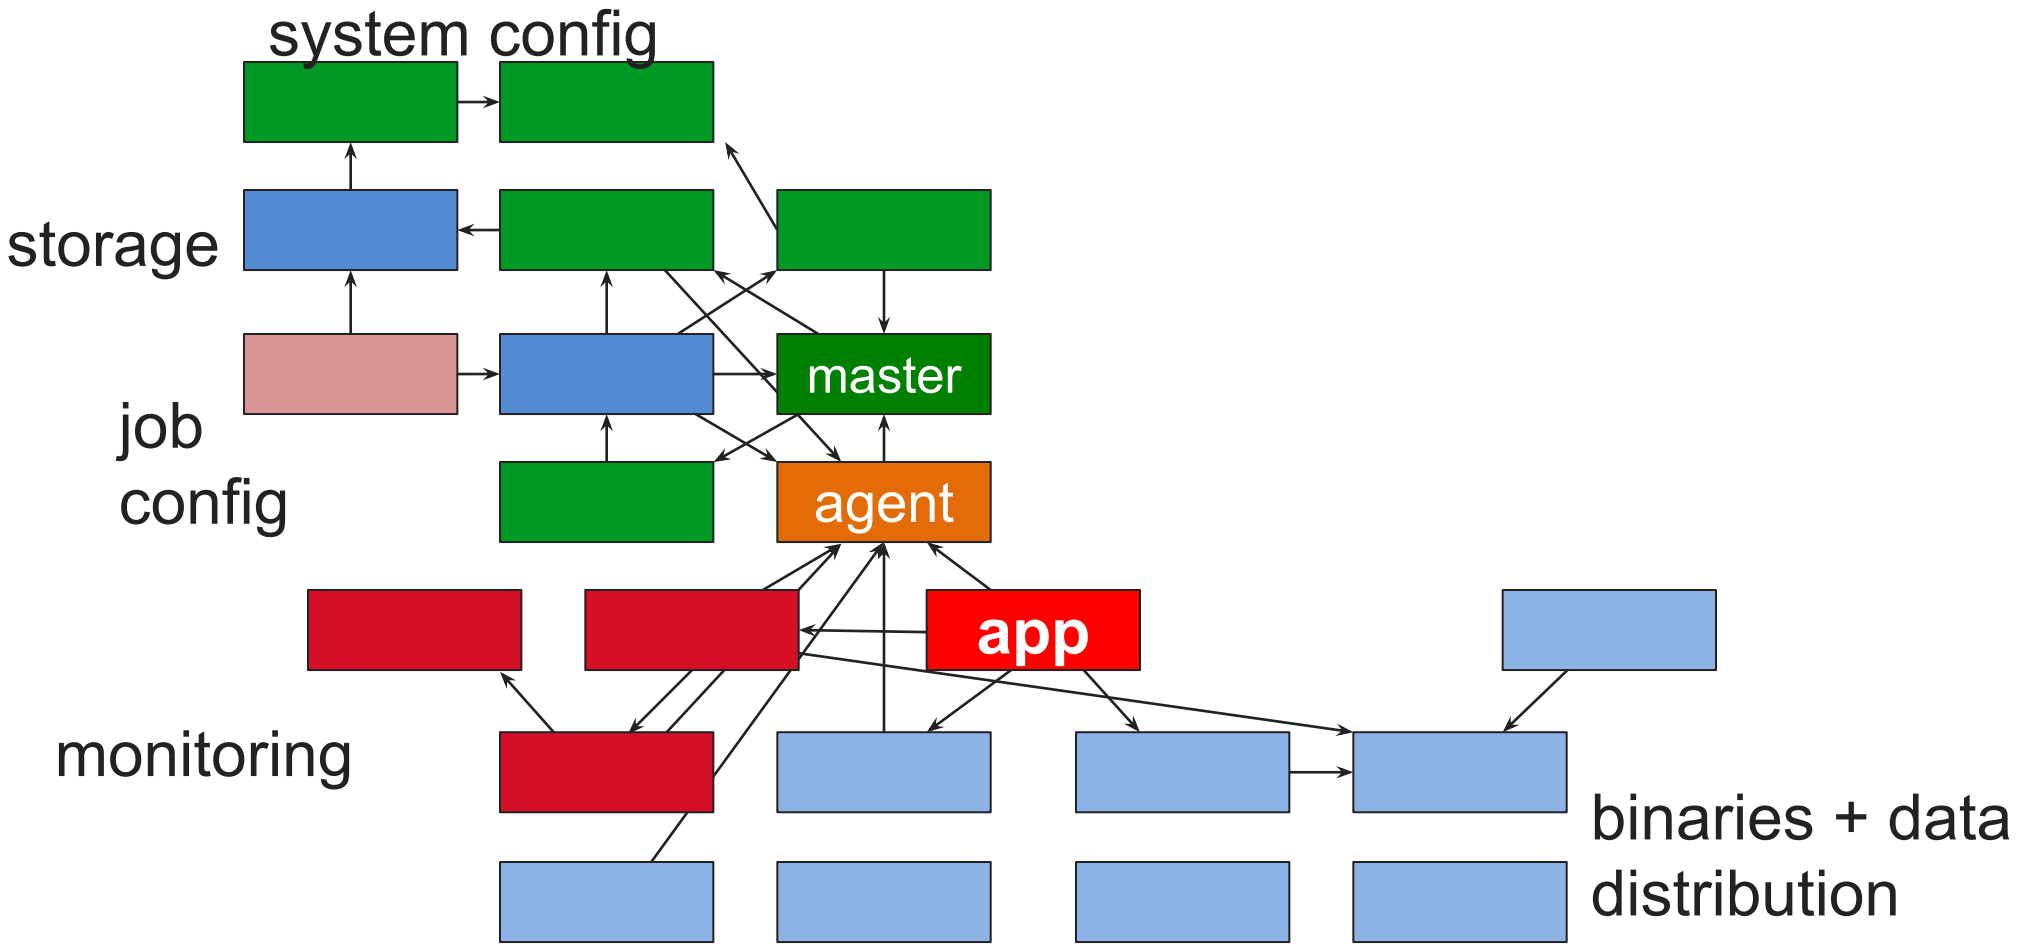

A few other moving parts

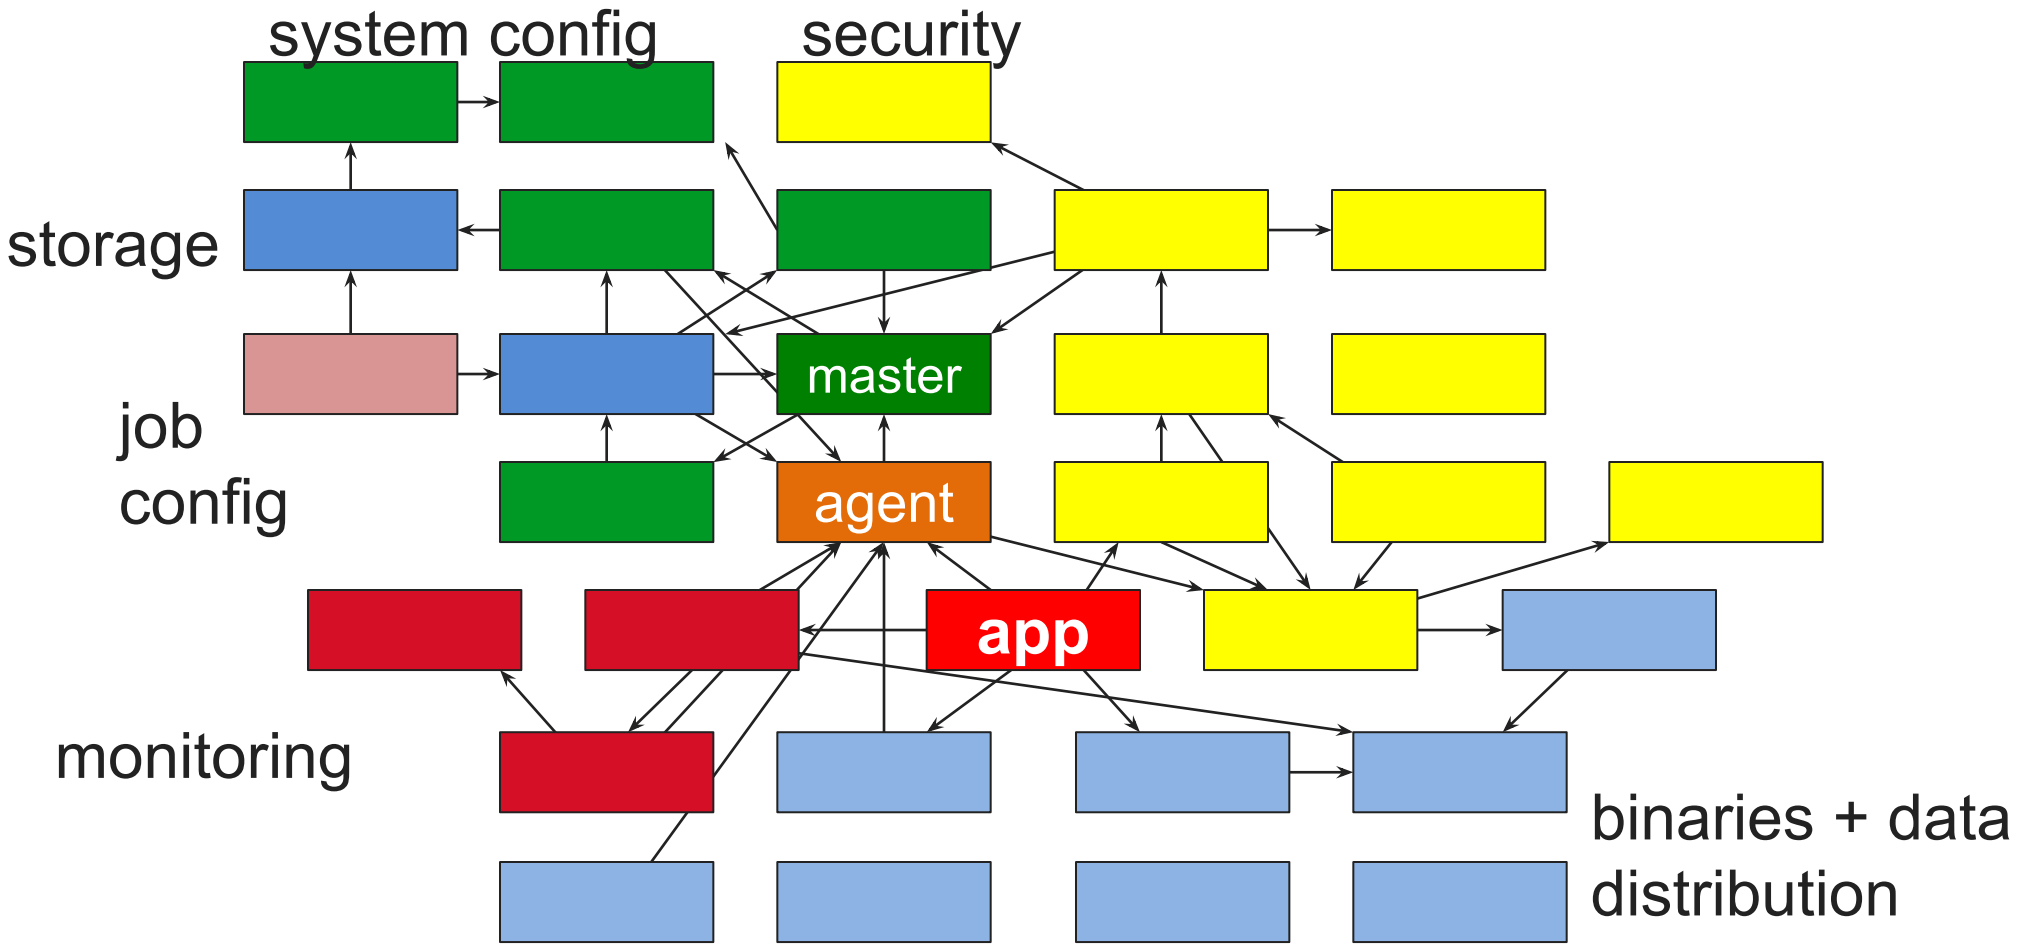

A few other moving parts

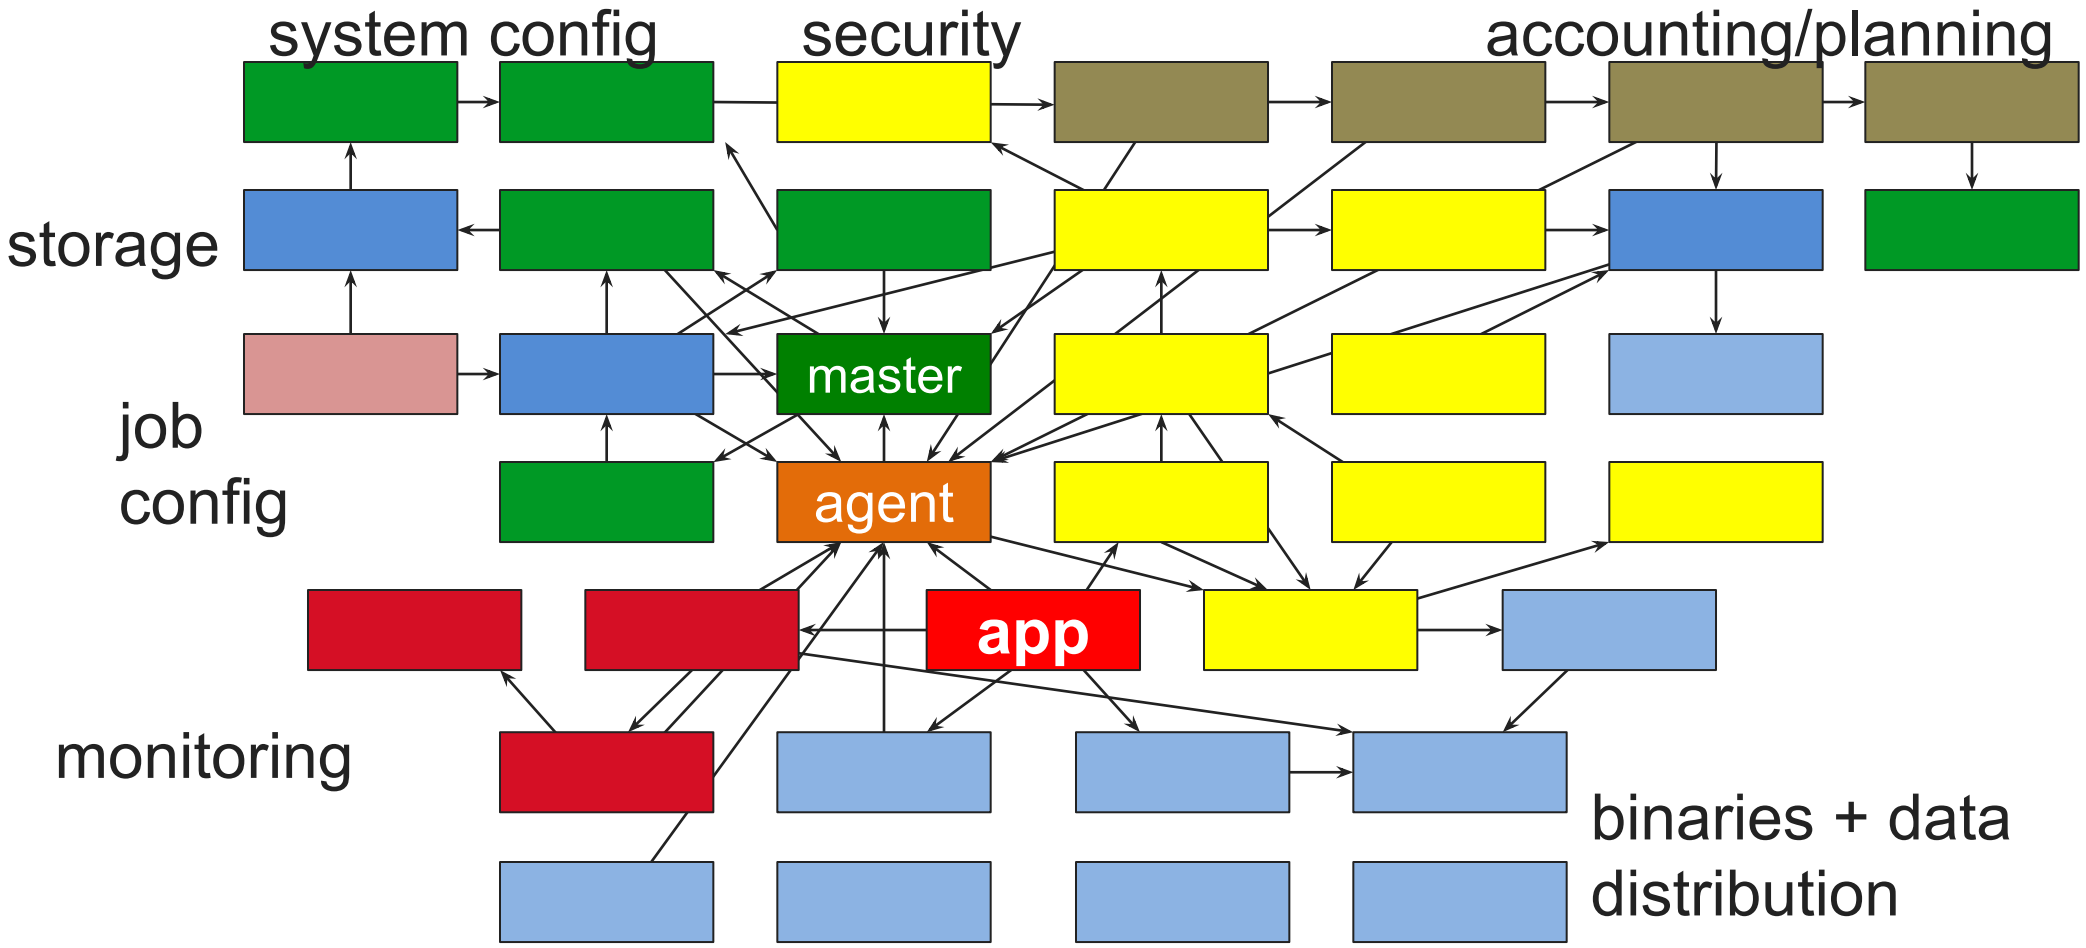

Diagram from an original by Cody Smith.

A few other moving parts

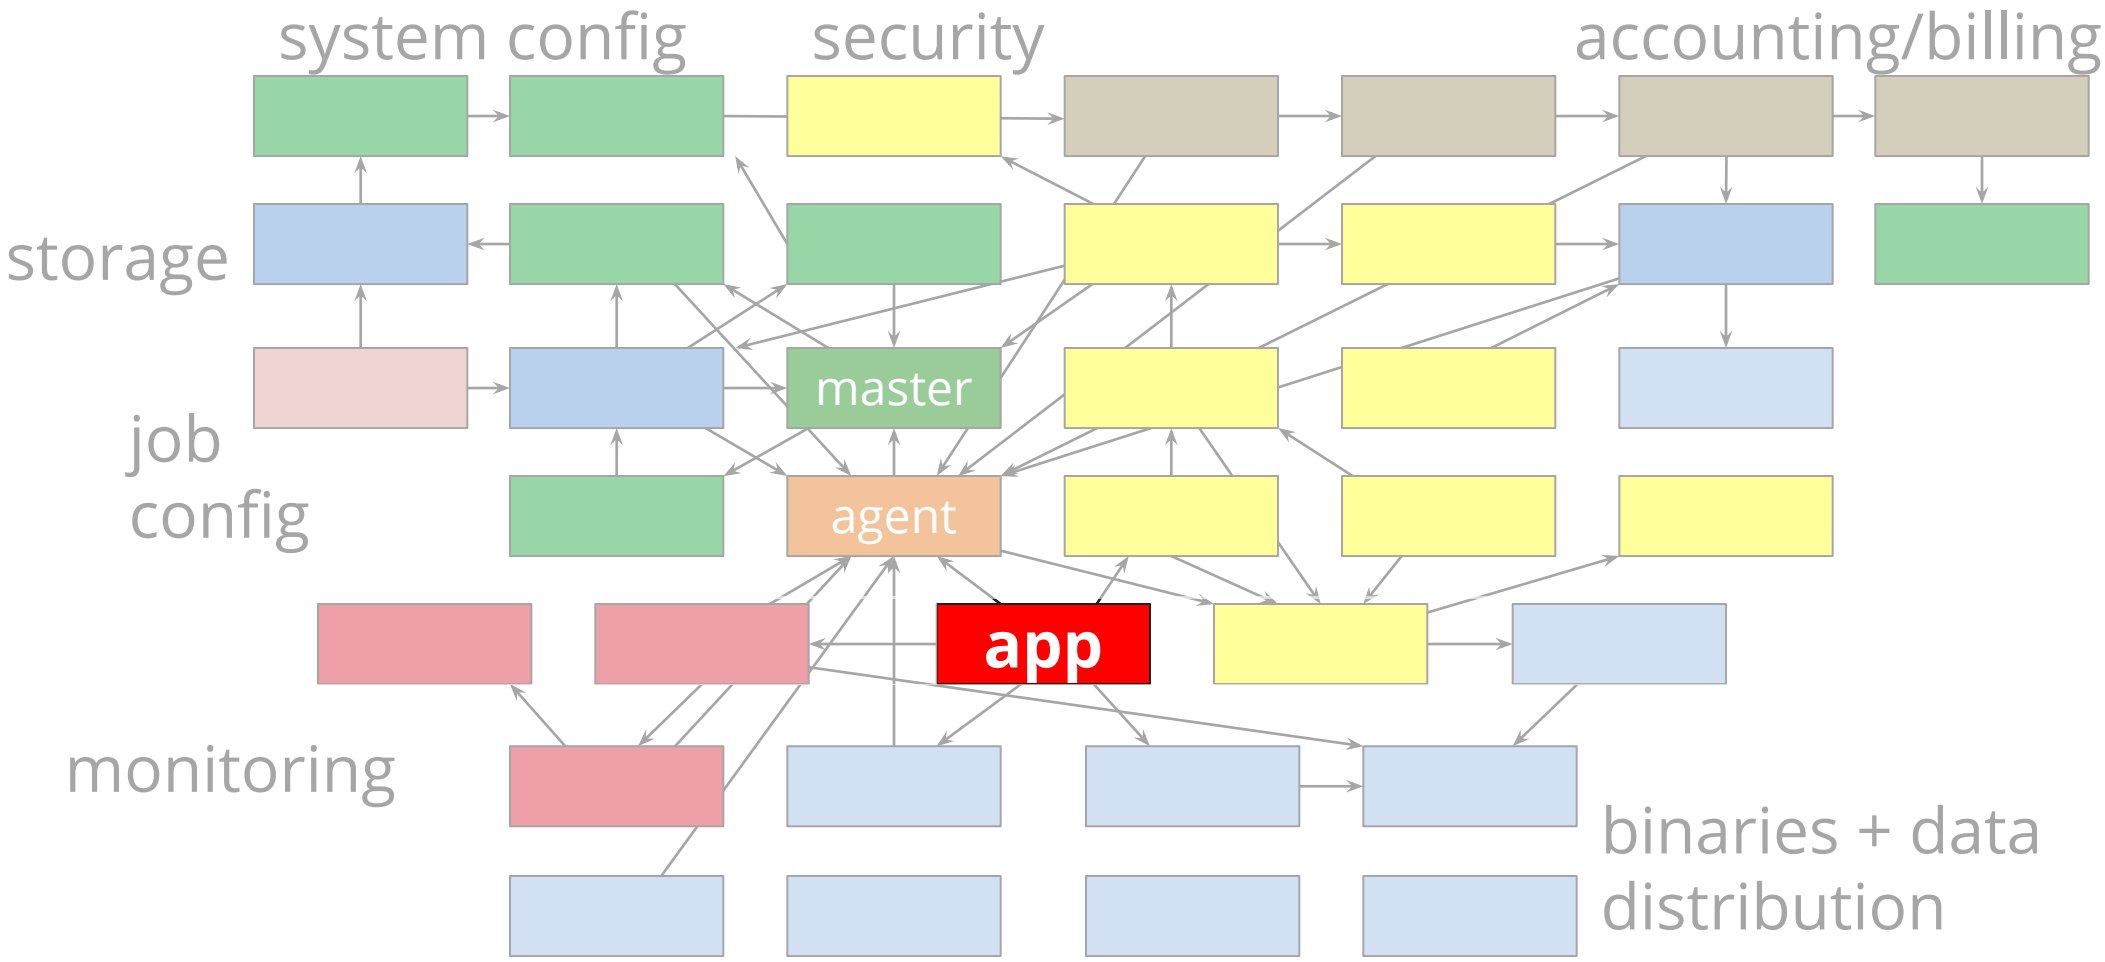

Diagram from an original by Cody Smith.

Containers

Everything at Google runs in a container -- including our VMs

Containers give us:

- resource isolation
- execution isolation
- CPU QoS

We start over 2 billion containers per week.

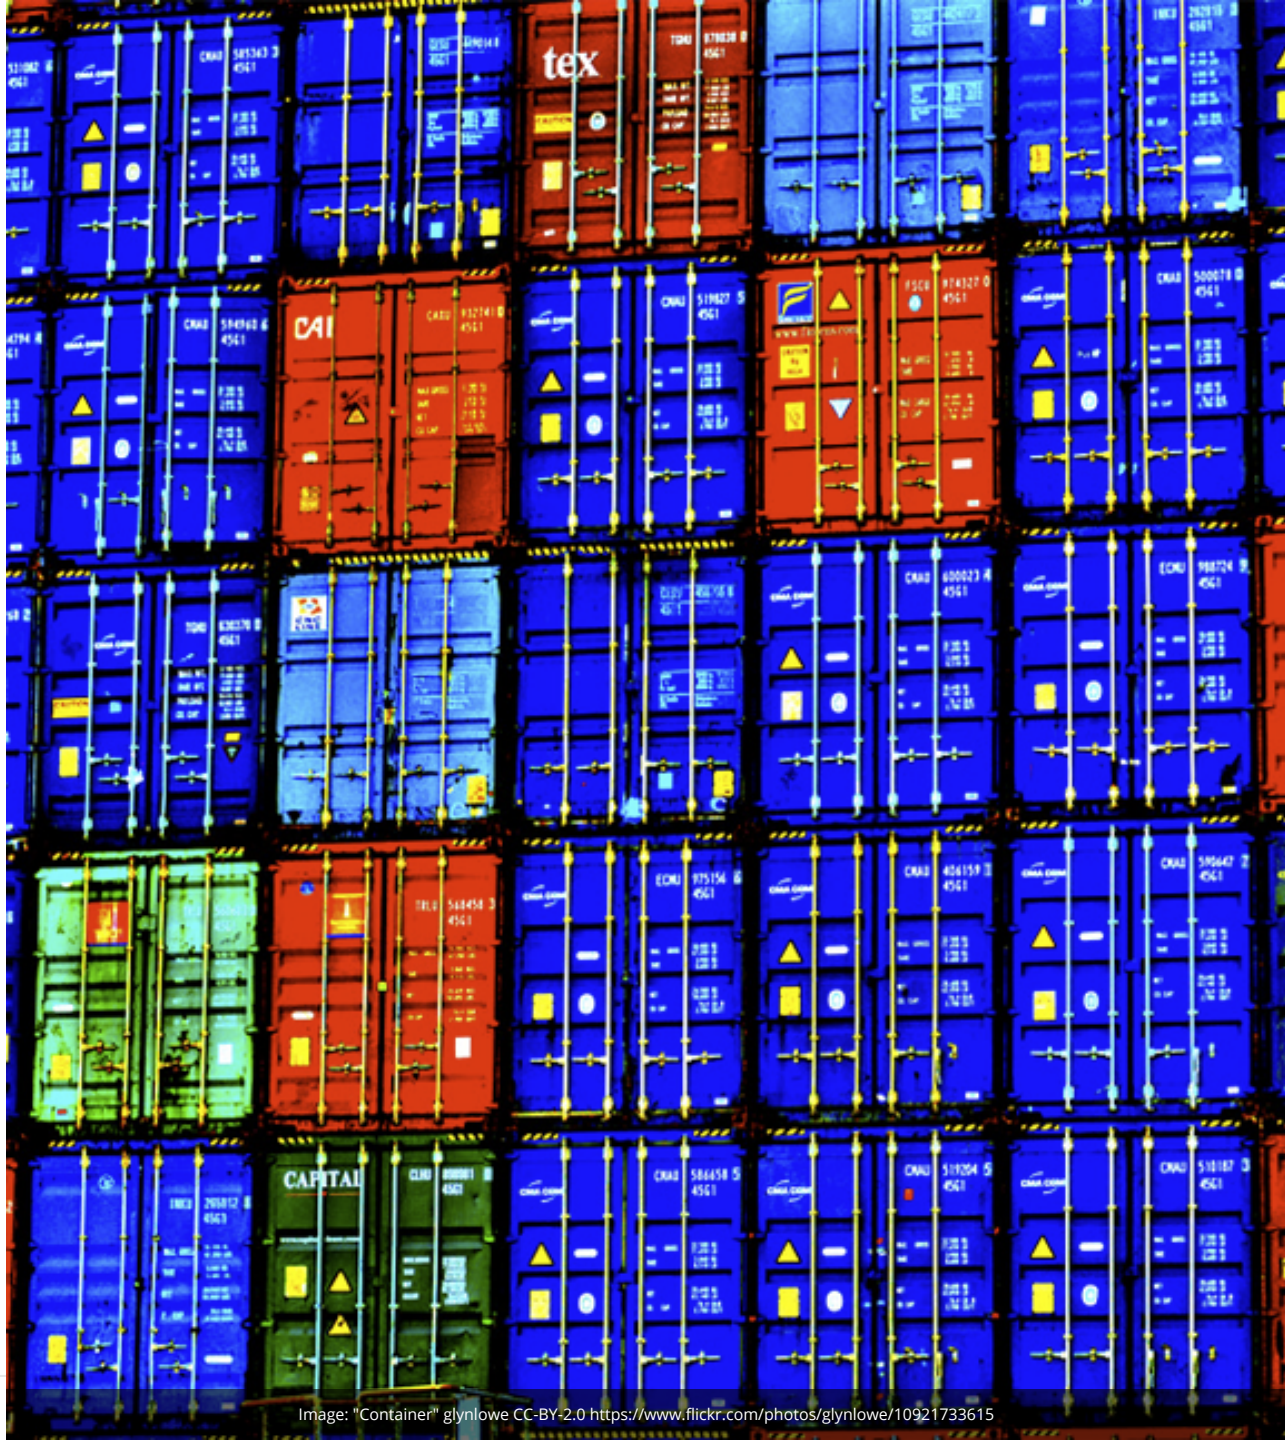

Image: "Container" glynlowe CC-BY-2.0 https://www.flickr.com/photos/glynlowe/10921733615

kubernetes

Home | Getting Started | Community | News | Events | Blog

Kubernetes

κυβερνήτης:
Greek for "pilot" or "helmsman of a ship"

The open source cluster manager from Google.

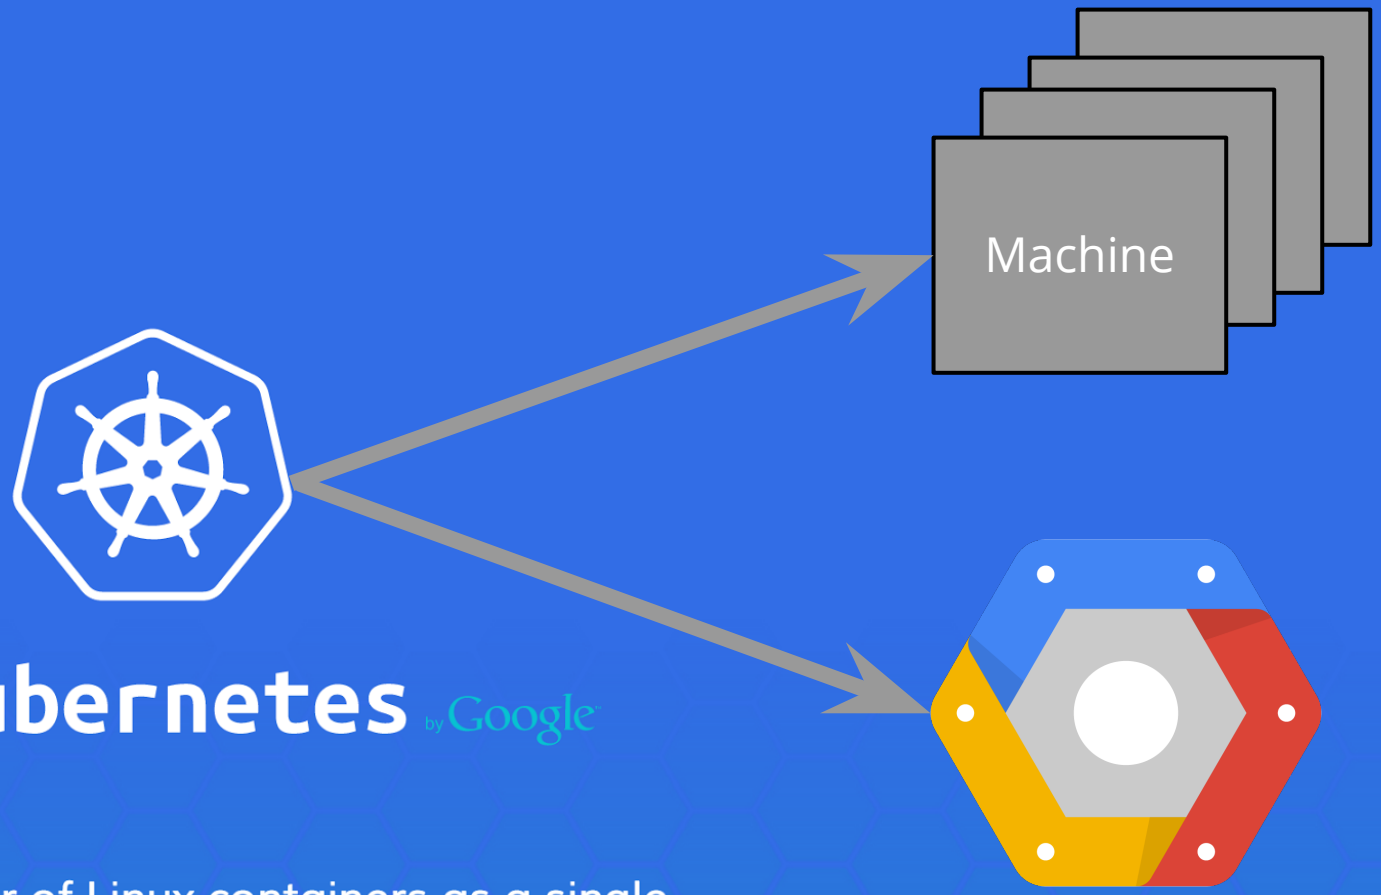

Manage a cluster of Linux containers as a single system to accelerate Dev and simplify Ops.

View on GitHub | Try Kubernetes

Kubernetes

Web server

Log roller

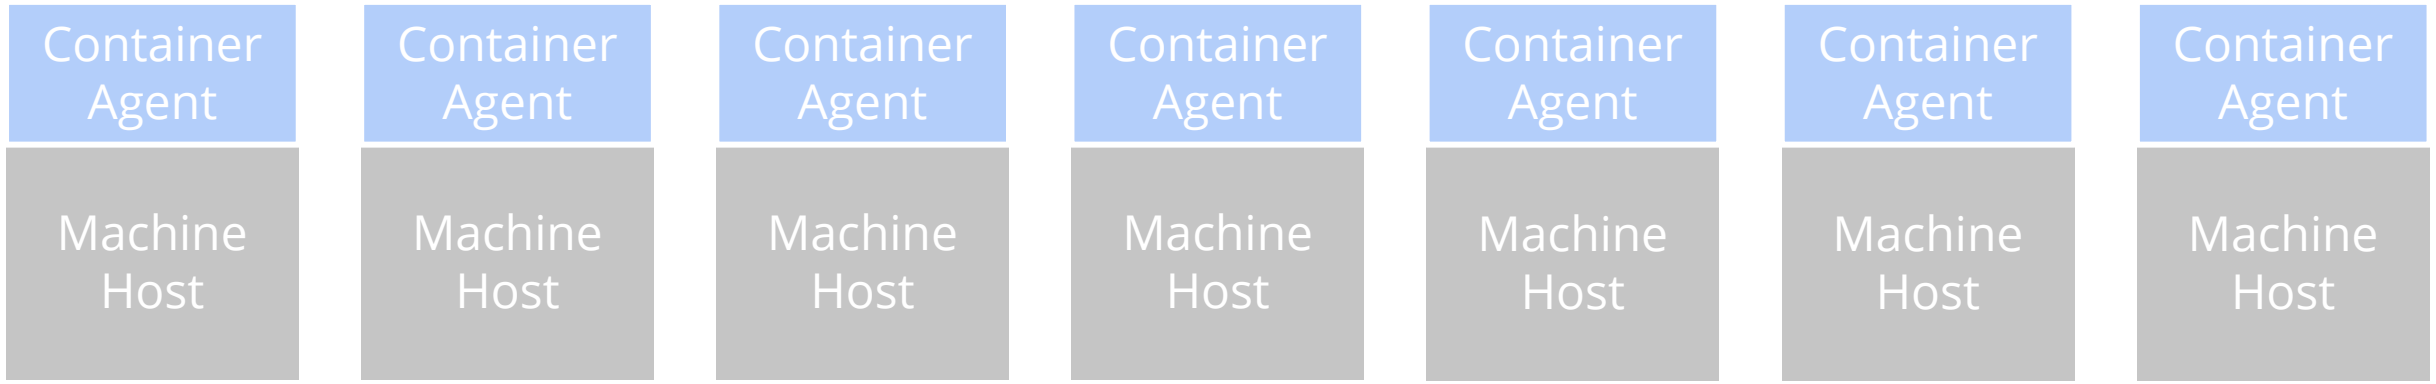

Pods

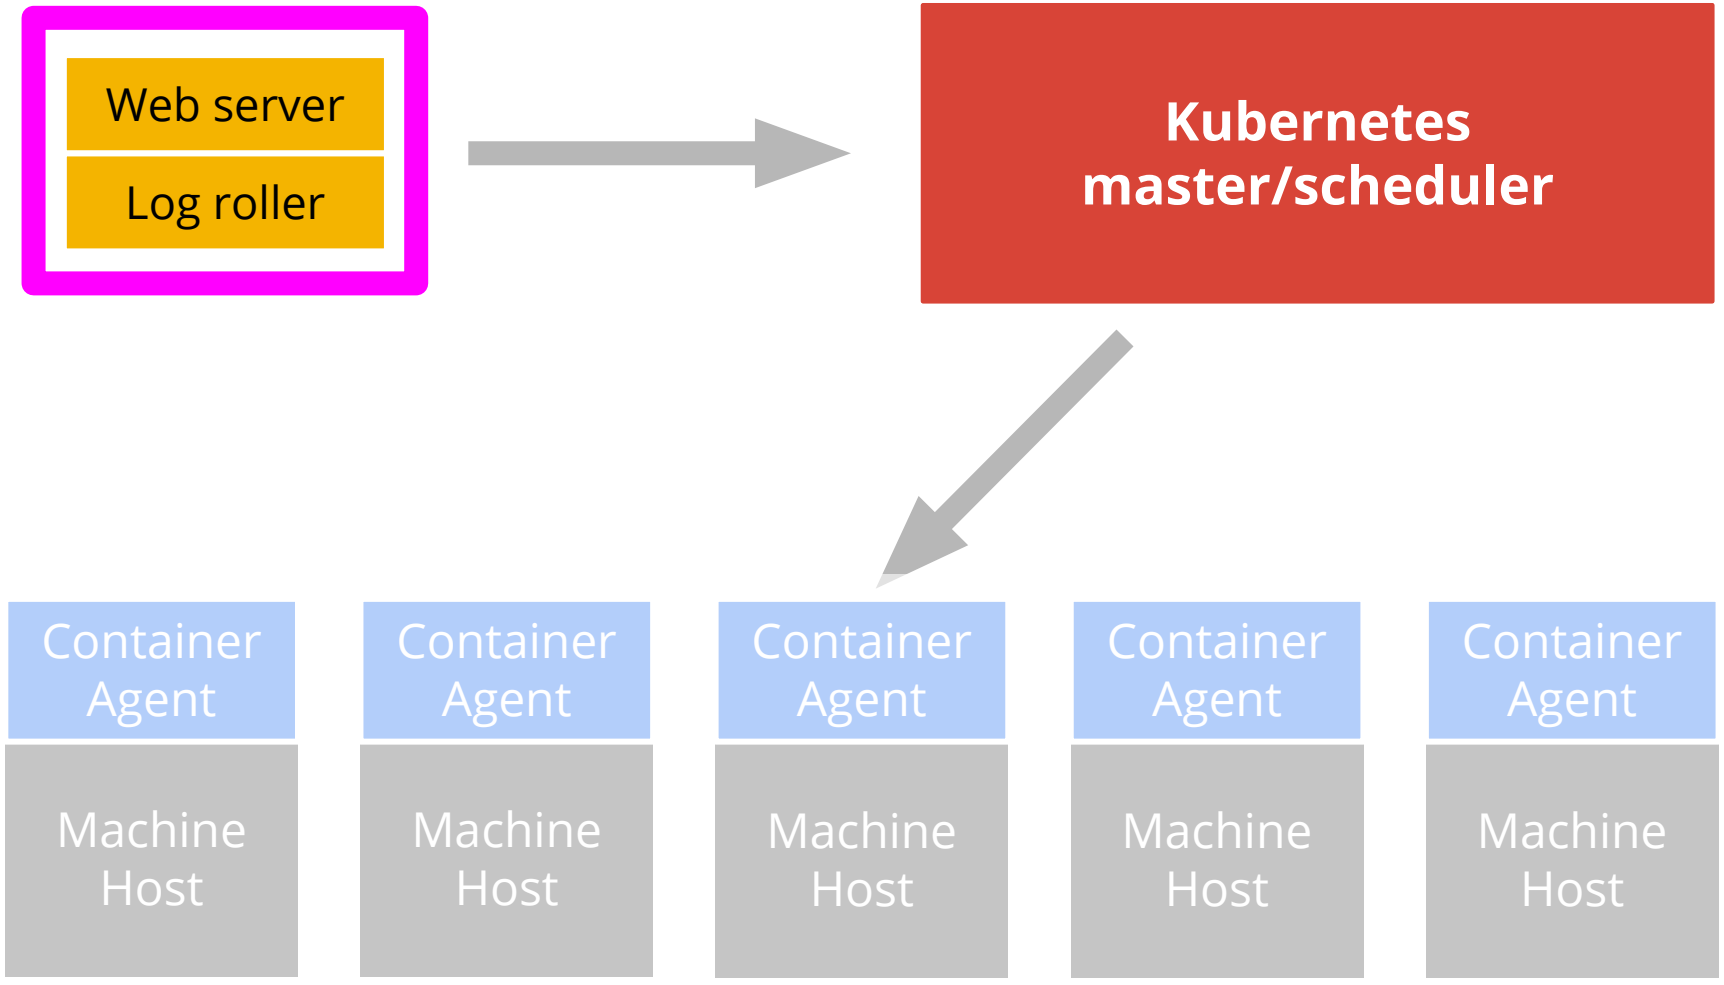
Web server
Log roller
Kubernetes master/scheduler
Container Agent
Machine Host
Container Agent
Machine Host
Container Agent
Machine Host
Container Agent
Machine Host
Container Agent
Machine Host
Container Agent
Machine Host
Container Agent
Machine Host

Labels

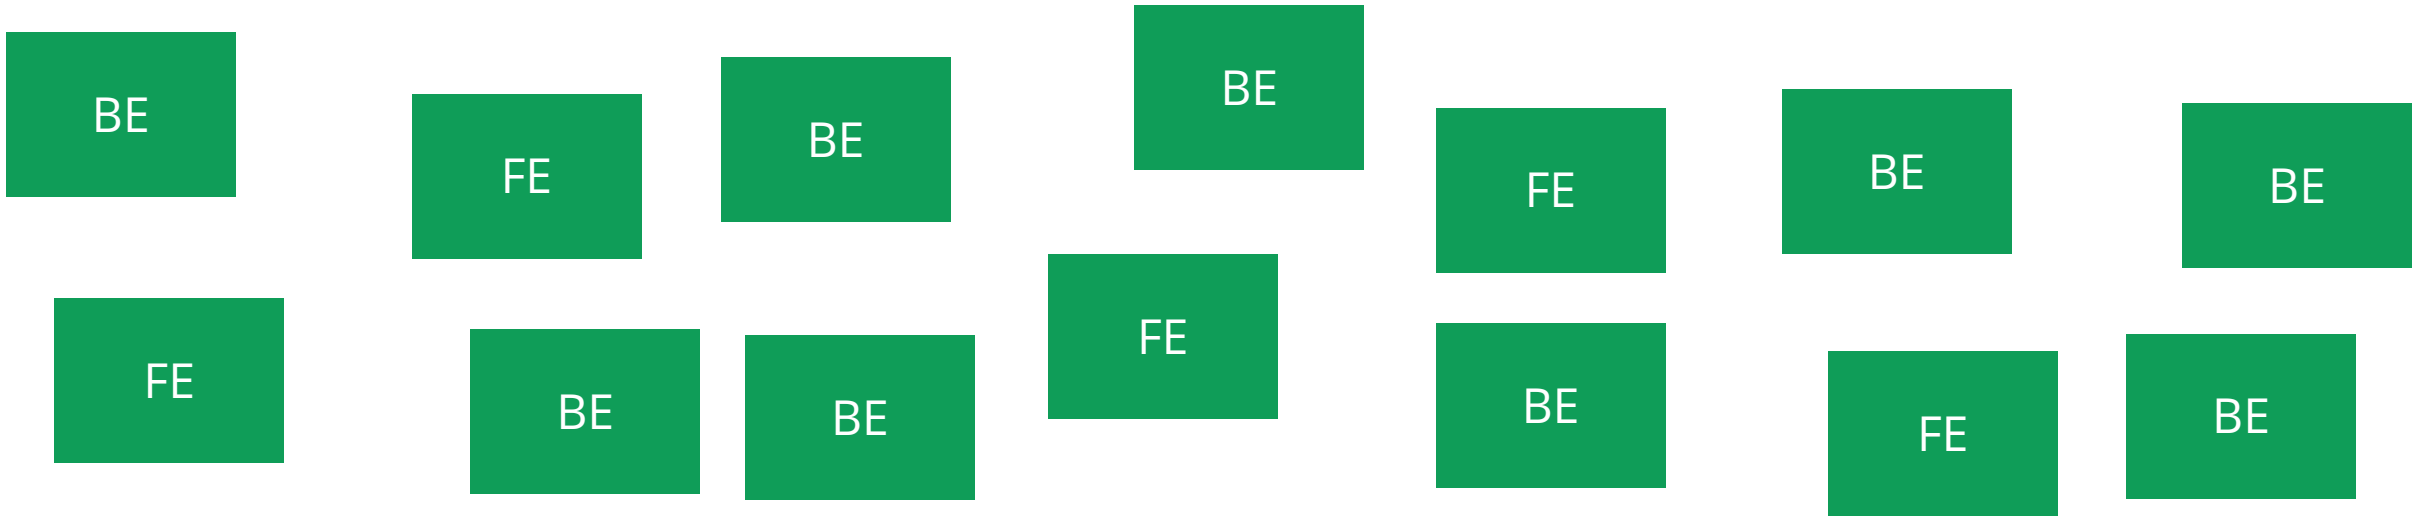

Kubernetes master/scheduler

Container Agent	Container Agent	Container Agent	Container Agent	Container Agent	Container Agent	Container Agent
Machine Host	Machine Host	Machine Host	Machine Host	Machine Host	Machine Host	Machine Host

Label selectors

labels:
role: frontend

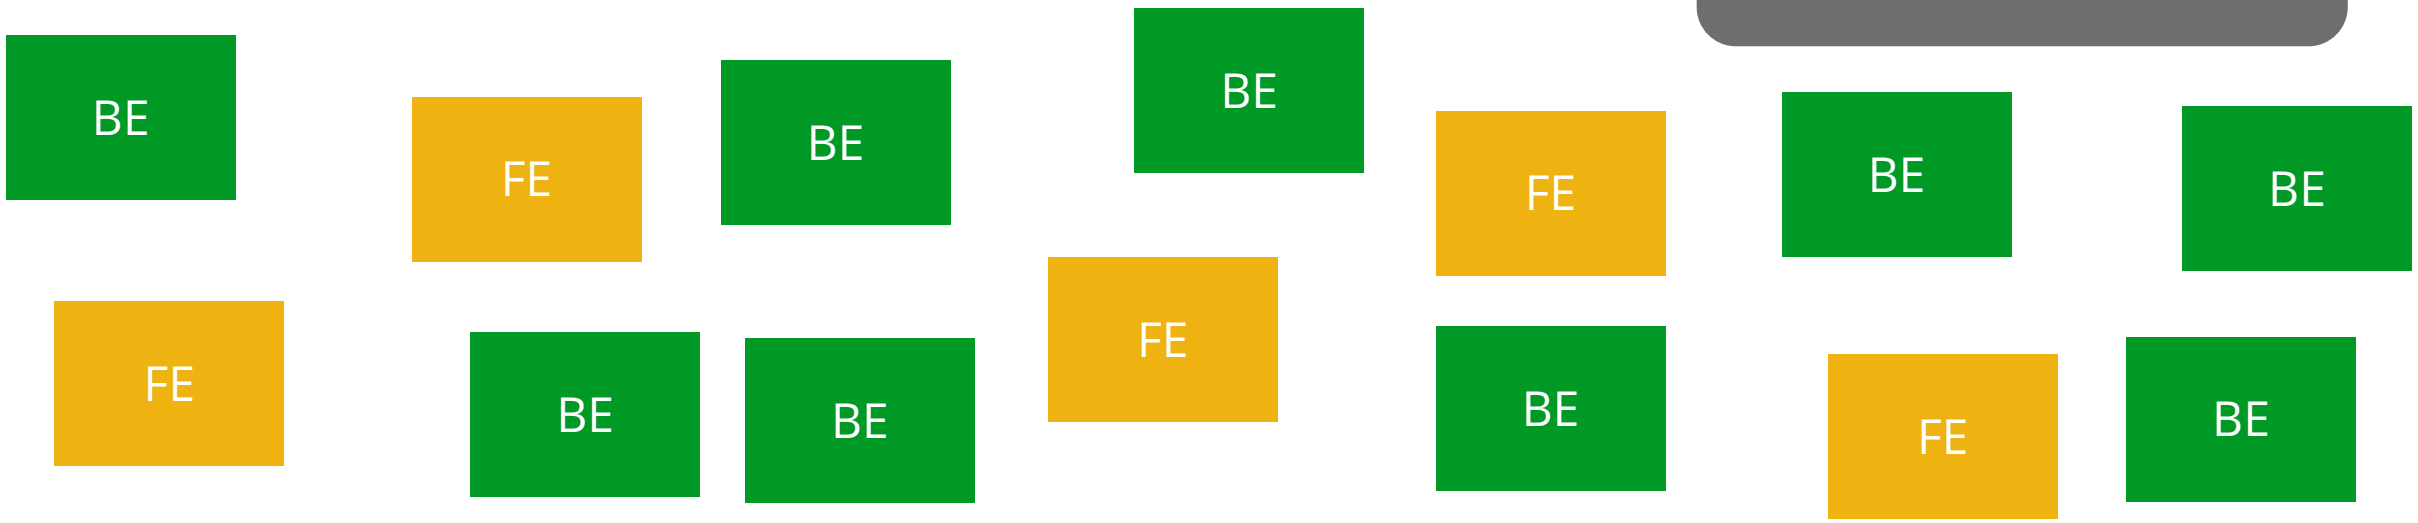

Kubernetes master/scheduler

Container Agent	Container Agent	Container Agent	Container Agent	Container Agent	Container Agent	Container Agent
Machine Host	Machine Host	Machine Host	Machine Host	Machine Host	Machine Host	Machine Host

Label selectors

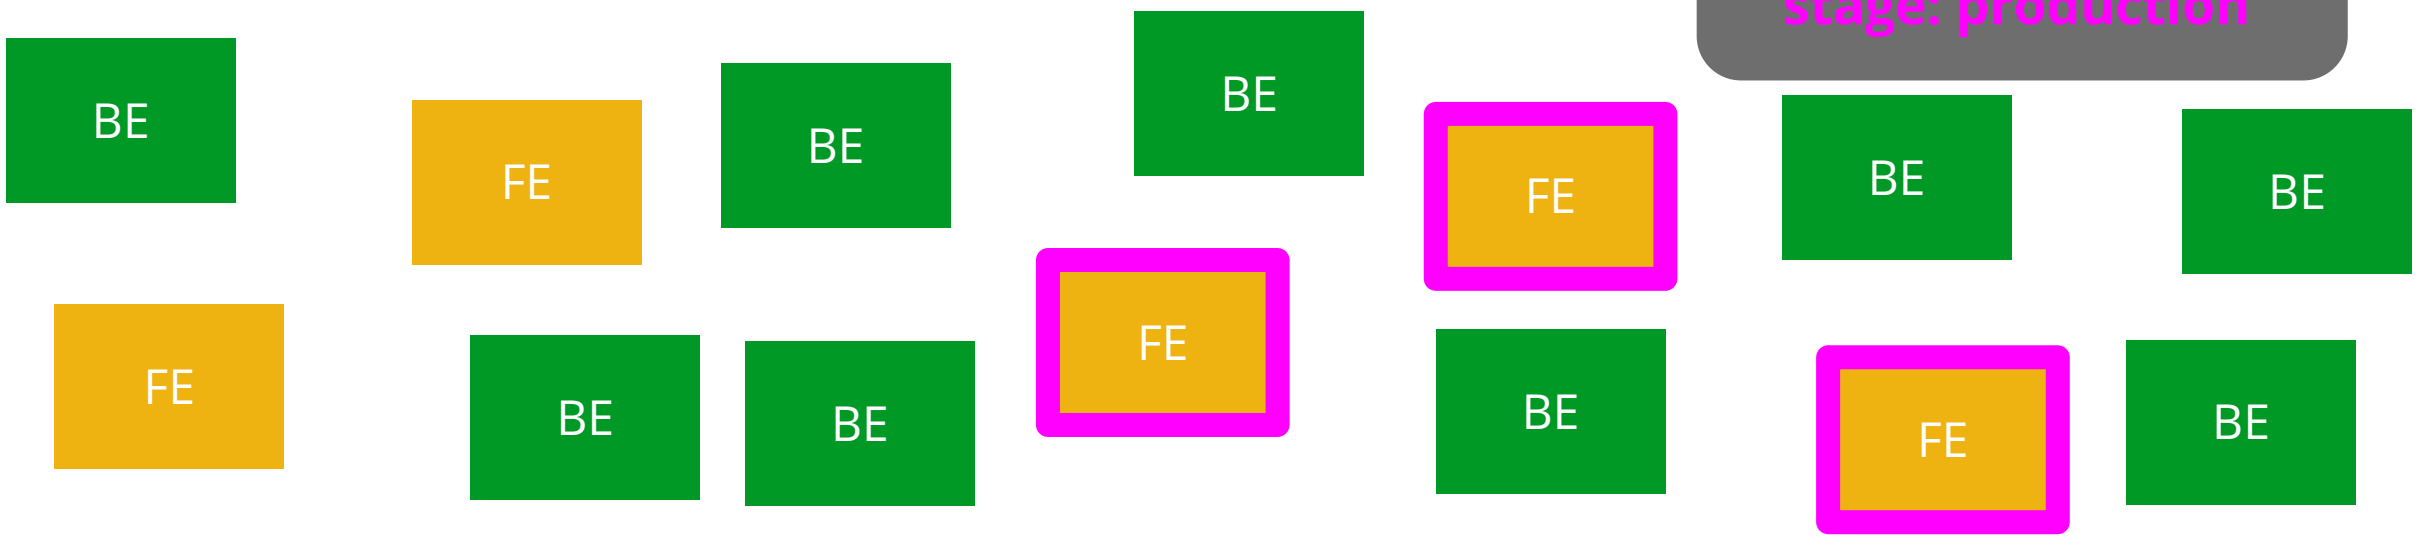

Replica controller

FE FE FE

replicas: 3
template:
  ...
labels:
  role: frontend

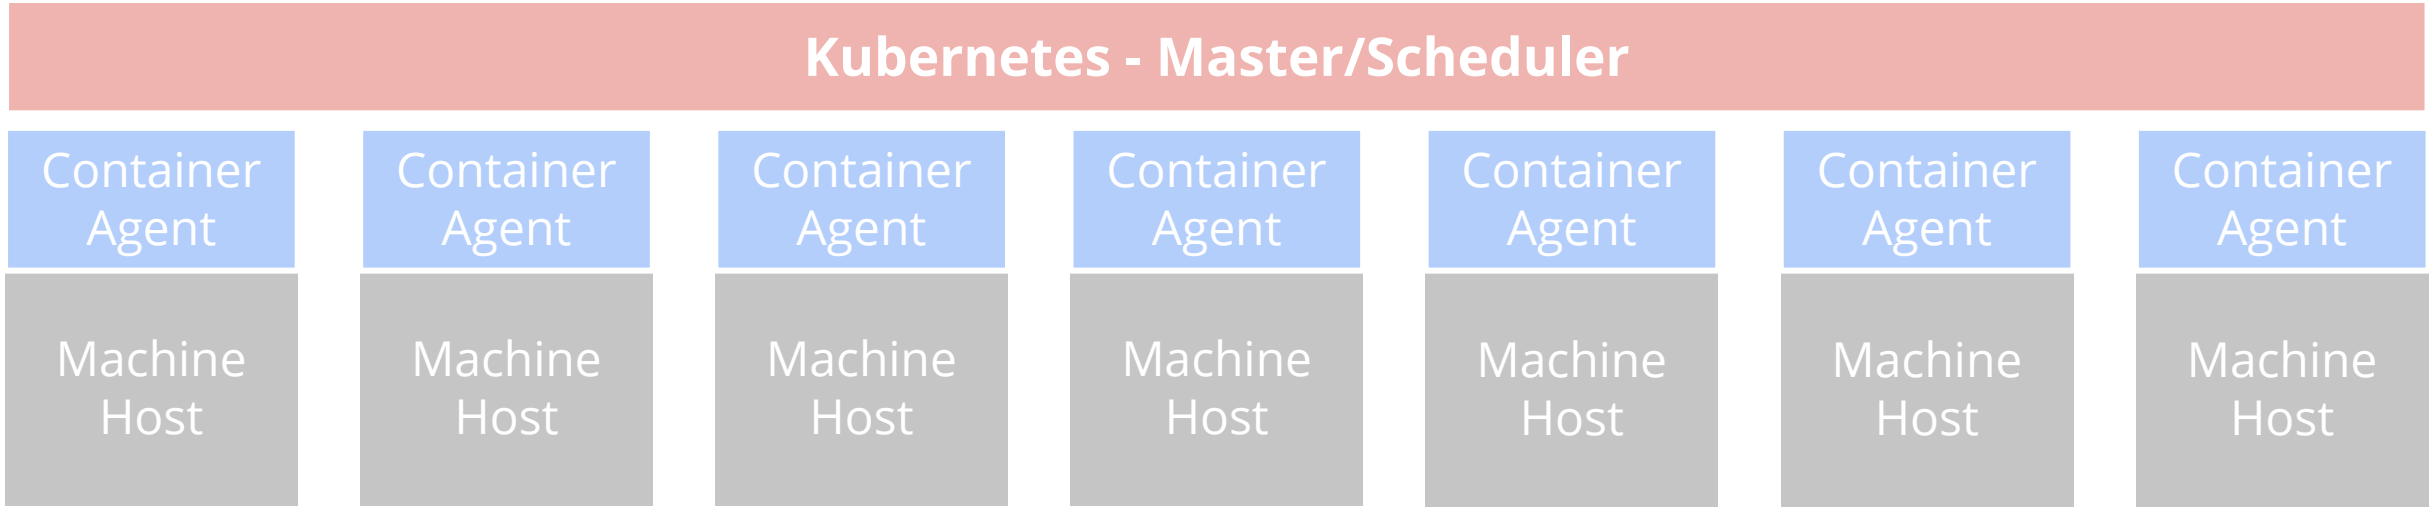

Replica controller

FE FE FE FE

replicas: 4
template:
  ...
labels:
  role: frontend

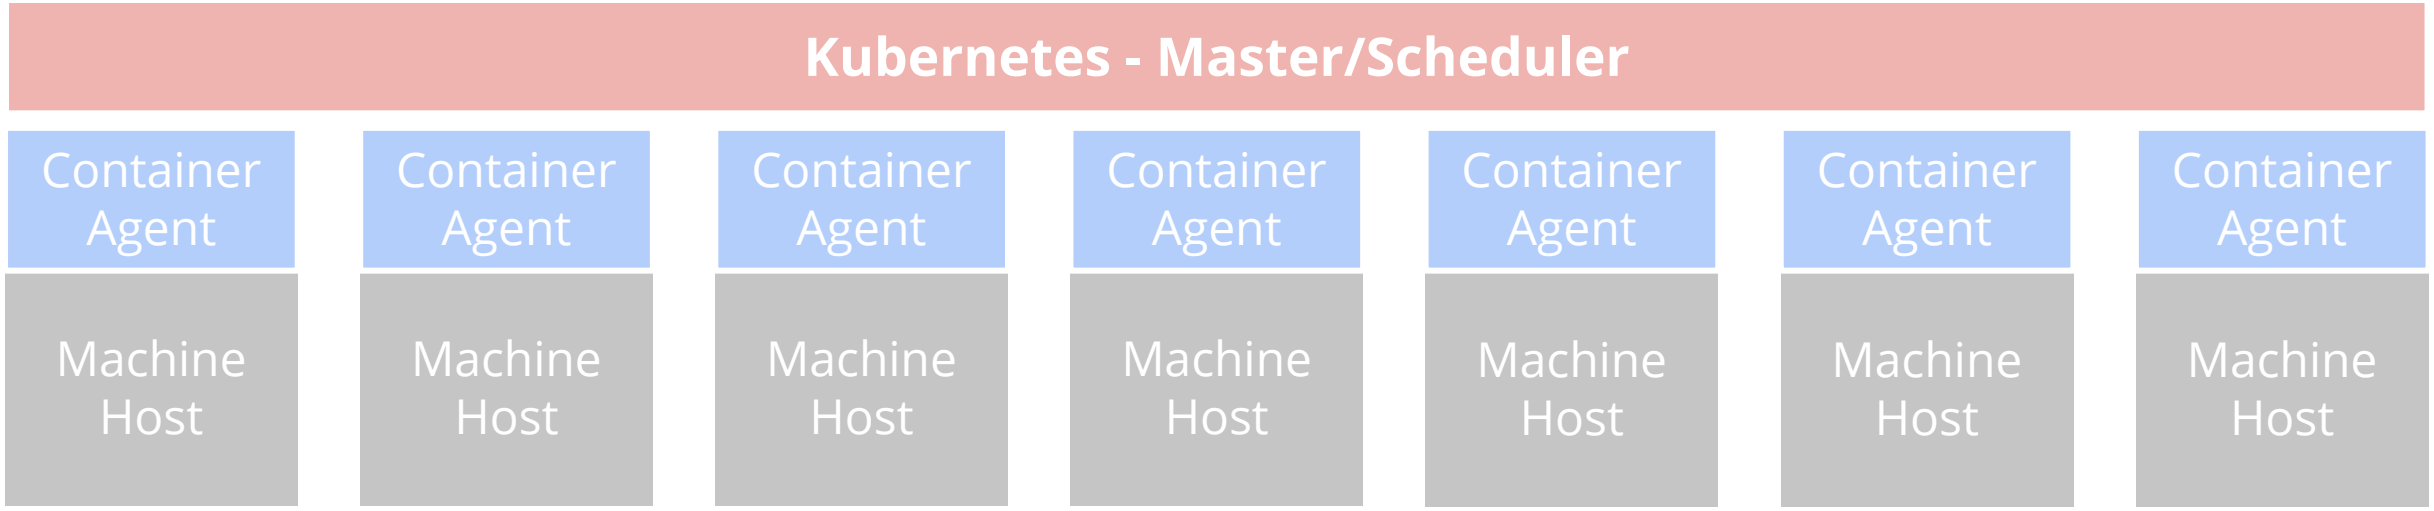

Service

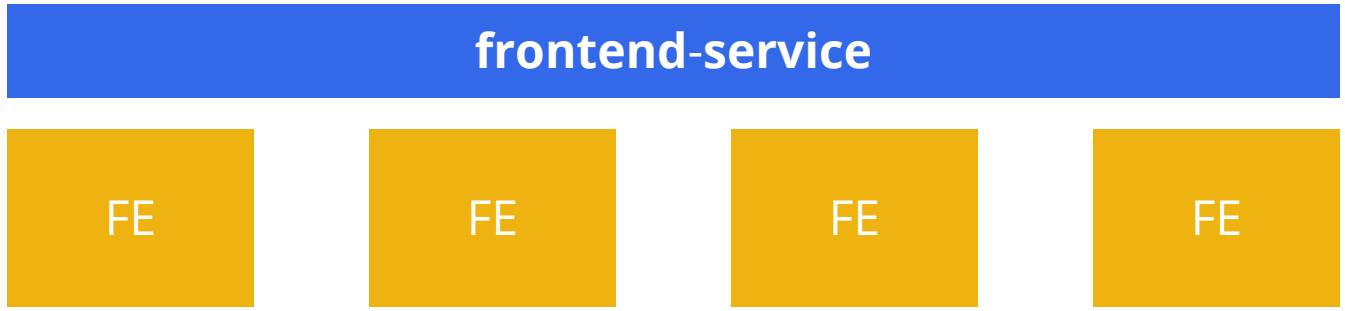

id: **frontend-service**
port: 9000
labels:
  role: **frontend**

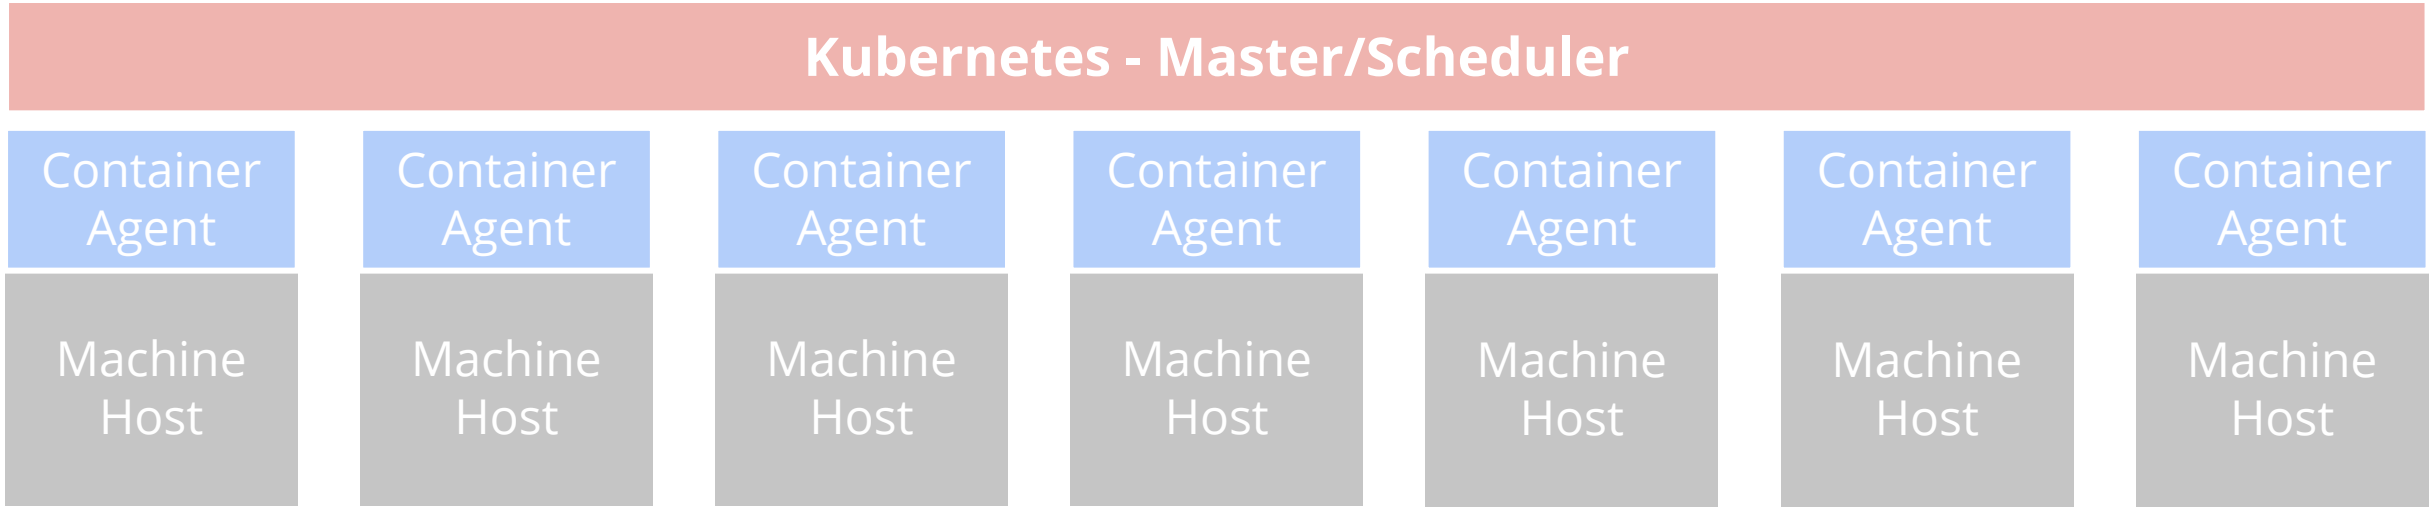

kubernetes Home Getting Started Community News Events Blog

Kubernetes

The open source cluster manager from Google.

- Pods: groups of containers
- Labels
- Replica controller
- Services

http://kubernetes.io

kubernetes by Google

Manage a cluster of Linux containers as a single system to accelerate Dev and simplify Ops.

Pulling it all together

Do it yourself? Sure.

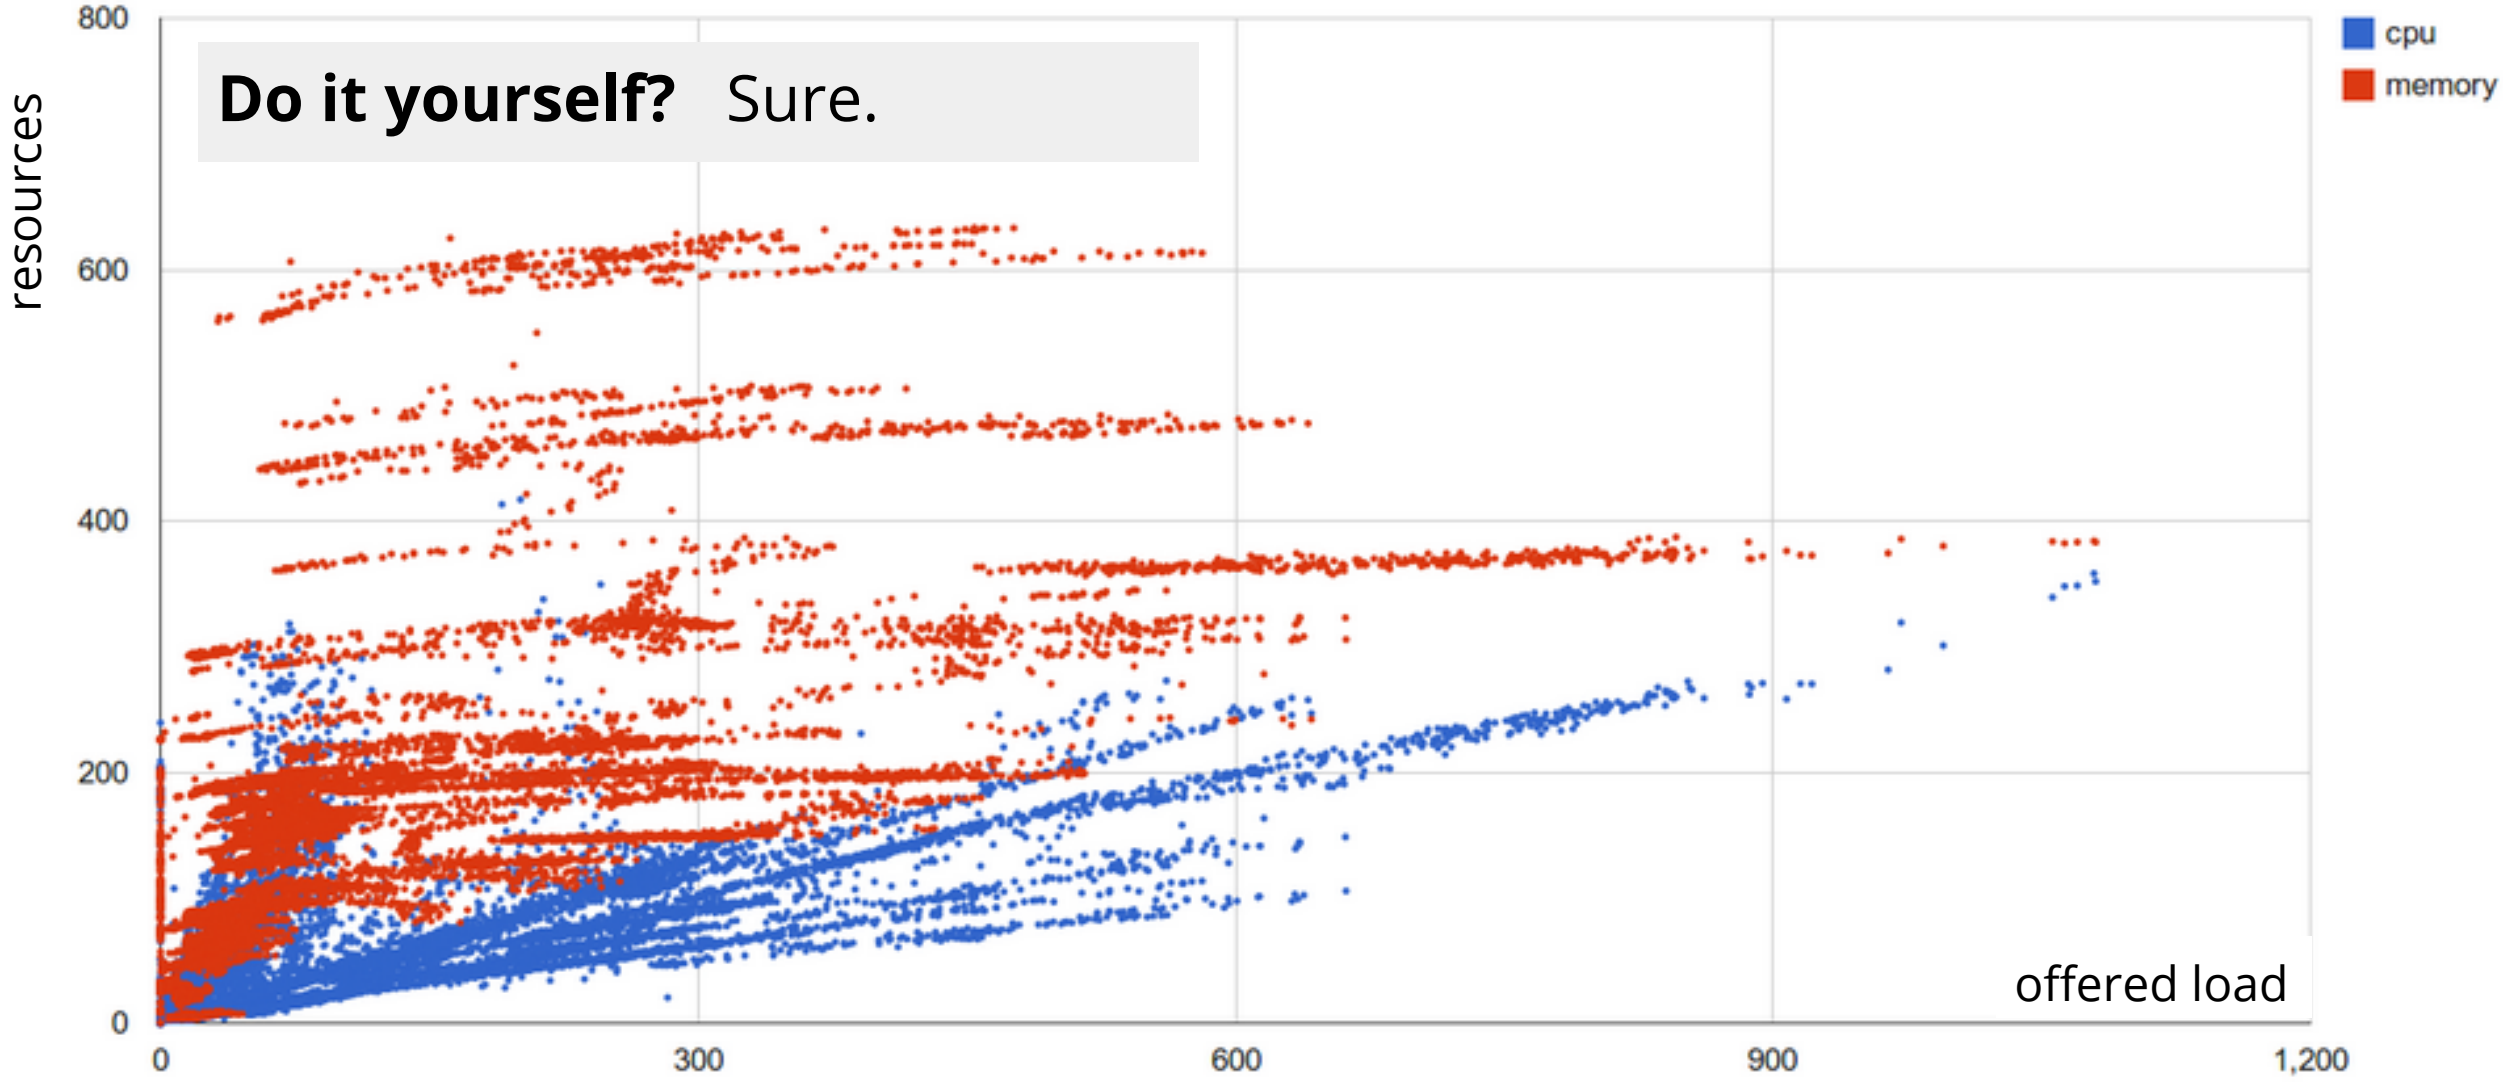

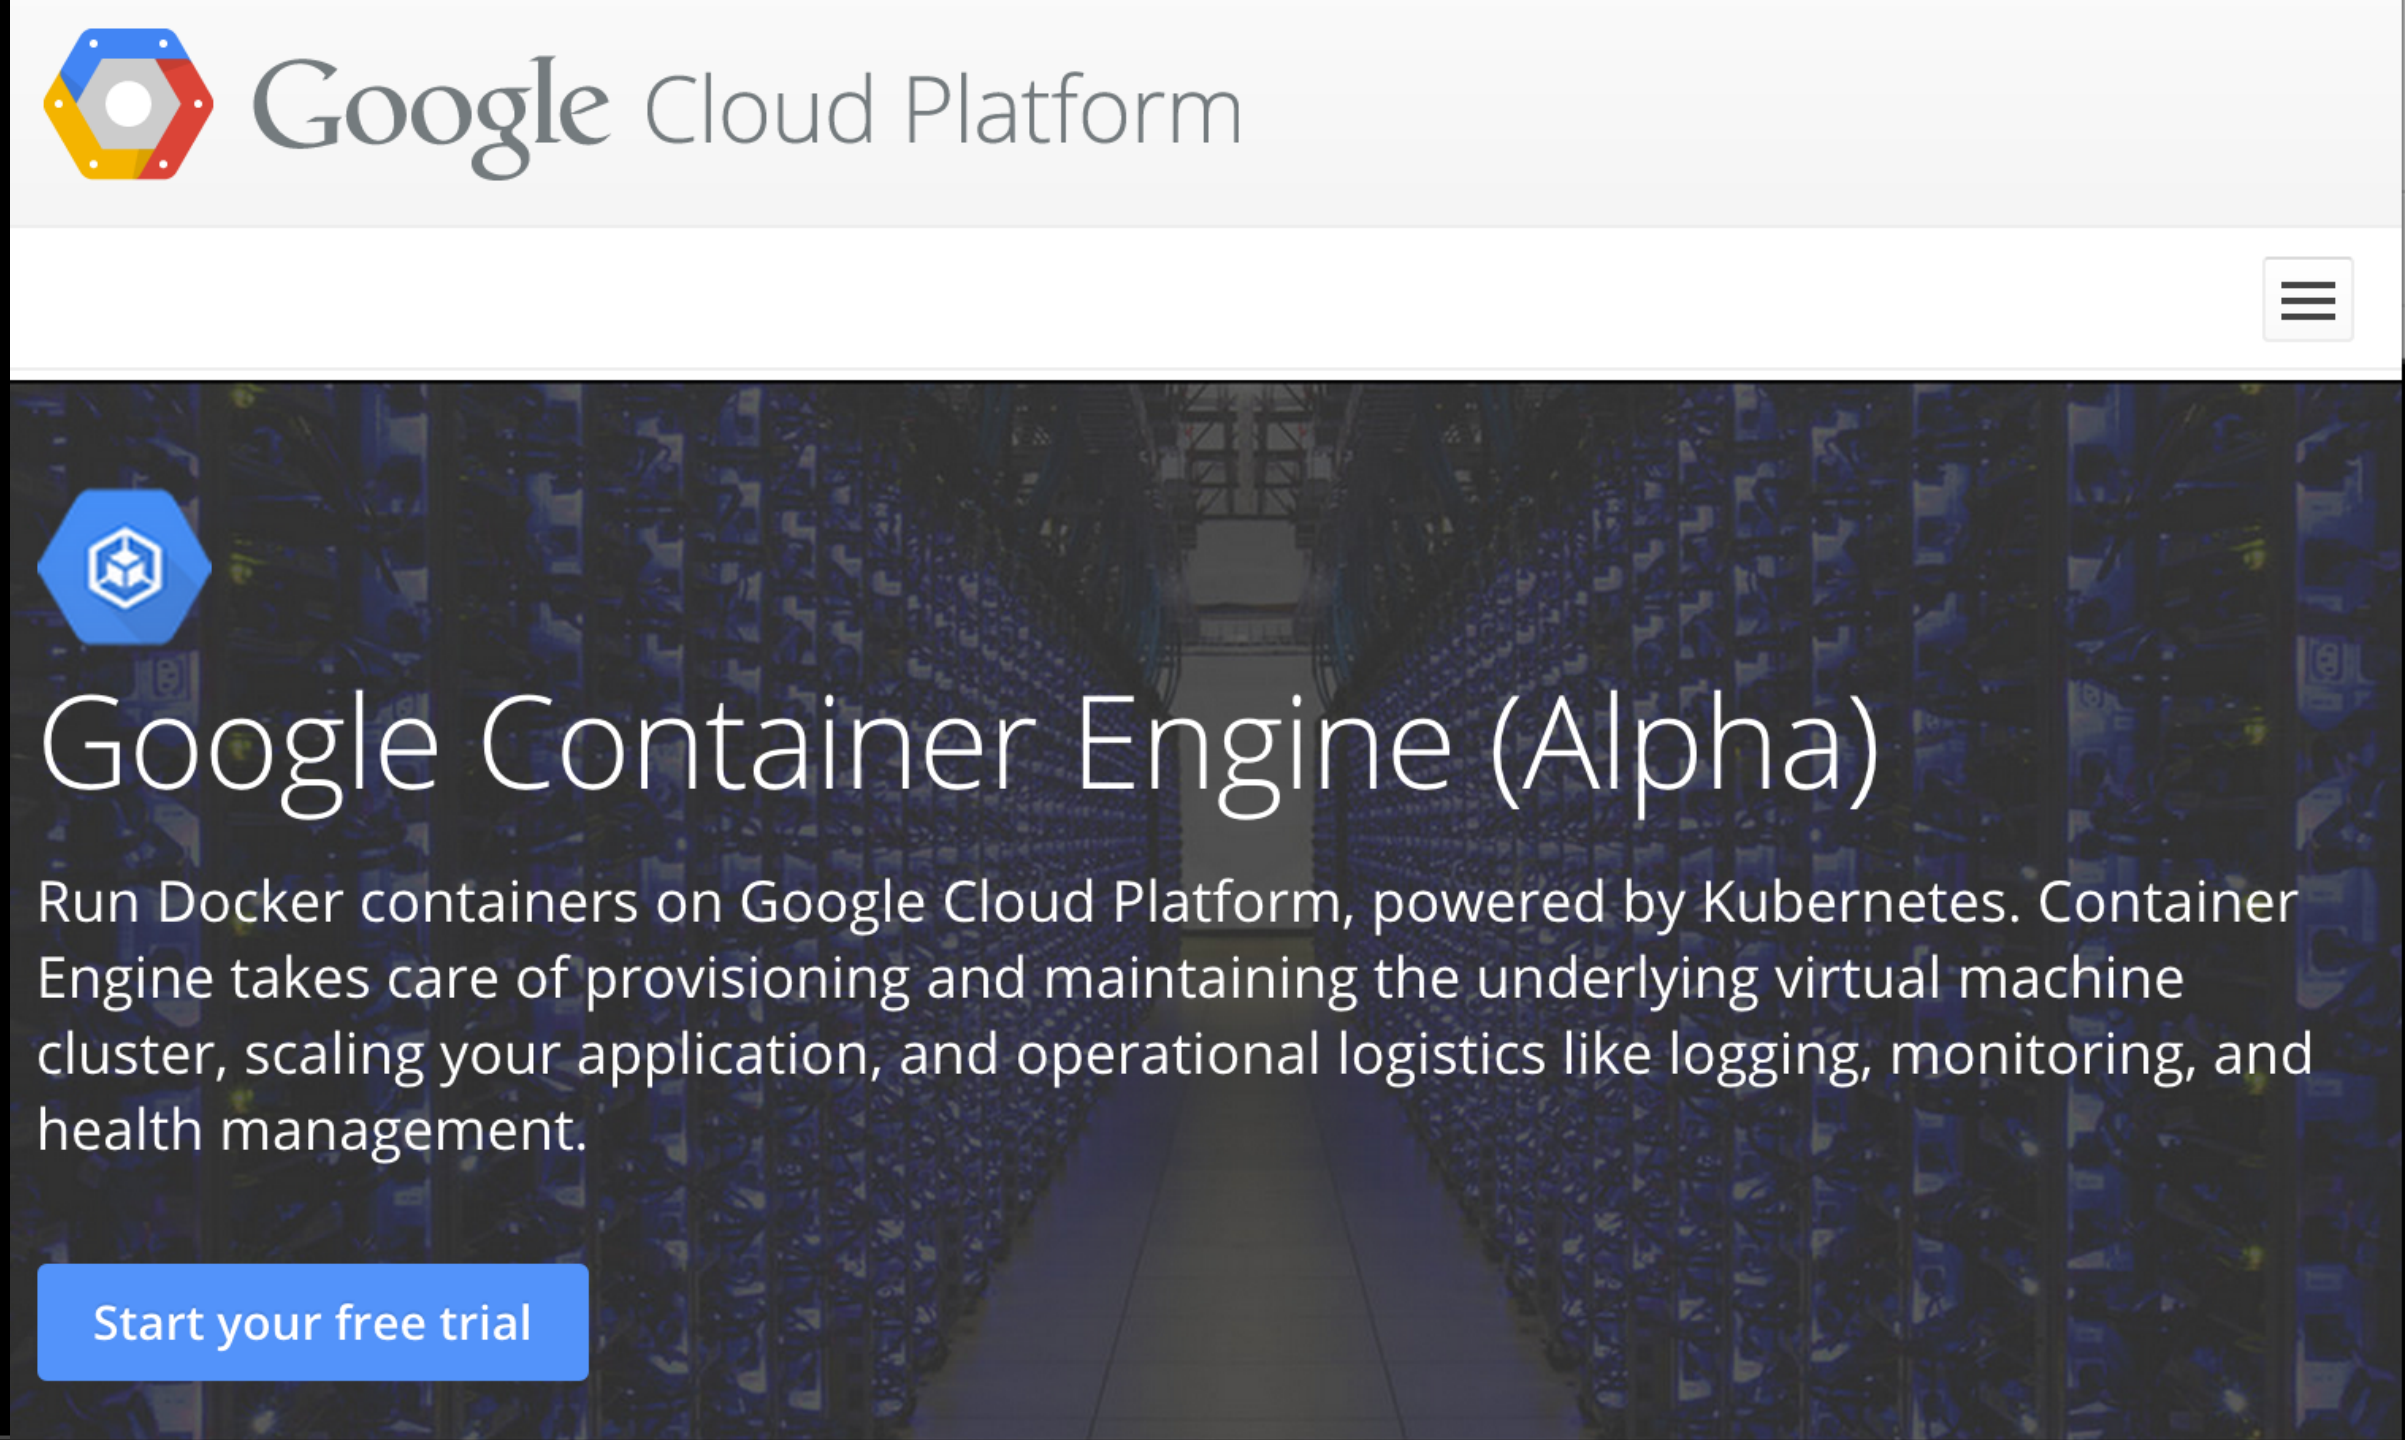
Google Cloud Platform
Google Container Engine (Alpha)
Run Docker containers on Google Cloud Platform, powered by Kubernetes. Container Engine takes care of provisioning and maintaining the underlying virtual machine cluster, scaling your application, and operational logistics like logging, monitoring, and health management.
Start your free trial

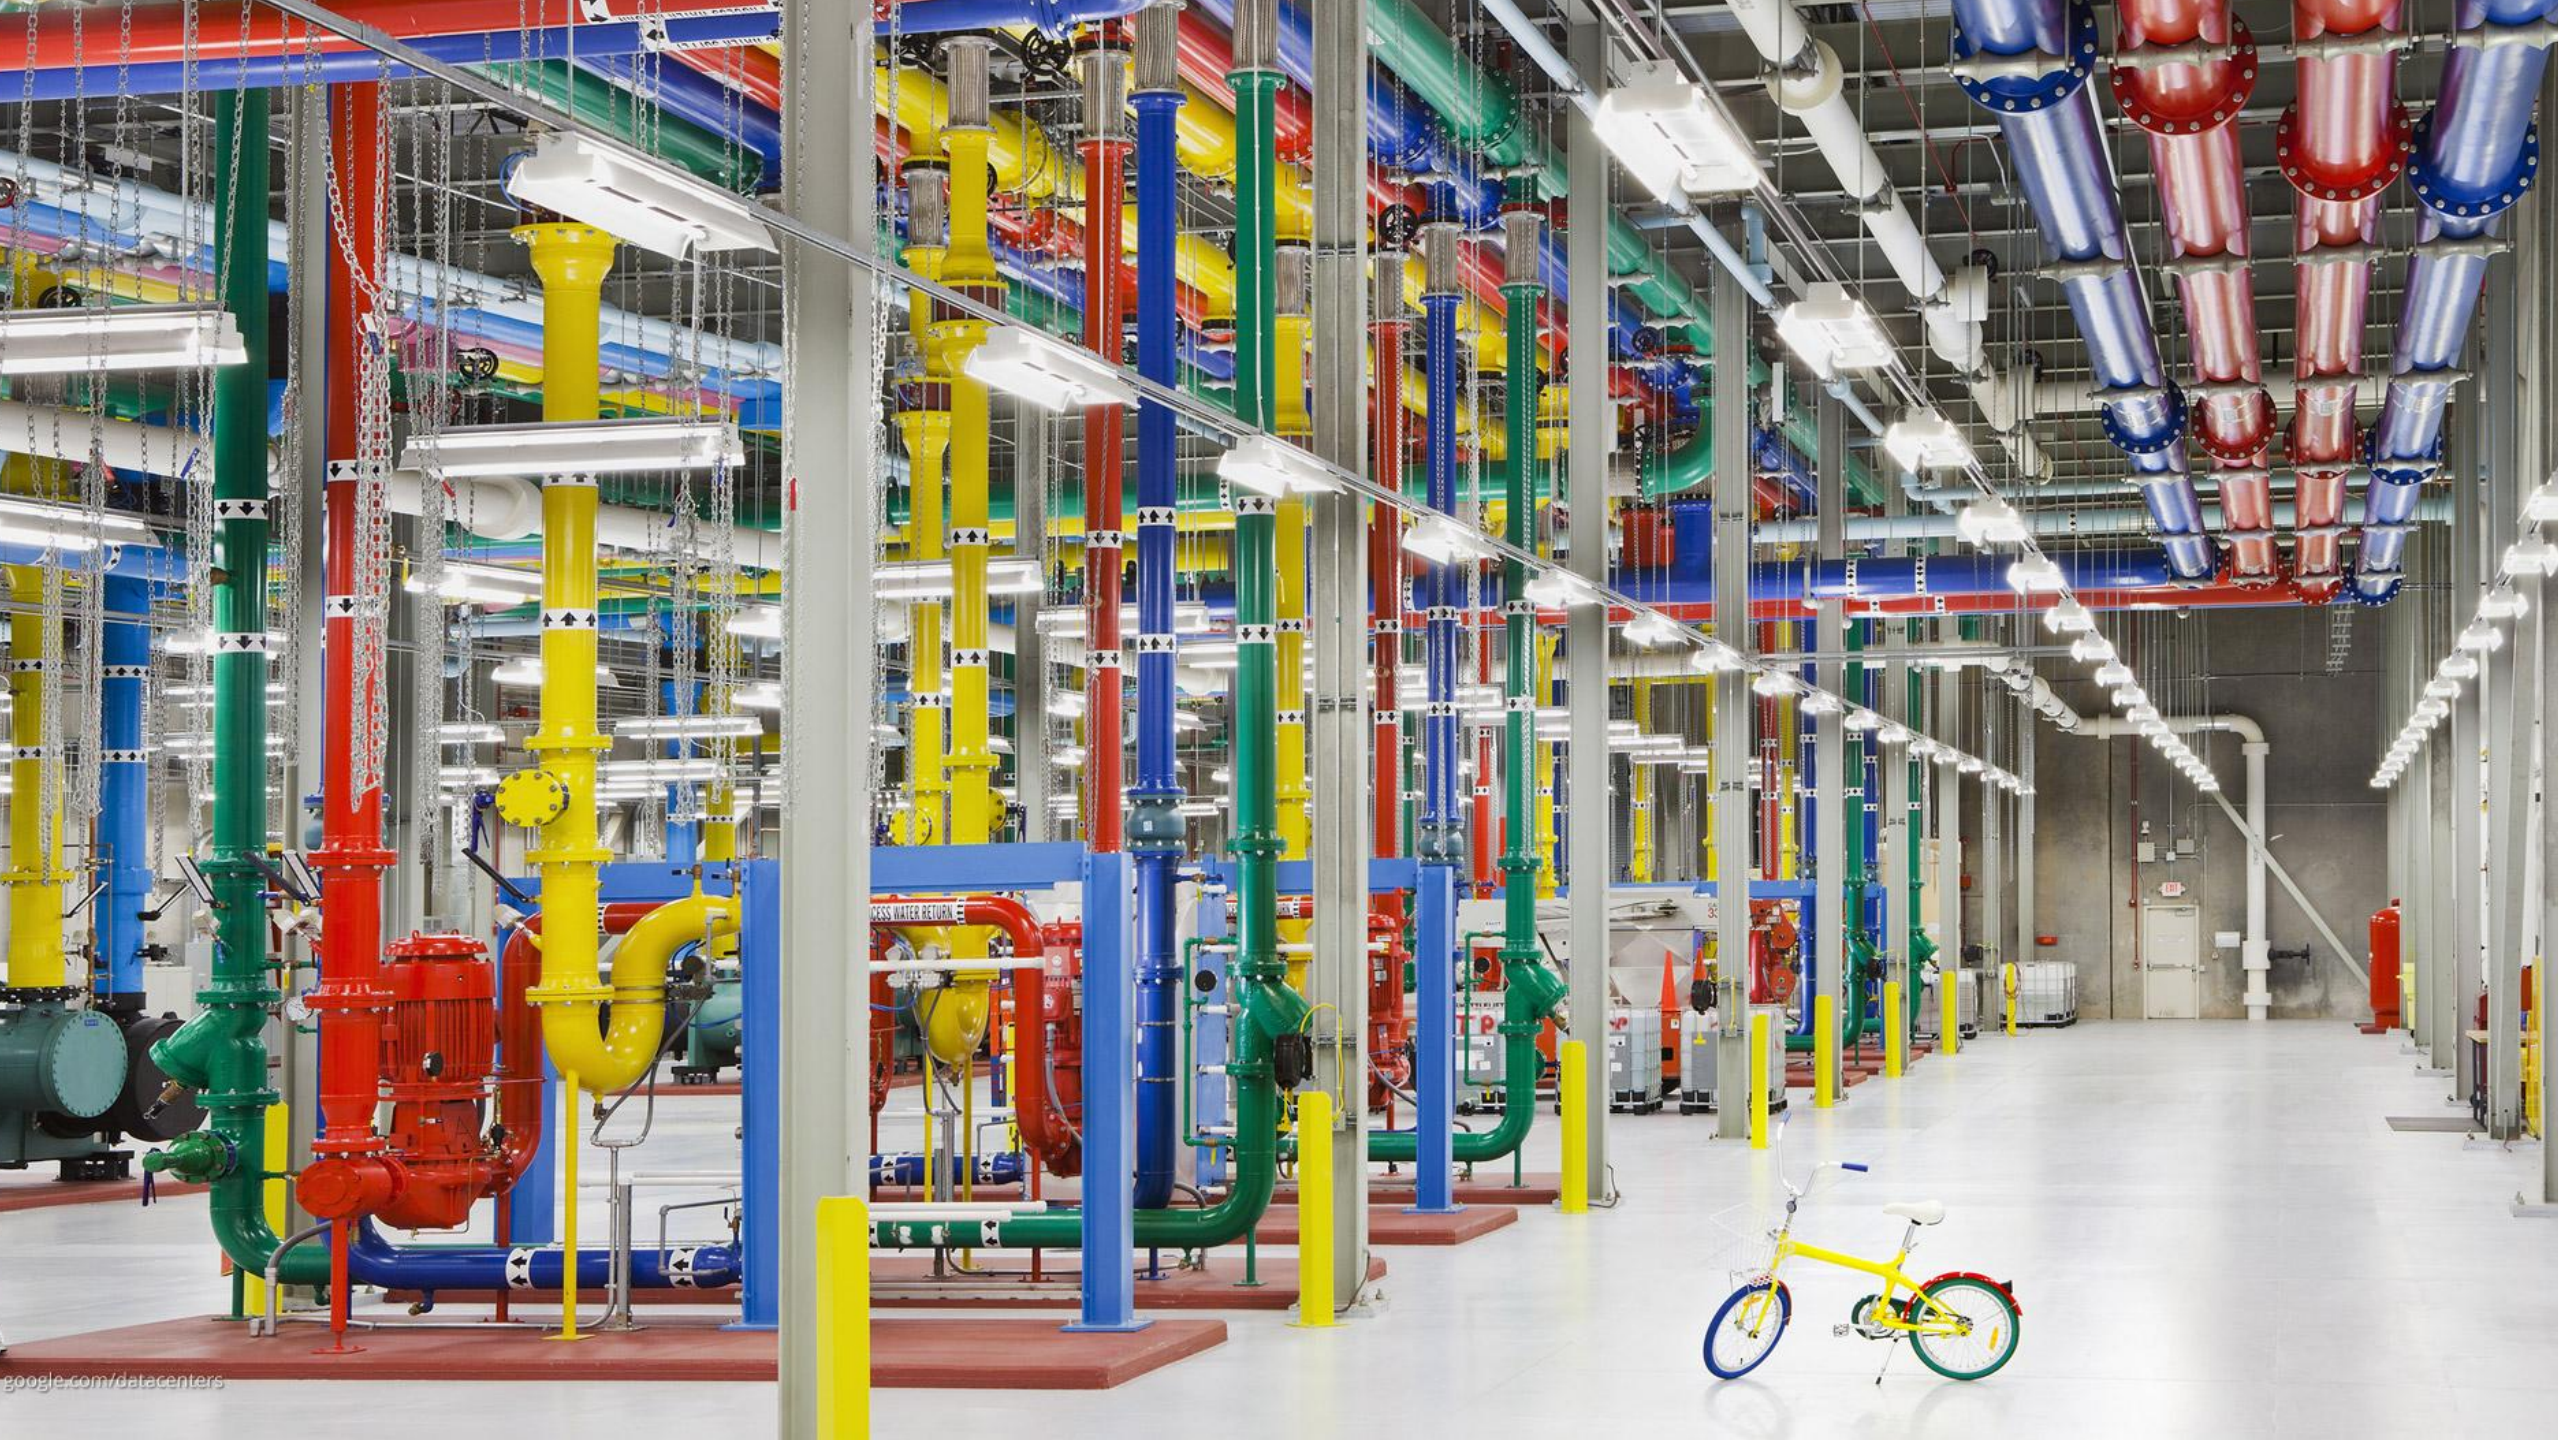
google.com/datacenters

Pulling it all together

We choose to go to the roof not because it is glamorous, but because it is right there!

> ... the bulk of our success is the result of the methodical, relentless, persistent pursuit of 1.3-2x opportunities -- what I have come to call "**roofshots**".
>
> -- Luiz Barroso

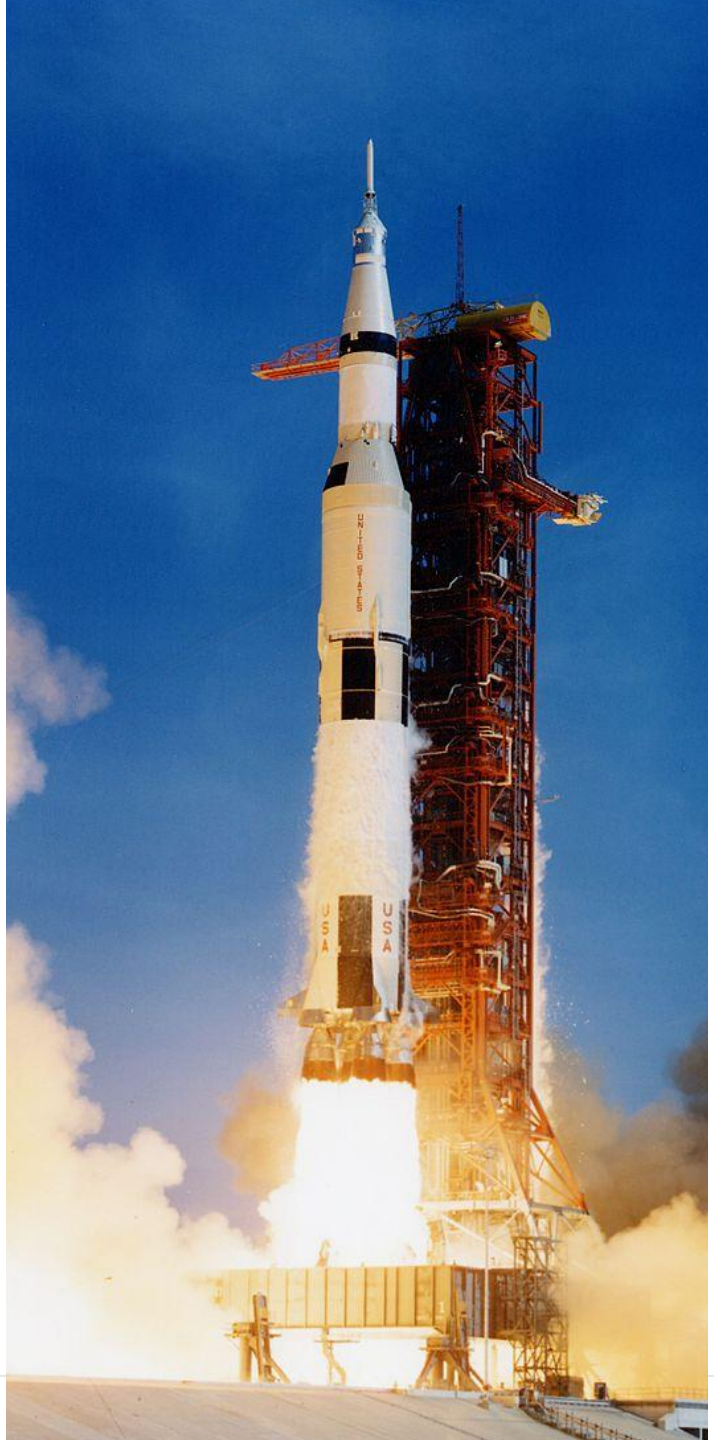

Pulling it all together

Data: Volkswagen, 2014-07-31
Image: john wilkes

Porsche doesn't **make** cars:
it designs and assembles them

1H2014:

- **1.7%** (89k) of VW group's vehicles
- **23% (€1.4b) of its profits**

Pulling it all together

Cloud system providers are getting better at *everything* ...

- capacity management
- monitoring
- storage + networking
- reliability
- software development tooling
- ...

Wouldn't you like to stand on others' shoulders?

Three rules of thumb:

1. *Resiliency* is more important than performance.

2. Relax. Let go. *Build on* what others have done.

3. Do more *monitoring*.

johnwilkes@google.com

http://kubernetes.io

Images by Connie Zhou

Achieving desired behavior

Service level agreement (SLA)

SLA = SLOs + consequences of achieving or missing them

Example:

- if *availability* > *99.95%* (SLO) user pays £xx/CPU-week
- else gets a 30% refund

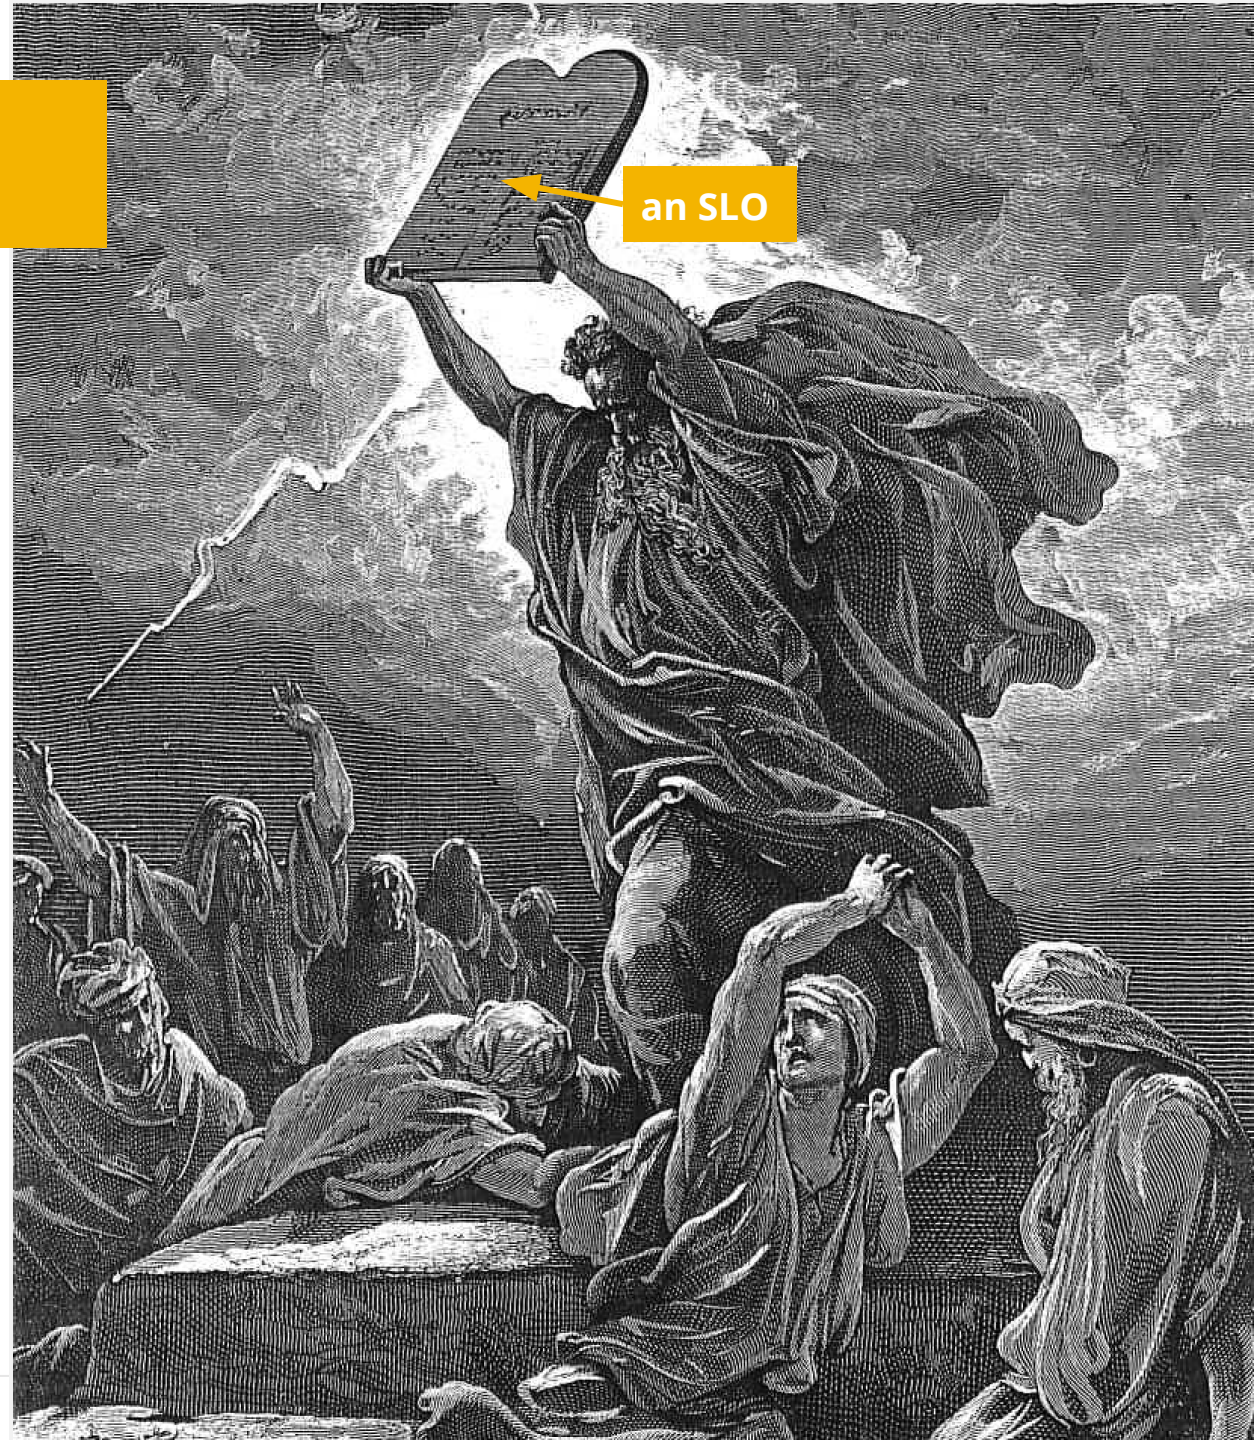

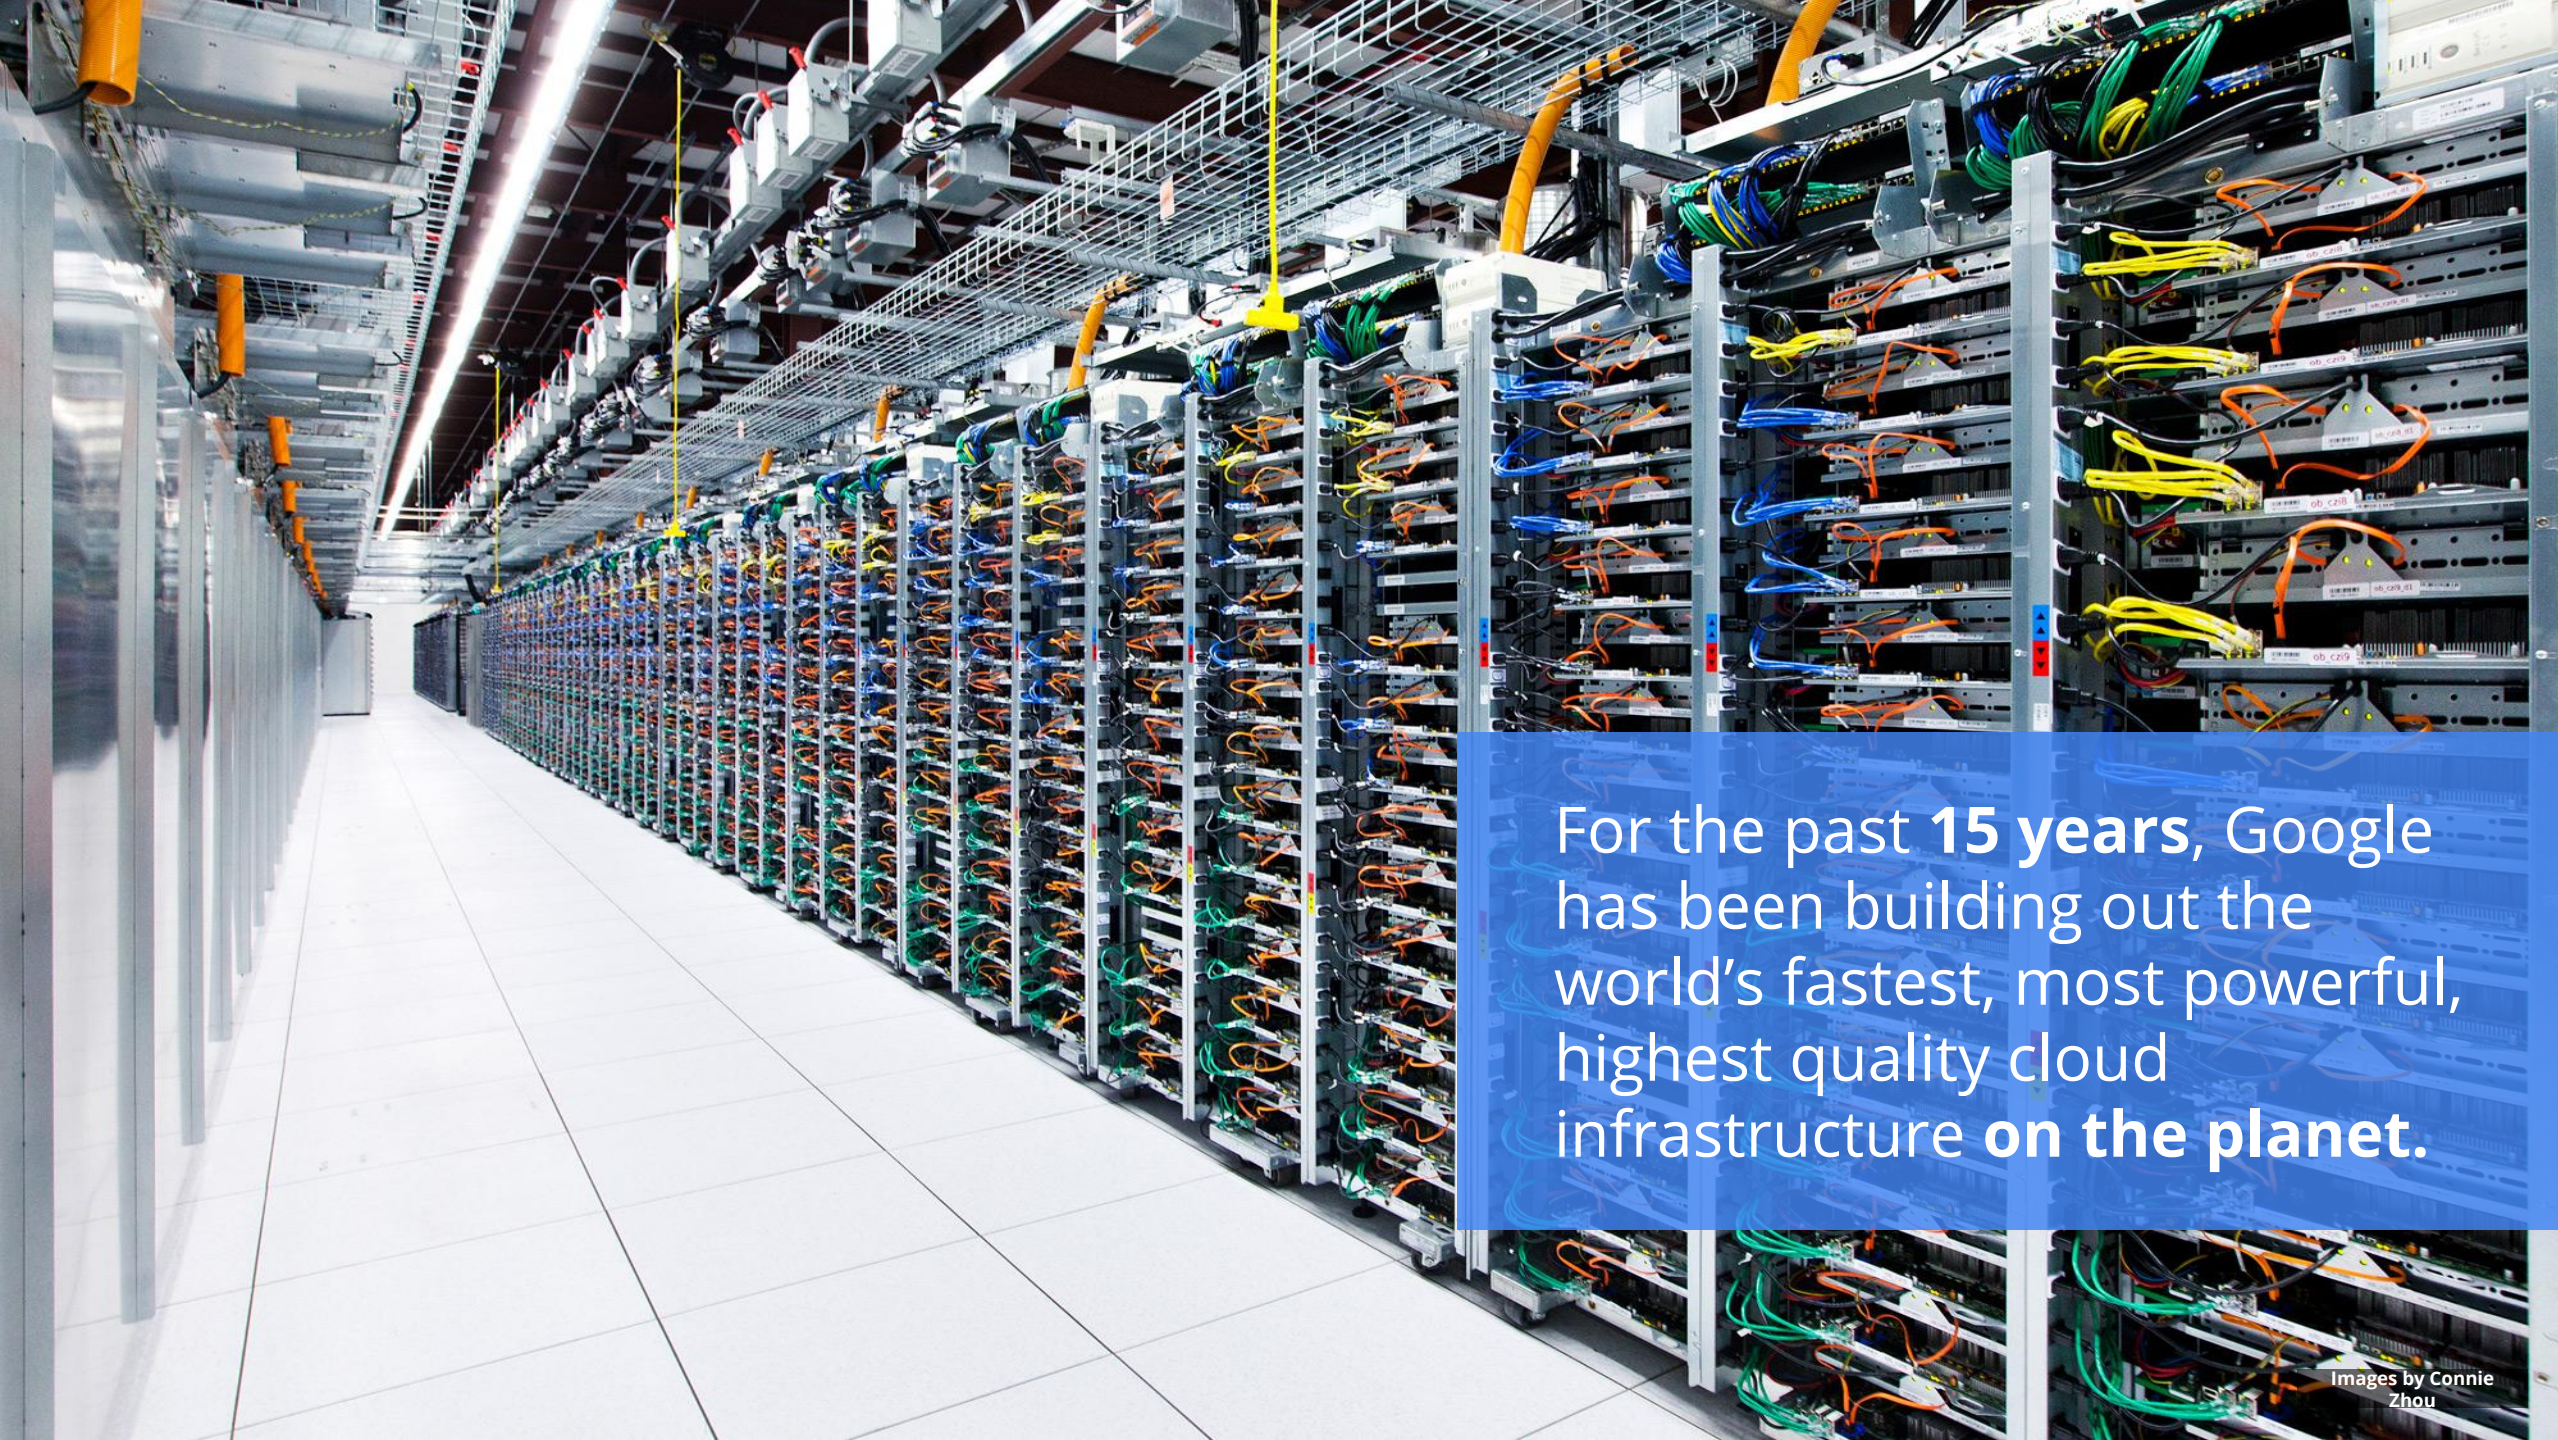

For the past **15 years**, Google has been building out the world's fastest, most powerful, highest quality cloud infrastructure **on the planet.**

Images by Connie Zhou