The Truth about Docker Container Lifecycles

KEVIN MCGUIRE
AUGUST 17/2015

A story begins…

Servicification of New Relic

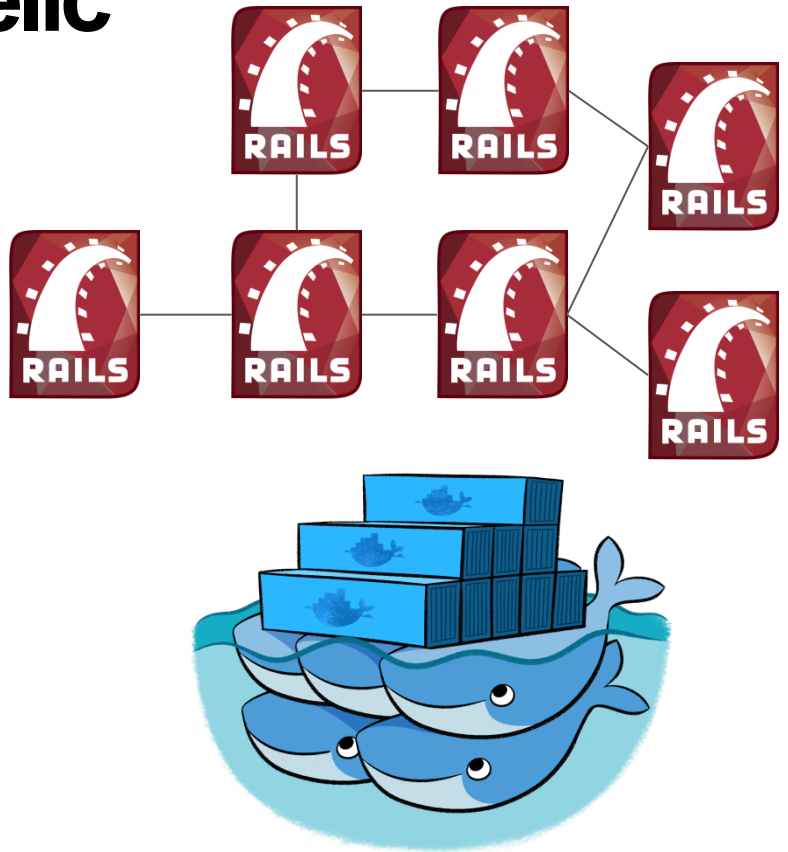

Service container life cycle

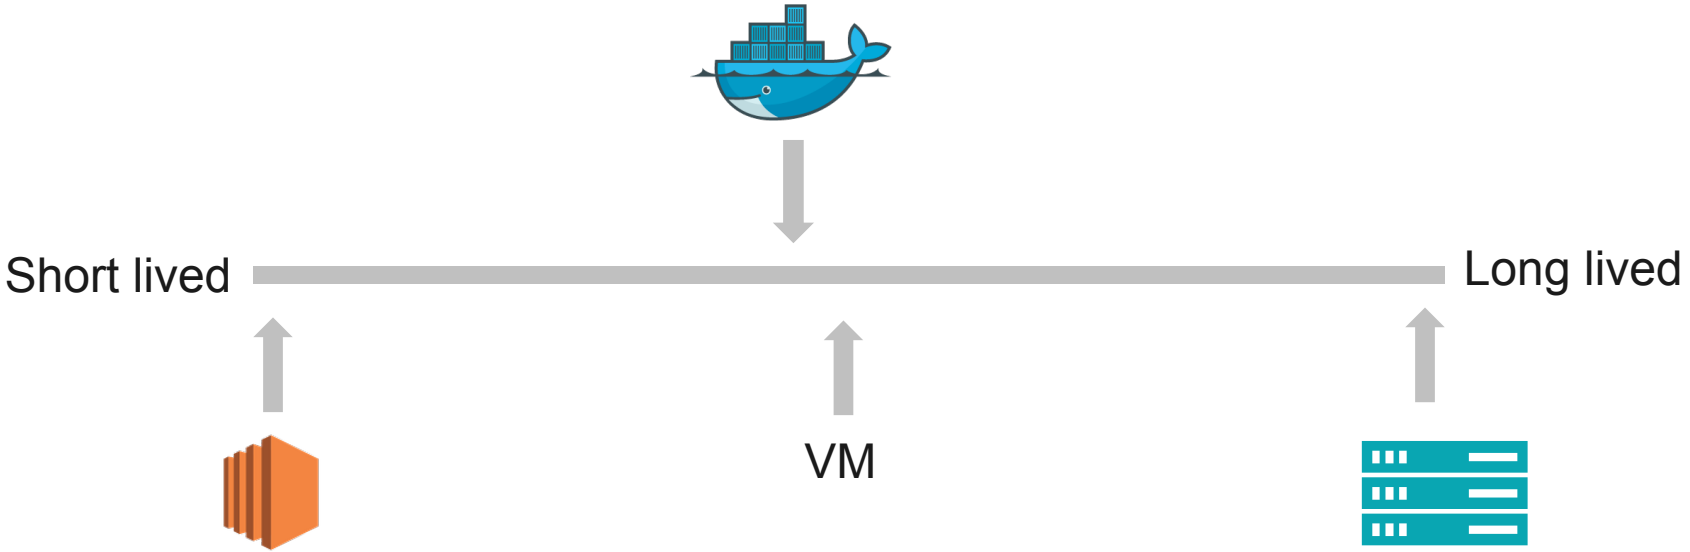

Theory

Docker is a lightweight VM

Let's monitor it!

Community Forum

Log In

How to try out New Relic's beta support for Docker

Server Monitoring Docker coreos ecs

adam NR Engineering 17 Mar 27

The following will guide you through the process to try out the latest beta release of New Relic's Docker support. If you have any questions or confusion, please do not hesitate to reach out or reply to this post directly. Happy monitoring!

Currently the New Relic Docker Beta will enable the following features:

- Navigation between APM and Servers when applications are hosted within Docker containers
- Granular visibility about containers and hosts that are running your APM-instrumented applications
- CPU and Memory metrics rolled up by Docker image for a given host
- A historical view of number of running containers of a given Docker image type

Well that was surprising

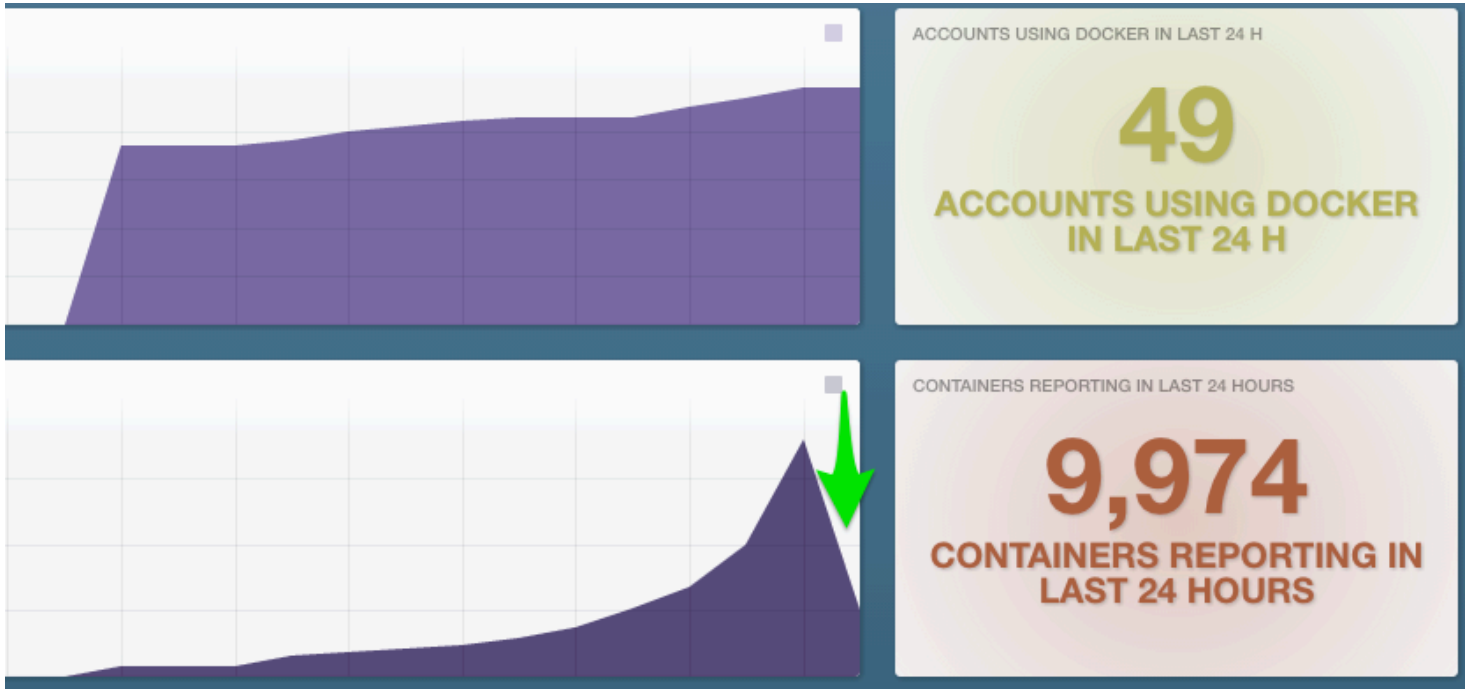

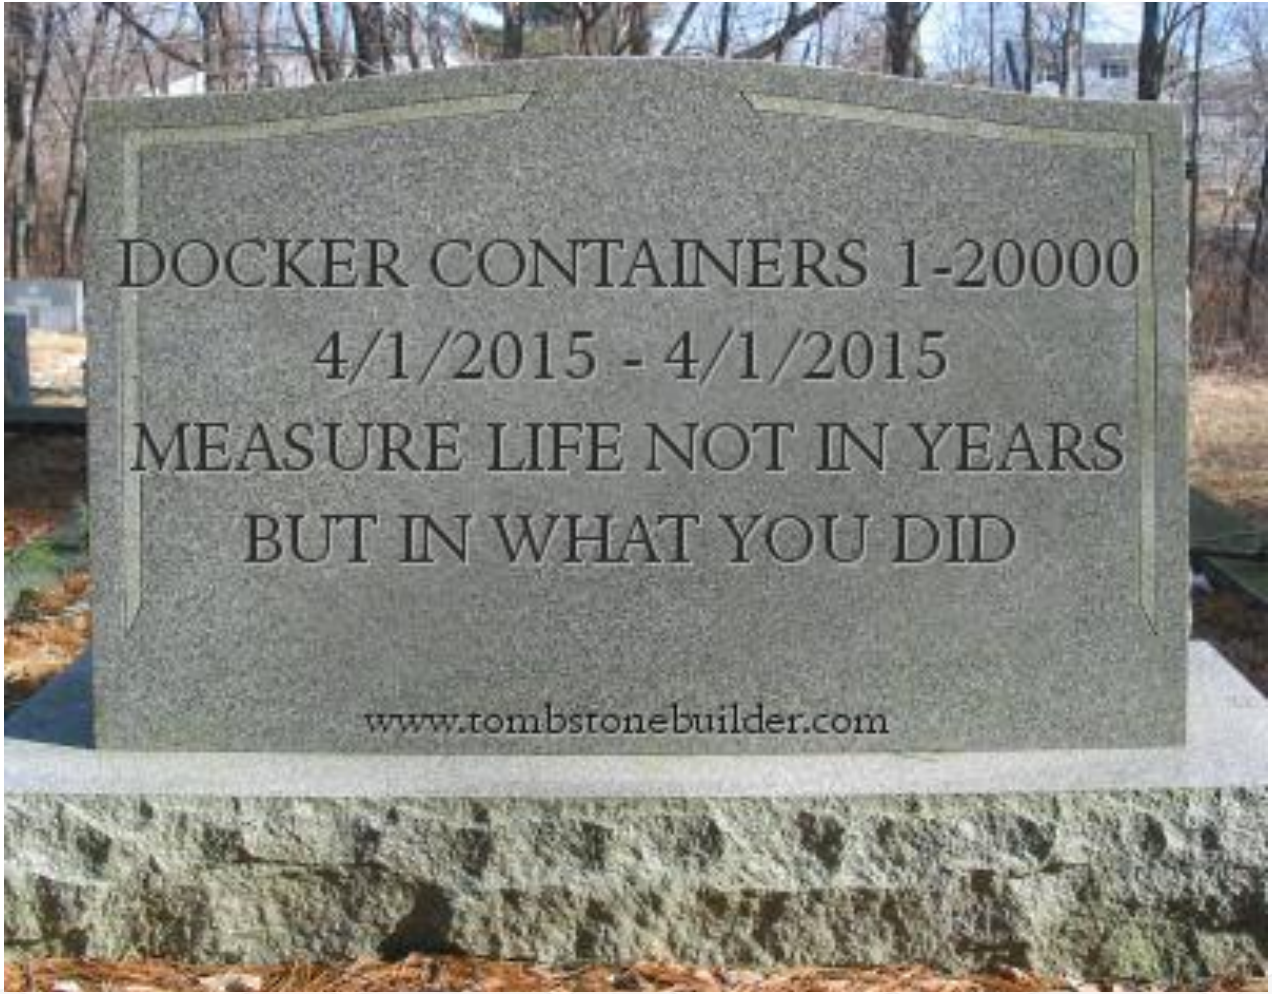
DOCKER CONTAINERS 1-20000
4/1/2015 - 4/1/2015
MEASURE LIFE NOT IN YEARS
BUT IN WHAT YOU DID
www.tombstonebuilder.com

Apparent customer usage

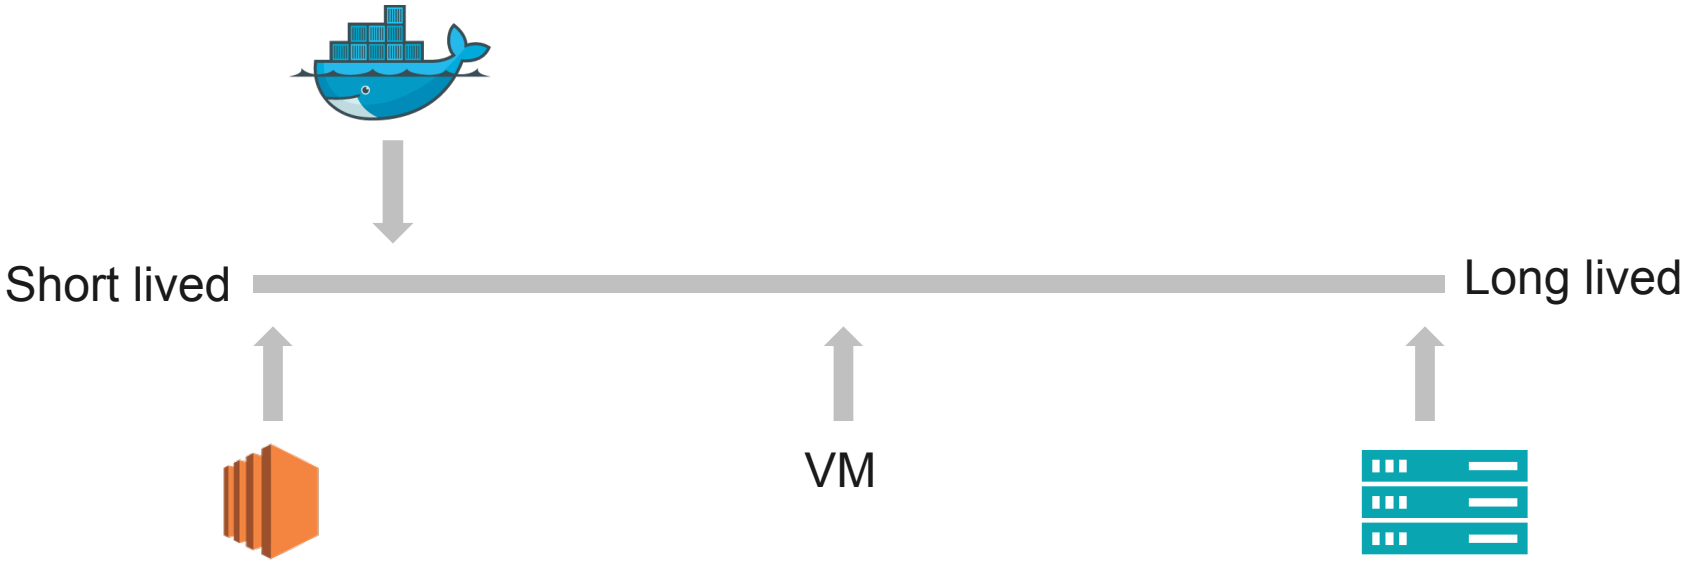

Theory

~~Docker is a lightweight VM~~
Docker is a cloud compute container

Pets vs Cattle

Along came New Relic Synthetics…

- User authored selenium scripts run in our data center
- Each run in its own container for security isolation
- Mostly run for under a minute

Along came New Relic Synthetics…

- Test external availability and performance
- User authored selenium scripts run in our data center
- Each run in its own container for security isolation
- Most run for under a minute

A lighter weight usage

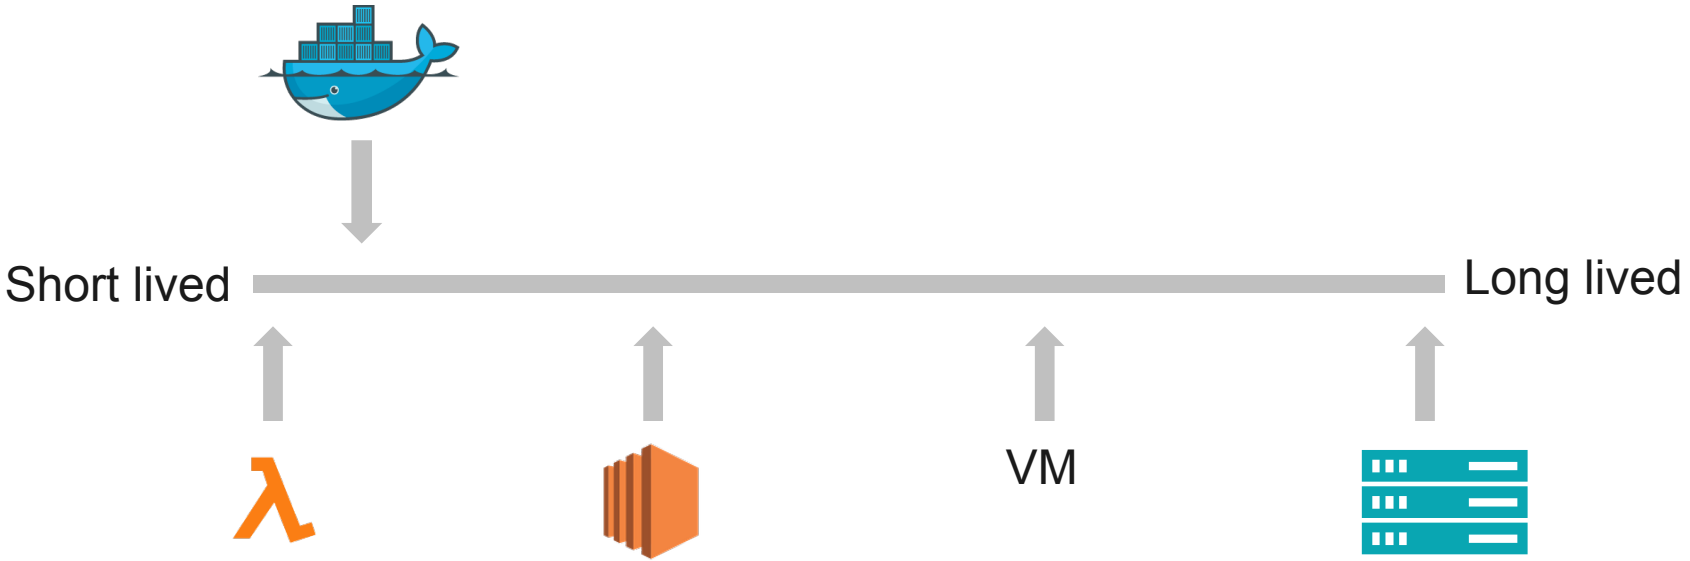

Theory

~~Docker is a lightweight VM~~
~~Docker is a cloud compute container~~
Docker is a short lived compute engine

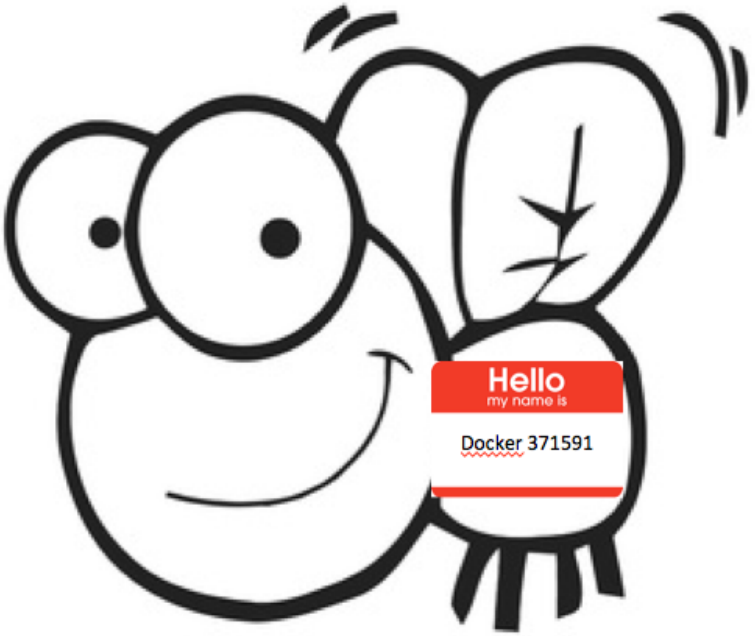

What the heck's going on?

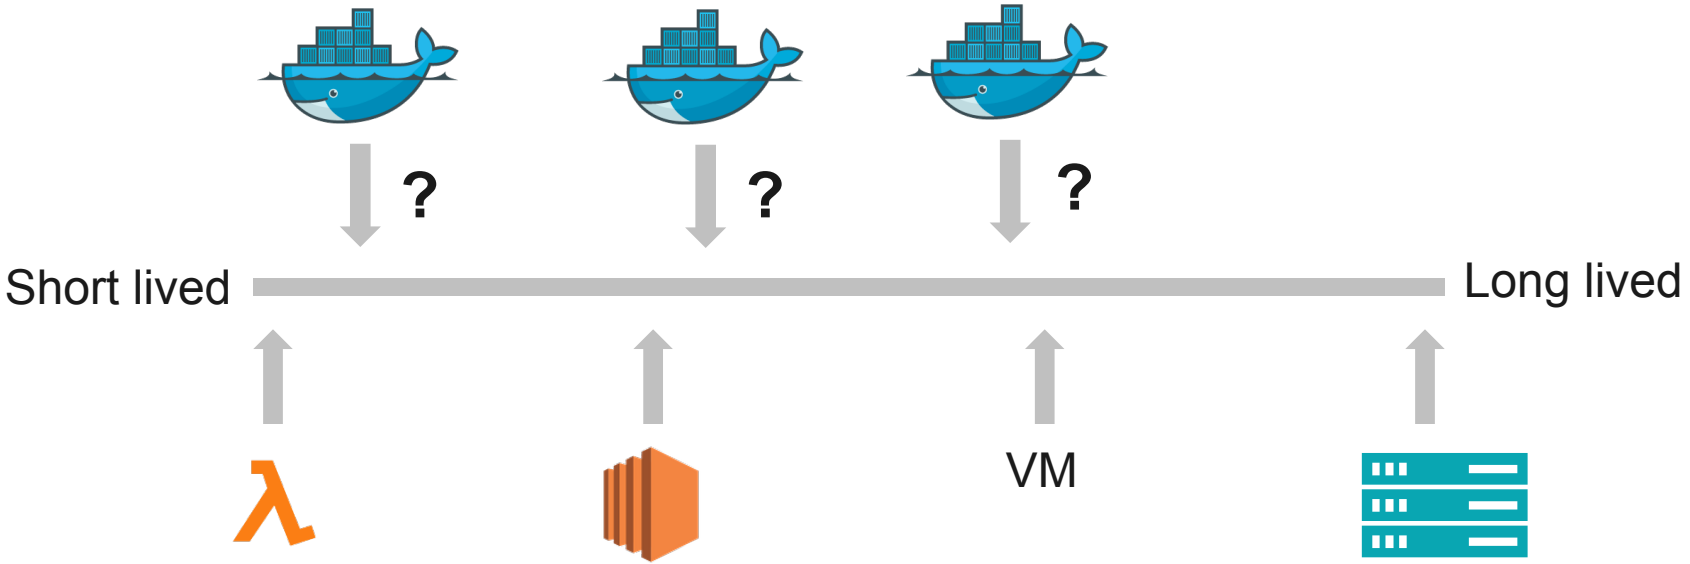

1: Hey we're monitoring it!

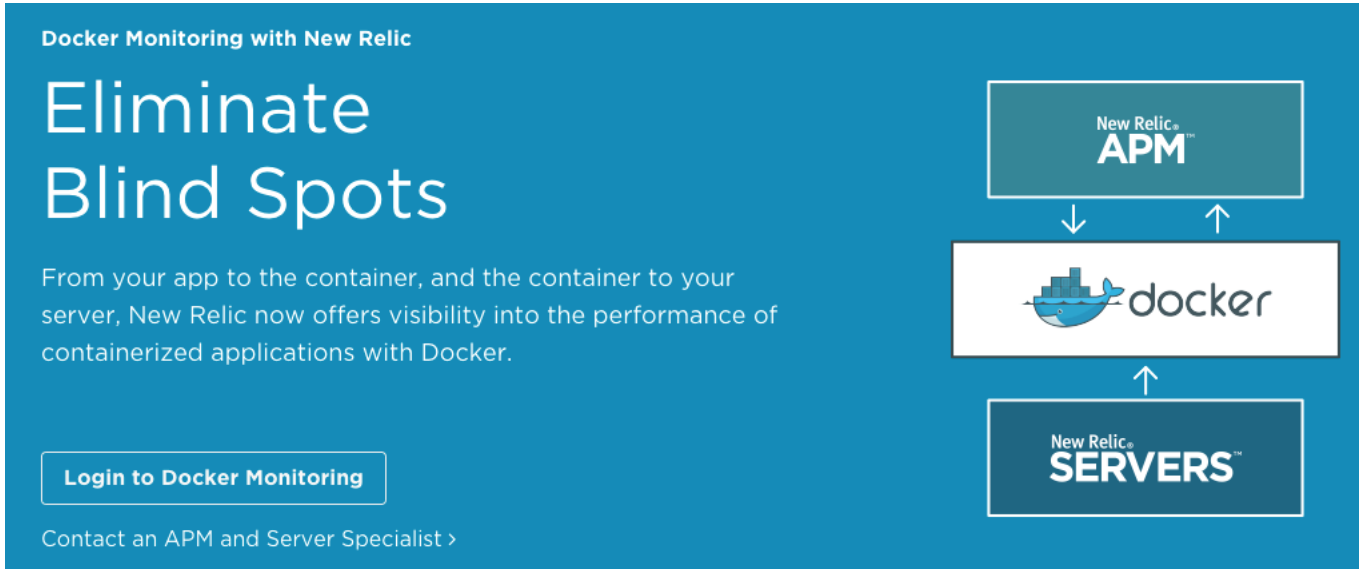

2: We're Data Nerds!

New Relic.

Data helps me build the right product for my customers.

We are *all* data nerds.

Kevin McGuire
Director of Engineering
New Relic

Data set analyzed

TOTAL CONTAINERS

8+ million

CUSTOMERS

Approx. 1000

AVERAGE 24 HR CONTAINERS

300,000+

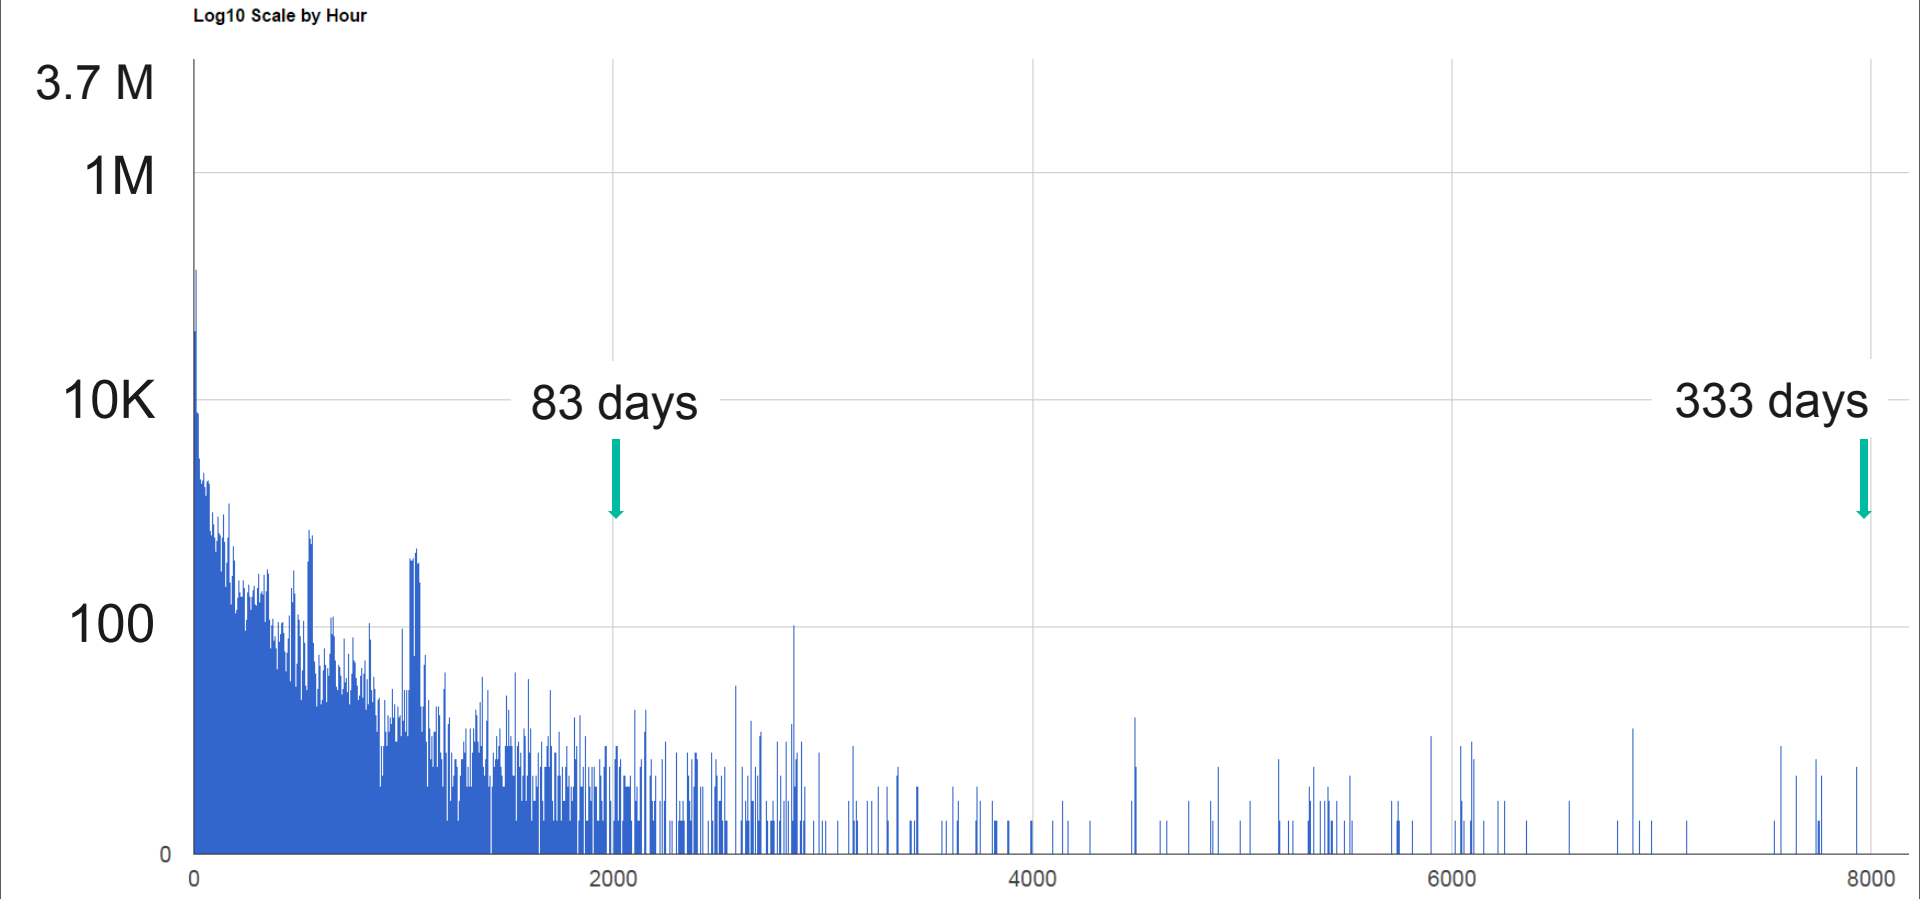
Log10 Scale by Hour
3.7 M
1M
10K
100
0
83 days
333 days
0
2000
4000
6000
8000

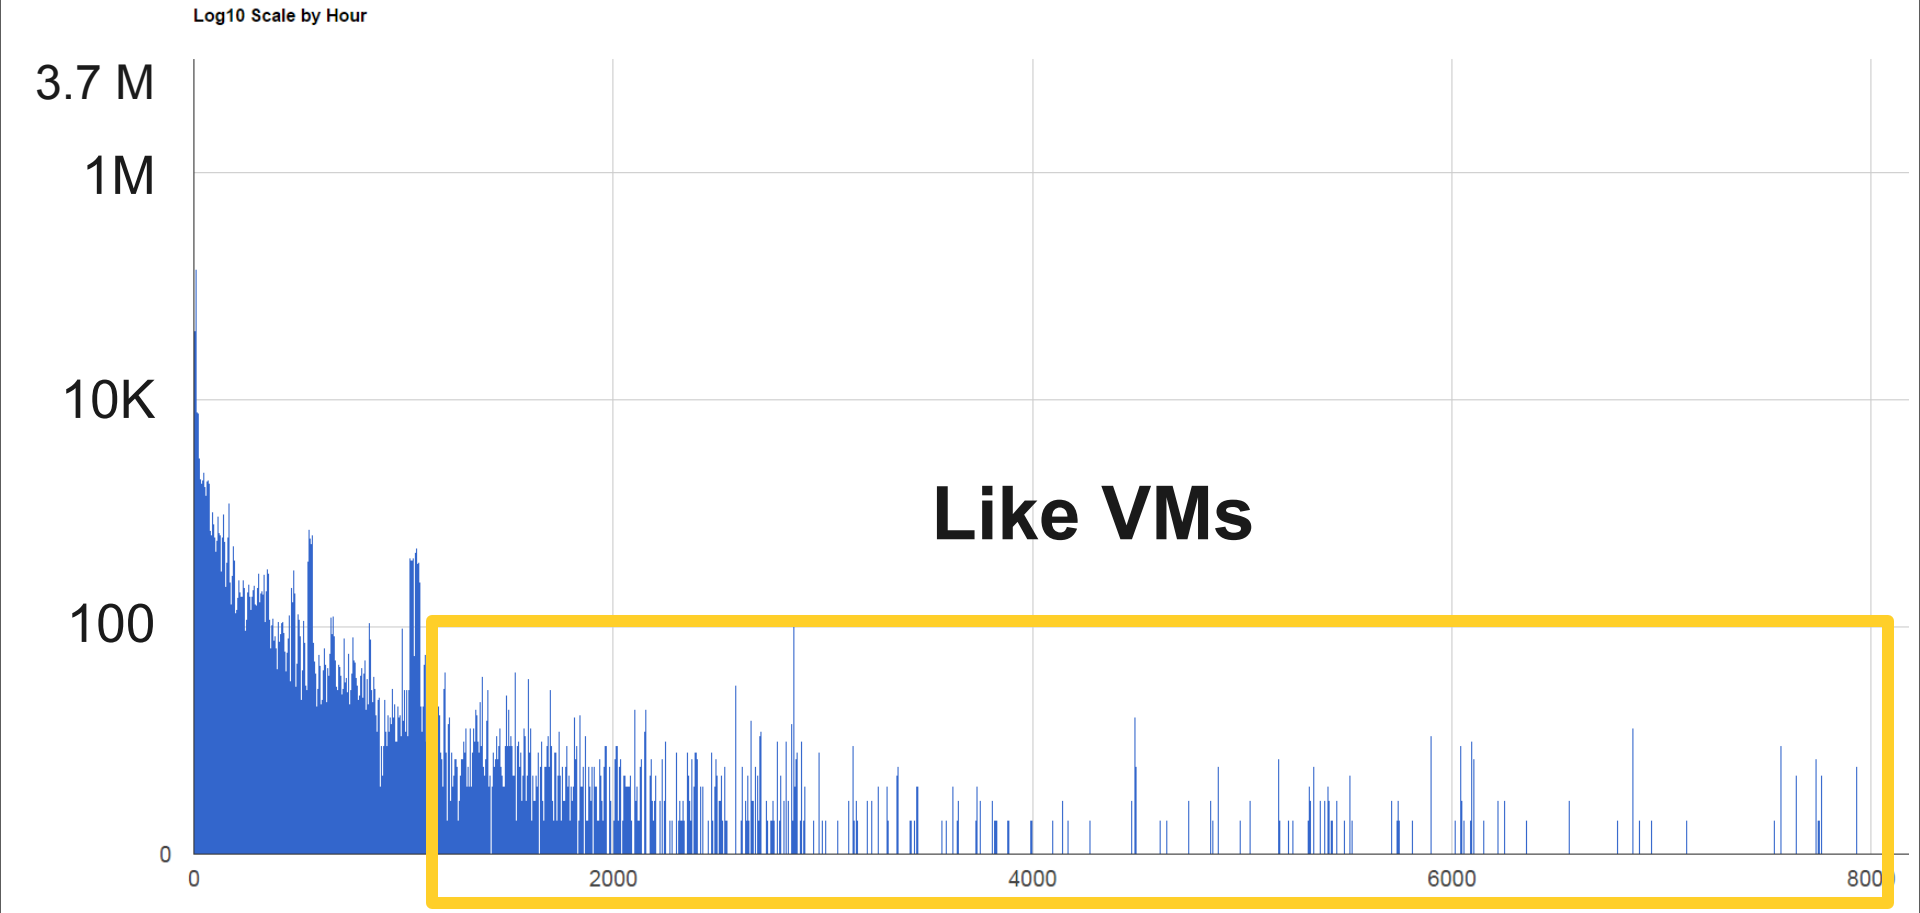
Log10 Scale by Hour
3.7 M
1M
10K
100
0
0
2000
4000
6000
800
Like VMs

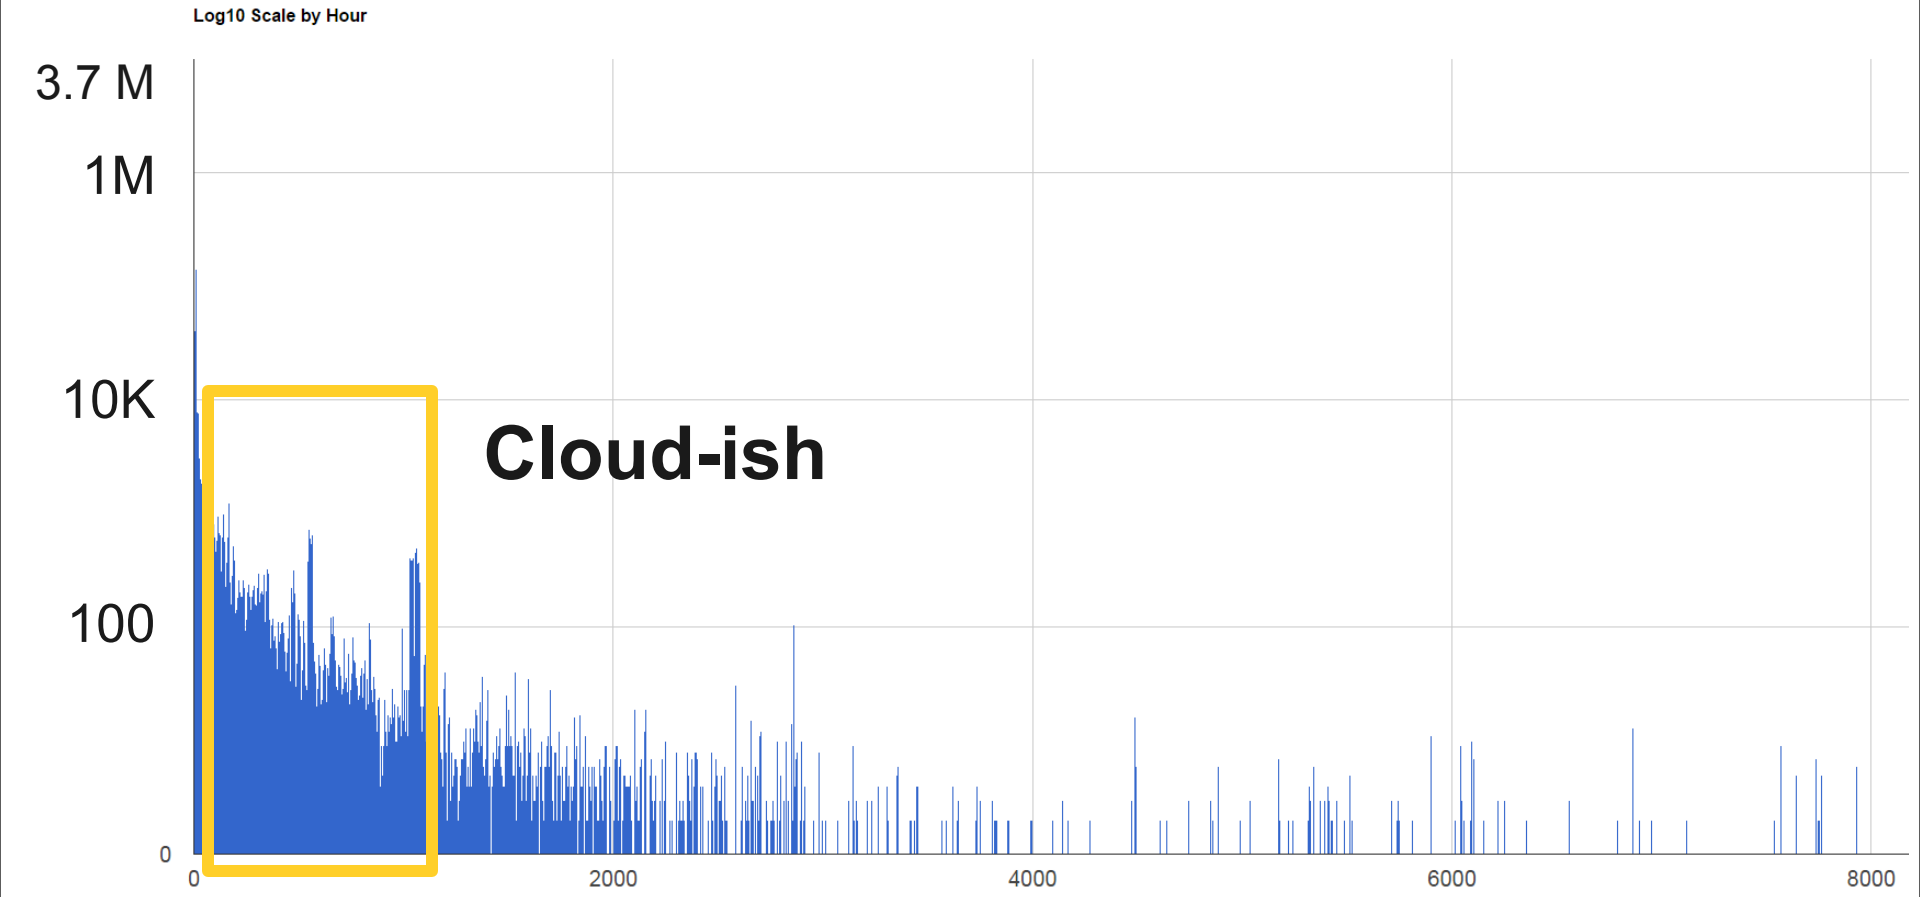
Log10 Scale by Hour
3.7 M
1M
10K
100
0
0
2000
4000
6000
8000
Cloud-ish

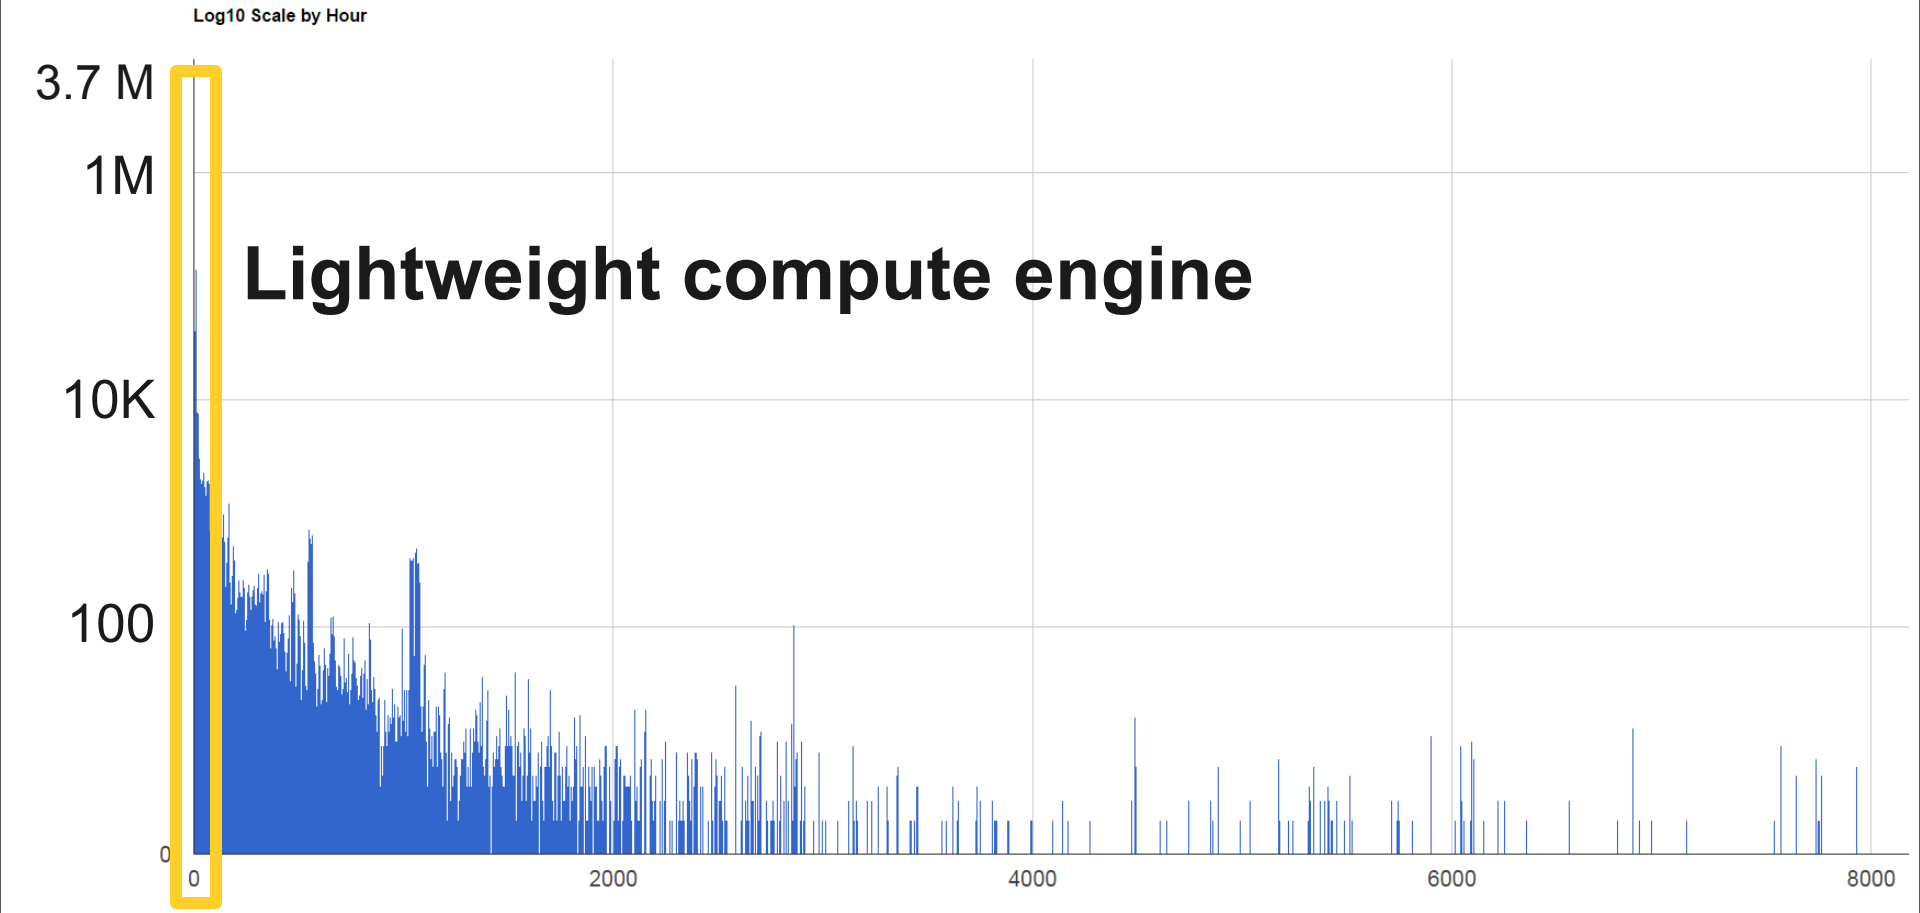

Lightweight compute engine

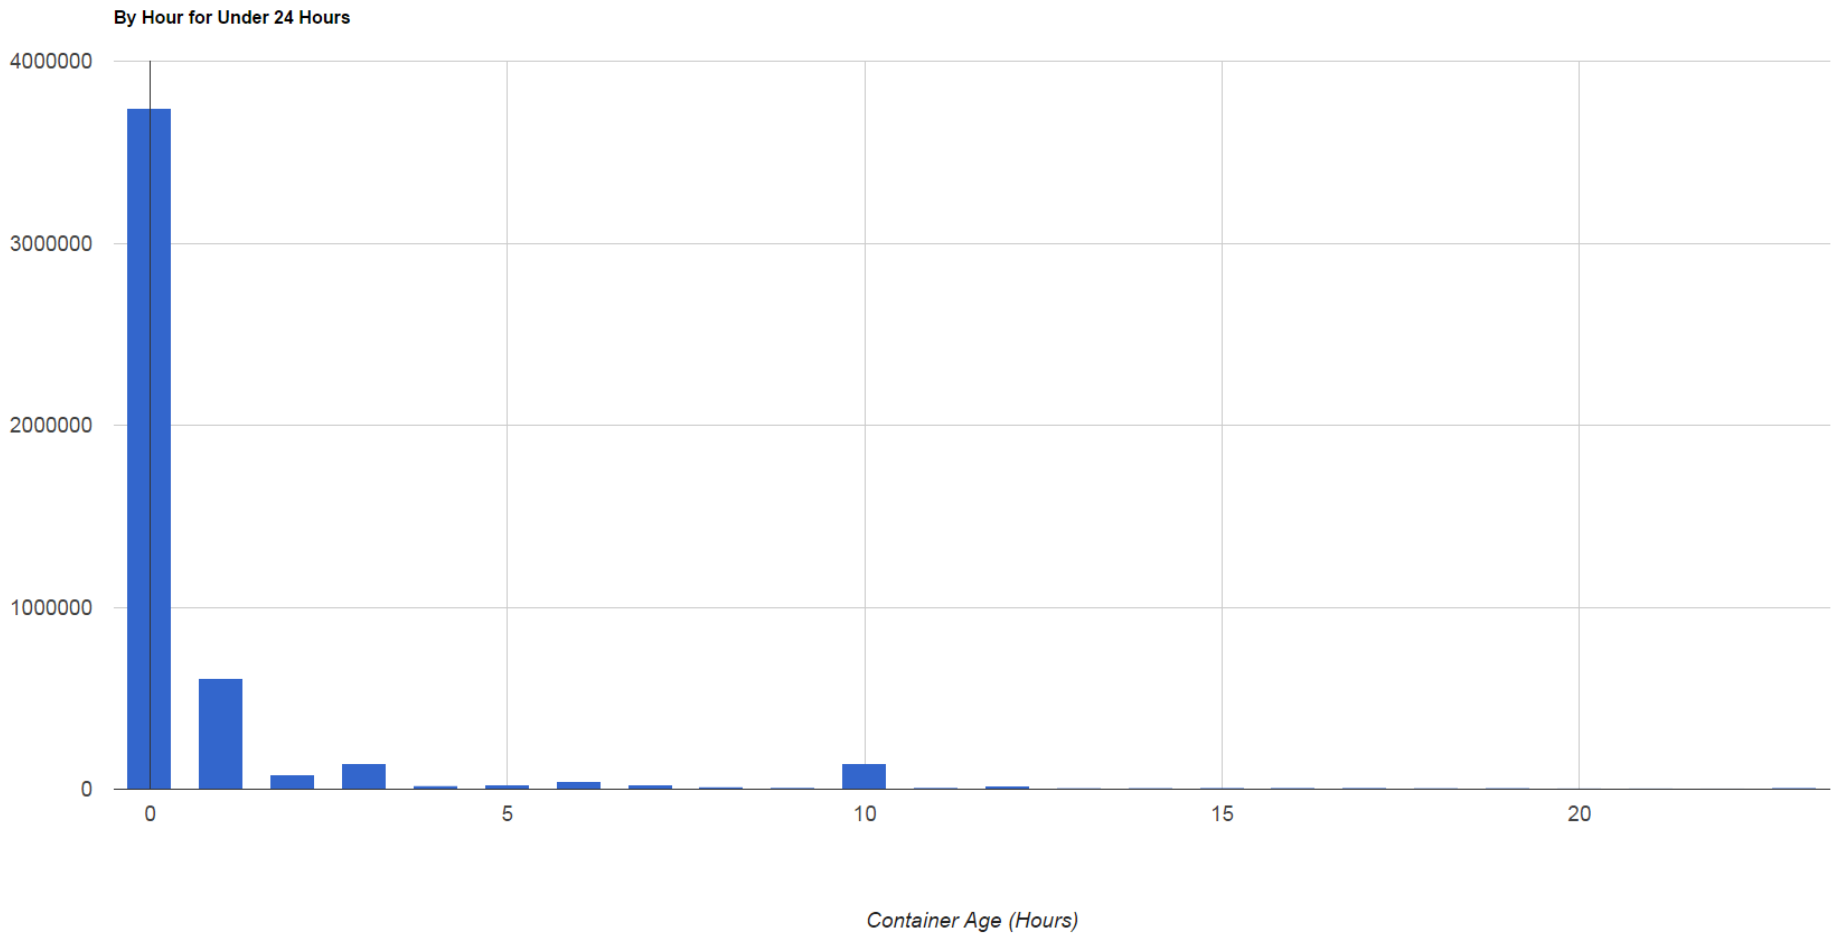
By Hour for Under 24 Hours
4000000
3000000
2000000
1000000
0
0
5
10
15
20
Container Age (Hours)

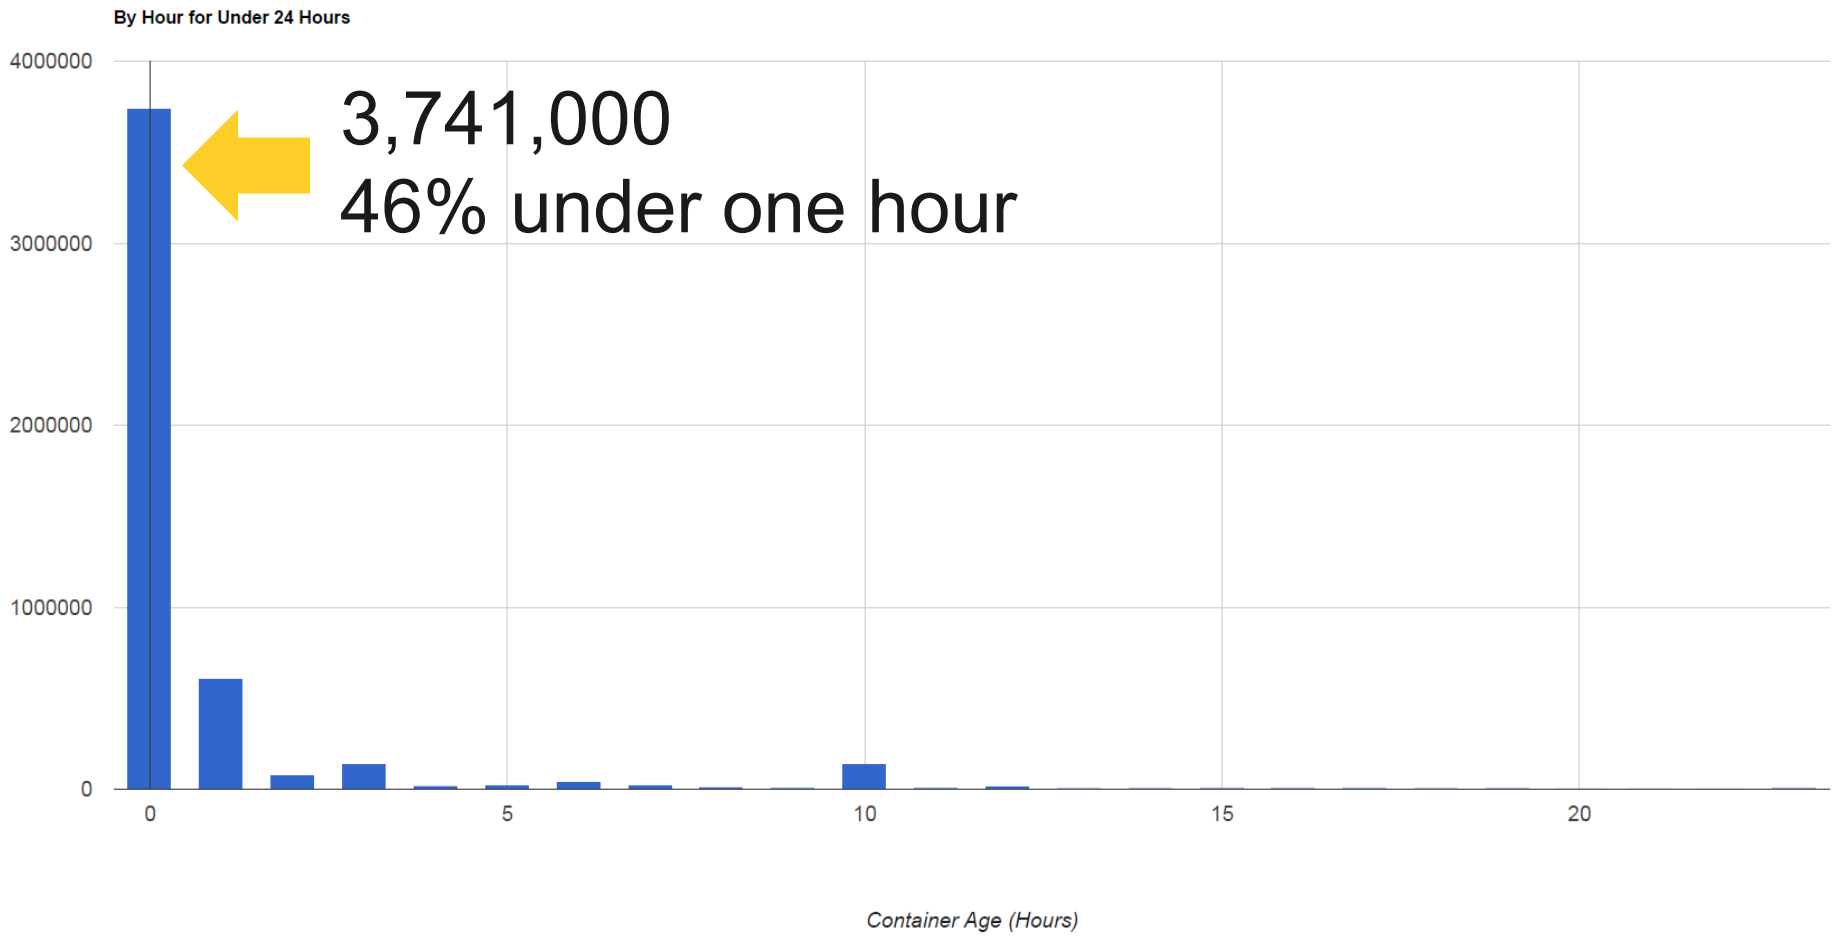
By Hour for Under 24 Hours
4000000
3000000
2000000
1000000
0
0
5
10
15
20
Container Age (Hours)
3,741,000
46% under one hour

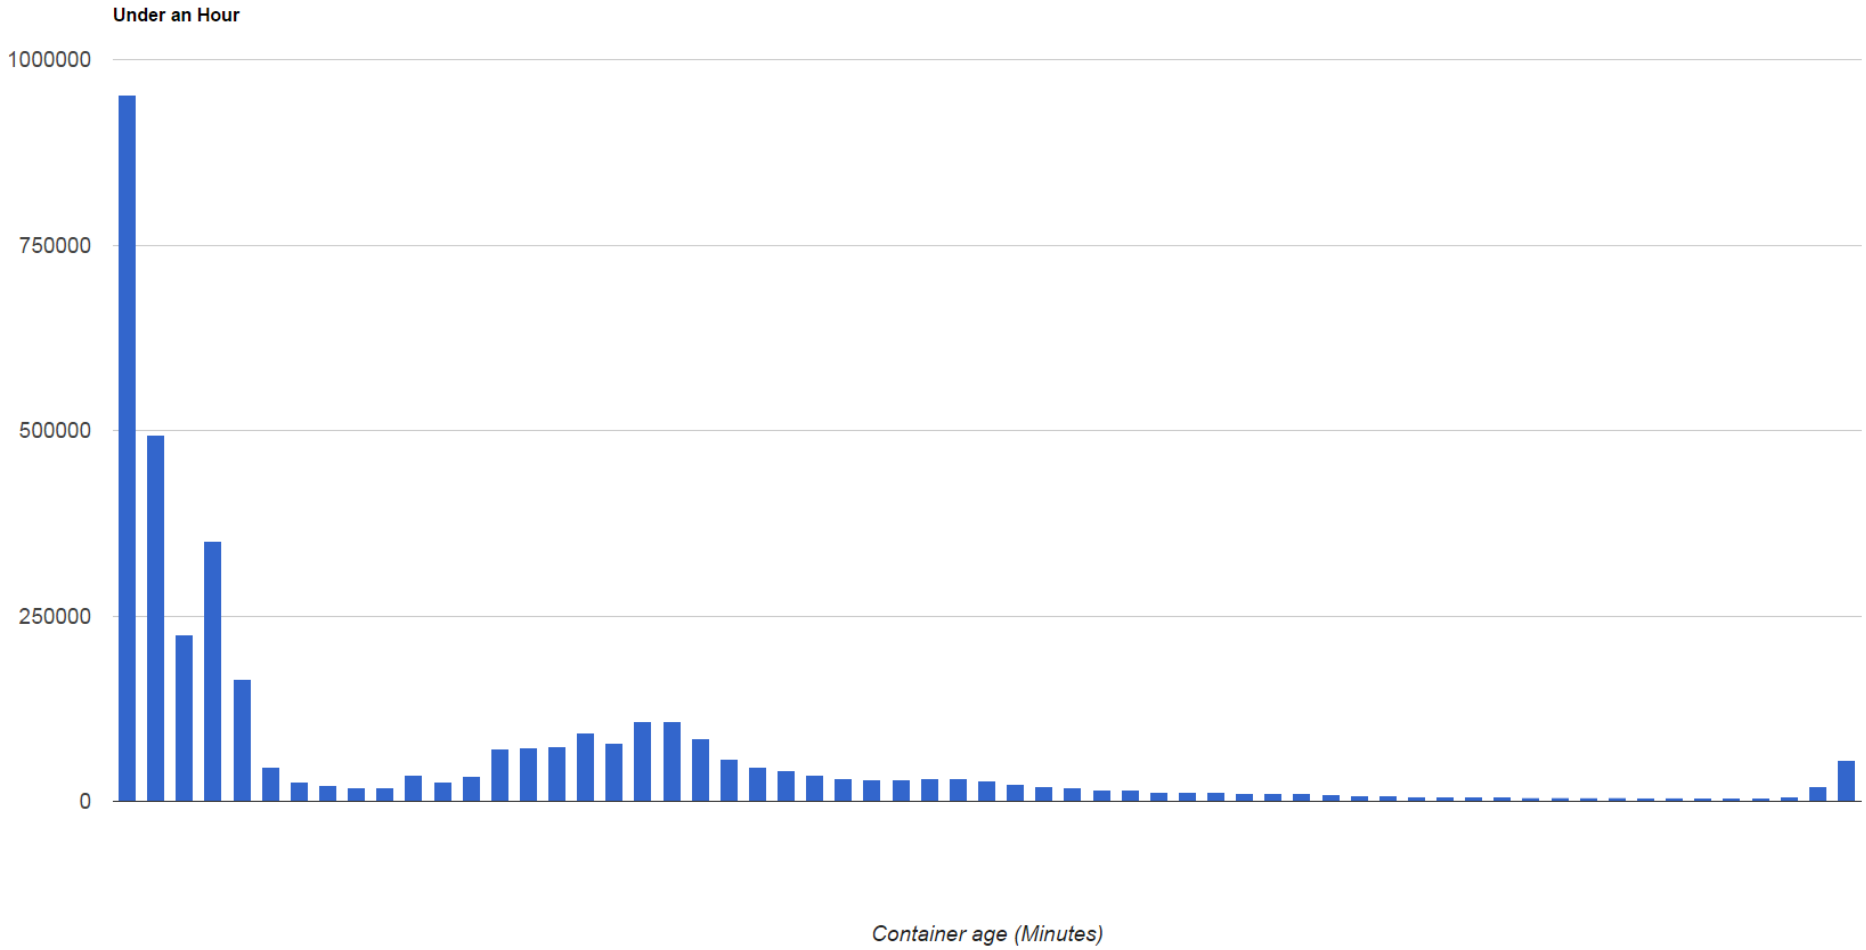
Under an Hour
1000000
750000
500000
250000
0
Container age (Minutes)

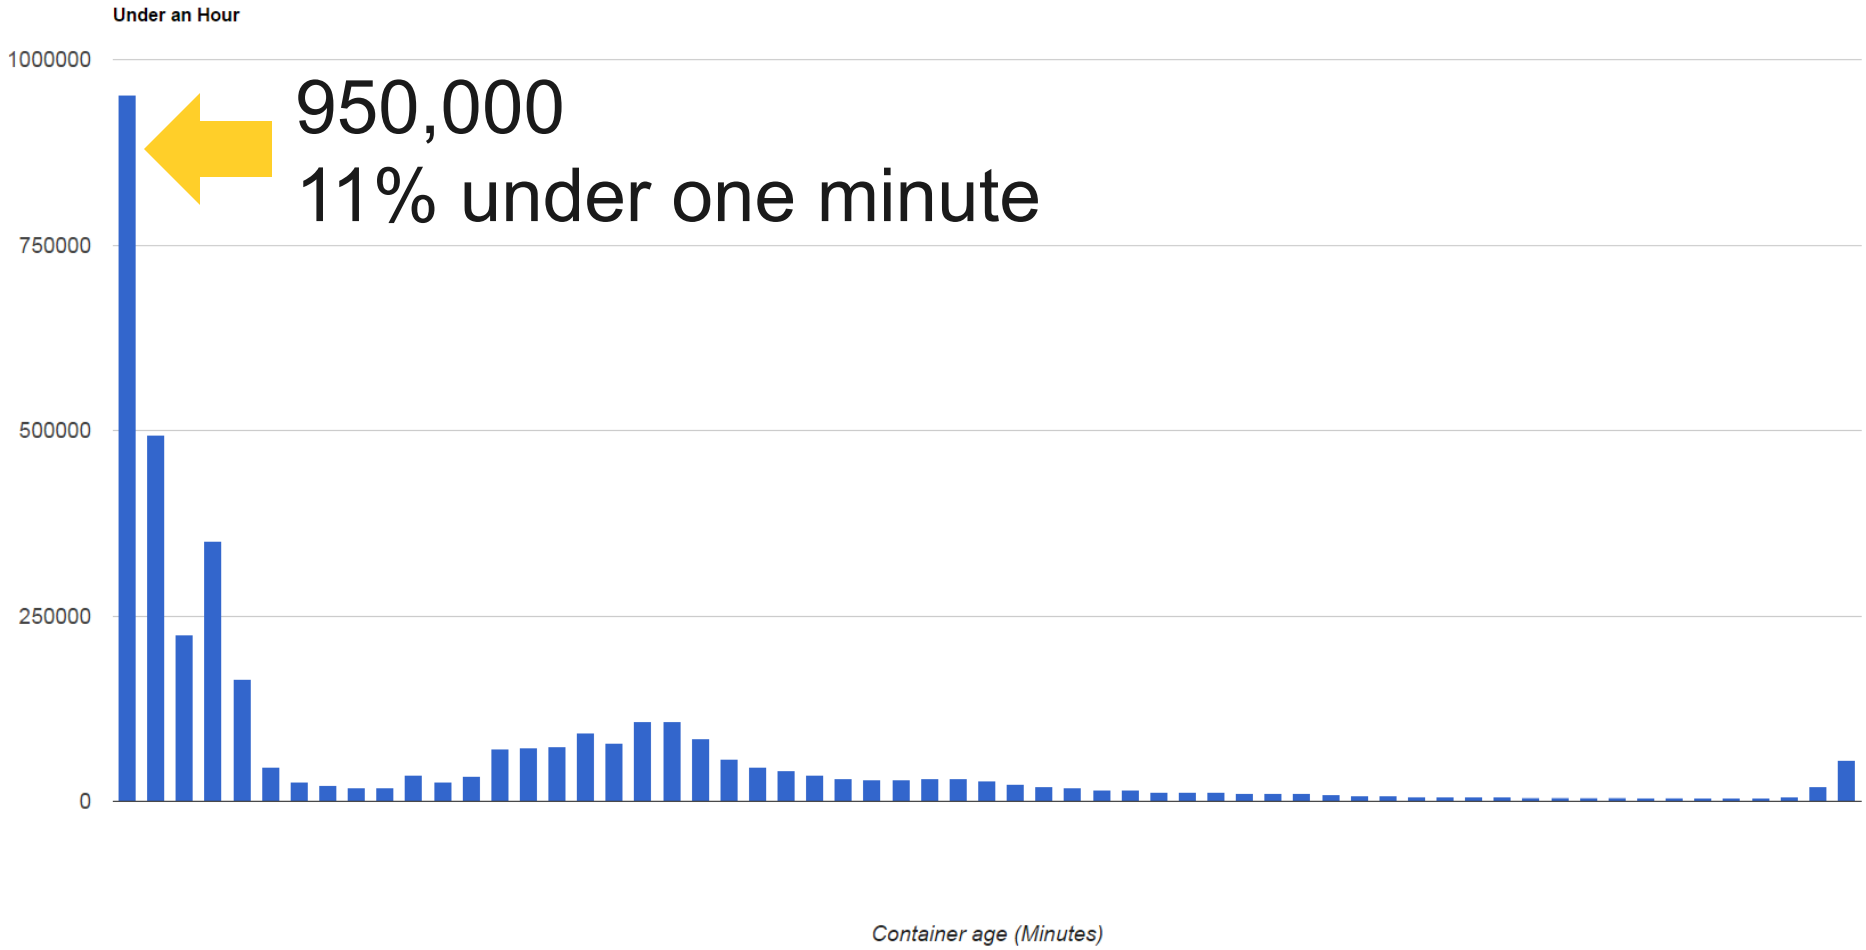
Under an Hour
1000000
750000
500000
250000
0
950,000
11% under one minute
Container age (Minutes)

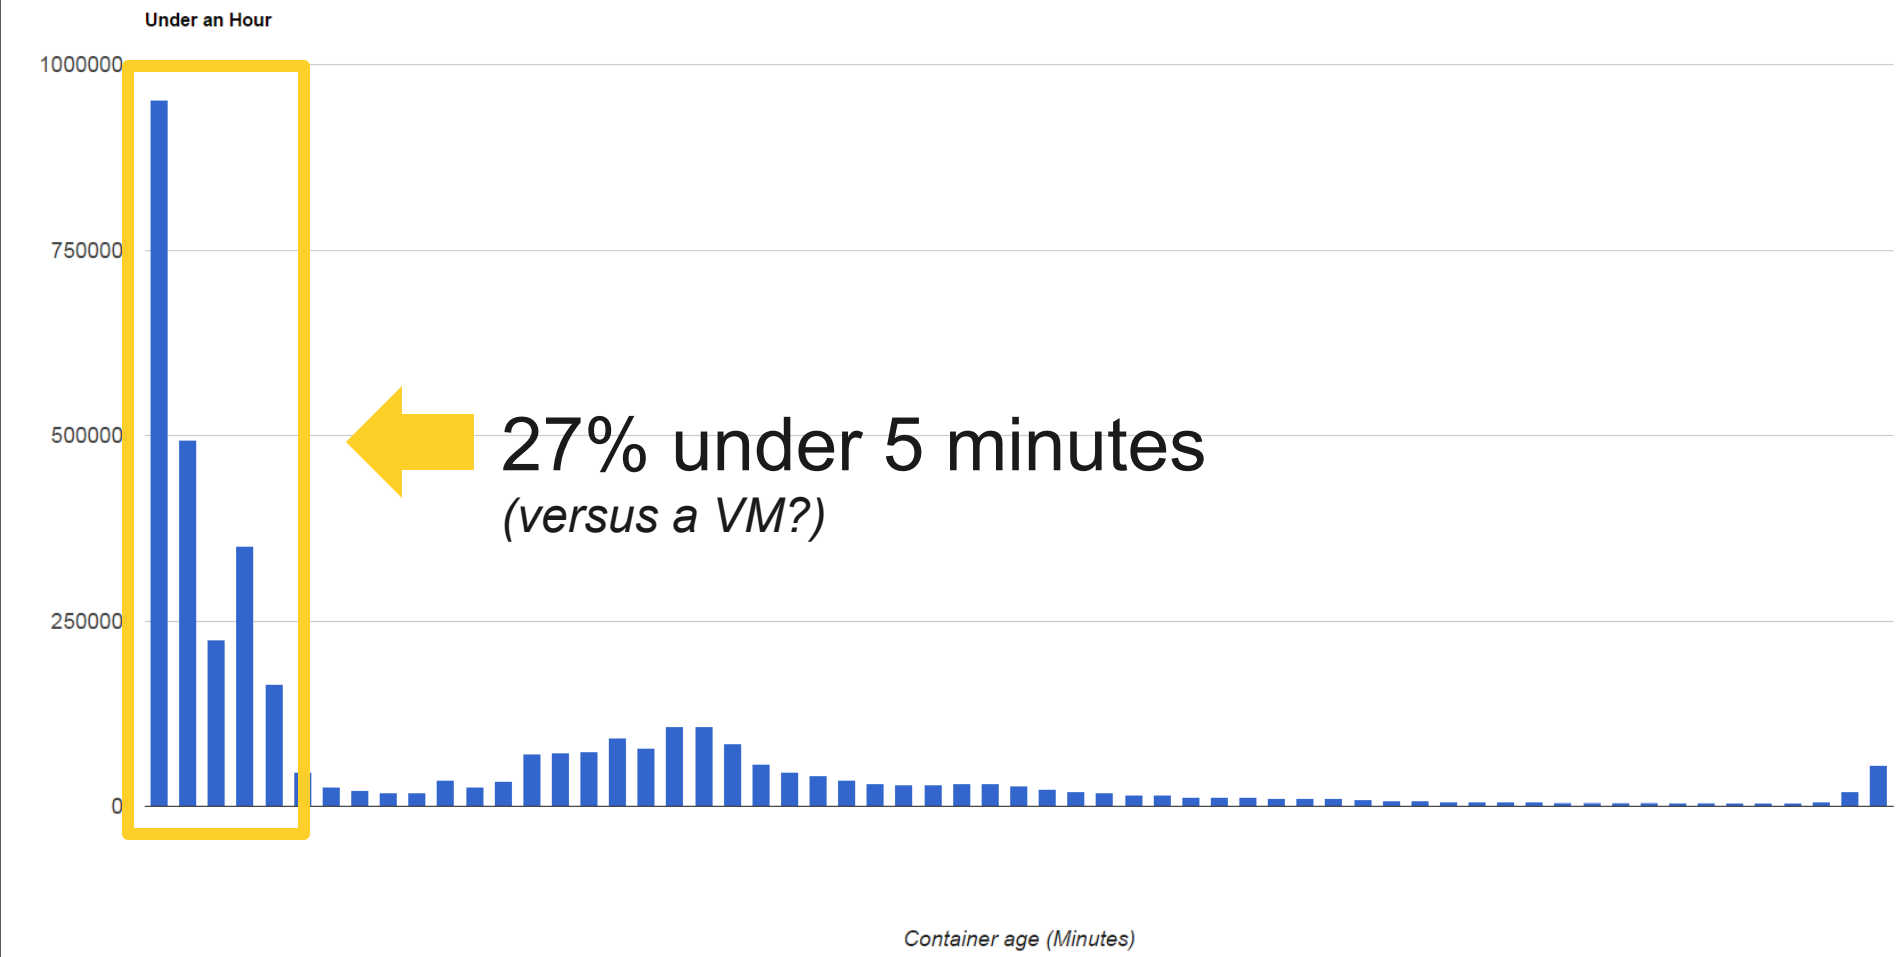
Under an Hour
1000000
750000
500000
250000
0
27% under 5 minutes
(versus a VM?)
Container age (Minutes)

June versus now: 5x data, same shape

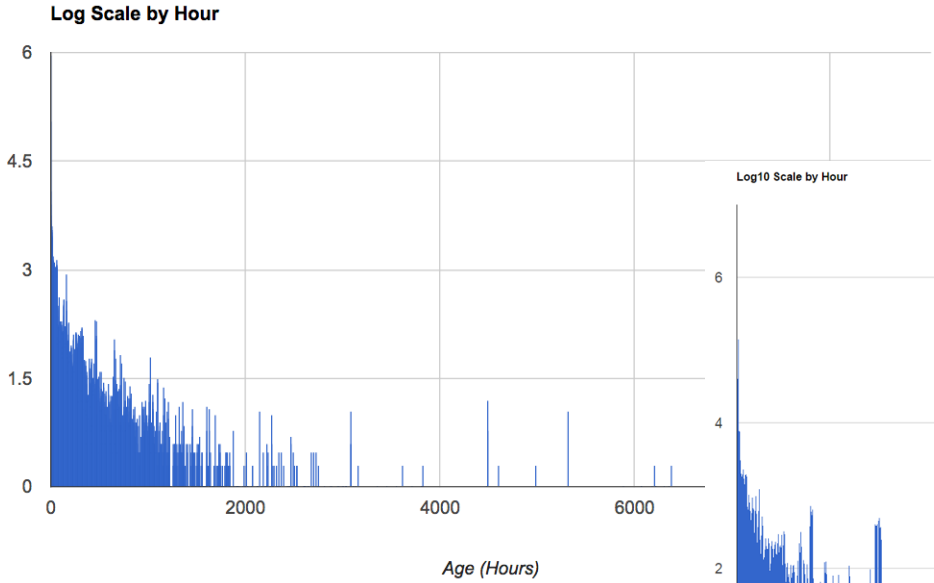

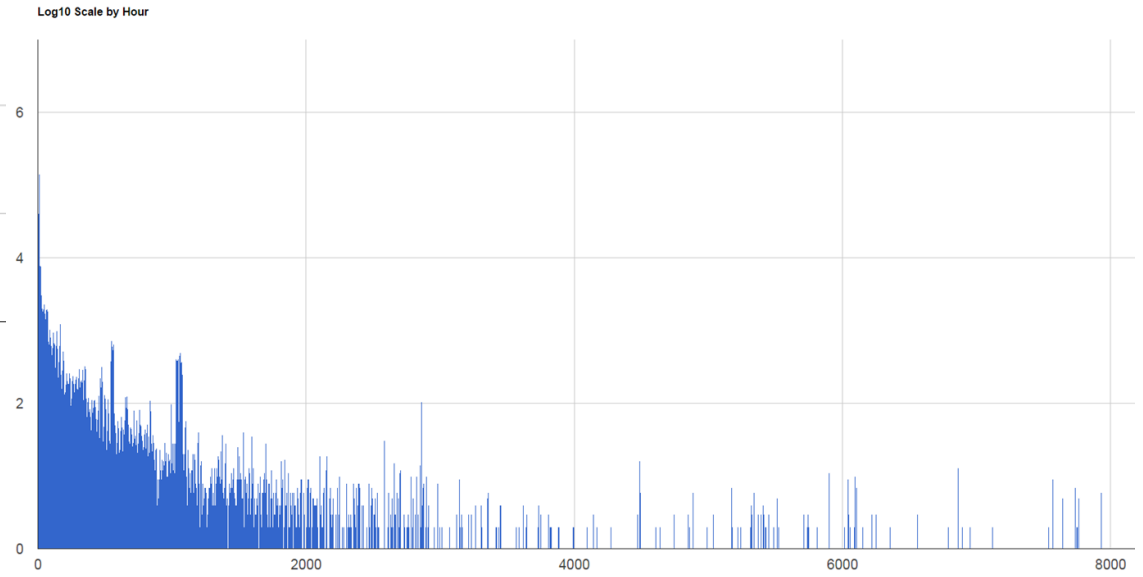

A surprising result

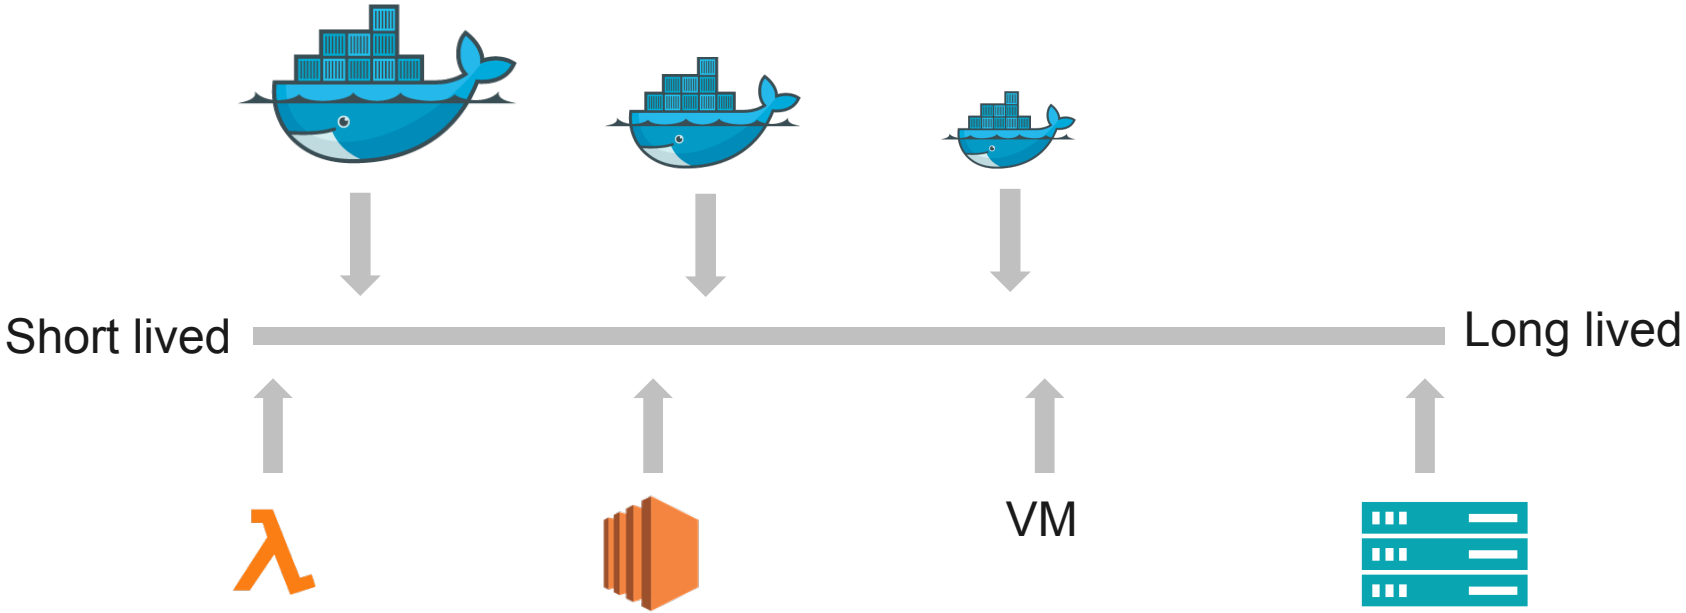

Why this matters

1. Usage is evolving in (fascinating) unexpected ways
2. Single technology that can span such wide usage is a game changer
3. Monitoring tools need to fit the lifecycle

Takeaways

- Same technology spans a *huge* usage scope
- What explains all this?
  - Batch jobs?
  - “Microservices”?
- The missing metric: computational work
  - There will be a lot more short lived *anything*

The evolution of computation as a service

- Short startup time (orders mag.) allows very short lived computing
  - Containers are created
  - Do their work
  - Go away

- Containers only exist, and only for as long, as they provide value. Full stop.

Implications of computation as a service?

- What does it mean to network them together?
- What does it mean for orchestration of work?
- What does it mean for CI? Does it increase agility?
- What does it mean for provisioning, load balance, availability?
- How do we know what they're doing?
  (And, what is "what"?)

Monitoring servers

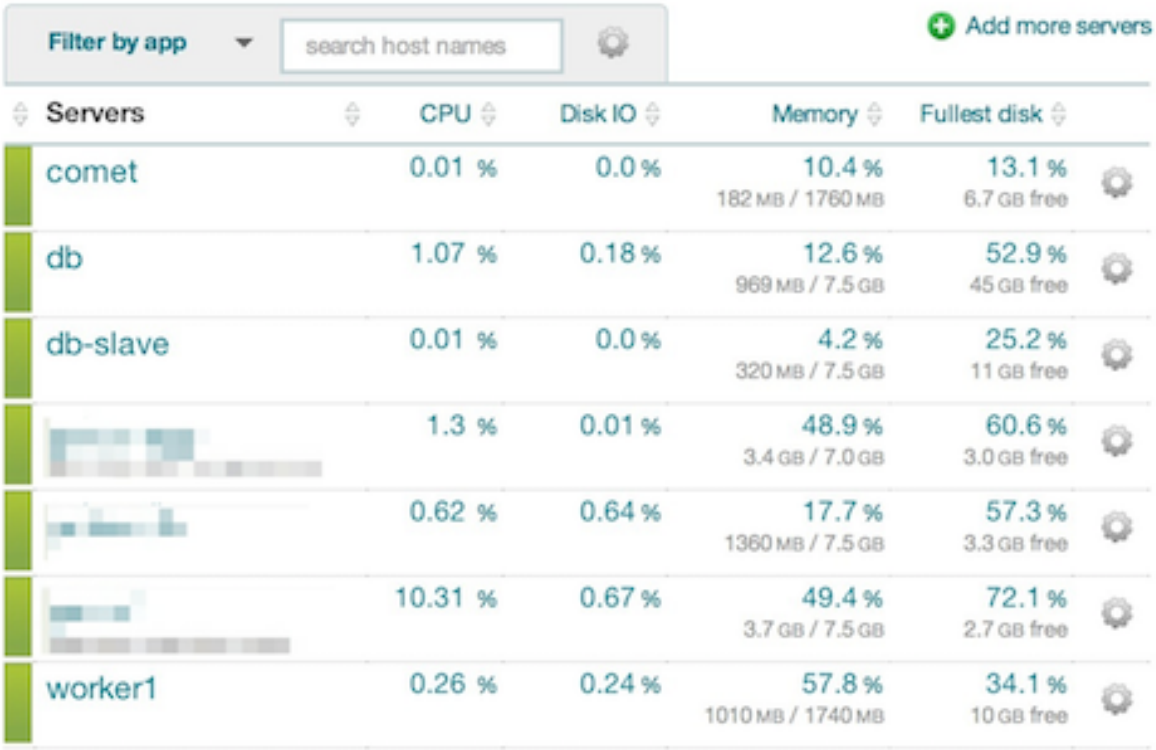

Monitoring computation

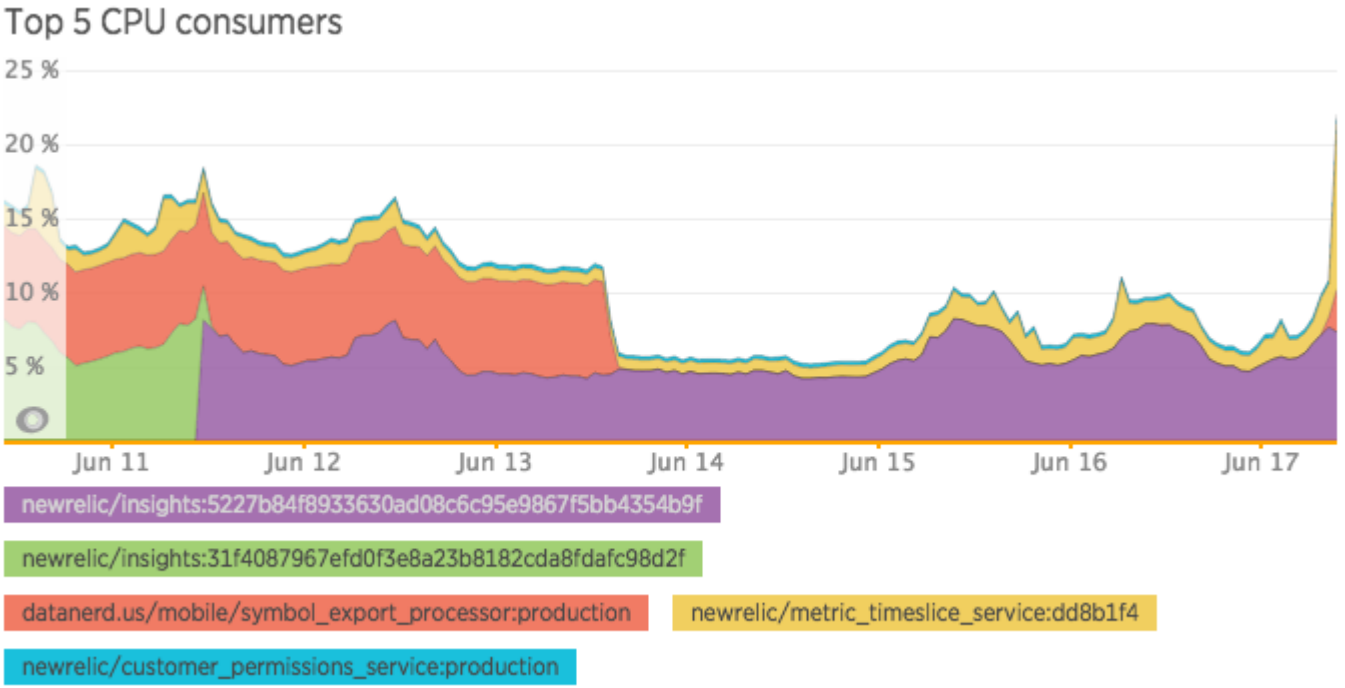

Thank you

Kevin McGuire
kevin@newrelic.com
@KevinMcGUI