

MONITORING K8S: FOLLOW THE DATA

DANIEL MAHER, DATADOG

@phrawzty

1. TRENDS IN CONTAINERISATION

2. THE QUALITIES OF GOOD METRICS

3. SCALING KUBERNETES (W/ METRICS)



DANIEL MAHER

DOCS & TALKS

DEVOPSDAYS

GREAT OUTDOORS

GOURMAND

SUITS FOR NO REASON

@phrawzty

DATADOG

SAAS-BASED MONITORING

TRILLIONS OF POINTS/DAY

WE'RE HIRING:

www.datadoghq.com/careers

@datadoghq



DATADOG

Docker support #844

 Merged

remh merged 14 commits into `DataDog:master` from `steeve:master` on Mar 8, 2014

 Conversation **23**

 Commits **14**

 Files changed **2**



steeve commented on Feb 25, 2014

This PR enables instrumentation on Docker (LXC/cgroups) containers for:

- Memory
- CPU
- Disk

It uses the native Unix socket to communicate with `dockerd`, and thus doesn't need any

Updated June 2018

8 SURPRISING FACTS ABOUT REAL
DOCKER
ADOPTION

<https://www.datadoghq.com/docker-adoption/>



8 EMERGING TRENDS IN CONTAINER ORCHESTRATION

<https://www.datadoghq.com/container-orchestration/>

Docker Adoption

Dabbler: used Docker during a given month, but hadn't reached significant use as defined by Adopter.

Docker Adoption

Dabbler: used Docker during a given month, but hadn't reached significant use as defined by Adopter.

Adopter: the average number of containers running during the month was at least 50% the number of distinct hosts run, or there were at least as many distinct containers as distinct hosts run during the month.


Docker Adoption

Dabbler: used Docker during a given month, but hadn't reached significant use as defined by Adopter.

Adopter: the average number of containers running during the month was at least 50% the number of distinct hosts run, or there were at least as many distinct containers as distinct hosts run during the month.

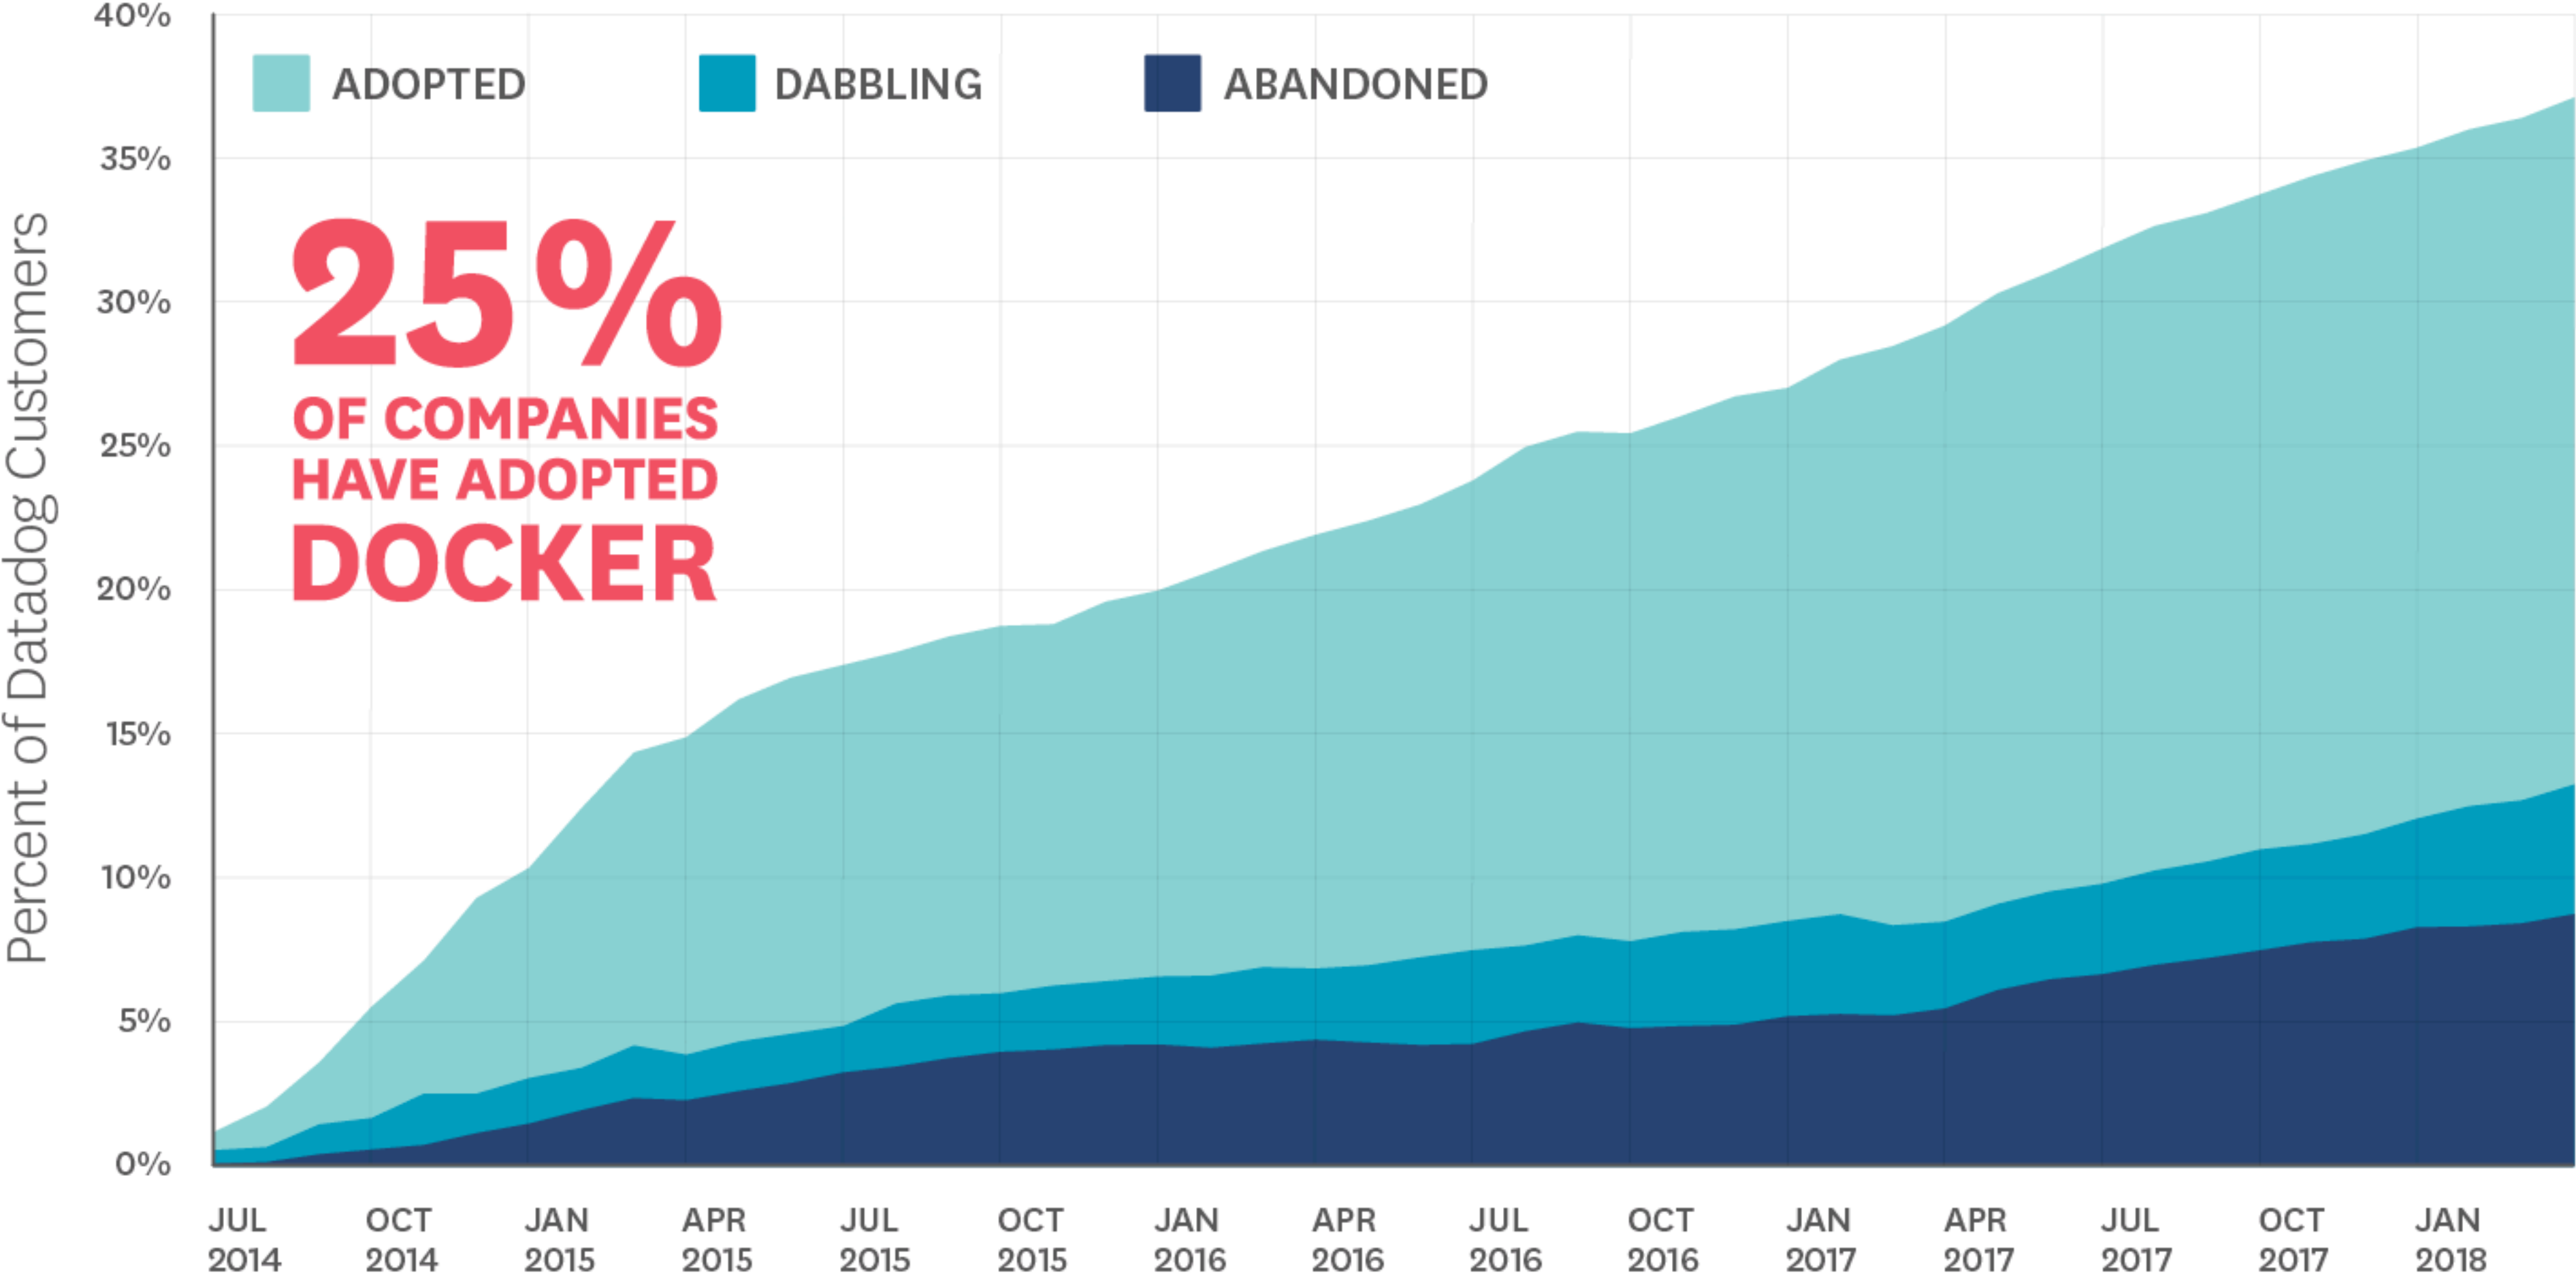
Abandoner: a currently active company that used Docker in the past, but hasn't used it at all in the last month.

Docker Adoption Behavior



Source: Datadog


Docker Adoption Behavior



Month (segmentation based on end-of-month snapshot)


Source: Datadog


Portion of Hosts Running Docker



Source: Datadog


WHAT'S RUNNING?






ENTER: ORCHESTRATION

Usage of Orchestrators




Source: Datadog

ECS PREVAILS IN AWS BUT KUBERNETES LEADS ELSEWHERE



Source: Datadog


Kubernetes share of Docker environments



Source: Datadog


NEW CONTAINER: WHO DIS?

Average Container Lifetimes



Source: Datadog

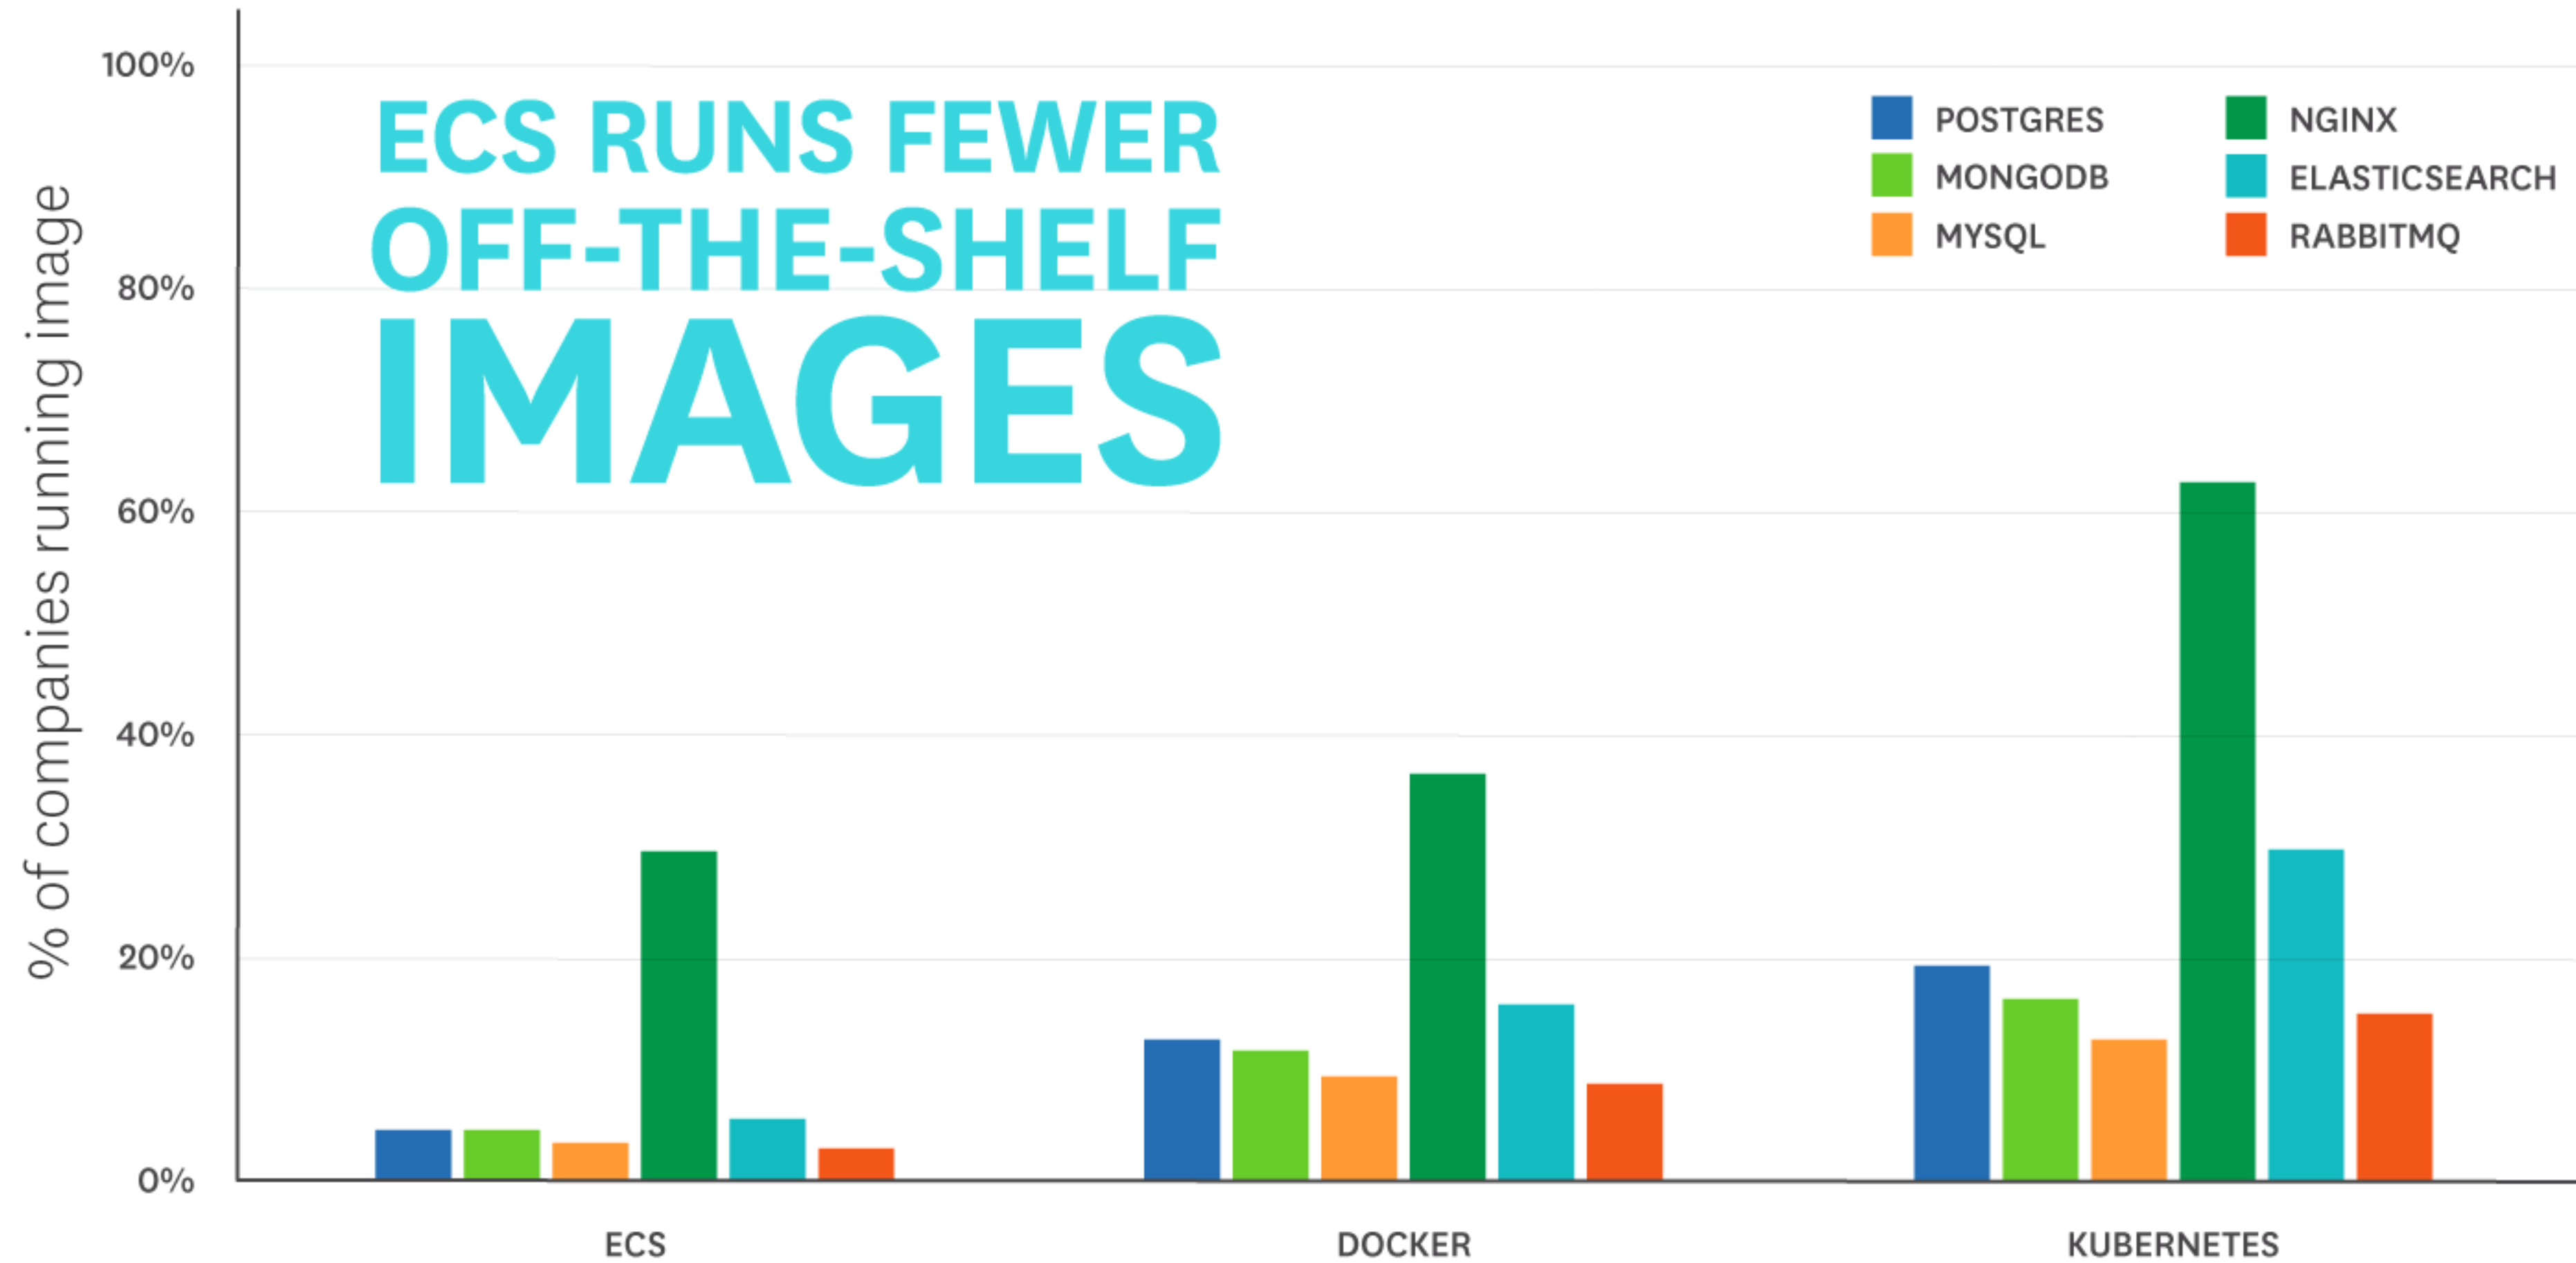
Average lifetimes of hosts



WHAT'S RUNNING?

PART 2: THE RUNNINATING


Common container images by orchestrator



Source: Datadog

Companies using tagged image versions

MOST COMPANIES MIX "LATEST" AND VERSIONED IMAGES



Source: Datadog

4 QUALITIES OF GOOD METRICS









NOT ALL METRICS ARE EQUAL

A large, bright orange and yellow fireball or explosion is shown against a dark, starry background. The fireball is the central focus, with a bright yellow core and a surrounding orange and red glow. The background is black with numerous small white stars. The text "1. MUST BE WELL UNDERSTOOD" is overlaid in white, bold, sans-serif font in the upper right quadrant.

1. MUST BE WELL UNDERSTOOD



2. SUFFICIENT GRANULARITY

RANK	PARTICIPANT	RESULT
G	 Anthony ERVIN USA	21.40
S	 Florent MANAUDOU FRA	21.41
B	 Nathan ADRIAN USA	21.49
4.	 Ben PROUD GBR	21.68
5.	 Andrii GOVOROV UKR	21.74
6.	 Bruno FRATUS BRA	21.79
6.	 Bradley Edward TANDY RSA	21.79
8.	 Simonas BILIS LTU	22.08




3. TAGGED & FILTERABLE

4. LONG-LIVED



CUSTOM METRICS & K8S



1. MUST BE WELL UNDERSTOOD

```
$ kubectl autoscale rs foo --min=5 --max=10 --cpu-percent=80
```

Questions:


What does cpu percent mean exactly?

Of user-space only? Does that even matter?


Is that on any one instance? Across the pod? Of a single core? Multi-core?

And what about spikes - do those count?


2. SUFFICIENT GRANULARITY



1 second
Peak 46%



1 minute
Peak 36%




5 minutes
Peak 12%

3. TAGGED & FILTERABLE

“Add more web servers when the frontend is getting busy.”

metric: nginx.net.request_per_s
scope: kube_container_name: frontend

“oh no”



4. LONG-LIVED


“There always seems to be ten web servers running. Should we increase the maximum number in the set?”


Questions:

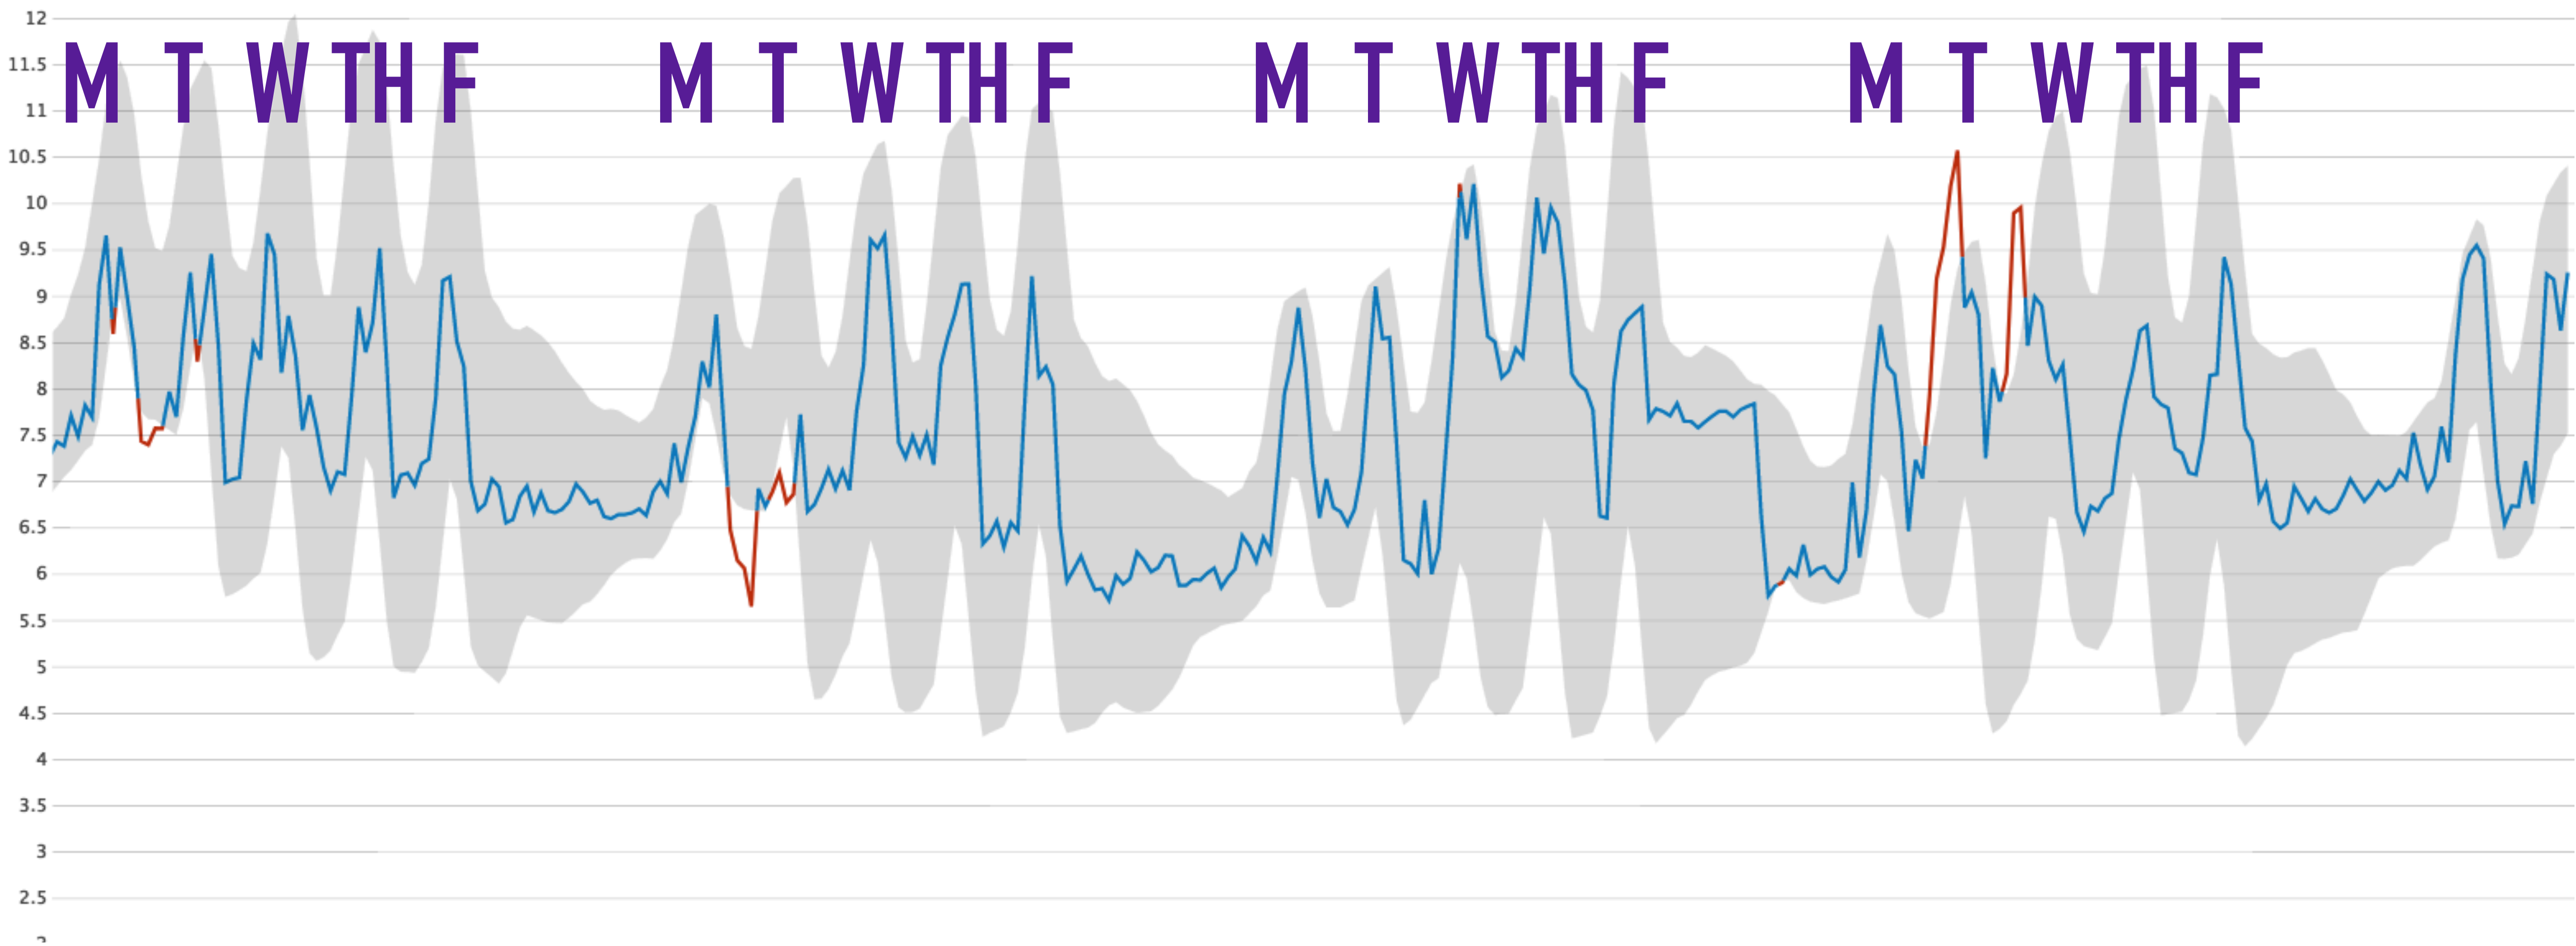
What does normal even look like?


Is normal stable over time? What about week-ends?

Did you just notice now? Seriously, how long has it been this way?









ALSO, WHAT HAPPENED HERE?

1. TRENDS IN CONTAINERISATION

2. THE QUALITIES OF GOOD METRICS

3. SCALING KUBERNETES (W/ METRICS)

MERCI !

DANIEL MAHER

@phrawzty

daniel.maher@datadoghq.com