

Switching horses midstream: the challenge of migrating 150+ services to kubernetes

Sarah Wells

Technical Director for Operations and Reliability, Financial Times

@sarahjwells







FINANCIAL TIMES

myFT

HOME WORLD UK COMPANIES MARKETS OPINION WORK & CAREERS LIFE & ARTS

Portfolio My Account

fastFT

Swiss franc heading for further declines - UBS 49M AGO

Govt appoints Credit Suisse banker to manage RBS stake ahead of sale 57M AGO

Au revoir, Arsène Wenger 1H AGO

BoE's Saunders: 'gradual' policy tightening needn't mean 'glacial' 1H AGO

Banks

Barclays CEO fined in whistleblowing probes


Jes Staley to keep job after seeking to uncover source of allegations against colleague

- Govt appoints Credit Suisse banker to manage RBS stake ahead of sale
- AMP chief quits amid allegations group overcharged clients
- China's bad loan bank shares tumble after chairman investigated



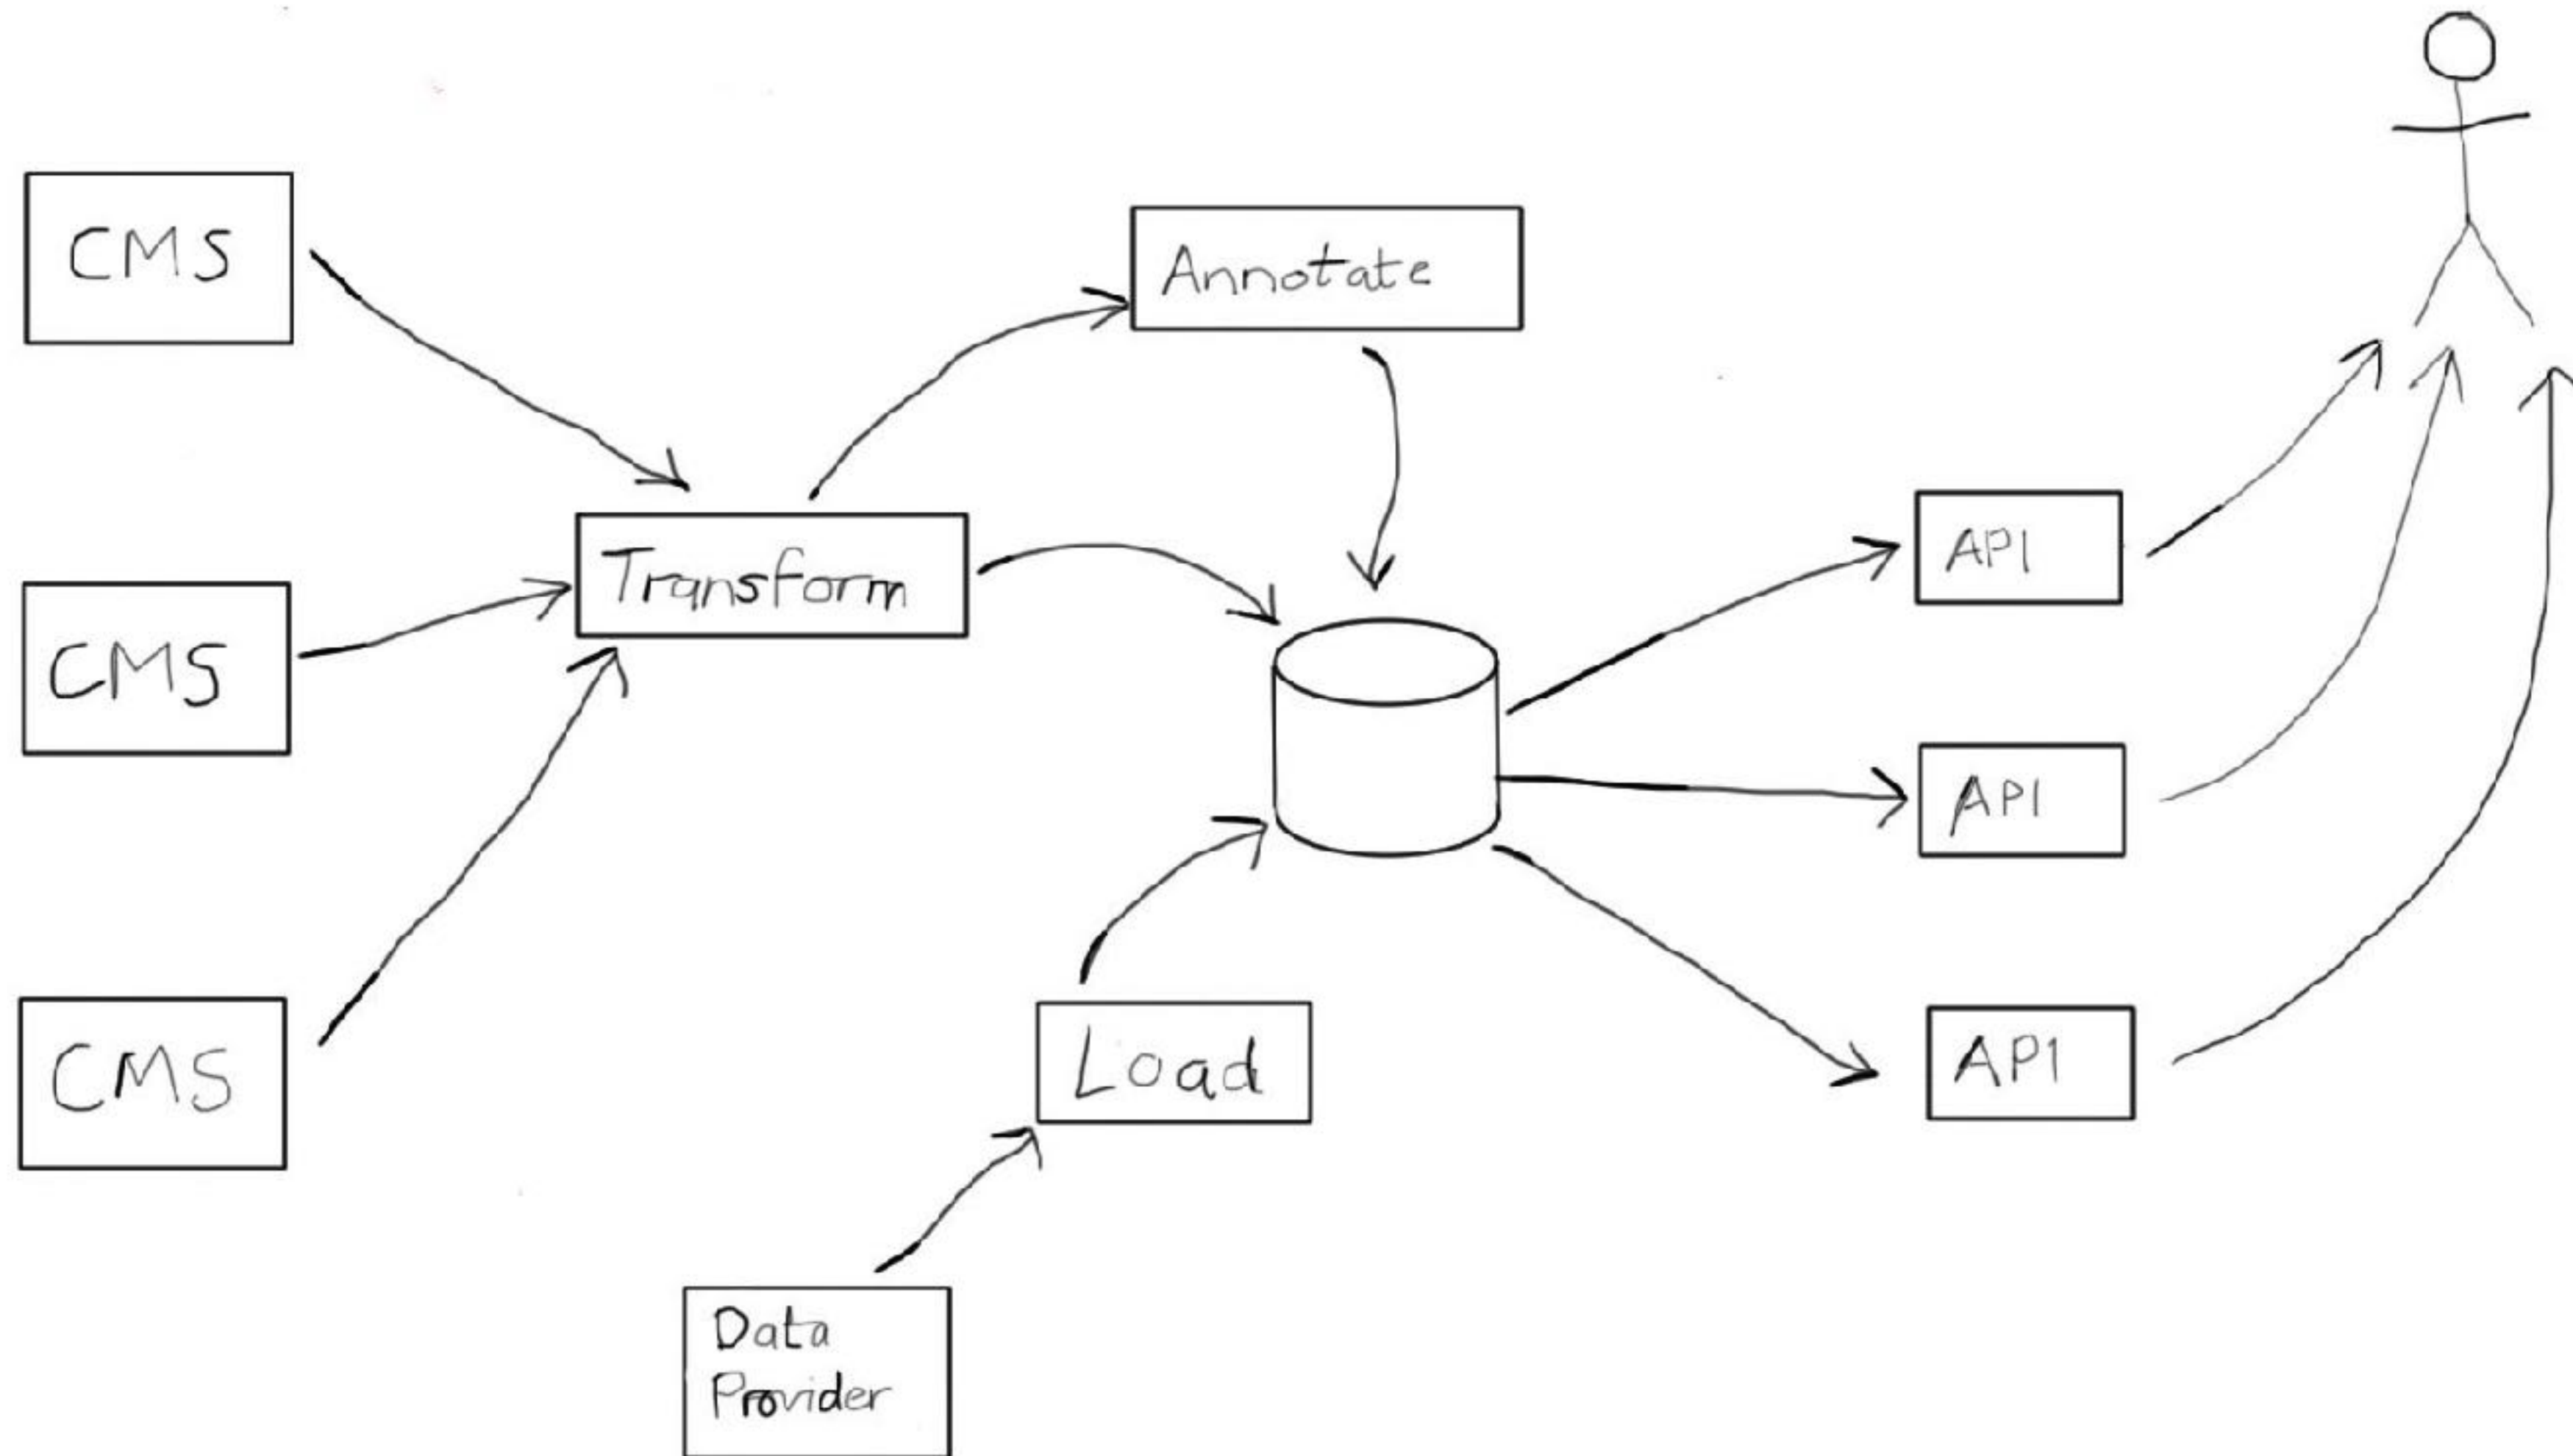
UK politics & policy

The roots of Windrush scandal: 'They saw enemies everywhere'

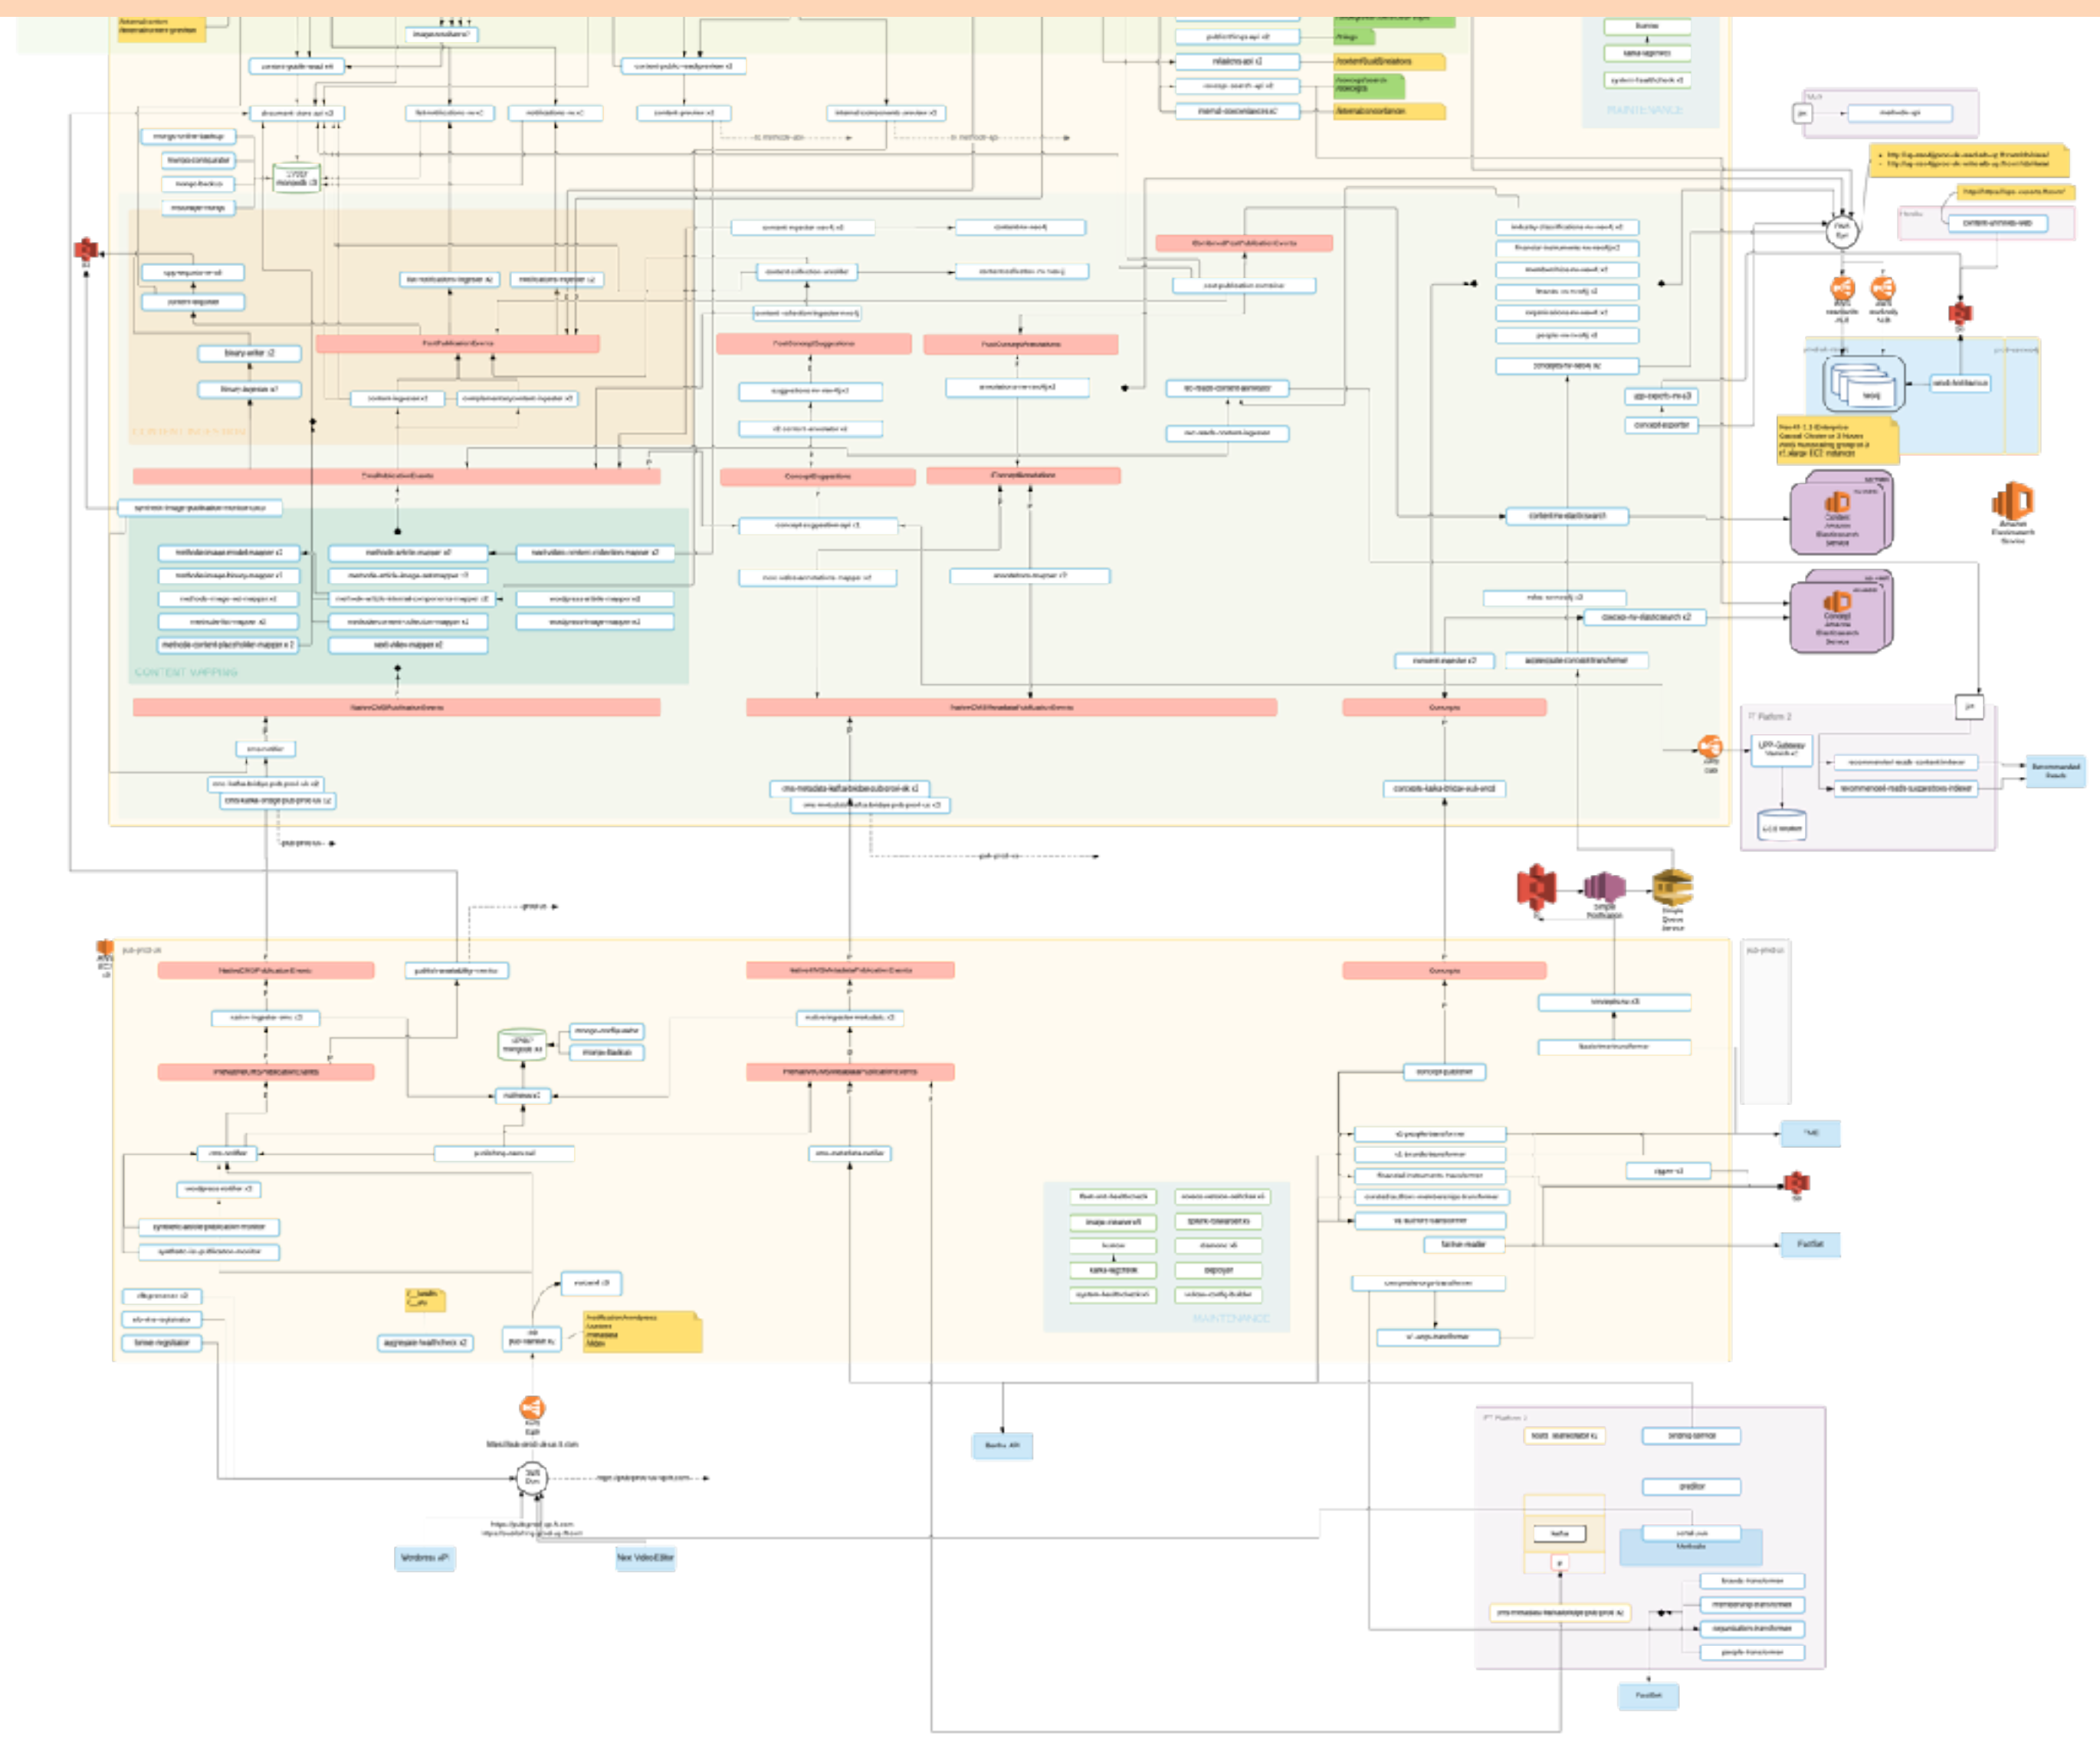




The FT's Content platform




This is what it really looks like...



Why **did** we migrate to k8s?

Mid 2015: a hand-rolled container stack



<https://medium.com/wardleymaps>

@sarahjwells

Spend your innovation tokens wisely

@sarahjwells

~80% reduction in EC2 costs

Many fewer steps to start running a new service
in production

@sarahjwells

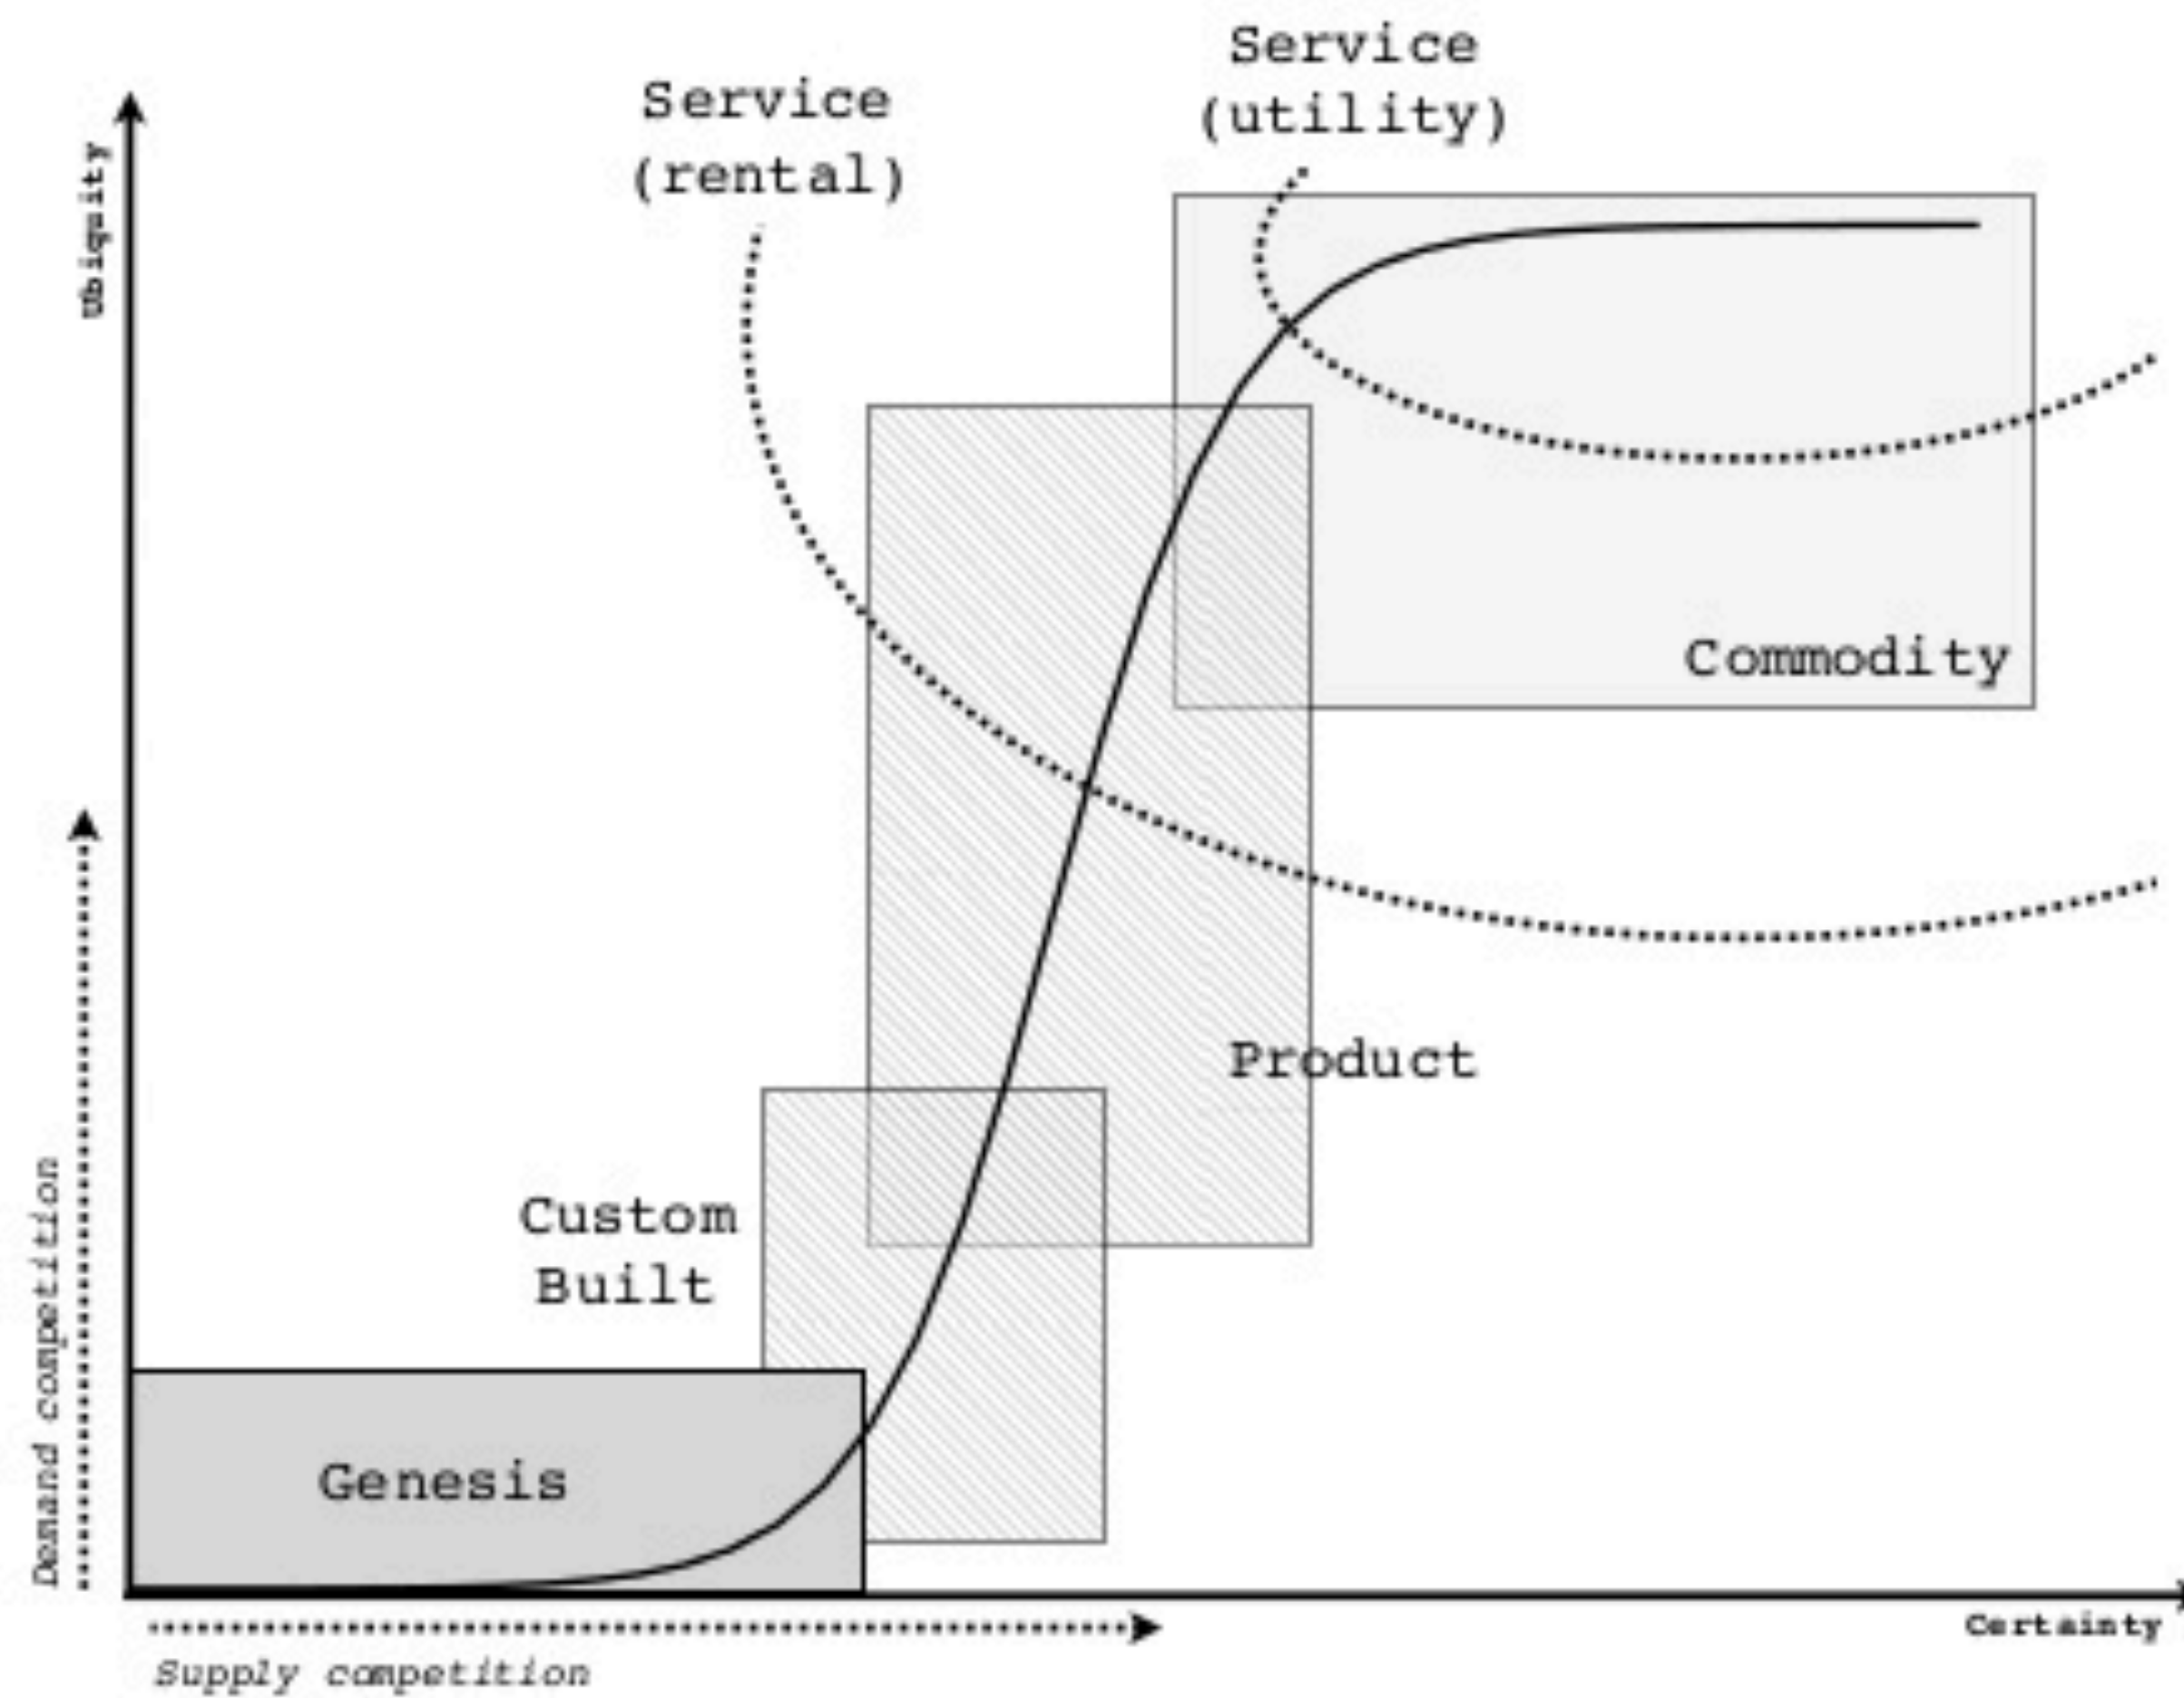
But: supportability of an in-house platform is a challenge

Choose boring technology

<http://mcfunley.com/choose-boring-technology>

@sarahjwells

By late 2016, tools were
maturing



<https://medium.com/wardleymaps>

@sarahjwells

The FT is not a cluster orchestration company

@sarahjwells

Late 2016: Consider the alternatives

@sarahjwells

Metrics for success:

- amount of time spent keeping cluster healthy
 - number of sarcastic comments on slack
-

Opted for kubernetes

@sarahjwells



Liz Rice

@lizrice

Following

Which orchestrator are you using in
production? Please RT!

#containers #docker #Kubernetes #mesos
#hashicorp

25% Docker Swarm

61% Kubernetes

10% Mesos / Marathon

4% Nomad

367 votes • Final results

Using leading edge
technologies requires you to
be comfortable with change

Shouldn't be (too) scared about making the wrong decision

<http://uk.businessinsider.com/jeff-bezos-on-type-1-and-type-2-decisions-2016-4>

@sarahjwells

Switching horses midstream

At the start of this migration we had 150 services

Complications of running in parallel



We had well over 2000 code releases while running at least part of the stack in parallel

@sarahjwells

Decisions, decisions, decisions...



@sarahjwells

Separate branches vs if/else in code



Separate deployment mechanisms vs a single deployment mechanism

Risk-based approach to testing



@sarahjwells

Doing anything 150 times
takes time

@sarahjwells

Changes per service weren't **that** big

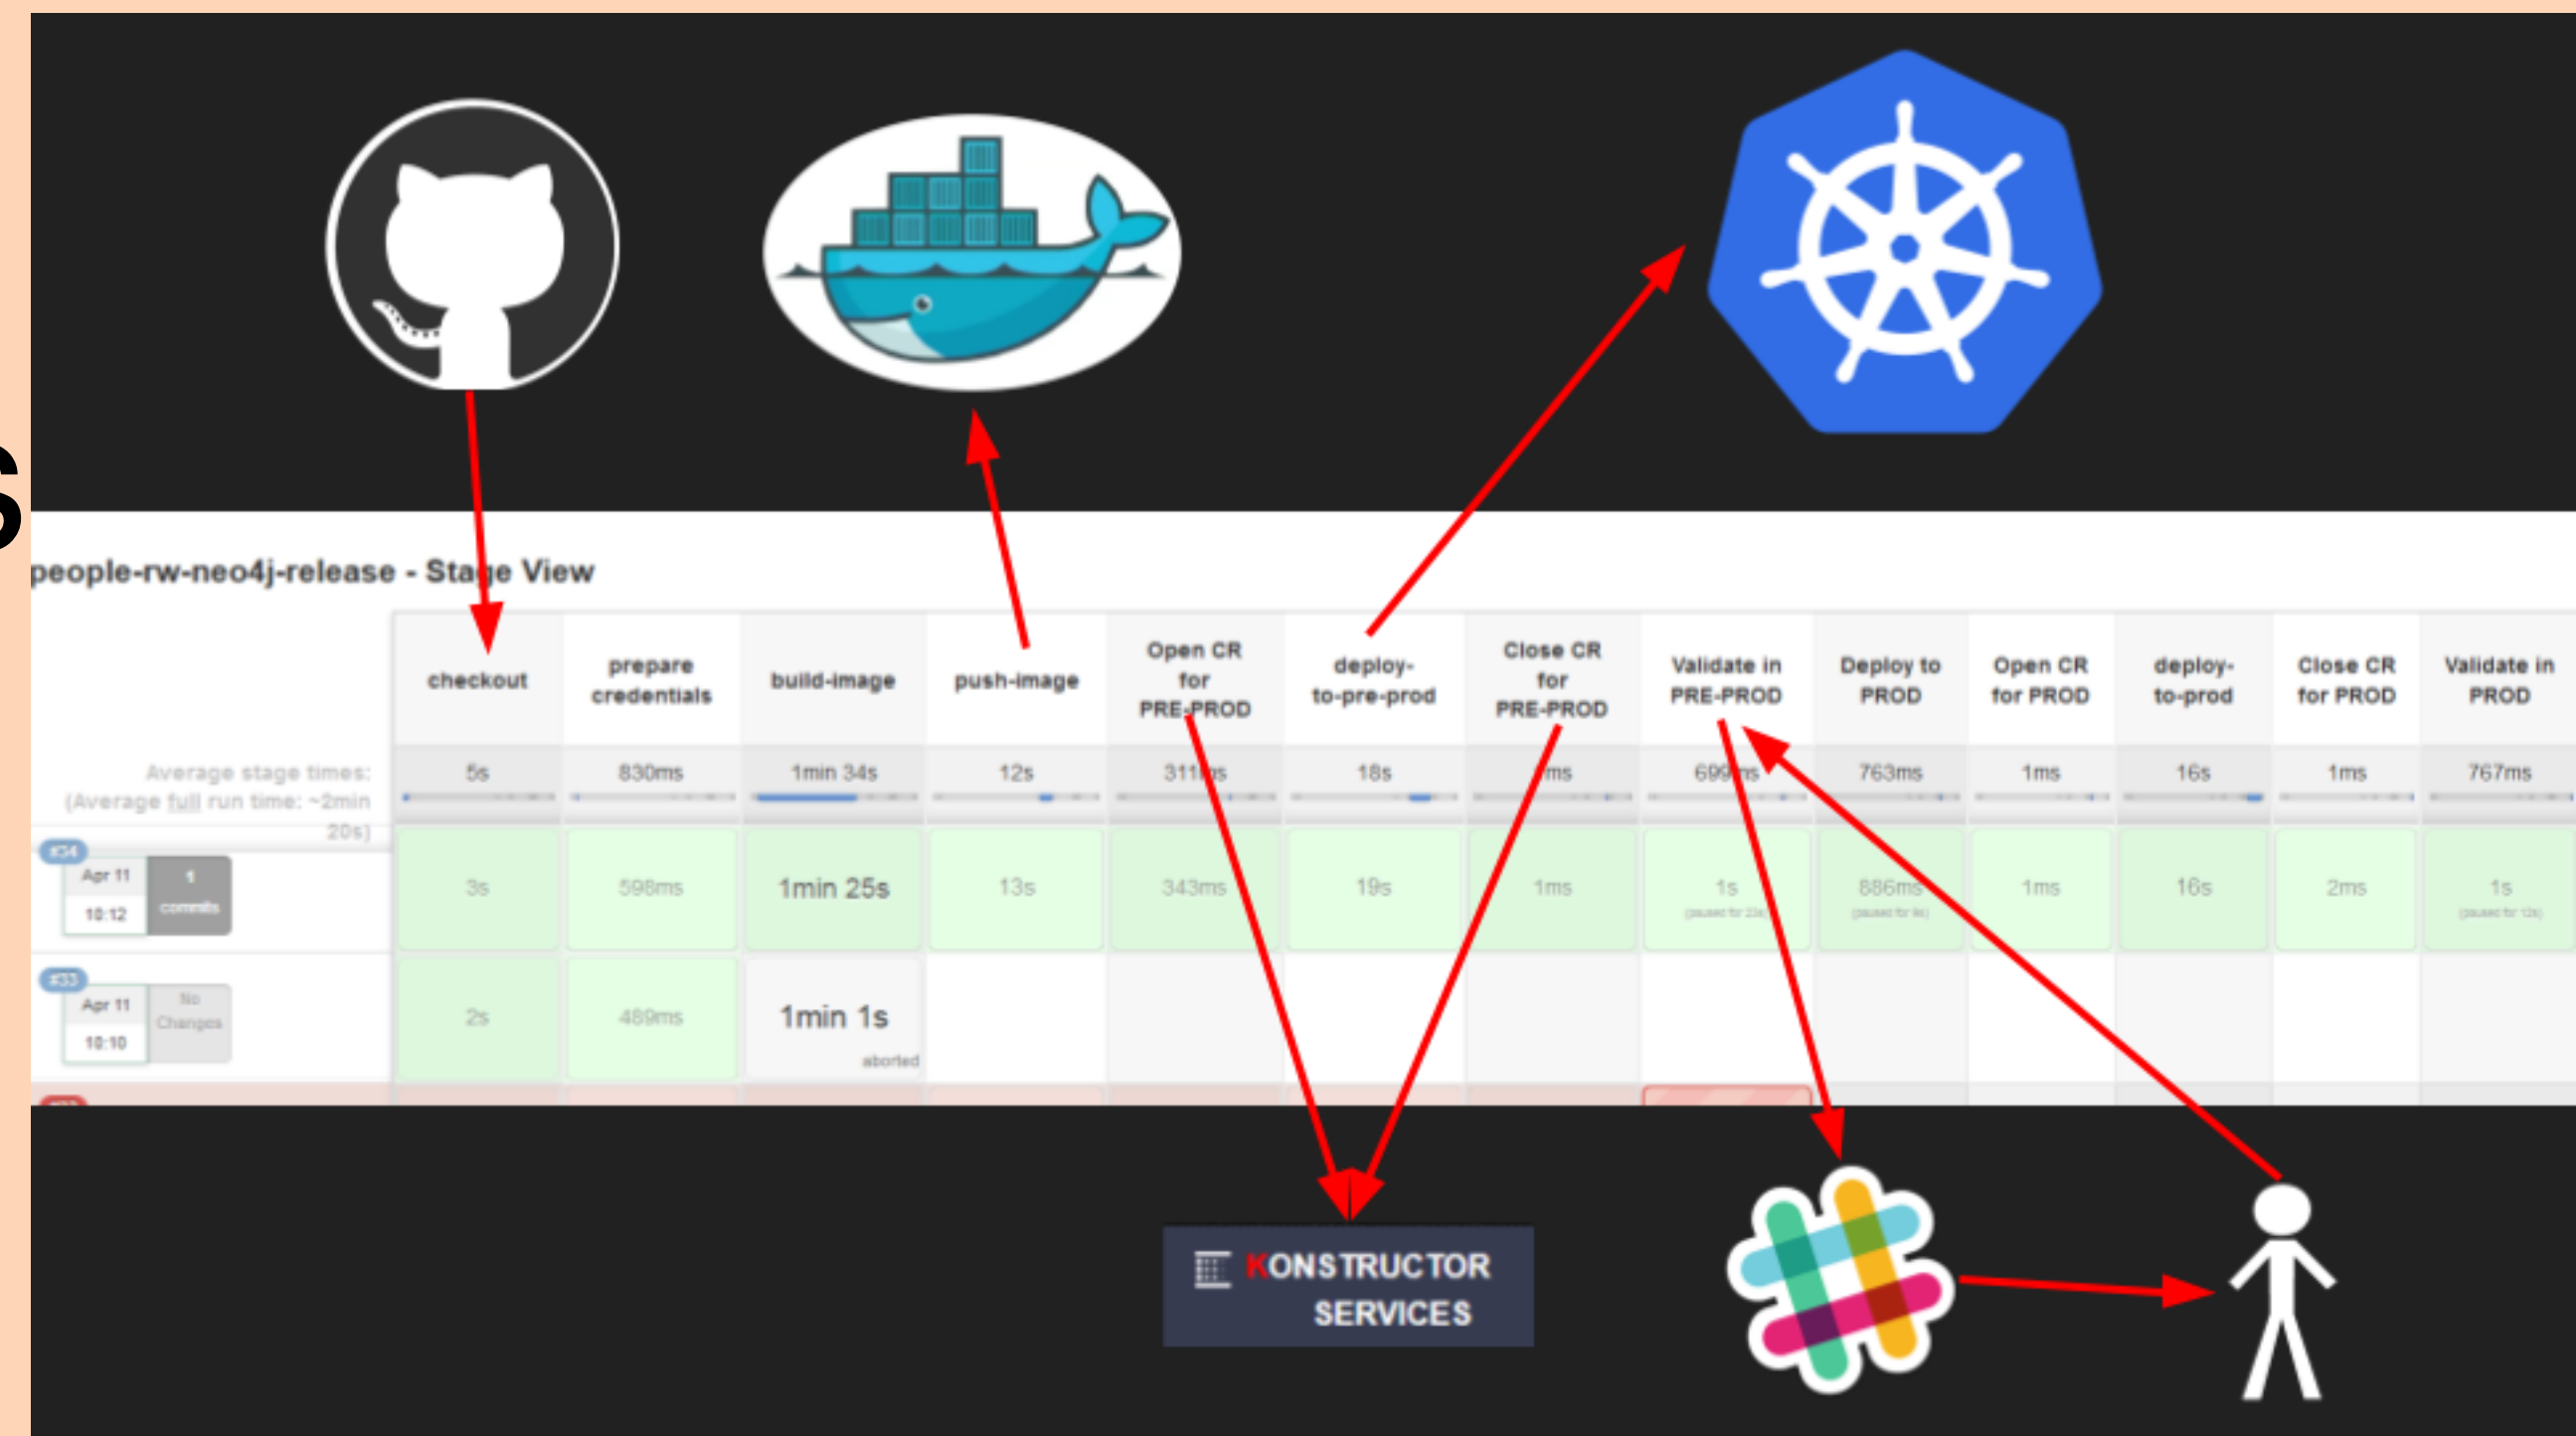
@sarahjwells

Migrating from systemd service files to helm charts



@sarahjwells

Integrating the service into a templated Jenkins pipeline



@sarahjwells

Good to get everyone involved - “Helm days”

@sarahjwells

Discovered a lot of 'broken'
things

Services that hadn't been built for a long time

@sarahjwells

A standard that isn't enforced may will not be complied with:

- healthcheck timeouts

@sarahjwells

UPP upp-prod-delivery-eu cluster's services from categories default (**healthy**)

Search:

Name	Health status	Output	Last Updated	Ack msg	Action
aggregate-concept-transformer	ok	2/2 pods available	16:04:53 UTC		Ack service
annotations-mapper	ok	2/2 pods available	16:04:53 UTC		Ack service
annotations-monitoring-service	ok	1/1 pods available	16:04:53 UTC		Ack service
annotations-publish-healthchecker	ok	1/1 pods available	16:04:53 UTC		Ack service
annotations-rw-neo4j	ok	2/2 pods available	16:04:53 UTC		Ack service
api-policy-component	ok	2/2 pods available	16:04:53 UTC		Ack service
binary-ingester	ok	2/2 pods available	16:04:53 UTC		Ack service
binary-writer	ok	2/2 pods available	16:04:53 UTC		Ack service
cms-kafka-bridge-pub-prod-eu	ok	2/2 pods available	16:04:53 UTC		Ack service

- /__gtg endpoints

```
livenessProbe:  
  tcpSocket:  
    port: 8080  
  initialDelaySeconds: 5  
readinessProbe:  
  httpGet:  
    path: "/__gtg"  
    port: 8080  
  initialDelaySeconds: 15  
  periodSeconds: 30
```

Making sure a service will recover if k8s moves it elsewhere

Easy to get sucked into making things better

@sarahjwells

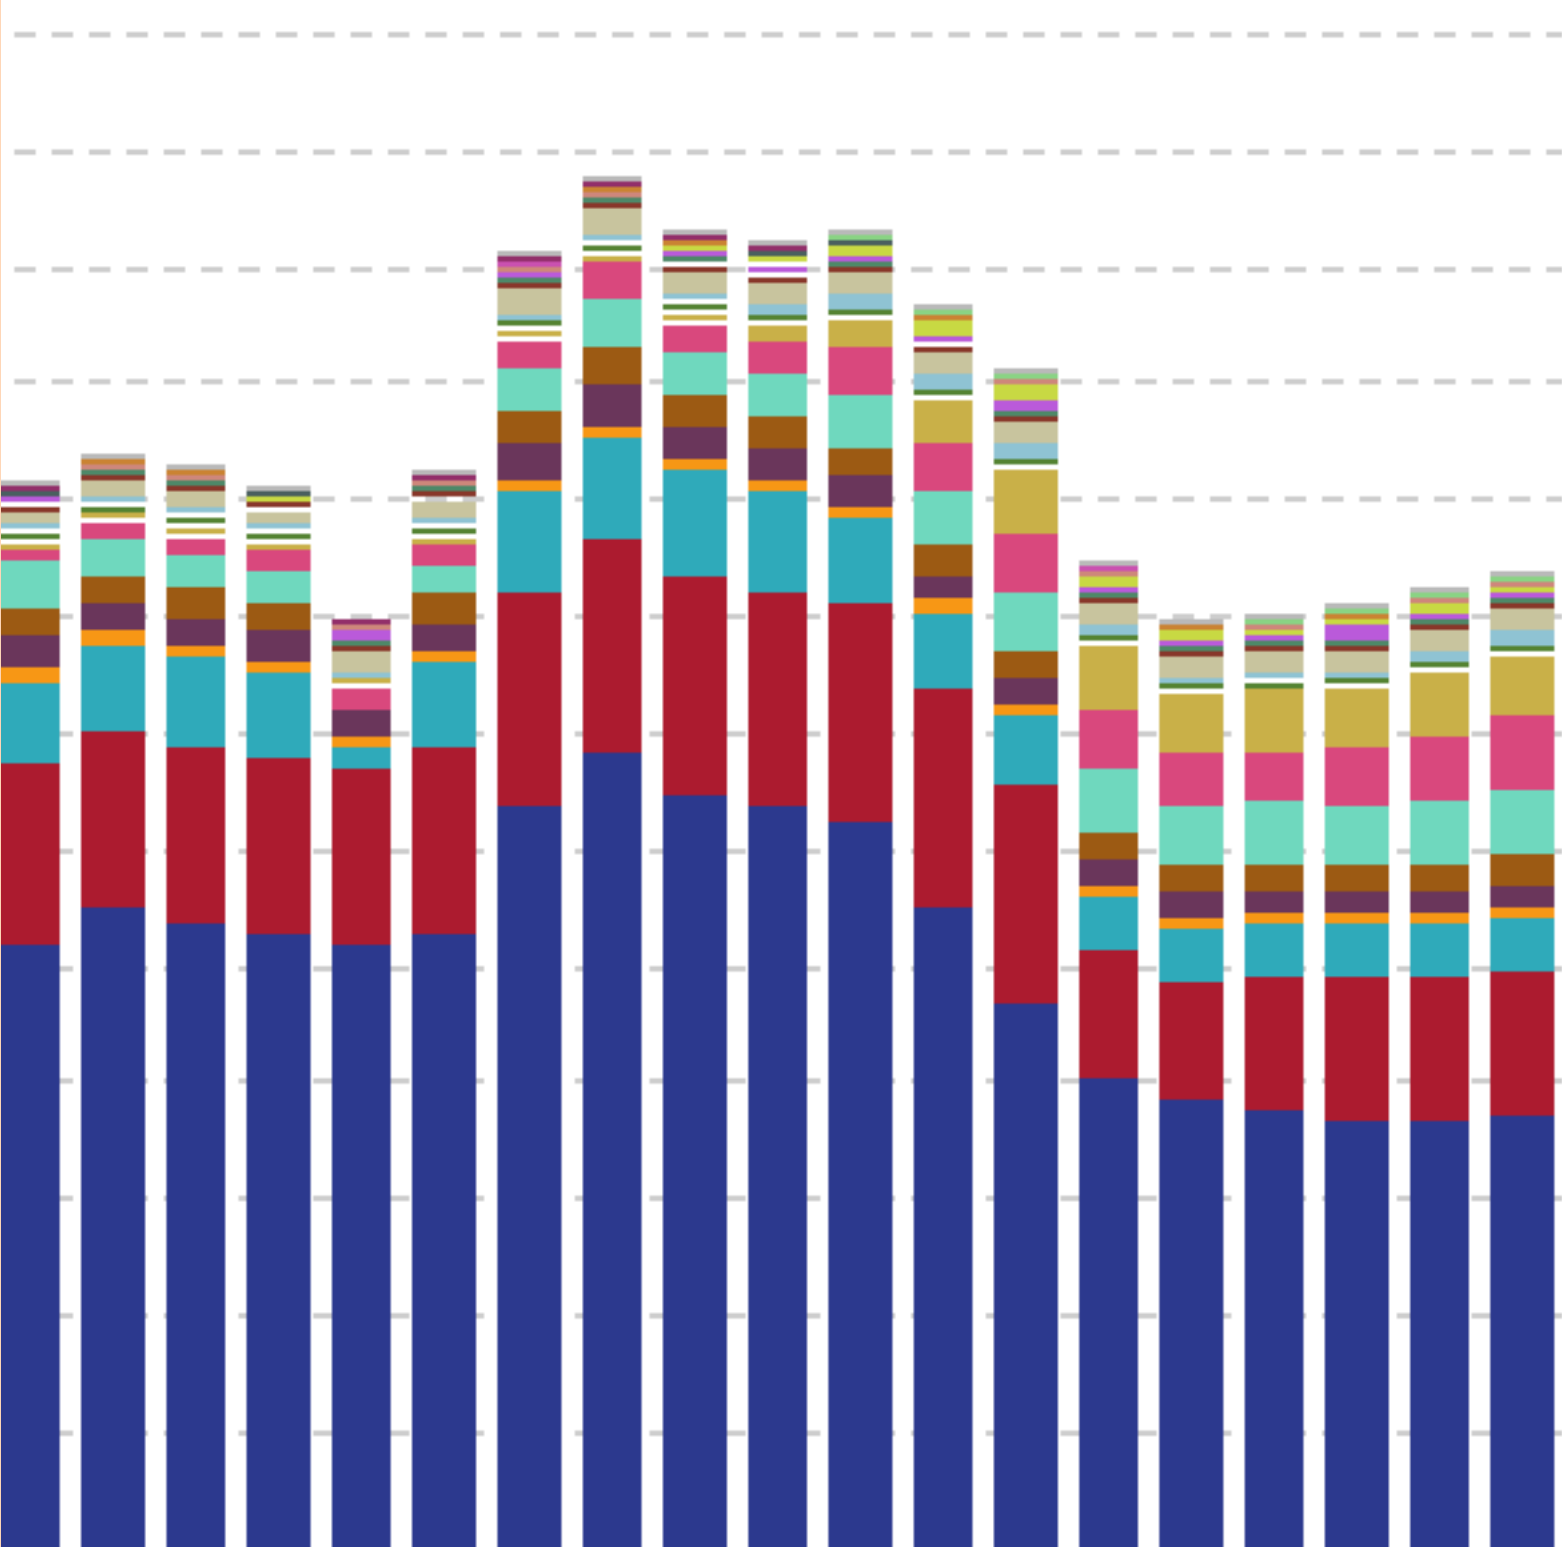
Would have been better if...

We'd swarmed on the work

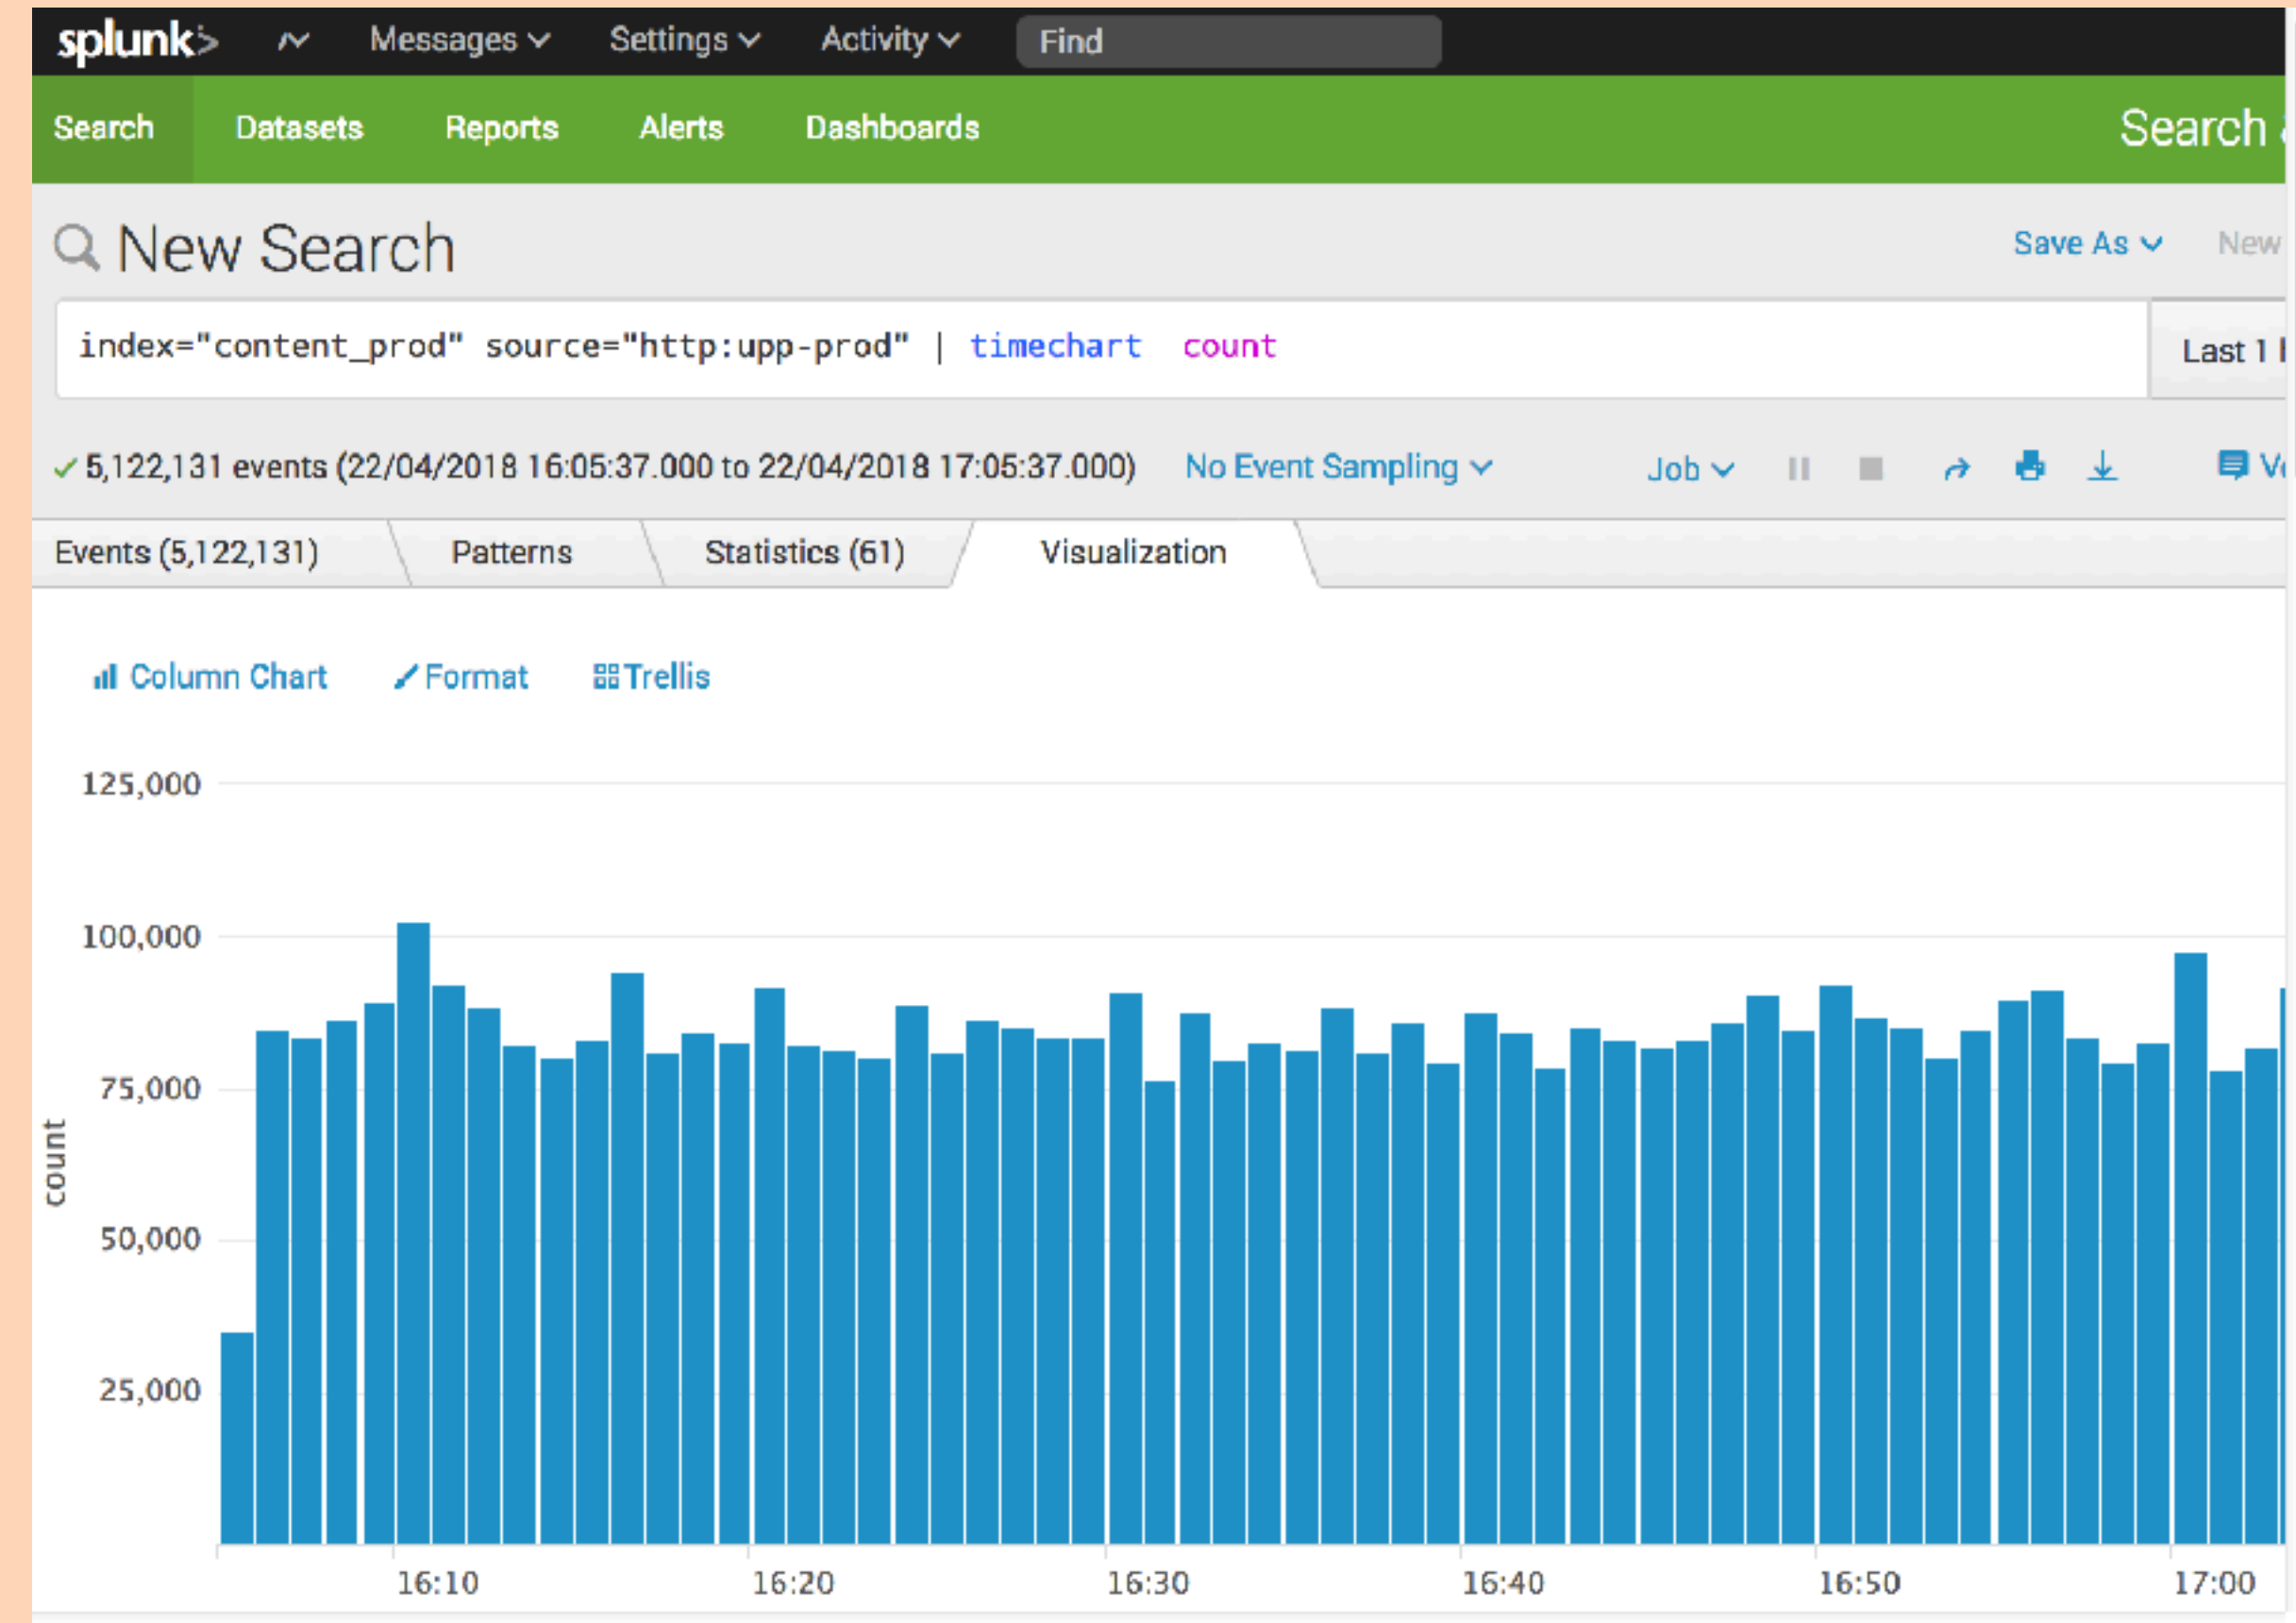
@sarahjwells

The longer you run in parallel, the more overhead for releasing code changes

and the higher the costs



Not just AWS costs
either



@sarahjwells

Going live

@sarahjwells

Doing the migration



@sarahjwells

The results

A more stable platform

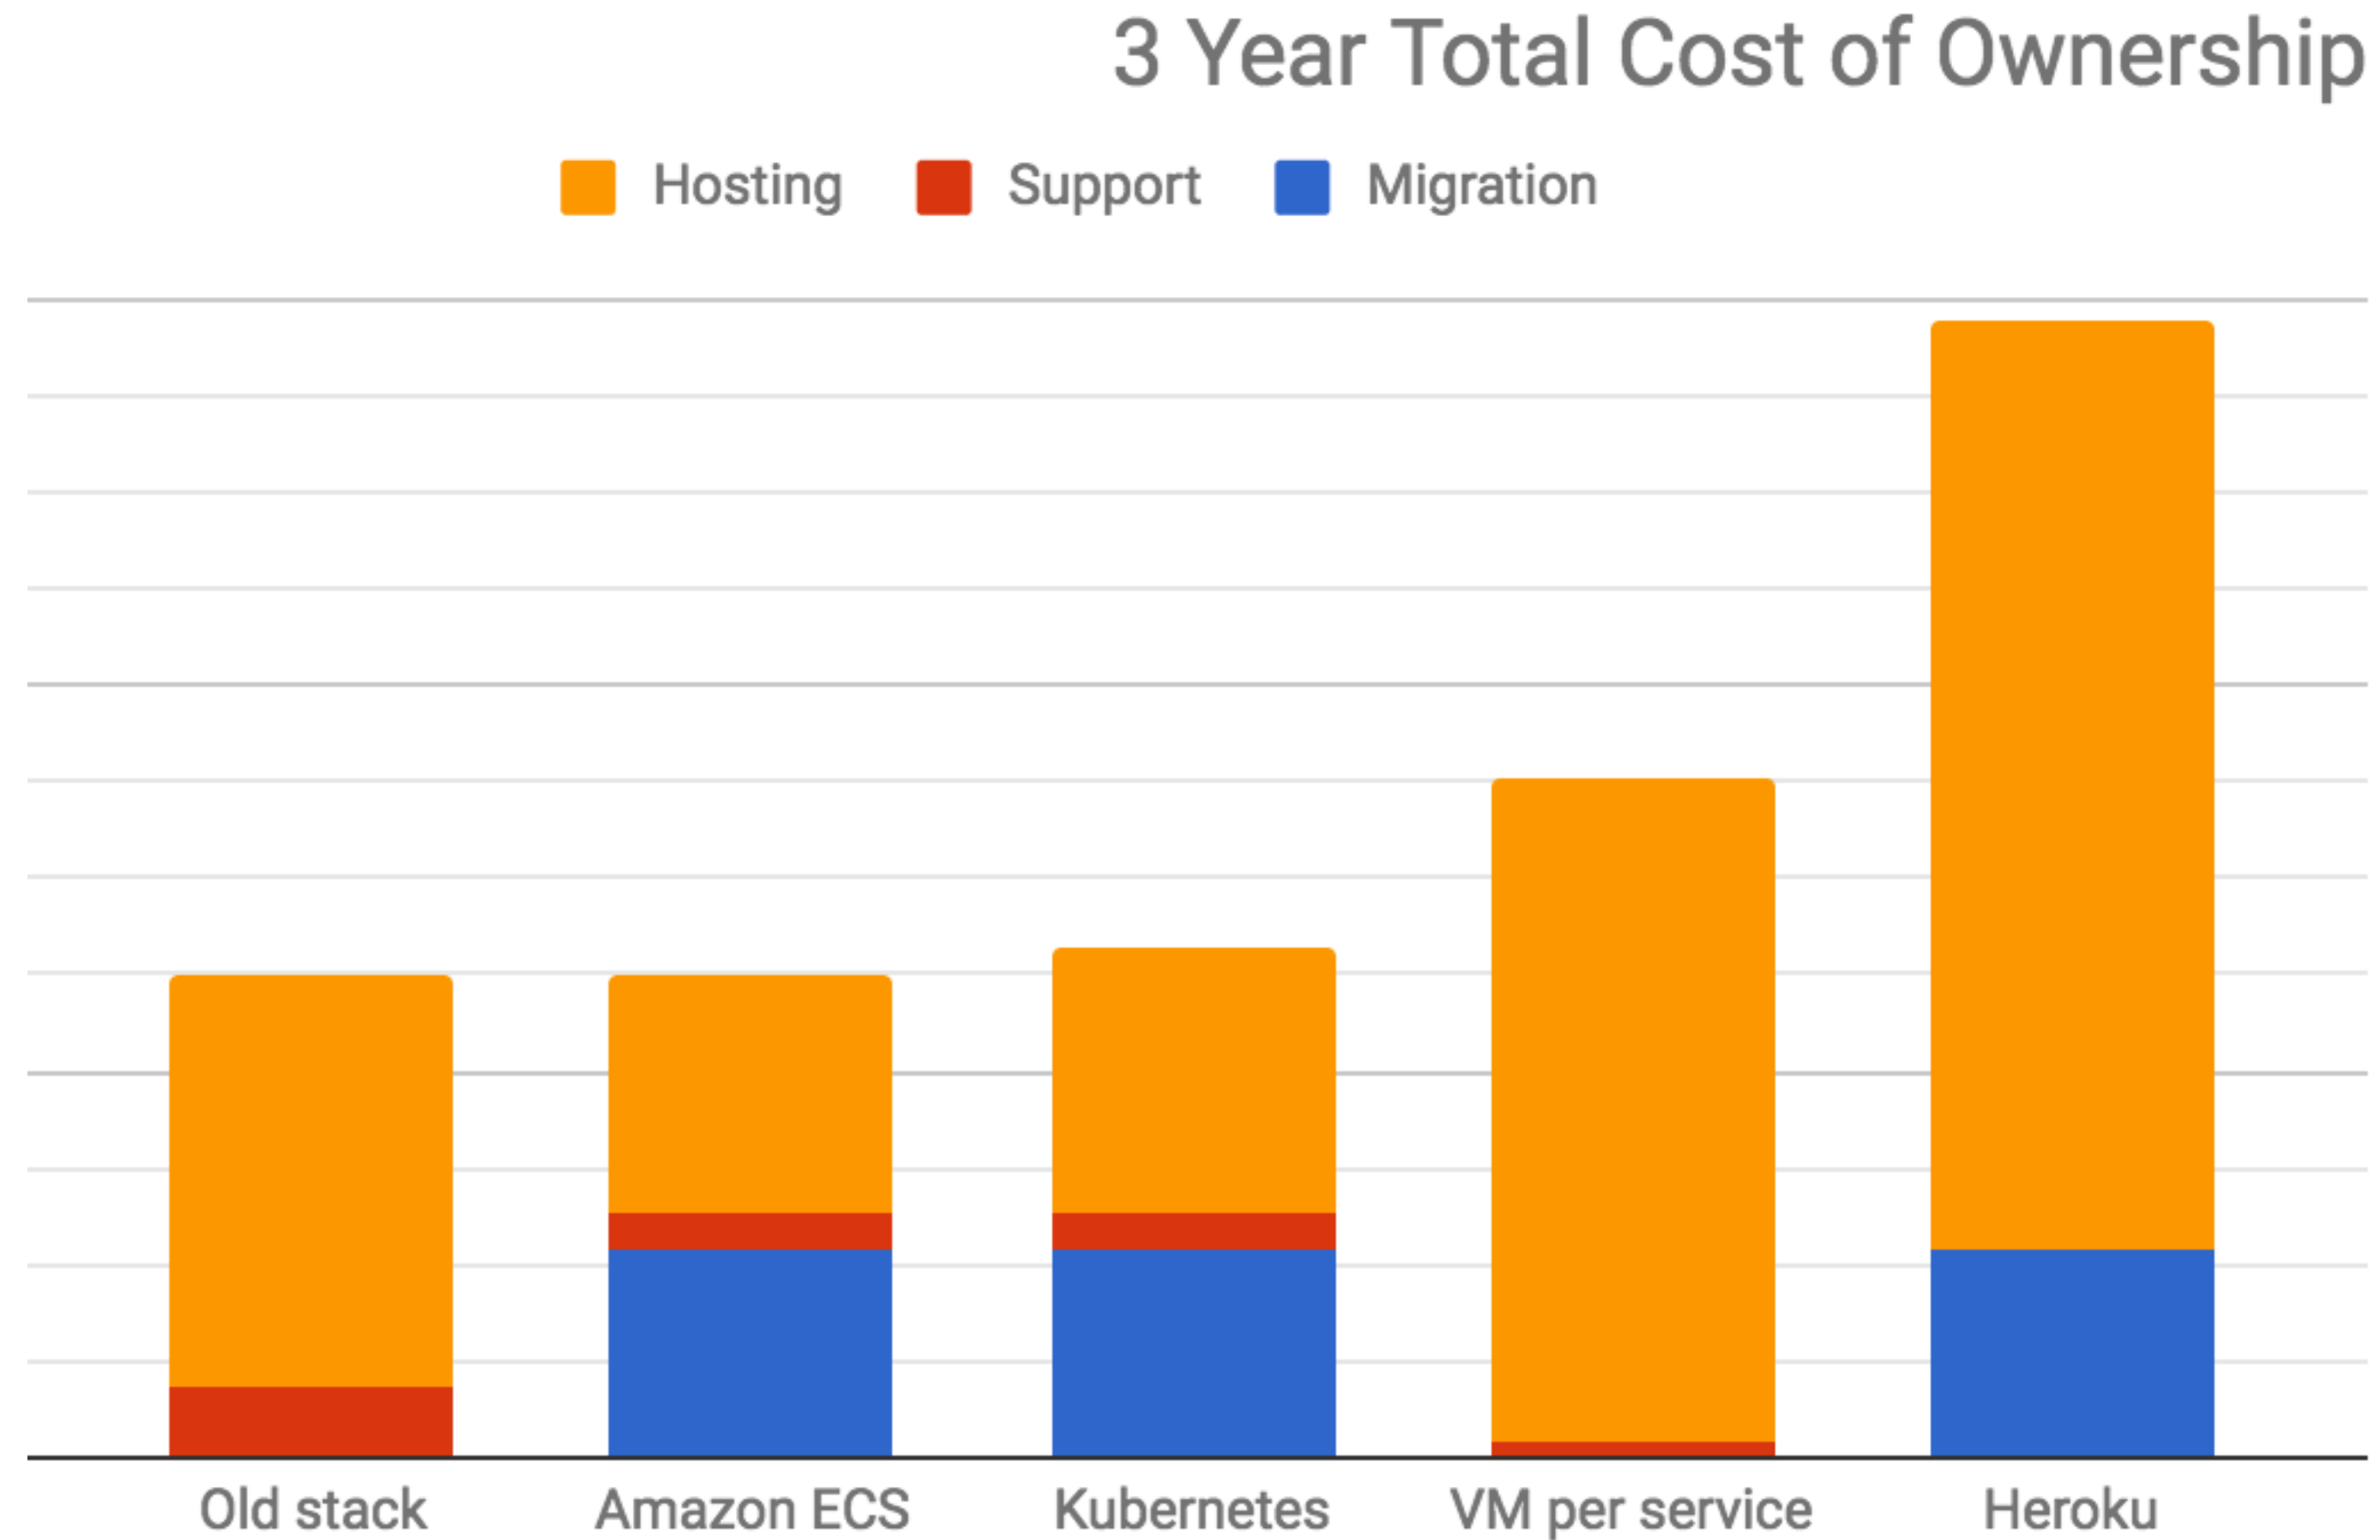


@sarahjwells

Something where we can learn from others

@sarahjwells

Reduction in hosting and support costs





Thank you!

We're hiring: <https://aboutus.ft.com/careers/>

@sarahjwells