

Avoiding alerts overload from microservices

Sarah Wells

Principal Engineer, Financial Times

@sarahjwells



More

1-100 of many

COMPOSE

Inbox (1,158 : 6,858)

Starred

Important

Sent Mail

Drafts (24)

Circles

Agile

Architecture

Content API

ApigeeAnalytics (...)

ApigeeDailyDigest...

BuildNotifications...

CodeReviews

Developer Portal Ac...

JIRAs

Monitoring (261)

svn commits

Dynamic Semantic ...

Alerting (splunk) (...)

BuildNotifications...

Monitoring (nagio...

Incidents (105)

Interesting Info (5)

Nagios recent (10,721)

Neo notifications (2...

Pingdom Monitorin...

Prod Alerts to Conte...

Publish Failures (2)

More

<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios (9)	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** PROBLEM Service Alert: ftweb39407-law1c-eu-p.osb.ft.com/TCP: blogs-varnish-03.labs.p... - Nagios Notification Type	23 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios (9)	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** PROBLEM Service Alert: ftweb21626-law1a-eu-p.osb.ft.com/TCP: blogs-varnish-01.labs.p... - Nagios Notification Type	23 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios (9)	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** PROBLEM Service Alert: ftweb21627-law1b-eu-p.osb.ft.com/TCP: blogs-varnish-02.labs.p... - Nagios Notification Type	23 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios (100)	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** PROBLEM Service Alert: ftweb39407-law1c-eu-p.osb.ft.com/TCP: blogs-varnish-03.labs.p... - Nagios Notification Type	23 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios (100)	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** PROBLEM Service Alert: ftweb21626-law1a-eu-p.osb.ft.com/TCP: blogs-varnish-01.labs.p... - Nagios Notification Type	23 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios (100)	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** PROBLEM Service Alert: ftweb21627-law1b-eu-p.osb.ft.com/TCP: blogs-varnish-02.labs.p... - Nagios Notification Type	23 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** RECOVERY Service Alert: ftweb21626-law1a-eu-p.osb.ft.com/Disk: /var is OK ** - Nagios Notification Type: RECOVER	22 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** RECOVERY Service Alert: ftweb21626-law1a-eu-p.osb.ft.com/Disk: /boot is OK ** - Nagios Notification Type: RECOVE	22 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** RECOVERY Service Alert: ftweb21626-law1a-eu-p.osb.ft.com/Disk: / is OK ** - Nagios Notification Type: RECOVERY S	22 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios (8)	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** PROBLEM Service Alert: ftweb21626-law1a-eu-p.osb.ft.com/Disk: /var is WARNING ** - Nagios Notification Type: PROI	22 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios (8)	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** PROBLEM Service Alert: ftweb21626-law1a-eu-p.osb.ft.com/Disk: /boot is WARNING ** - Nagios Notification Type: PR	22 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** RECOVERY Service Alert: ftaps21899-law1b-eu-p.osb.ft.com/Disk: / is OK ** - Nagios Notification Type: RECOVERY S	22 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** RECOVERY Service Alert: ftaps21899-law1b-eu-p.osb.ft.com/Disk: /var is OK ** - Nagios Notification Type: RECOVER	22 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** RECOVERY Service Alert: ftweb21627-law1b-eu-p.osb.ft.com/Disk: / is OK ** - Nagios Notification Type: RECOVERY S	22 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** RECOVERY Service Alert: ftaps21899-law1b-eu-p.osb.ft.com/Disk: /boot is OK ** - Nagios Notification Type: RECOVEI	22 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** RECOVERY Service Alert: ftweb21627-law1b-eu-p.osb.ft.com/Disk: /boot is OK ** - Nagios Notification Type: RECOVE	22 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** RECOVERY Service Alert: ftweb21627-law1b-eu-p.osb.ft.com/Disk: /var is OK ** - Nagios Notification Type: RECOVER	22 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios (2)	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** PROBLEM Service Alert: ftaps21899-law1b-eu-p.osb.ft.com/Disk: / is WARNING ** - Nagios Notification Type: PROBLE	22 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios (2)	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** PROBLEM Service Alert: ftaps21899-law1b-eu-p.osb.ft.com/Disk: /var is WARNING ** - Nagios Notification Type: PROE	22 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** PROBLEM Service Alert: ftweb21627-law1b-eu-p.osb.ft.com/Disk: / is CRITICAL ** - Nagios Notification Type: PROBLE	22 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios (2)	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** PROBLEM Service Alert: ftaps21899-law1b-eu-p.osb.ft.com/Disk: /boot is WARNING ** - Nagios Notification Type: PR	22 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** PROBLEM Service Alert: ftweb21627-law1b-eu-p.osb.ft.com/Disk: /boot is CRITICAL ** - Nagios Notification Type: PR	22 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** PROBLEM Service Alert: ftweb21627-law1b-eu-p.osb.ft.com/Disk: /var is CRITICAL ** - Nagios Notification Type: PROI	22 Sep
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	nagios	Dyn.../Monitoring (nagios)	Nagios recent	Prod Alerts to Content ...	** RECOVERY Service Alert: dummy ftaps21901-law1a-eu-p.osb.ft.com/ftaps21901-law1a-eu-p... - Nagios Notification Ty	22 Sep



Content API V2 Content - PROD UK backend is down

[prod-uk-up.ft.com](#) • [View details](#)

UPP Prod Delivery UK: Image Publish Aggregate Healthcheck is down

[prod-uk-up.ft.com](#) • [View details](#)

Content API V2 Content - PROD UK backend is up

[prod-uk-up.ft.com](#) • [View details](#)

UPP Prod Publishing Active: Aggregate Healthcheck is down (Incident #55002)

[pub-prod-up.ft.com](#) • [View details](#)

UPP Prod Delivery UK: Image Publish Aggregate Healthcheck is down (Incident #55004)

[prod-uk-up.ft.com](#) • [View details](#)

Content API V2 Content - PROD UK backend is down

[prod-uk-up.ft.com](#) • [View details](#)

Content API V2 Enriched Content - PROD UK backend is down

[prod-uk-up.ft.com](#) • [View details](#)

Content API V2 Content - PROD UK backend is up

[prod-uk-up.ft.com](#) • [View details](#)

UPP Prod Delivery UK: Concordances Read Aggregate Healthcheck is down

[prod-uk-up.ft.com](#) • [View details](#)

UPP Prod Delivery UK: List Read Aggregate Healthcheck is down

[prod-uk-up.ft.com](#) • [View details](#)

UPP Prod Delivery UK: Concordances Read Aggregate Healthcheck is up

[prod-uk-up.ft.com](#) • [View details](#)

UPP Prod Delivery UK: List Read Aggregate Healthcheck is up

[prod-uk-up.ft.com](#) • [View details](#)

Content API V2 Enriched Content - PROD UK backend is up

[prod-uk-up.ft.com](#) • [View details](#)

[prod-uk-up.ft.com](#) • [View details](#)

Fri 21 Oct, 19:46

ALRT #2117, #2119 on UPP Platinum. Reply 4:Ack all, 6:Resolv all

4

2 (of 2) acknowledged. You are now assigned 0 triggered incidents

ALRT #2123, #2124 on UPP Platinum. Reply 14:Ack all, 16:Resolv all

14

2 (of 2) acknowledged. You are now assigned 0 triggered incidents

ALRT #2129 on UPP Platinum: Pingdom Alert: incident #49747 is open for Content API


Knowing when there's a problem isn't enough

@sarahjwells

You only want an alert when you
need to take action

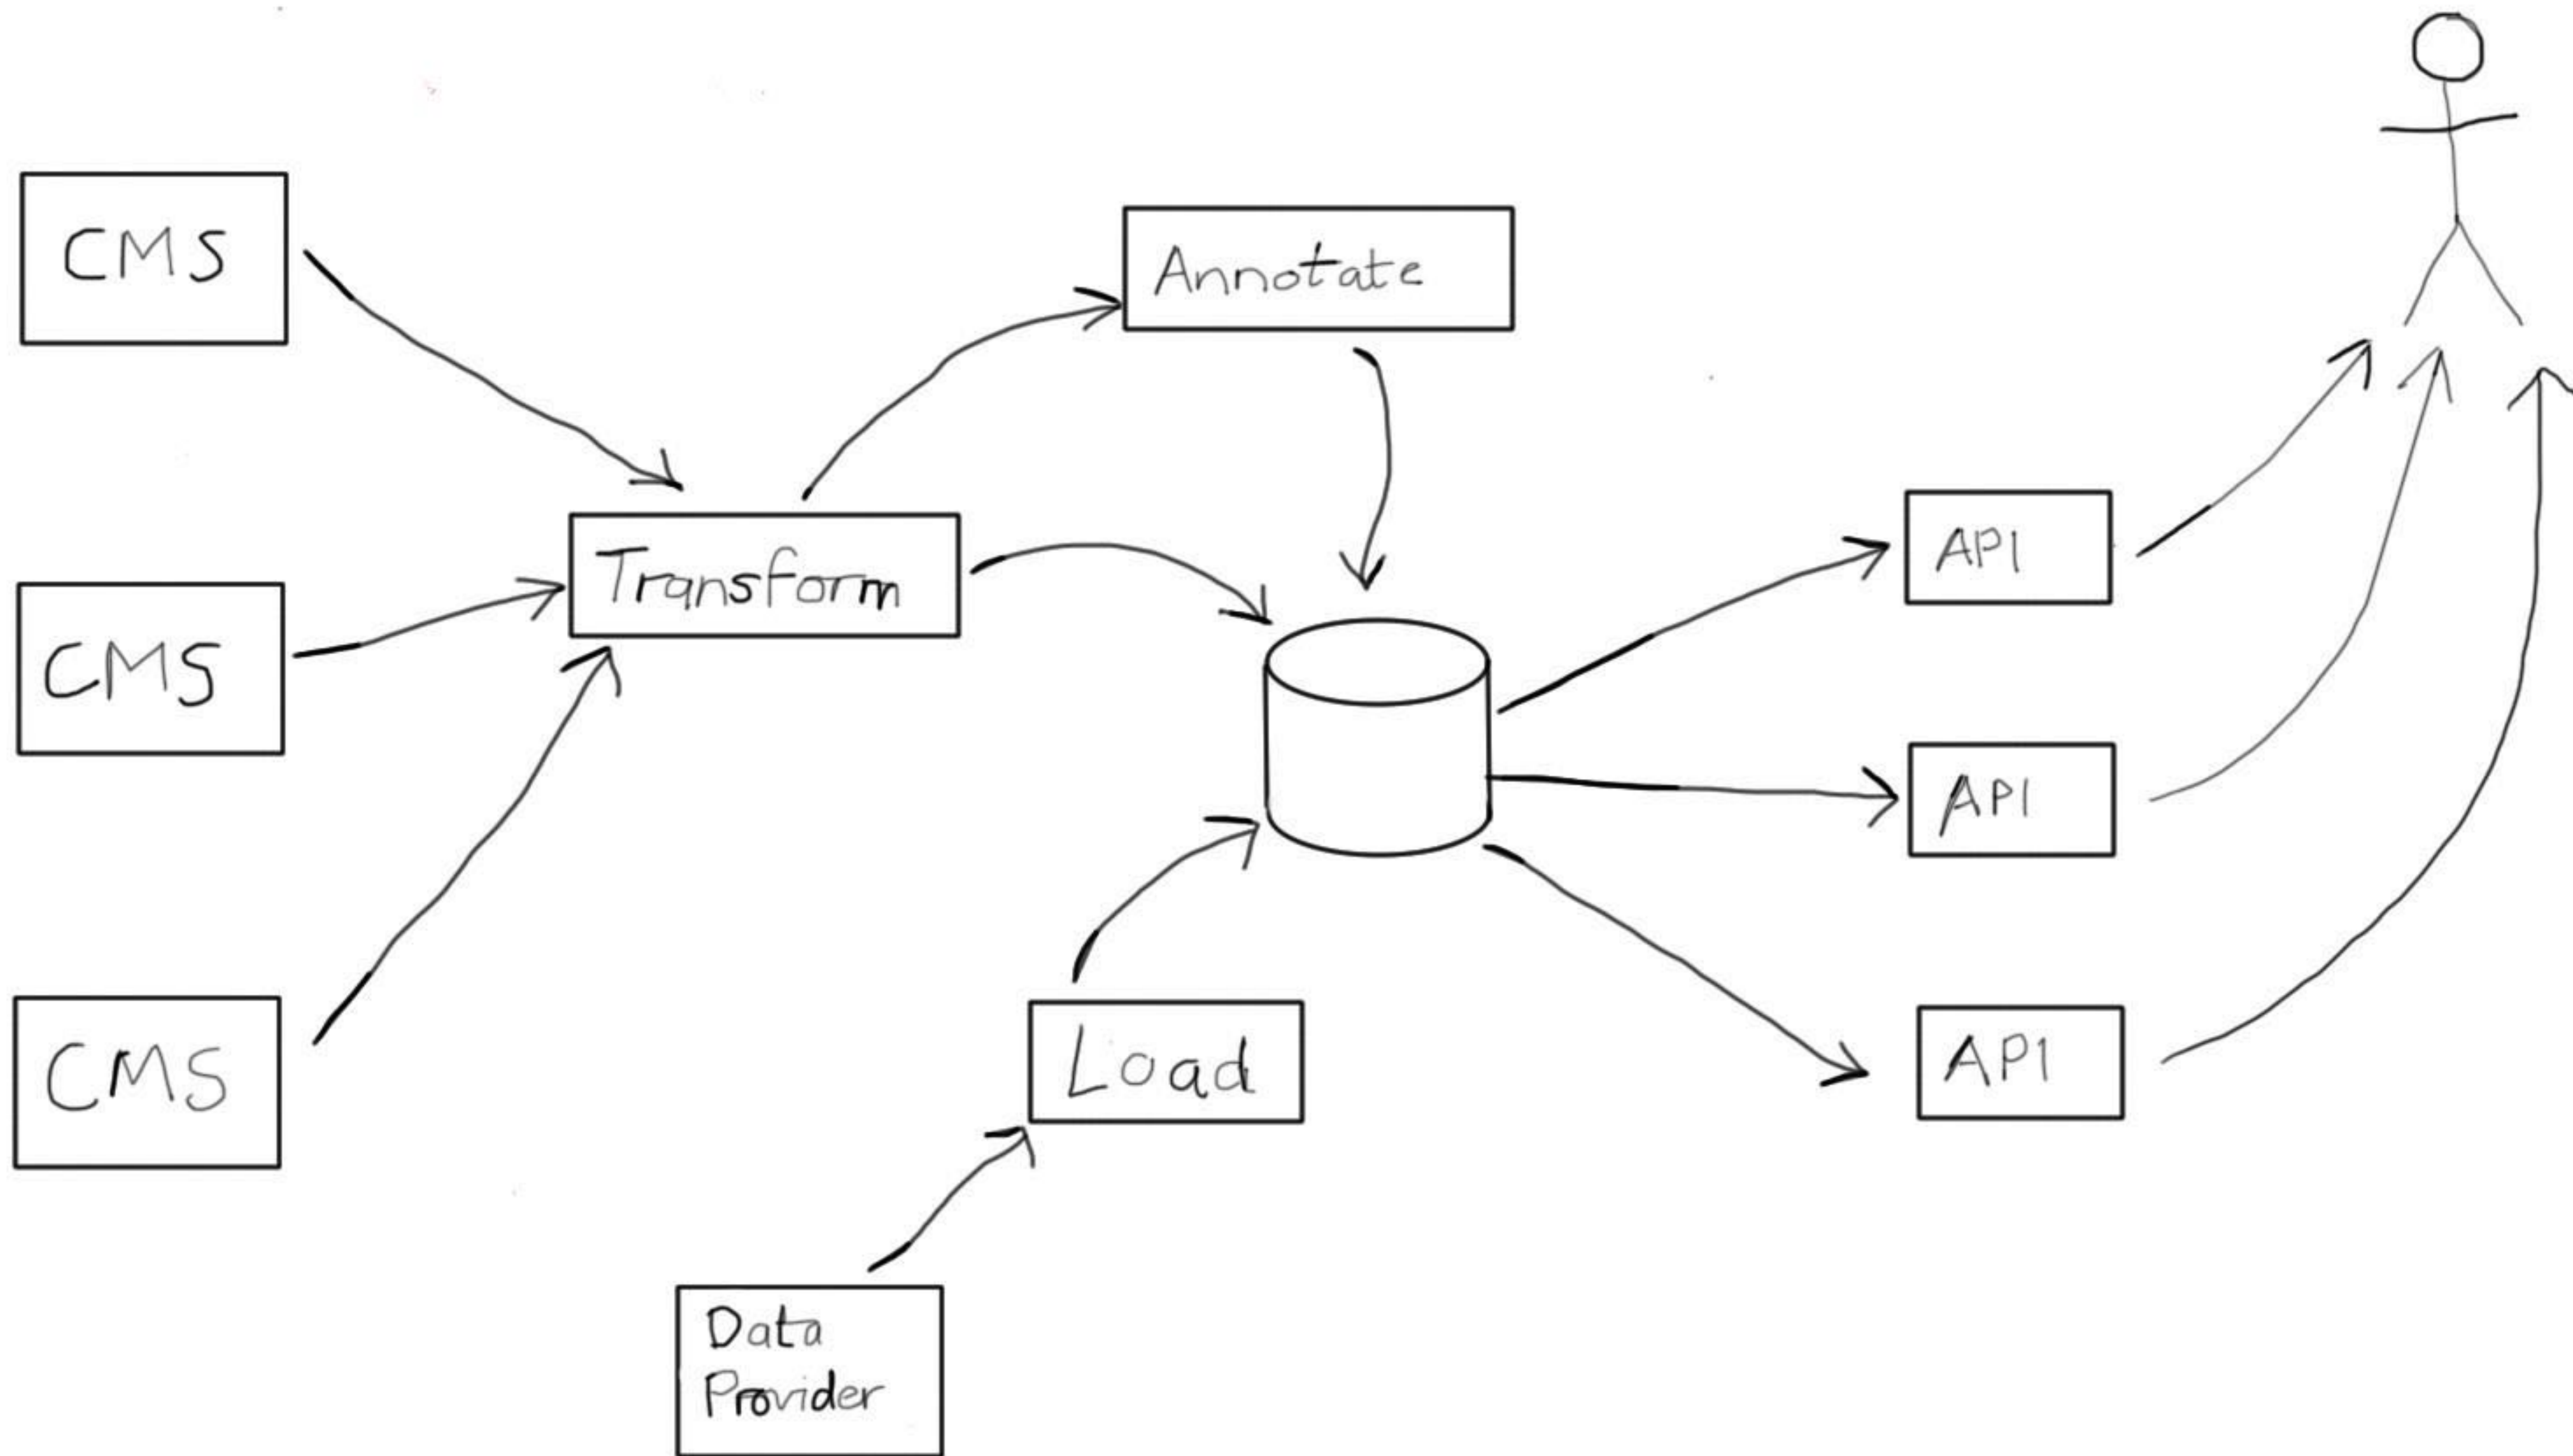
Hello

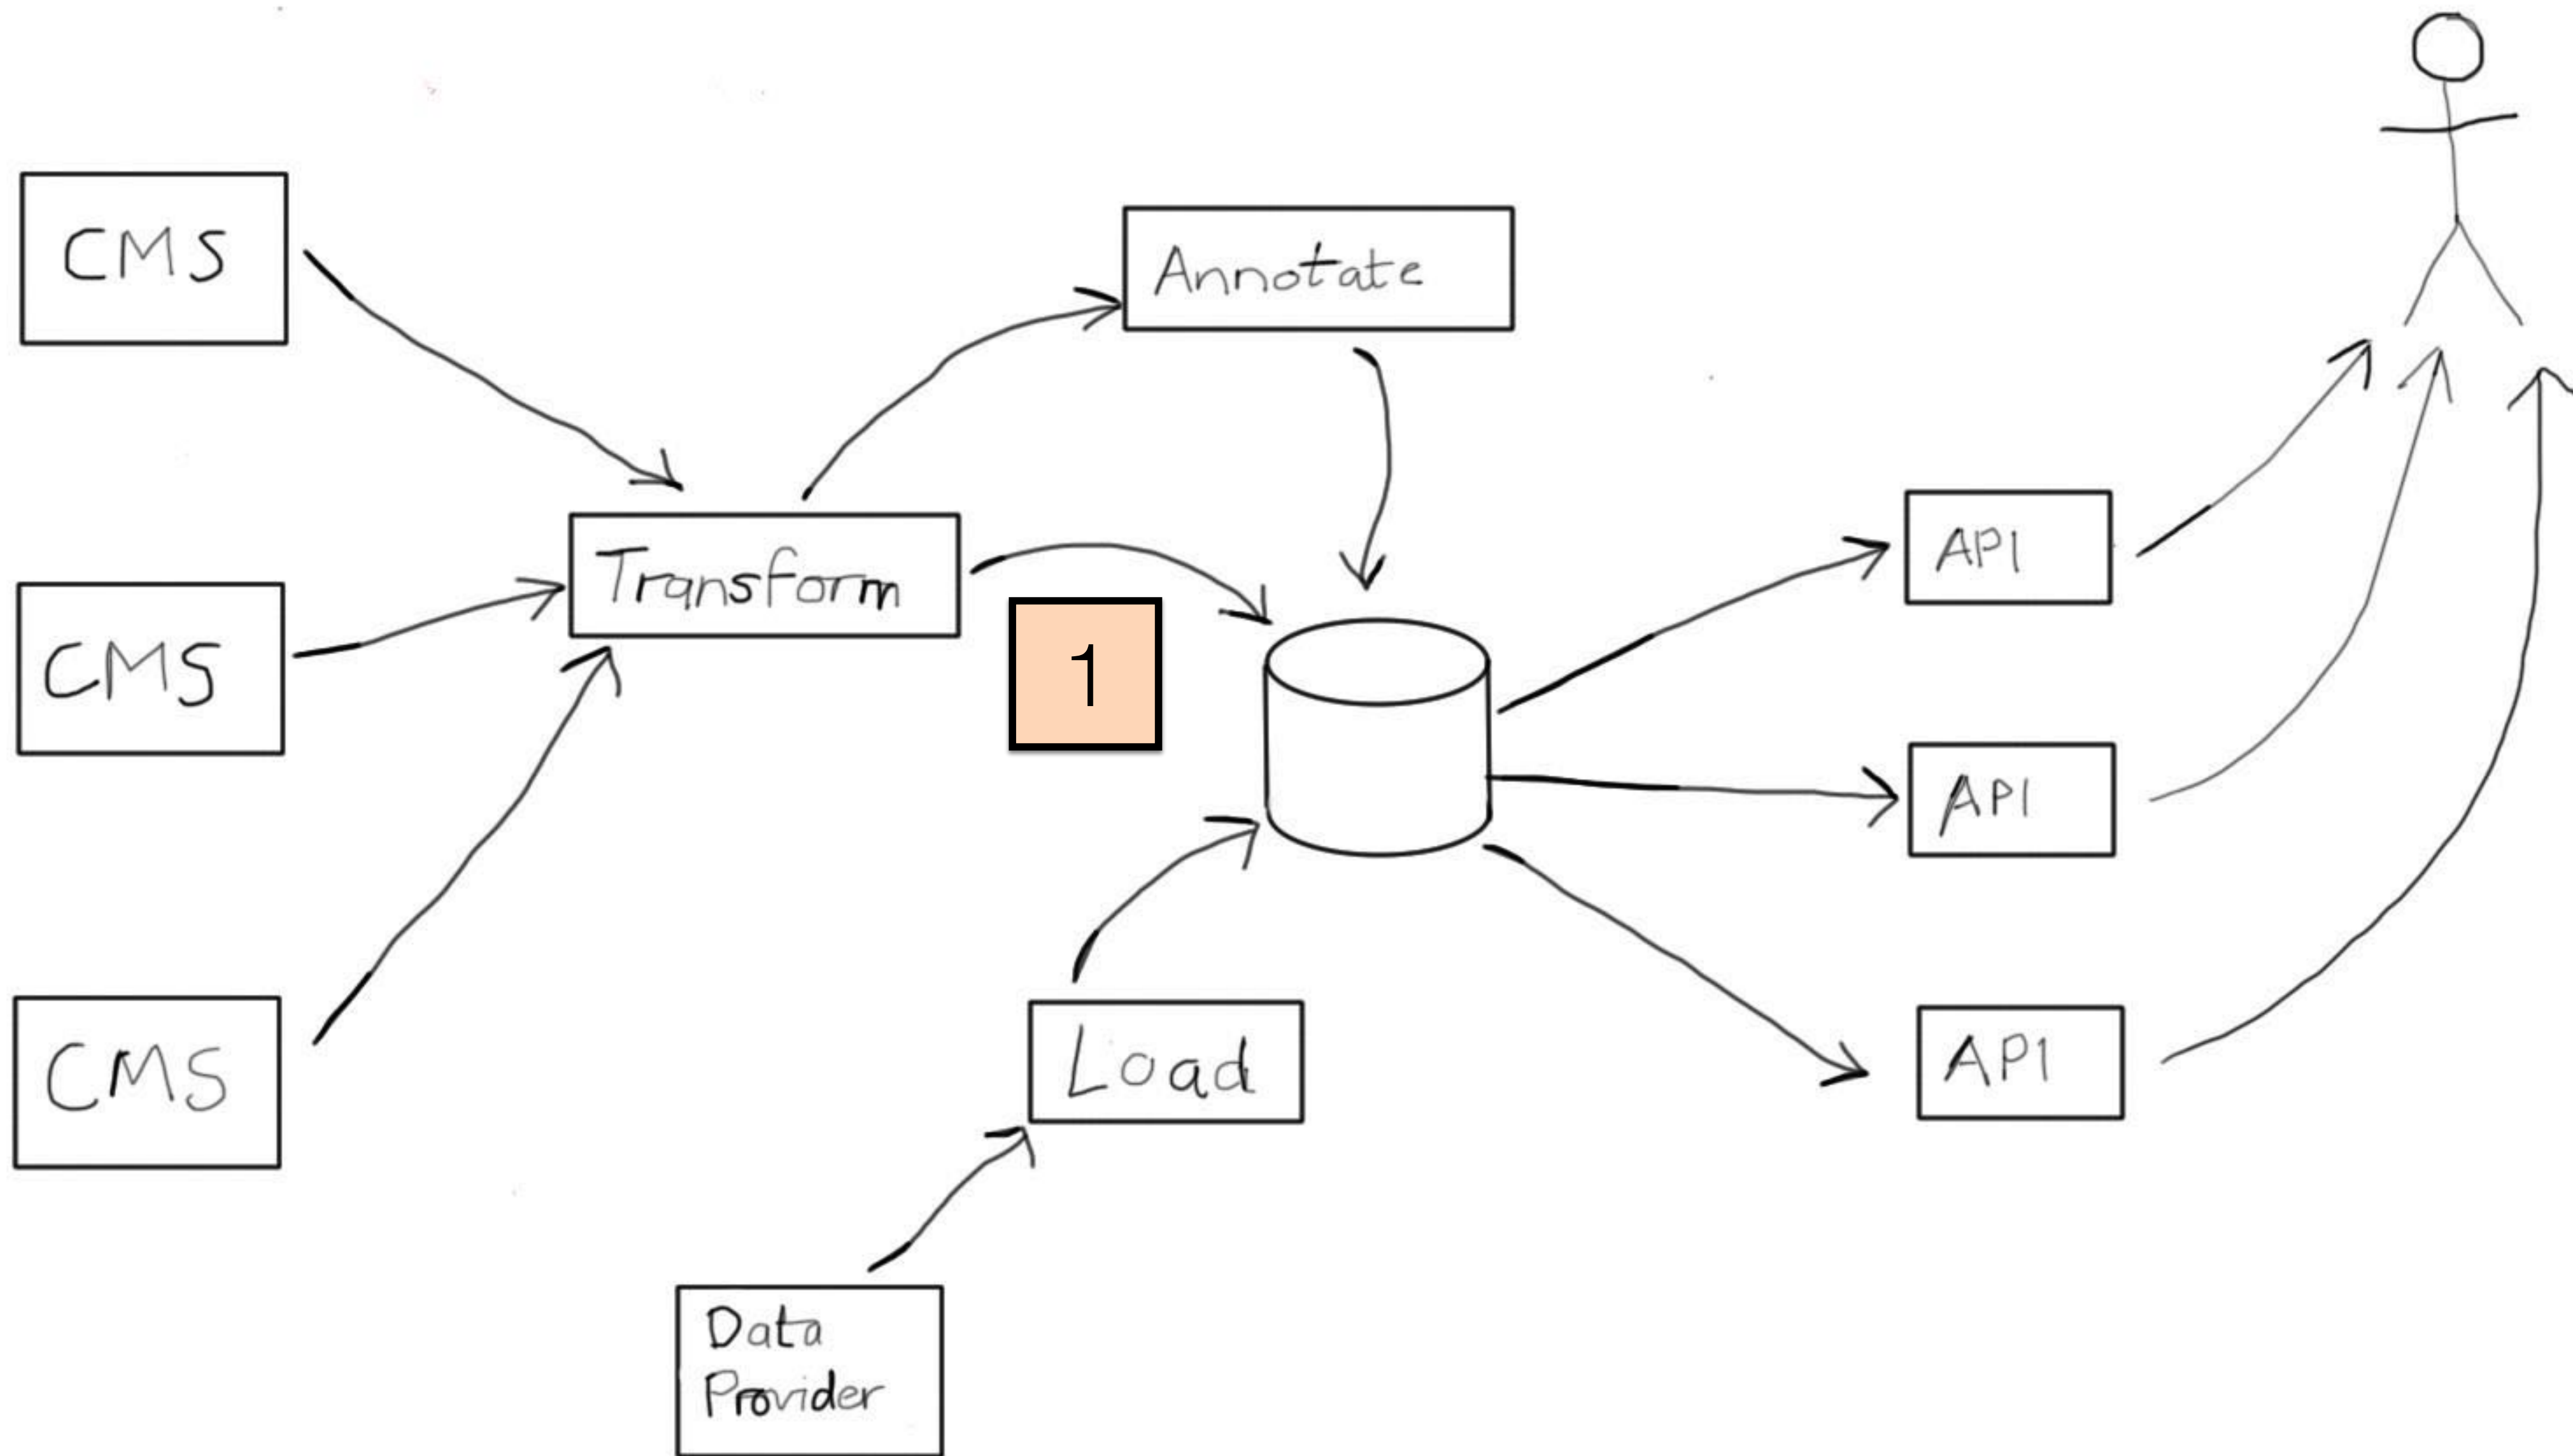
@sarahjwells





elastic

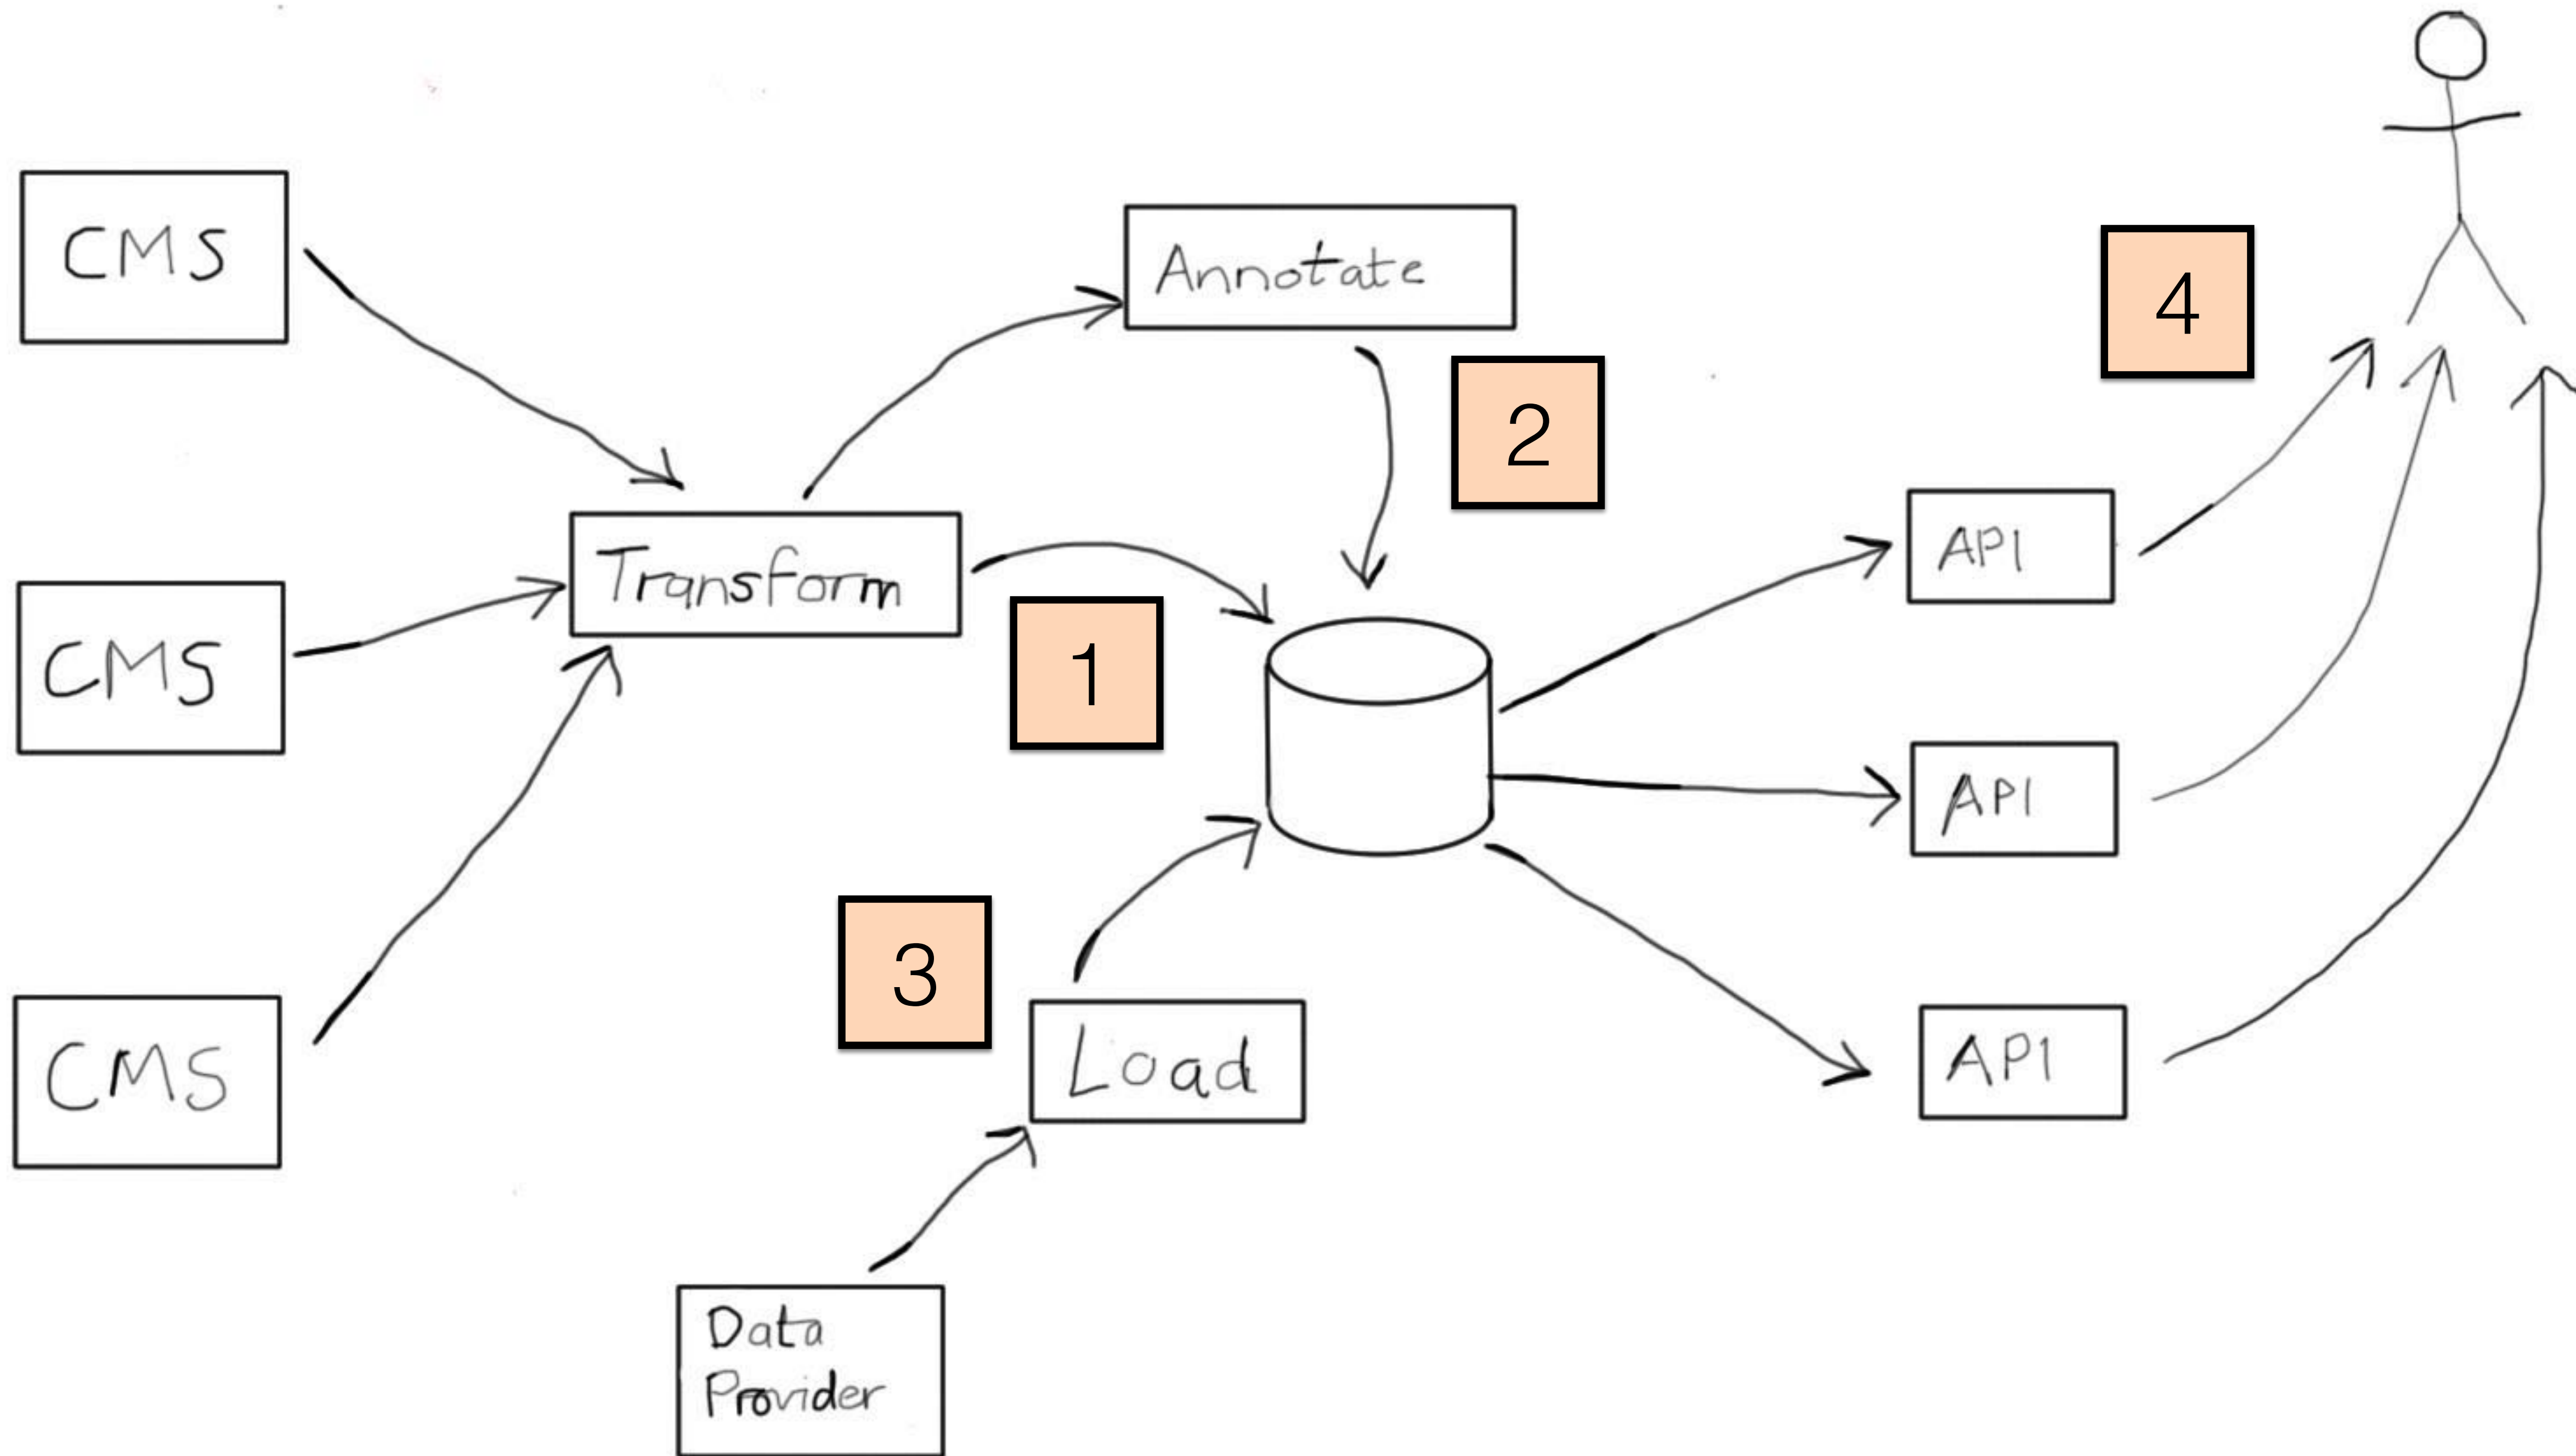






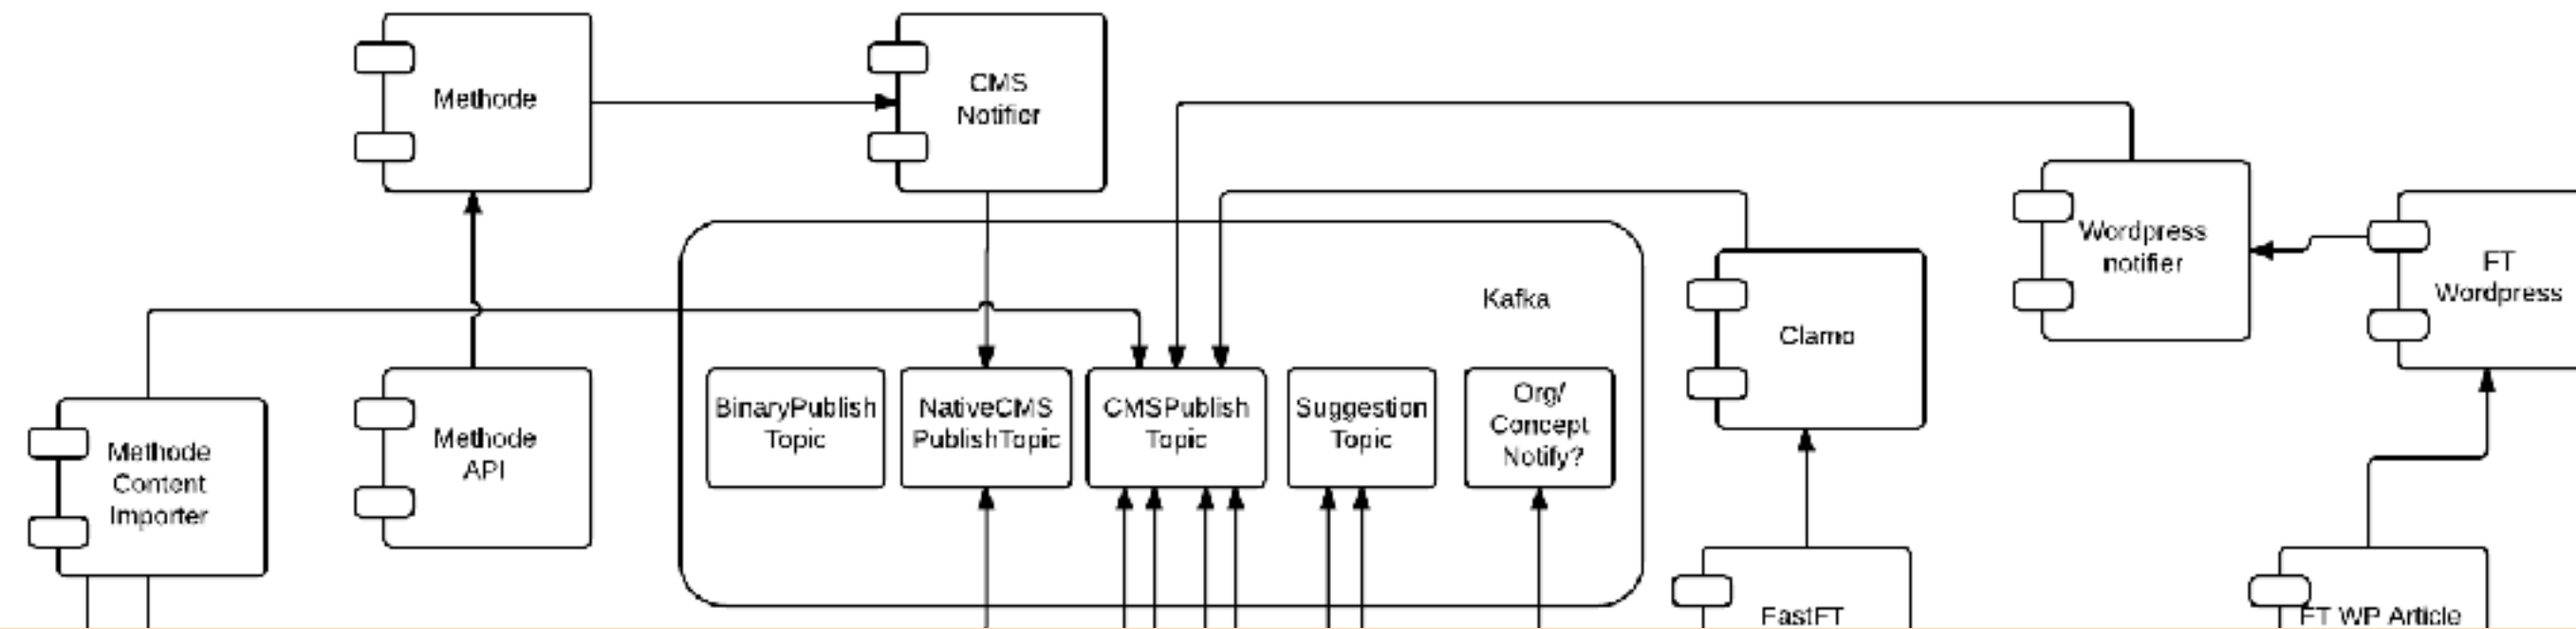




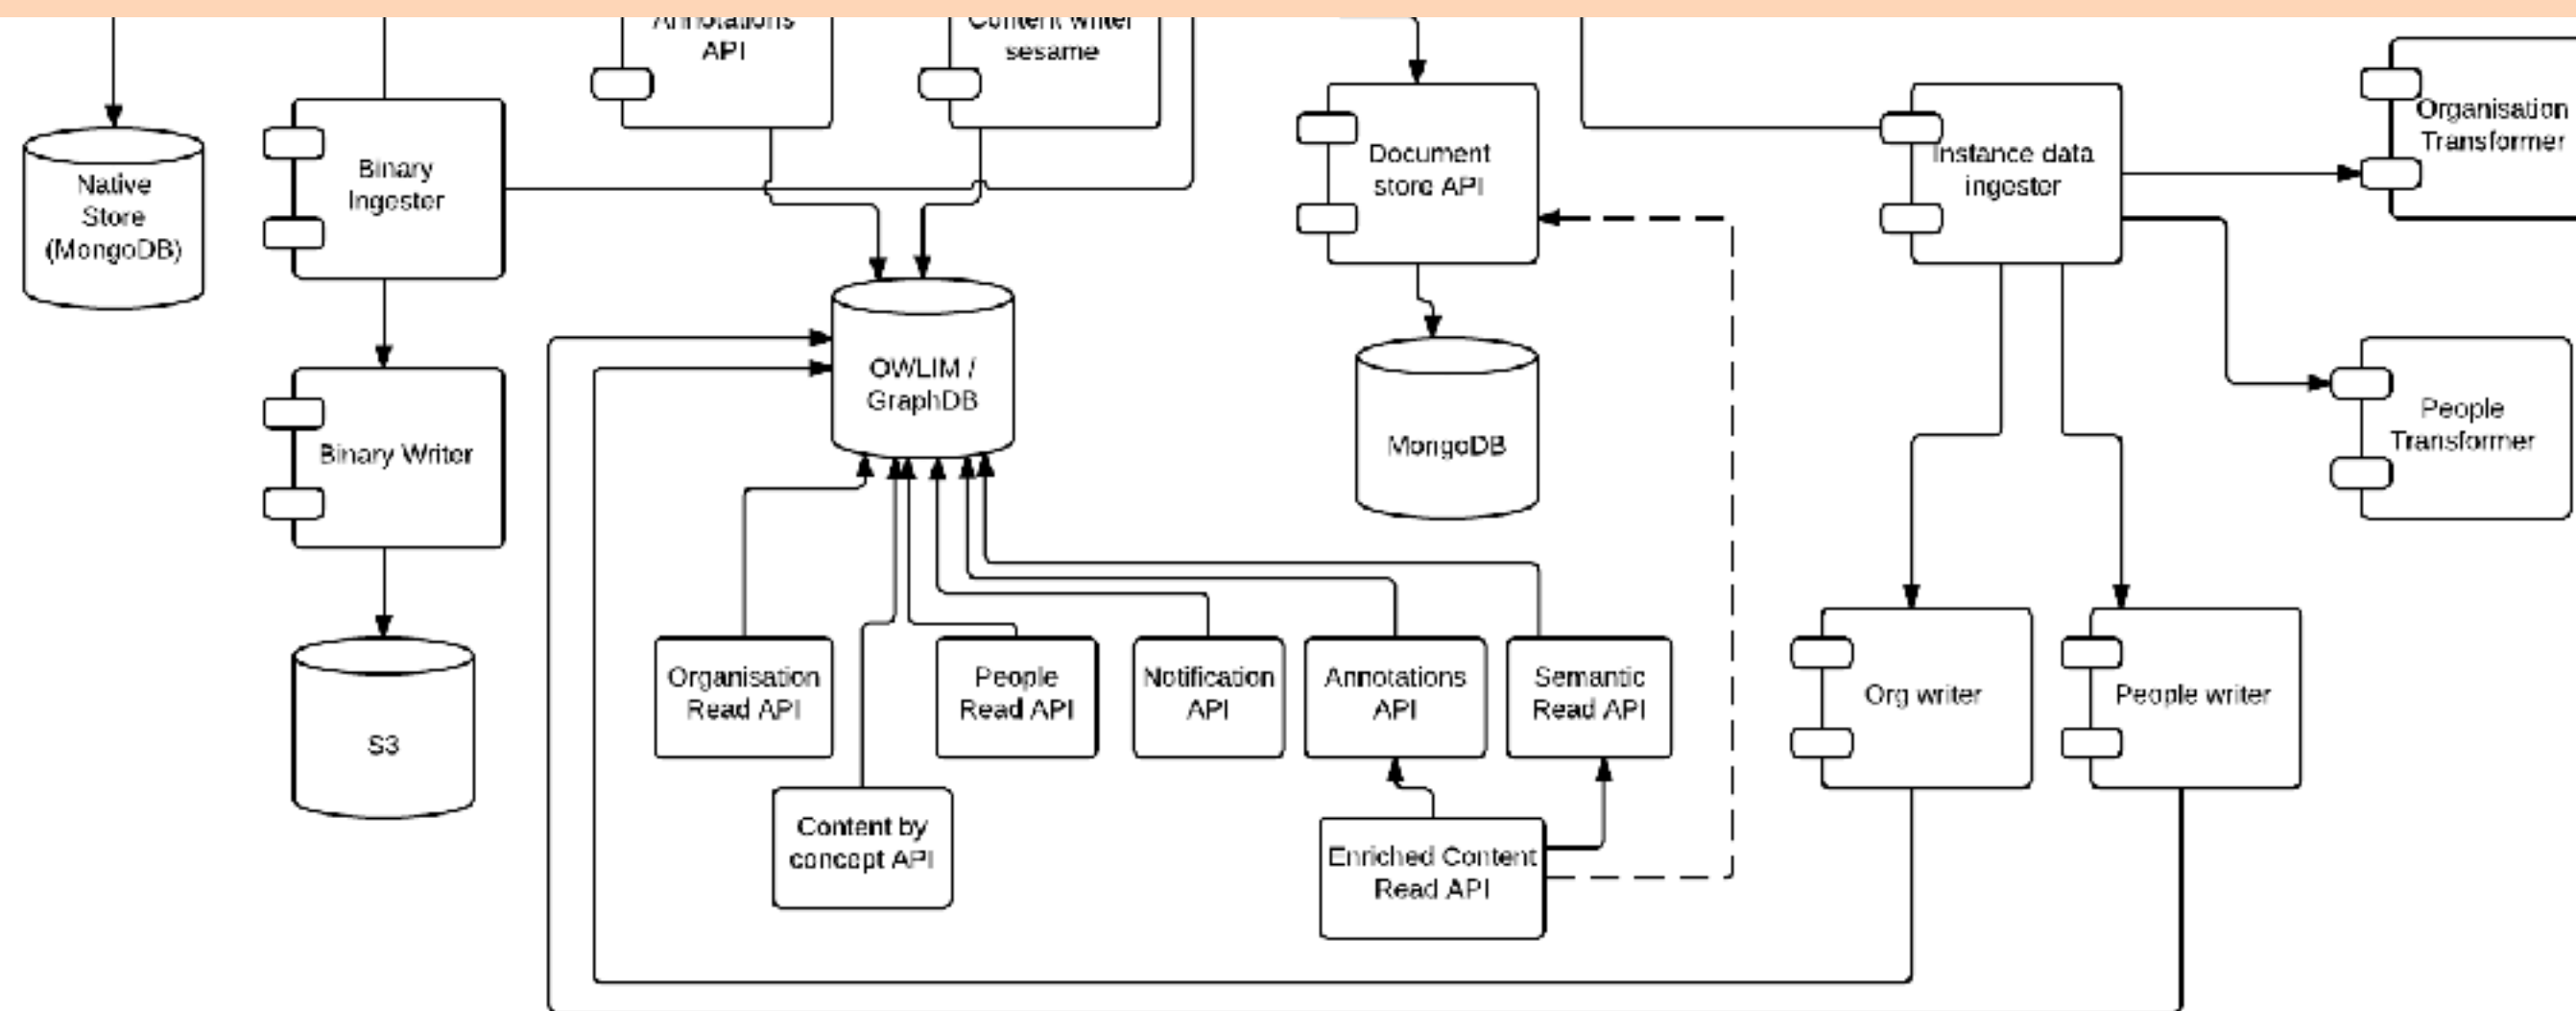


Monitoring this system...

@sarahjwells



Microservices make it worse



@sarahjwells

“microservices (n,pl): an efficient device for transforming **business** problems into **distributed transaction** problems”



@drsnooks

The services *themselves* are simple...

@sarahjwells

There's a lot of complexity around them

@sarahjwells

Why do they make monitoring harder?

@sarahjwells

You have a lot more services

@sarahjwells

99 functional microservices

350 running instances

52 non functional services

218 running instances

That's 568 separate services

@sarahjwells

If we checked each service every minute...

@sarahjwells

817,920 checks per day

@sarahjwells

What about *system* checks?

@sarahjwells

16,358,400 checks per day

@sarahjwells

“One-in-a-million” issues would hit us **16** times
every day

@sarahjwells

Running containers on shared VMs reduces this
to 92,160 system checks per day

@sarahjwells

For a total of 910,080 checks per day


@sarahjwells


It's a distributed system


@sarahjwells


Services are not independent

@sarahjwells











<http://devopsreactions.tumblr.com/post/122408751191/alerts-when-an-outage-starts>

You have to change how you think about
monitoring

@sarahjwells

How can you make it better?



1. Build a system you can support



The basic tools you need

@sarahjwells

Log aggregation

@sarahjwells

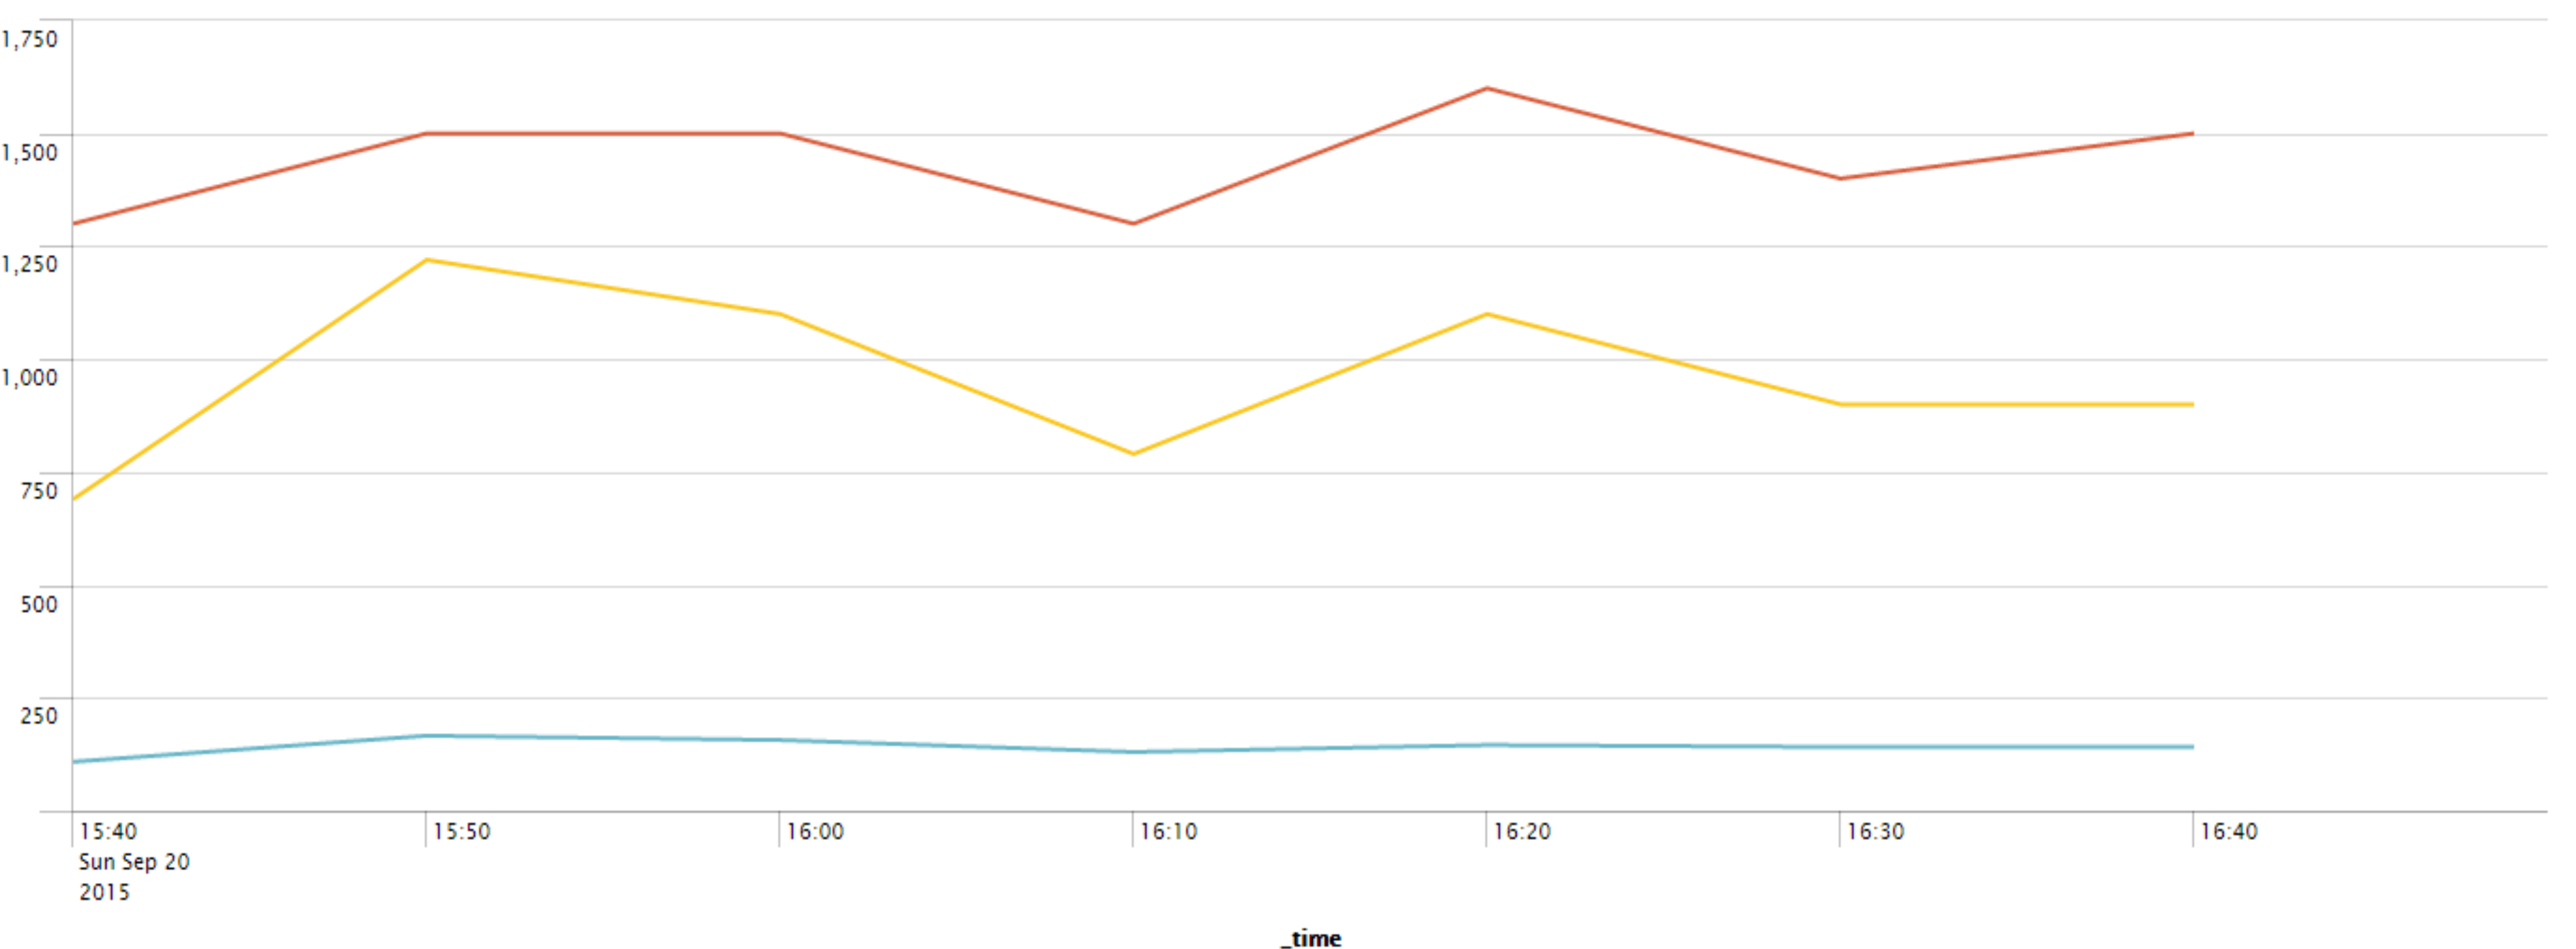
New Search

```
source="/var/log/apps/api-policy-component-dw-access.log" | timechart span=10m avg(dw_resptime), perc95(dw_resptime), perc99(dw_resptime)
```

✓ 39,304 events (20/09/2015 15:47:32.000 to 20/09/2015 16:47:32.000) Job [] [] [] [] []

Events (39,304) Patterns Statistics (7) Visualization

Line Format



Logs go missing or get delayed more now

@sarahjwells

Which means log based alerts may miss stuff

@sarahjwells

Monitoring

@sarahjwells

Service tier: **Bronze**
 Serving live traffic? **No**

Current Network Status
 Last Updated: Sun Sep 20 16:26:49 GMT 2015
 Updated every 90 seconds
 Nagios® Core™ 3.3.1 - www.nagios.org
 Logged in as nagiosadmin

Contact:
contentplatform.nagios.aler

[View History For all hosts](#)
[View Notifications For All Hosts](#)
[View Host Status Detail For All Hosts](#)

Host Status Totals

Up	Down	Unreachable	Pending
48	0	0	0
All Problems		All Types	
0		48	

Service Status Totals

Ok	Warning	Unknown	Critical	Pending
980	0	0	2	0
All Problems		All Types		
2		982		

General

- Home
- Documentation

Current Status

- Tactical Overview
- Map
- Hosts
- Services
- Host Groups
 - Summary
 - Grid
- Service Groups
 - Summary
 - Grid
- Problems
 - Services (Unhandled)
 - Hosts (Unhandled)
 - Network Outages

Quick Search:

Reports

- Availability
- Trends
- Alerts
 - History
 - Summary
 - Histogram
- Notifications
- Event Log

System

- Comments
- Downtime
- Process Info
- Performance Info
- Scheduling Queue
- Configuration

Service Status Details For All Hosts

Host	Service	Status	Last Check	Duration	Attempt	Status Information
app01-api-policy-component-iv-uk-p.content-platform.prod.cloud.ft.com	CPU	OK	20-09-2015 16:21:57	88d 13h 31m 44s	1/3	OK: CPU Idle = 98.56%
	DNS: resolve ntp.svc.ft.com on app01-api-policy-component-iv-uk-p.content-platform.prod.cloud.ft.com	OK	20-09-2015 16:24:22	88d 13h 29m 16s	1/3	DNS OK: 0.008 seconds response time. ntp.svc.ft.com: 10.117.101.254, 10.118.101.254, 10.160.70.4
	Disk: /	OK	20-09-2015 16:26:45	79d 17h 18m 3s	1/3	DISK OK - free space: / 7876 MB (83% inode=91%)
	Disk: /boot	OK	20-09-2015 16:23:23	88d 13h 31m 33s	1/3	DISK OK - free space: /boot 360 MB (77% inode=91%)
	Disk: /var	OK	20-09-2015 16:25:51	52d 16h 16m 47s	1/3	DISK OK - free space: /var 8744 MB (92% inode=91%)
	HTTP: ftapp06096-iv-uk-p:8081/healthcheck:OK:200:1:2.0:2.0	OK	20-09-2015 16:26:34	2d 6h 12m 26s	1/3	HTTP OK: Status line output matched "200" - 2521ms response time
	Load	OK	20-09-2015 16:25:54	79d 17h 18m 30s	1/3	OK - load average: 0.00, 0.00, 0.00
	NTP: ntp.svc.ft.com	OK	20-09-2015 16:23:47	21d 18h 13m 26s	1/3	NTP OK: Offset -0.0002371072769 secs
	Service: crond	OK	20-09-2015 16:24:03	88d 13h 31m 29s	1/3	crond (pid 1375) is running...
	Service: nrpe	OK	20-09-2015 16:26:35	88d 13h 29m 39s	1/2	nrpe (pid 10586 10585 1206) is running...
	Service: nscd	OK	20-09-2015 16:24:17	88d 13h 30m 10s	1/3	nscd (pid 1149) is running...
	Service: ntpd	OK	20-09-2015 16:22:36	88d 13h 29m 9s	1/3	ntpd (pid 1195) is running...
	Service: pe-mcollective	OK	20-09-2015 16:26:28	88d 13h 30m 9s	1/3	mcollectived (pid 6225) is running...
	Service: pe-puppet	OK	20-09-2015 16:24:19	53d 3h 17m 29s	1/3	puppet (pid 30204) is running...
	Service: rsyslog	OK	20-09-2015 16:24:29	88d 13h 30m 12s	1/3	rsyslogd (pid 1012) is running...
Service: splunk	OK	20-09-2015 16:26:32	88d 13h 29m 6s	1/3	Splunk status:	
Service: sshd	OK	20-09-2015 16:21:50	88d 13h 29m 9s	1/3	openssh-daemon (pid 1184) is running...	
Swap	OK	20-09-2015 16:25:41	79d 17h 18m 4s	1/3	SWAP OK - 99% free (2013 MB out of 2047 MB)	
app01-binary-ingester-iv-uk.content-platform.prod.cloud.ft.com	CPU	OK	20-09-2015 16:24:30	79d 17h 18m 21s	1/3	OK: CPU Idle = 99.06%
	DNS: resolve ntp.svc.ft.com on app01-binary-ingester-iv-uk.content-platform.prod.cloud.ft.com	OK	20-09-2015 16:23:58	88d 13h 29m 14s	1/3	DNS OK: 0.010 seconds response time. ntp.svc.ft.com: 10.117.101.254, 10.118.101.254, 10.160.70.4
	Disk: /	OK	20-09-2015 16:26:32	88d 13h 29m 34s	1/3	DISK OK - free space: / 7666 MB (81% inode=91%)
	Disk: /boot	OK	20-09-2015 16:26:41	88d 13h 31m 51s	1/3	DISK OK - free space: /boot 360 MB (77% inode=91%)
	Disk: /var	OK	20-09-2015 16:25:51	88d 13h 29m 38s	1/3	DISK OK - free space: /var 9109 MB (96% inode=91%)
	HTTP: ftapp06144-iv-uk-p:8081/healthcheck:OK:200:1:1.0:1.5	OK	20-09-2015 16:22:22	4d 2h 9m 26s	1/3	HTTP OK: Status line output matched "200" - 5391ms response time
	HTTP: localhost:8080/build-info:binary-ingester-service:200:1:1.0:1.5	OK	20-09-2015 16:23:23	88d 13h 29m 17s	1/3	HTTP OK: Status line output matched "200" - 6611ms response time
	HTTP: localhost:8081/ping:pong:200:1:1.0:1.5	OK	20-09-2015 16:26:29	79d 17h 18m 8s	1/3	HTTP OK: Status line output matched "200" - 1761ms response time
	Load	OK	20-09-2015 16:26:34	79d 17h 18m 31s	1/3	OK - load average: 0.00, 0.00, 0.00
	NTP: ntp.svc.ft.com	OK	20-09-2015 16:24:28	22d 23h 28m 44s	1/3	NTP OK: Offset -0.0004309415817 secs
	Service: crond	OK	20-09-2015 16:24:20	88d 13h 29m 9s	1/3	crond (pid 1354) is running...

Limitations of our nagios integration...

@sarahjwells

No 'service-level' view

@sarahjwells

Default checks included things we couldn't fix

@sarahjwells

A new approach for our container stack

@sarahjwells

We care about each service

@sarahjwells

CoCo semantic cluster's default services are **unhealthy** (1 services acked)

Search: public

Fleet Name	Is Healthy	Last Updated	Ack msg
content-public-read-preview@1	OK	13:17:57 UTC	
content-public-read-preview@2	OK	13:17:57 UTC	
content-public-read@1	OK	13:17:57 UTC	
content-public-read@2	OK	13:17:57 UTC	
public-annotations-api@1	OK	13:17:57 UTC	
public-annotations-api@2	OK	13:17:57 UTC	
public-brands-api@1	OK	13:17:57 UTC	
public-brands-api@2	OK	13:17:57 UTC	
public-concordances-api@1	OK	13:17:57 UTC	
public-concordances-api@2	OK	13:17:57 UTC	
public-content-by-concept-api@1	OK	13:17:57 UTC	
public-content-by-concept-api@2	OK	13:17:57 UTC	
public-organisations-api@1	OK	13:17:57 UTC	
public-organisations-api@2	OK	13:17:57 UTC	
public-people-api@1	OK	13:17:57 UTC	

We care about each VM

@sarahjwells

Fleet Name ▲	Is Healthy ◆	Last Updated ◆	Ack msg ◆
system-healthcheck-ip-172-24-149-166.eu-west-1.compute.internal	OK	13:17:57 UTC	
system-healthcheck-ip-172-24-161-130.eu-west-1.compute.internal	OK	13:17:57 UTC	
system-healthcheck-ip-172-24-23-157.eu-west-1.compute.internal	OK	13:17:57 UTC	
system-healthcheck-ip-172-24-75-37.eu-west-1.compute.internal	OK	13:17:57 UTC	
system-healthcheck-ip-172-24-84-0.eu-west-1.compute.internal	OK	13:17:57 UTC	

```
app := cli.App("System-healthcheck", "A service that report on current VM status at __health")

hostPath = app.String(cli.StringOpt{
    Name:  "hostPath",
    Value: "",
    Desc:  "The dir path of the mounted host fs (in the container)",
    EnvVar: "SYS_HC_HOST_PATH",
})

checks = append(checks, diskFreeChecker{20}.Checks()...)
checks = append(checks, memoryChecker{20}.Checks()...)
checks = append(checks, loadAverageChecker{}.Checks()...)
checks = append(checks, ntpChecker{}.Checks()...)
checks = append(checks, tcpChecker{}.Checks()...)
checks = append(checks, versionChecker{}.Checks()...)
```

We care about unhealthy instances

@sarahjwells

Monitoring needs aggregating somehow

@sarahjwells

SAWS

@sarahjwells



Built by Silvano Dossan

See our Engine room blog: <http://bit.ly/1GATHLy>

"I imagine most people do exactly what I do -
create a google filter to send **all Nagios emails**
straight to the bin"

@sarahjwells

"Our screens have a viewing angle of about
10 degrees"

@sarahjwells

"It never seems to show the page I want"

@sarahjwells



Code at: <https://github.com/muce/SAWS>

@sarahjwells

Dashing

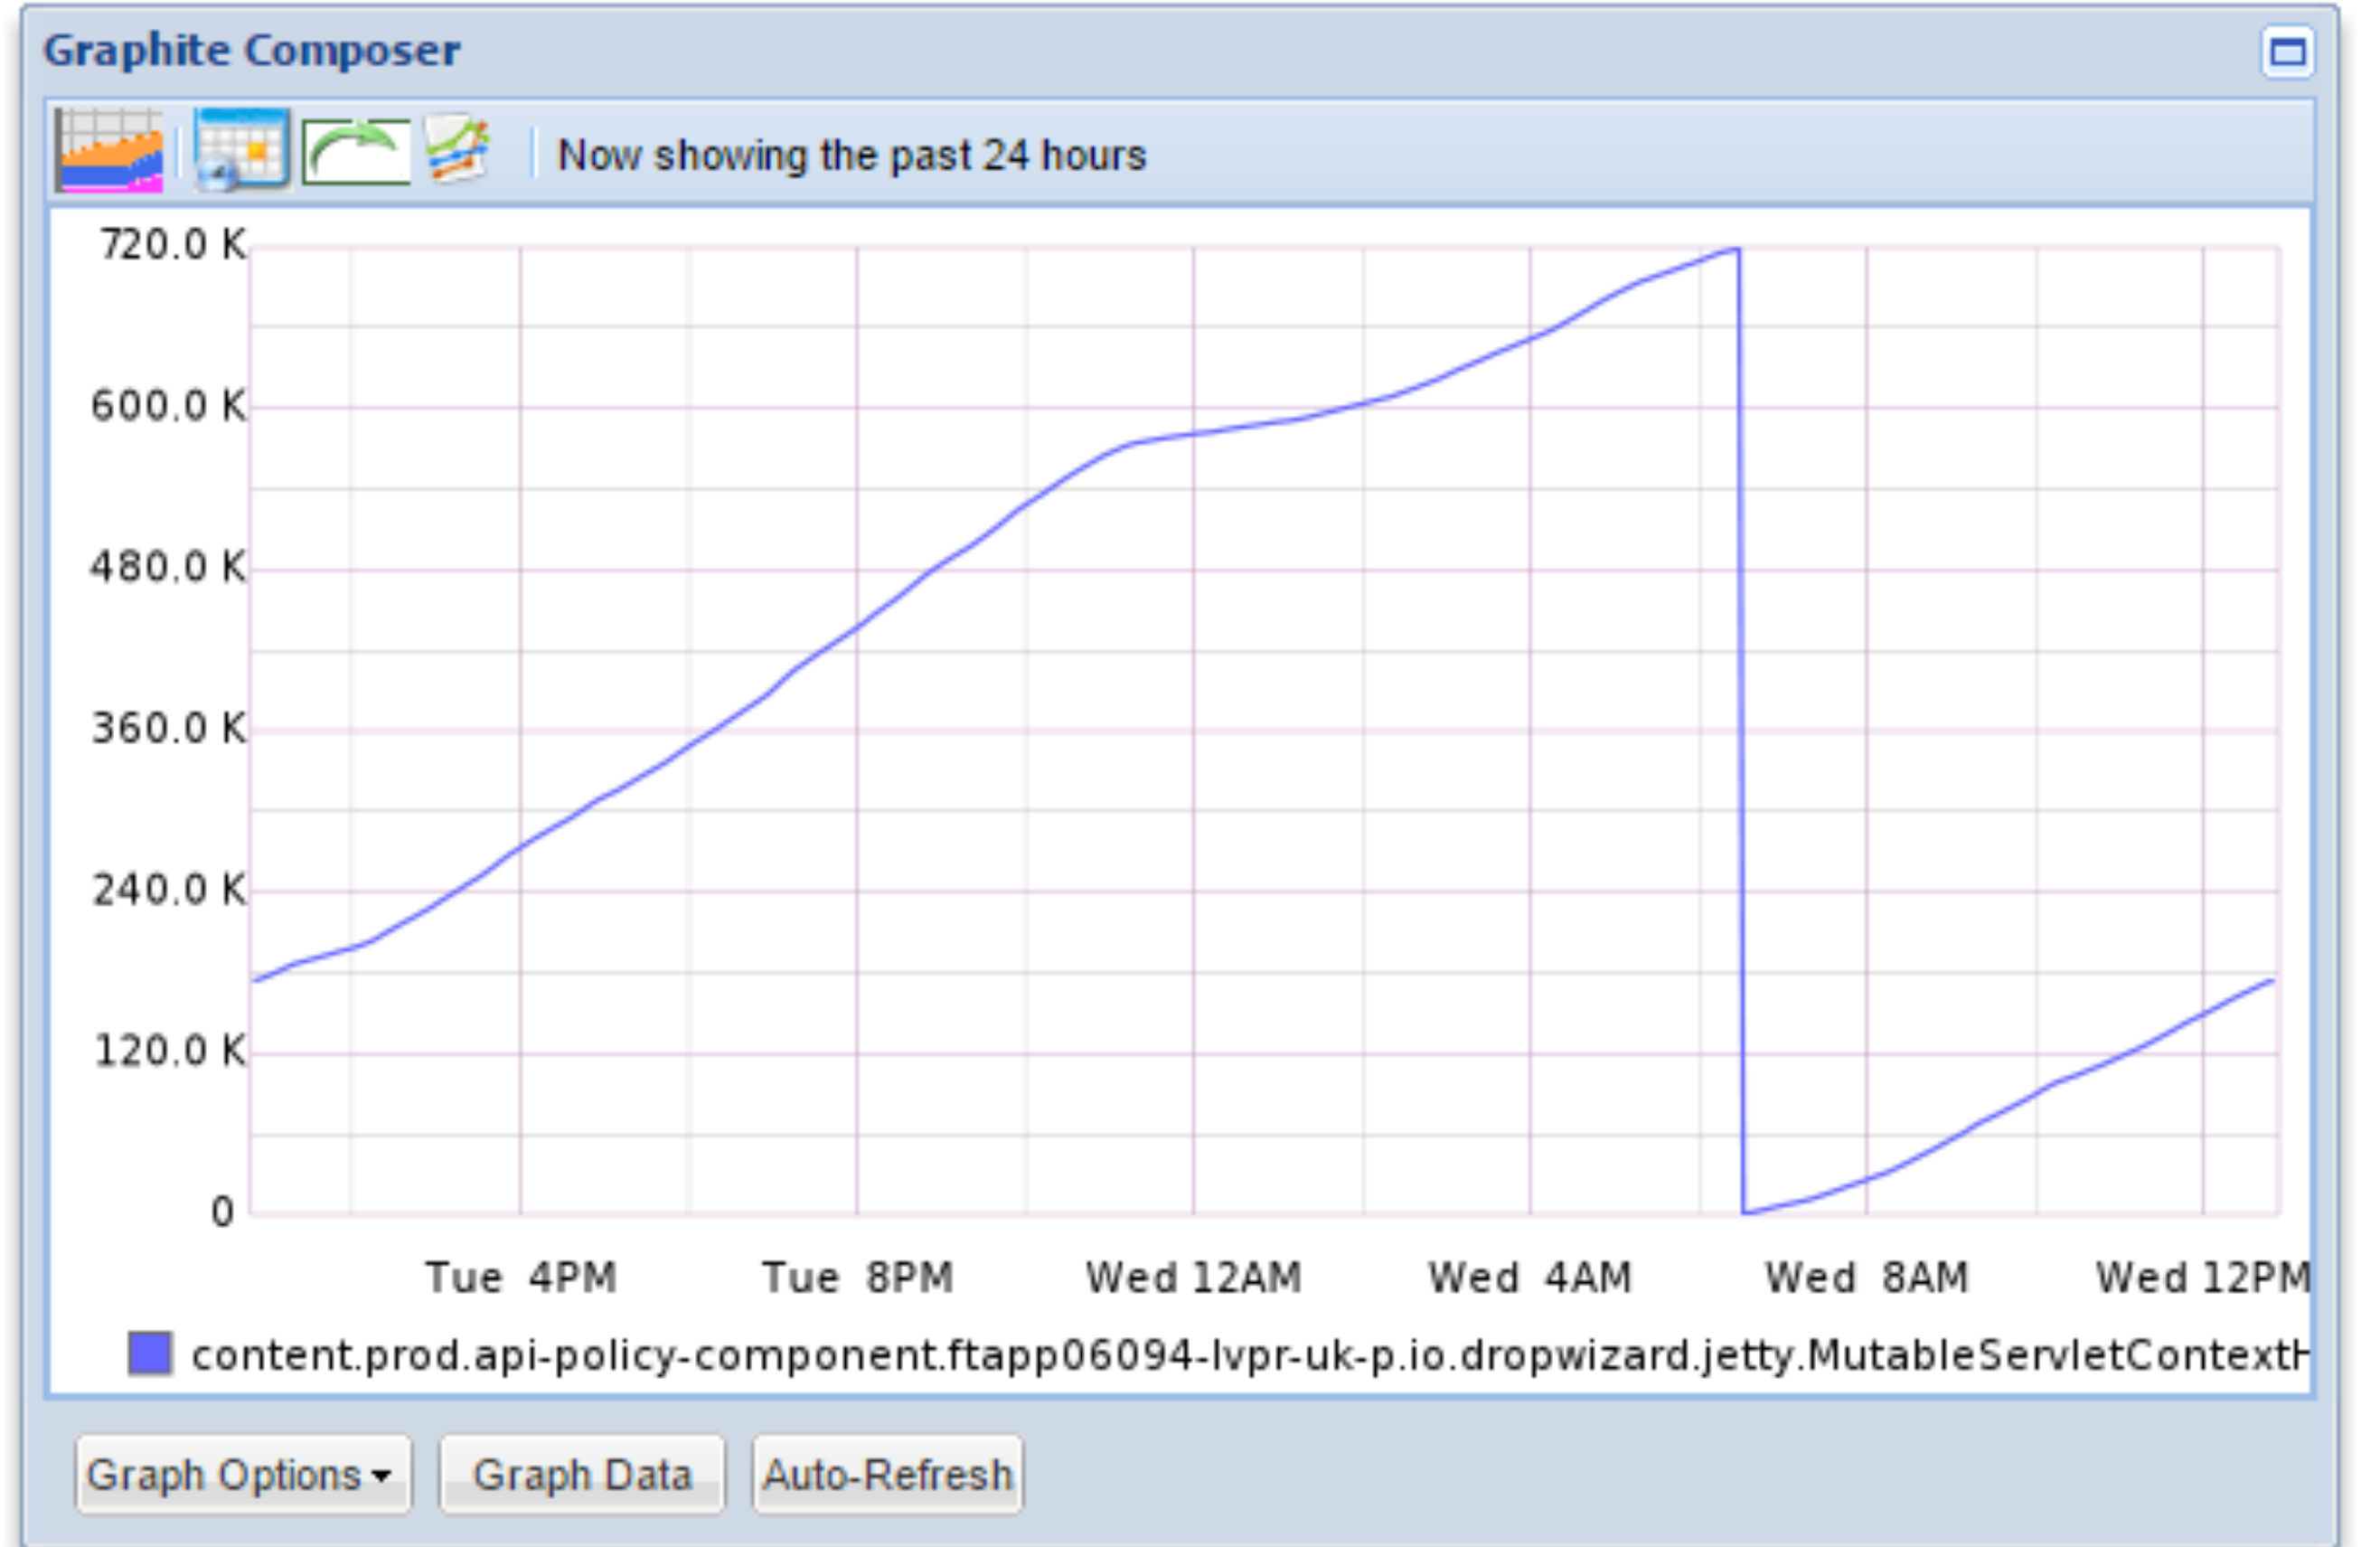
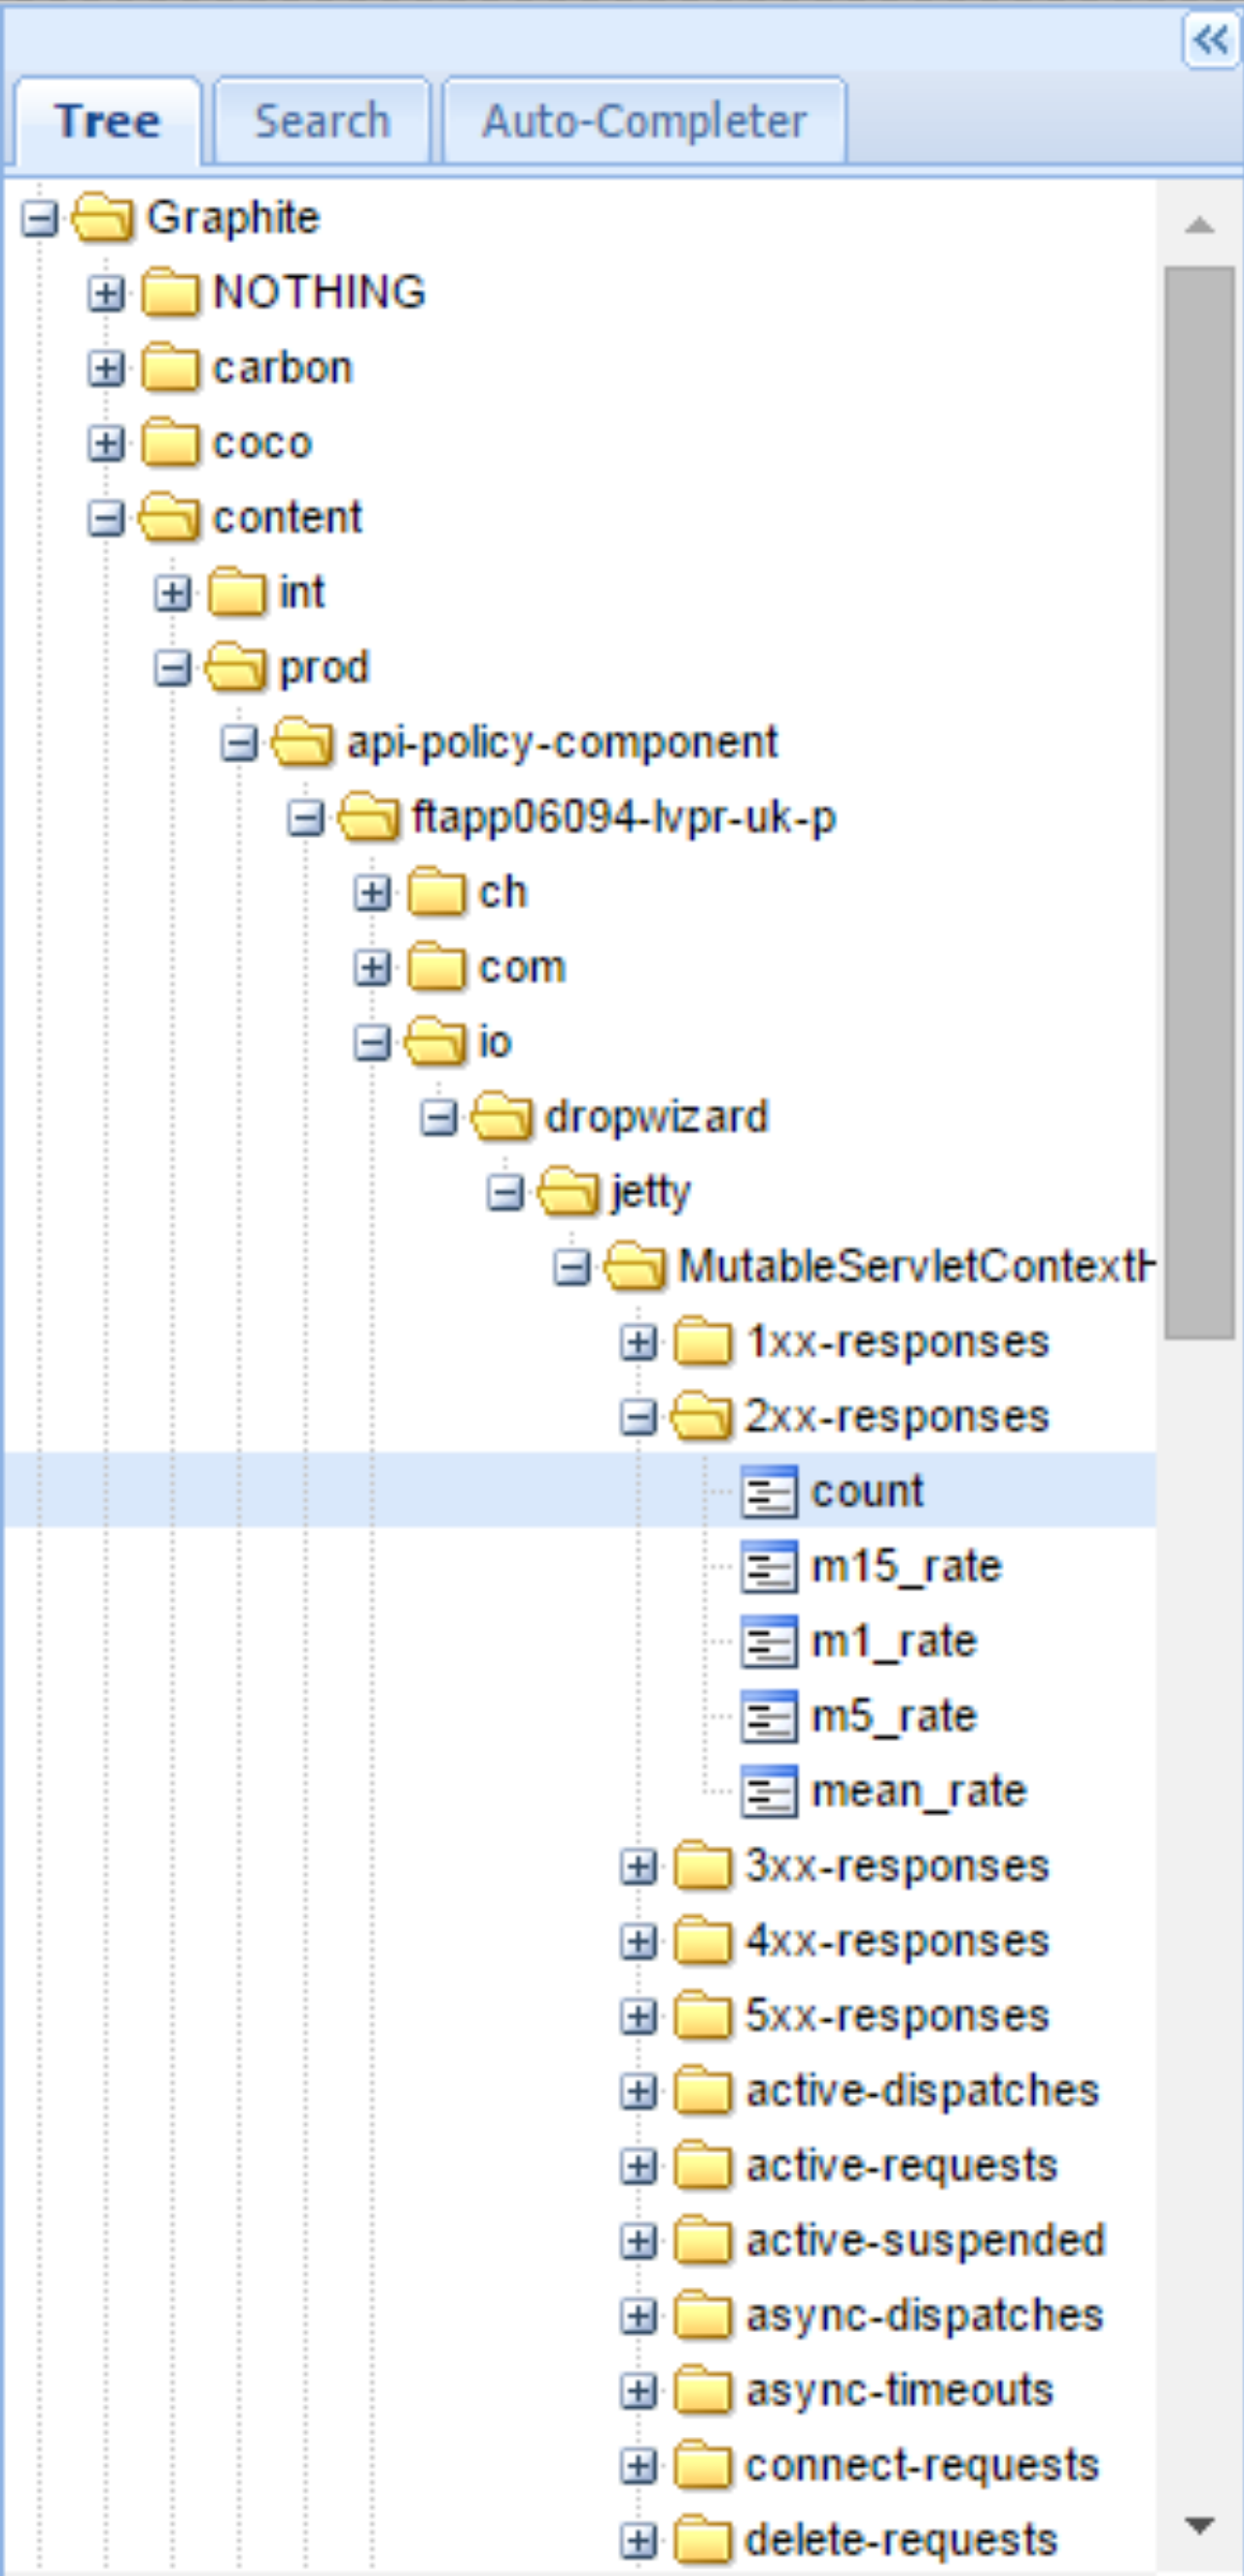
@sarahjwells

<p>CO-CO Prod UK</p> <p>OK Healthy</p> <p>Last updated at 12:42</p>	<p>CO-CO Prod US</p> <p>OK Healthy</p> <p>Last updated at 12:42</p>	<p>CO-CO Prod UK lists publish</p> <p>OK Healthy</p> <p>Last updated at 12:43</p>	<p>CO-CO Prod US lists publish</p> <p>OK Healthy</p> <p>Last updated at 12:43</p>	<p>CO-CO Prod UK lists read</p> <p>OK Healthy</p> <p>Last updated at 12:43</p>	<p>CO-CO Prod US lists read</p> <p>OK Healthy</p> <p>Last updated at 12:43</p>	<p>CO-CO Prod UK image publish</p> <p>OK Healthy</p> <p>Last updated at 12:43</p>
<p>CO-CO Prod US image publish</p> <p>OK Healthy</p> <p>Last updated at 12:43</p>	<p>CO-CO Prod UK content publish</p> <p>OK Healthy</p> <p>Last updated at 12:43</p>	<p>CO-CO Prod US content publish</p> <p>OK Healthy</p> <p>Last updated at 12:43</p>	<p>CO-CO Prod UK content read</p> <p>OK Healthy</p> <p>Last updated at 12:43</p>	<p>CO-CO Prod US content read</p> <p>OK Healthy</p> <p>Last updated at 12:43</p>	<p>CO-CO Prod UK enriched-content read</p> <p>OK Healthy</p> <p>Last updated at 12:43</p>	<p>CO-CO Prod US enriched-content read</p> <p>OK Healthy</p> <p>Last updated at 12:43</p>
<p>CO-CO Prod UK concordances read</p> <p>OK Healthy</p> <p>Last updated at 12:43</p>	<p>CO-CO Prod US concordances read</p> <p>OK Healthy</p> <p>Last updated at 12:43</p>	<p>CO-CO Publishing Prod UK</p> <p>OK Healthy</p> <p>Last updated at 12:43</p>	<p>CO-CO Publishing Prod US</p> <p>OK Healthy</p> <p>Last updated at 12:43</p>	<p>CO-CO Pre-Prod UK</p> <p>WARNING Severity 2 Issue.</p> <p>Last updated at 11:02</p>	<p>CO-CO Pre-Prod US</p> <p>WARNING Severity 2 Issue.</p> <p>Last updated at 11:02</p>	<p>CO-CO Publishing Pre-Prod UK</p> <p>OK Healthy</p> <p>Last updated at 12:43</p>
<p>CO-CO Publishing Pre-Prod US</p> <p>OK Healthy</p> <p>Last updated at 12:43</p>	<p>CO-CO UP-Registry</p> <p>OK Healthy</p> <p>Last updated at 12:43</p>	<p>CO-CO DynPub</p> <p>WARNING Severity 2 Issue.</p> <p>Last updated at 11:02</p>	<p>CO-CO XP</p> <p>CRITICAL Error while retrieving cluster health</p> <p>Last updated at 12:43</p>	<p>CO-CO Publishing XP</p> <p>CRITICAL Error while retrieving cluster health</p> <p>Last updated at 12:43</p>	<p>CO-CO Semantic</p> <p>WARNING Severity 2 Issue.</p> <p>Last updated at 11:02</p>	

CO-CO Prod UK llists publish	CO-CO Prod US llists publish	CO-CO Prod UK llists read OK	CO-CO Prod US llists read OK	CO-CO Prod UK Image publish	CO-CO Prod US Image publish	CO-CO Prod UK content publish	CO-CO Prod US content publish	CO-CO Prod UK content read	CO-CO Prod US content read	CO-CO Prod UK enriched-content read	CO-CO Prod US enriched-content read	CO-CO Prod UK concordance read	CO-CO Prod US concordance read	CO-CO Publishing Prod Active	CO-CO Publishing Prod Active	Method- Staging- Prod- (PR/IW) Editorial	Method- Staging- Prod- (PR/IW) Editorial	Web-App- API- (EU/US) FT.COM	Web-App- API- (EU/US) FT.COM
Session-API- (EU) Membership	Session-API- (US) Membership	Cassandra- EU Membership	Cassandra- US Membership	Access-EU Membership	Access-US Membership	Session- Service-UK Membership	Session- Service-US Membership	Content(WT) Content	Content(NJ) Content	Content- (PR) Content	Render- DB-(WT)	SearchFeeder (WT) Content	SearchService (WT) Content	Content- NFS- (MariaDB- Backups) Content	Nipa/Press- cuttings Content	Preditor- (PR) Content	Render	SearchFeeder Content	SearchService Content
Preditor- (WT) Content	API-Gateway- (EU) Environments	API-Gateway- (US) Environments	Graphite Environments	Grafana Environments	Print-slte	APIG PLATINUM OK Healthy	build- service	Fastly API Gateway	FT App SSR & UI	Next FT.com apl in EU	Next FT.com classification- api in EU	Next FT.com classification- api in US	Next FT.com es- interface In EU	Next FT.com es- interface In US	Next FT.com front- page In EU	Next FT.com front- page In US	Next FT.com origami- navigation- v1 in	Next FT.com origami- navigation- v1 in	user- profile- svc
user- profile- svc	UsrProductSv	UsrProductSv	Web App Api	Wires- Services Editorial Technology	Wires- Services Editorial Technology	Method- msrender Editorial Technology	Method- msrender Editorial Technology	Claro- Prod-(WT) Editorial Technology	Claro- Prod-(PR) Editorial Technology	MPS/FTAdviser login-barrier- Specialist Ti	MPS/FTAdviser login-barrier- Specialist Ti	Federated- SSO Membership	Blogs ft.com	Quick- Metadata- Interface- (WT)	Quick- Metadata- Interface- (PR)	Syndication- Services- (WT) Content	OM-Apps- nagios-(EU- West) Membership	OM-Apps- nagios-(US- East) Membership	Dashing-IV Environment
Dashing-PR Environments	s3o Environments	s3o Environments	Talend-Prod Environments Critical.. 2	Tooling- Team-Nagios (PROD) Environments	SAML- Authenticator Environments	SMTP- Service-(WT) Environments	SMTP- Service-(NJ) Environments	SMTP- Service-(OSB) Environments	SMTP- Service-(PR) Environments	APIG GOLD OK Healthy	Email Platform	OfferApi	OfferApi	origami- image- service	origami- image- service	Web App Admin	Method- Archive Editorial Technology	Method- Swing Editorial Technology	Wires- Method- IW-or-PR Editorial Technology
Method- Archive Editorial Technology	Method- Swing Editorial Technology	Wires- Method- IW-or-PR Editorial Technology	IC-Mobile Speciallist Titles	Analyse- Africa-AWS Specialist Titles	FTLive Speciallist Titles	origami- imageservice	rankingspubli prod-cl Legacy deb	WPT-PR Environments	WPT-IW Environments	Infra-Unix- Servers Environments	APIG SILVER OK Healthy	Contact Organiser	ft-app- access	System Registry	Method- Toolbox Editorial Technology	EZ-Sites Speciallist Titles	Minisites Speciallist Titles	Web-App- testing-& deployments ft.com	Federated- SSO Membership
Guest- Pass-/-Gift- Article- Membership	Email- Provider- Facade Membership	Email- Provider- Facade Membership	SSO-De- provisioning- (UK) Membership	sso- acceptance- tests- prod- uk- Membership	Reglstration- nagios-(UK) Membership	Reglstration- nagios-(US) Membership	SSO-Email- Service Membership	Ads- Nagios Ads	Feeds-(IW)	Feeds	Dynamic- Publishing- IW Content	Dynamic- Publishing- PR Content	Recommend- Reads-AWS Content	Data- Platform- Prod Data	S3- Unload- Prod Digital	Healthcheck- Aggregator Environments	CI-Tools	dba-monyog- prod Environments	OI-testing-& deployment environment
splunkforward prod-cloud Environment	splunkforward prod-cloud Environment	FTP1-Prod Environments	CO-CO Prod UK OK Healthy	CO-CO Prod US OK Healthy	CO-CO Publishing Prod UK OK	CO-CO Publishing Prod US OK	APIG BRONZE OK Healthy	AnonEmailCo	AnonEmailLis	AnonEmailSv	CMDB Key Manager	Content Comparator	CustomerSso	CustomerSso	EmailFacade	EmailFacade	ftalphaville	ftalphaville- es- interface- service	google- amp
Hui API	Next FT.com affinity- api in	Next FT.com barrier- guru in EU	Next FT.com encryption- api in	Next FT.com rfv- api in EU	Next FT.com video- page in	PaymentMeth	PaymentMeth	Secure Upload	tagbot	UserDetailsS	UserDetailsS	UserSubsStat	usersvc- usersvc_euw prod-cloud	usersvc- usersvc_use prod-cloud	userprofilesvc prod-cloud	userprofilesvc prod-cloud	admintools_iv prod-cloud	admintools_pi prod-cloud	cofup_app_iv prod-cloud
cofup_app_pr	uk-prod- omans	us-prod- omans	lp_ftcom_elast	lp_ftcom_elast	apig-prod- aws-ami-in	modules- prod-uk	modules- prod-us	awslogin_s3o	pwm_proden	armz	login-app- acceptance-	login-app- acceptance-	login-api- acceptance-	login-api- acceptance-	External Authentication	api-authz- svc-	api-authz- svc-	Authorization	Authorization

Graphing of metrics

@sarahjwells






environment: prod host: All

ERROR RATES

5XX responses (1 minute rate)




min max avg

ftapp06094-lvpr-uk-p	0	0	0
ftapp06095-lvpr-uk-p	0	0	0
ftapp06096-lviw-uk-p	0	0	0
ftapp06097-lviw-uk-p	0	0	0

5XX over 1 hour

0 req

4XX responses (1 min rate)




min max avg

ftapp06094-lvpr-uk-p	0	29	7
ftapp06095-lvpr-uk-p	0	24	7
ftapp06096-lviw-uk-p	1	33	15
ftapp06097-lviw-uk-p	2	36	15

4XX over 1 hour

1302 req

200XX responses (1 min rate)



min max avg

ftapp06094-lvpr-uk-p	103	243	153
ftapp06095-lvpr-uk-p	104	247	154
ftapp06096-lviw-uk-p	101	249	179
ftapp06097-lviw-uk-p	103	237	179


2XX over 1 hour

9997 req

io.dropwizard.jetty.MutableServletContextHandler

REQUEST RATE BY HOST (REQ/SEC)


notifications



min max avg

ftapp06094-lvpr-uk-p	0.13	0.45	0.26
ftapp06095-lvpr-uk-p	0.18	0.43	0.26
ftapp06096-lviw-uk-p	0.10	0.35	0.23
ftapp06097-lviw-uk-p	0.10	0.39	0.24


content



min max avg

ftapp06094-lvpr-uk-p	0.07	1.05	0.37
ftapp06095-lvpr-uk-p	0.06	1.02	0.38
ftapp06096-lviw-uk-p	0.27	1.25	0.77
ftapp06097-lviw-uk-p	0.26	1.18	0.76


enriched content



min max avg

ftapp06094-lvpr-uk-p	0.93	2.36	1.45
ftapp06095-lvpr-uk-p	0.86	2.21	1.34
ftapp06096-lviw-uk-p	1.08	2.67	2.08
ftapp06097-lviw-uk-p	0.95	2.07	1.64

other




min max avg

ftapp06094-lvpr-uk-p	0.010	0.180	0.073
ftapp06095-lvpr-uk-p	0.010	0.150	0.061
ftapp06096-lviw-uk-p	0.020	0.190	0.117
ftapp06097-lviw-uk-p	0.040	0.190	0.116




environment: prod-uk services: All


Memory Available



CPU Idle Time Percent



Disk Space Free






environment: prod-uk + semantic

services: All


Unhealthiness of Services prod-uk




Aggregate Healthcheck Up prod-uk



Unhealthiness of Services semantic



Aggregate Healthcheck Up semantic



+ ADD ROW



<https://www.flickr.com/photos/davidmasters/2564786205/>

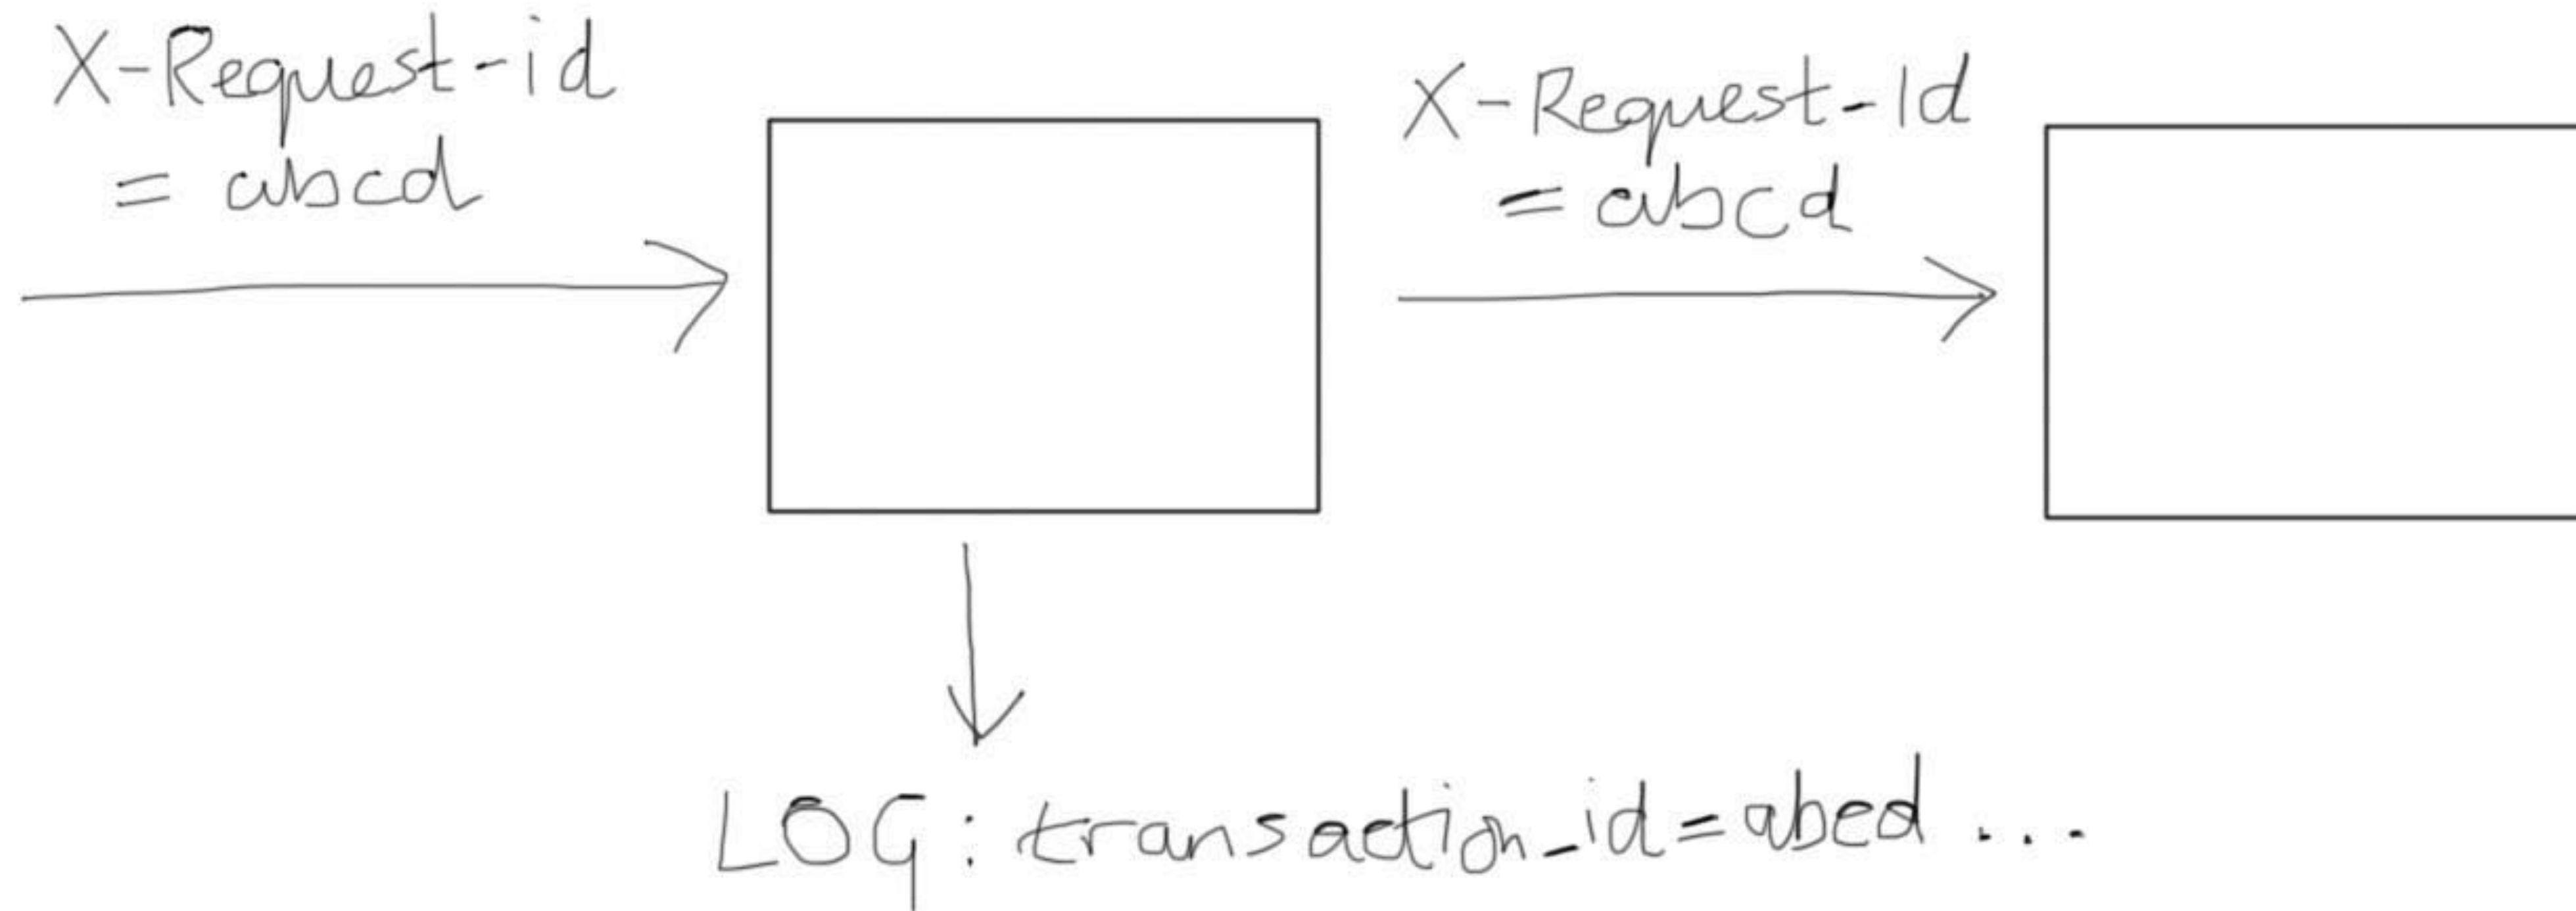
The things that make those tools **WORK**

@sarahjwells

Effective log aggregation needs a way to find all
related logs

@sarahjwells

Transaction ids tie all microservices together



Make it easy for any language you use

@sarahjwells

transaction_id=1C74FEC2-F7C5-89C3-4F57-0FB6A914D42D

Date time range ▼



✓ 22 events (23/10/2015 23:59:55.000 to 24/10/2015 00:00:00.000)

Job ▼ || ■ ↶ ⬇ 🖨 Verbose Mode ▼

Events (22) Patterns Statistics Visualization

Format Timeline ▼ — Zoom Out + Zoom to Selection × Deselect

100 milliseconds per column



List ▼ Format ▼ 50 Per Page ▼

< Hide Fields ☰ All Fields

Selected Fields

a host 4
a source 3

Interesting Fields

a category 6
date_hour 1
date_mday 1
date_minute 1
a date_month 1
date_second 3
a date_wday 1
date_year 1
a date_zone 1

<i>i</i>	Time	Event
>	23/10/2015 23:59:59.399	INFO [2015-10-23 23:59:59,399] com.ft.api.util.transactionid.TransactionIdFilter: transaction_id=1C74FEC2-F7C5-89C3-4F57-0FB6A914D42D [REQUEST HANDLED] uri=/enrichedcontent/96fbf7bc-702f-11e5-9b9e-690fdae72044 time_ms=3729 status=200 exception_was_thrown=false [dw-2053] host = ftapp06097-lviw-uk-p source = /var/log/apps/api-policy-component-dw-app.log
>	23/10/2015 23:59:59.395	INFO [2015-10-23 23:59:59,395] com.ft.jerseyhttpwrapper.ResilientClient: transaction_id=1C74FEC2-F7C5-89C3-4F57-0FB6A914D42D [REQUEST FINISHED] short_name=enrichedcontent, outcome=200, total_attempts=1, failed_attempts=0 [dw-2053 - GET /enrichedcontent/96fbf7bc-702f-11e5-9b9e-690fdae72044] host = ftapp06097-lviw-uk-p source = /var/log/apps/api-policy-component-dw-app.log
>	23/10/2015 23:59:59.385	INFO [2015-10-23 23:59:59,385] com.ft.api.util.transactionid.TransactionIdFilter: transaction_id=1C74FEC2-F7C5-89C3-4F57-0FB6A914D42D, [REQUEST HANDLED] uri=/enrichedcontent/96fbf7bc-702f-11e5-9b9e-690fdae72044 time_ms=3687 status=200 exception_was_thrown=false host = ftaps33815-law1a-eu-p source = /var/log/apps/enriched-content-read-api-dw-app.log
>	23/10/2015 23:59:59.384	INFO [2015-10-23 23:59:59,384] com.ft.api.util.transactionid.TransactionIdFilter: transaction_id=1C74FEC2-F7C5-89C3-4F57-0FB6A914D42D, [REQUEST HANDLED] uri=/content/96fbf7bc-702f-11e5-9b9e-690fdae72044/annotations time_ms=2642 status=200 exception_was_thrown=false host = ftaps32100-law1a-eu-p source = /var/log/apps/annotations-api-dw-app.log

Services need to report on their own health

@sarahjwells

The FT healthcheck standard

GET http://{service}/__health

The FT healthcheck standard

GET http://{service}/__health

returns 200 if the service can run the healthcheck

The FT healthcheck standard

GET http://{service}/__health

returns 200 if the service can run the healthcheck

each check will return "ok": true or "ok": false

```
{
  "schemaVersion": 1,
  "name": "IngesterApplication",
  "description": "IngesterApplication",
  "checks": [
    {
      "ok": true,
      "checkOutput": "Kafka is connected and topic is present",
      "panicGuide": "https://sites.google.com/a/ft.com/dynamic-publishing-team/ingester-panic-guide",
      "lastUpdated": "2015-10-25T17:46:11Z",
      "severity": 1,
      "businessImpact": "Content being published by journalists - for example, Methode articles and FastFT articles - will not be written to the Mongo Content Writer",
      "technicalSummary": "Tests that Kafka responds with correct Topic",
      "name": "Can connect to Kafka"
    },
    {
      "ok": true,
      "checkOutput": "OK",
      "panicGuide": "https://sites.google.com/a/ft.com/dynamic-publishing-team/ingester-panic-guide",
      "lastUpdated": "2015-10-25T17:46:11Z",
      "severity": 1,
      "businessImpact": "Content being published by journalists - for example, Methode articles and FastFT articles - will not be written to the Mongo Content Writer",
      "technicalSummary": "Tests that the ping endpoint for the Mongo Content Writer returns pong and a 200 response",
      "name": "Can connect to the Mongo Content Writer"
    },
    {
      "ok": true,
      "checkOutput": "ZooKeeper instance responded, at least 2 are up.",
      "panicGuide": "https://sites.google.com/a/ft.com/dynamic-publishing-team/ingester-panic-guide",
      "lastUpdated": "2015-10-25T17:46:12Z",
      "severity": 1,
      "businessImpact": "Content being published by journalists - for example, Methode articles and FastFT articles - will not be written to the Mongo Content Writer",
      "technicalSummary": "Tests that ZooKeeper is up and responding",
      "name": "ZooKeeper is up and responding"
    }
  ]
}
```

Service health status

IngesterApplication: IngesterApplication



All checks OK



Can connect to Kafka

a few seconds ago



Can connect to the Mongo Content Writer

a few seconds ago



CanConnectToZooKeeper

a few seconds ago




Kafka consumer health check for <http://cmdb.ft.com/systems/SemanticIngestor>

a few seconds ago

Knowing about problems before your clients do

@sarahjwells

Synthetic requests tell you about problems early



**TYRANNOSAURUS REX
EARLY WARNING
SYSTEM** JasonTheodor.com

[https://www.flickr.com/photos/jted/
5448635109](https://www.flickr.com/photos/jted/5448635109)



2. Concentrate on the stuff that matters



**It's the business functionality you should
care about**

@sarahjwells



US Politics & Policy >

Trump defiant as world leaders criticise travel ban

EU leaders see executive order as part of US offensive against world liberal order

2 HOURS AGO



More on this topic



Judge blocks Trump order to deport Muslim migrants

Apple, Google and Facebook condemn Trump's immigration ban

AN HOUR AGO

May does 'not agree' with Trump's refugee ban

AN HOUR AGO



We care about whether content got published successfully



UP Publish Availability

Edit More Info Download Print

Last 30 days Submit

Total Publish Availability



Total number of publishes

107,245

Publish Availability - Images



Publish Availability - Content



Publish Availability - Enrichedcontent



Publish Availability - List



Publish Availability - Notifications



Number of publishes - Images

7,931

Number of publishes - Content

38,255

Number of publishes - Enrichedcontent

38,228

Number of publishes - List

297


Number of publishes - Notifications

22,534

When people call our APIs, we care about speed



... we also care about errors




UP Read Availability

Edit More Info Download Print

Last 1 hour Submit


Total Read Availability




Total number of read requests

11,003


Read Availability - Content




Read Availability - EnrichedContent



Read Availability - Lists



Read Availability - Notifications



Number of reads - Content

5,992

Number of reads - EnrichedContent

725

Number of reads - Lists

334

Number of reads - Notifications


3,950

But it's the end-to-end that matters




<https://www.flickr.com/photos/robef/16537786315/>

If you just want information, create a dashboard or report




io.dropwizard.jetty.MutableServletContextHandler

REQUEST RATE BY HOST (REQ/SEC)



org.apache.http.client.HttpClient.*.get-requests.m1_rate

CUMULATIVE RESPONSE TIME (MS)



Checking the services involved in a business flow


```
[Unit]
```

```
Description=Method list mapper sidekick
```

```
Bindsto=methode-list-mapper@%i.service
```

```
After=methode-list-mapper@%i.service
```

```
[Service]
```

```
RemainAfterExit=yes
```

```
ExecStart=/bin/sh -c "\
```

```
export SERVICE=$(echo %p | sed 's/-sidekick//g'); \
```

```
etcdctl set /ft/services/$SERVICE/healthcheck true; \
```

```
etcdctl mkdir /ft/services/$SERVICE/servers; \
```

```
etcdctl set /ft/healthcheck/$SERVICE-%i/path /health; \
```

```
etcdctl set /ft/healthcheck/$SERVICE-%i/categories lists-publish; \
```

[/__health?categories=lists-publish](http://__health?categories=lists-publish)

CO-CO Prod UK

OK
Healthy

Last updated at 15:12

CO-CO Prod US

OK
Healthy

Last updated at 15:13

CO-CO Prod UK lists publish

OK
Healthy

Last updated at 15:13

CO-CO Prod US lists publish


OK
Healthy

CO-CO Prod UK lists read


OK
Healthy

CO-CO Prod US lists read

OK
Healthy



UPP Prod Delivery UK: List Publish Aggregate Healthcheck is down
[prod-uk-up.ft.com](#) • [View details](#)



UPP Prod Delivery UK: List Publish Aggregate Healthcheck is up
[prod-uk-up.ft.com](#) • [View details](#)



3. Cultivate your alerts



Make each alert great



<http://www.thestickerfactory.co.uk/>

Splunk Alert: PROD - **MethodAPIResponseTime5MAlert**

Business Impact

The methode api server is slow responding to requests. This might result in articles not getting published to the new content platform or publishing requests timing out.

...

@sarahjwells

Splunk Alert: PROD - MethodeAPIResponseTime5MAlert

Business Impact

The methode api server is slow responding to requests. This **might** result in articles not getting published to the new content platform or publishing requests timing out.

...

@sarahjwells

...

Technical Impact

The server is experiencing service degradation **because of network latency, high publishing load, high bandwidth utilization, excessive memory or cpu usage on the VM.** This might result in failure to publish articles to the new content platform.

Splunk Alert: PROD Content Platform Ingestor **Method** **Publish Failures Alert**

There has been one or more publish failures to the Universal Publishing Platform. The UUIDs are listed below.

Please see the [run book](#) for more information.

<code>_time</code>	<code>transaction_id</code>	<code>uuid</code>
Mon Oct 12 07:43:54 2015	tid_pbueyqnsqe	a56a2698-6e90-11e5-8608-a0853fb4e1fe

@sarahjwells

Splunk Alert: PROD Content Platform Ingestor Methode Publish Failures Alert

There has been one or more publish failures to the Universal Publishing Platform. The UUIDs are listed below.

Please see the [run book](#) for more information.

_time	transaction_id	uuid
Mon Oct 12 07:43:54 2015	tid_pbueyqnsqe	a56a2698-6e90-11e5-8608-a0853fb4e1fe

@sarahjwells

Splunk Alert: PROD Content Platform Ingestor Methode Publish Failures Alert

There has been one or more publish failures to the
Universal Publishing Platform. The UUIDs are listed below.

Please see the [run book](#) for more information.

<code>_time</code>	<code>transaction_id</code>	<code>uuid</code>
Mon Oct 12 07:43:54 2015	<code>tid_pbueyqnsqe</code>	a56a2698-6e90-11e5-8608-a0853fb4e1fe

@sarahjwells

Splunk Alert: PROD Content Platform Ingestor Methode Publish Failures Alert


There has been one or more publish failures to the Universal Publishing Platform. The UUIDs are listed below.

Please see the [run book](#) for more information.

_time	transaction_id	uuid
Mon Oct 12 07:43:54 2015	tid_pbueyqnsqe	a56a2698-6e90-11e5-8608-a0853fb4e1fe

@sarahjwells

Make sure you can't miss an alert



'Ops Cops' keep an eye on our systems

@sarahjwells

Use the right communication channel

@sarahjwells

It's not email

@sarahjwells

- Financial Times
- # semantic-data
- # six-degrees-workshop
- # tech-talks
- # universal-publishing
- # up-gp2-rollout
- # up-leadership
- # upp-escalation-drills
- # upp-incident-201701...
- # upp-prod-incidents
- # upp-pull-requests
- # upp-releases
- # upp-tech
- # upp-three-amigos
- Ammar Hassan
- Galia Rimon
- Hunor Kovács
- Luke Blaney
- Matt Andrews
- Nicky Wrightson
- Peter Schubert

Slack integration

Publish Availability Errors Severity 1 Business Impact: Most likely there is an issue in the pu...



Production alert APP 12:16 PM ☆

uploaded this email ▾

"financialtimes SplunkCloud" <alerts@splunkcloud.com>
Splunk Alert: Content Platform: Publish above SLA Notifications in Prod US
Content publish above SLA (Platinum) Severity 1 Business Impact: Most likely a piece of co...




Pingdom APP 12:22 PM

UPP Prod Publishing Active: Aggregate Healthcheck is down (Incident #55299)

[pub-prod-up.ft.com](#) • [View details](#)

UPP Prod Publishing Active: Aggregate Healthcheck is up (Incident #55299)

[pub-prod-up.ft.com](#) • [View details](#)



Scott Newton 12:23 PM

^due to content placeholder `aa8b00be-f401-11e6-95ee-f14e55513608` registering as false in PAM (edited)



Production alert APP 12:31 PM



Production alert APP 10:01 AM

uploaded this email ▾

"financialtimes SplunkCloud" <alerts@splunkcloud.com>

Splunk Alert: Content Platform: Publish above SLA Notifications in Prod US

Content publish above SLA (Platinum) Severity 1 Business Impact: Most likely a piece of co...



Production alert APP 10:01 AM

uploaded this email ▾

"financialtimes SplunkCloud" <alerts@splunkcloud.com>

Splunk Alert: Publish Availability Monitor Errors

Publish Availability Errors Severity 1 Business Impact: Most likely there is an issue in the pu...



Production alert APP 10:06 AM

uploaded this email ▾

"financialtimes SplunkCloud" <alerts@splunkcloud.com>

Splunk Alert: Content API V2 Concordances: V2ApiClientReceivingTooManyErrors 

<h1>V2 API prod-up-read.ft.com clients receiving too many errors</h1>
Severity 1<b...



Support isn't just getting the system fixed

@sarahjwells

Date:

Table of Contents

1. Summary
2. Health of our clusters
3. How was it fixed?
4. Impact
5. Actions
6. What happened?

‘You build it, you run it’?

@sarahjwells

Review the alerts you get


@sarahjwells

If it isn't helpful, make sure you don't get sent it again



See if you can improve it

ERRR...



**CAN'T STOP.
TOO BUSY!!**




www.workcompass.com/

When you didn't get an alert


@sarahjwells

What would have told you about this?




Line Format


Reset Zoom



Setting up an alert is part of fixing the problem



code



test



alerts

System boundaries are more difficult



Severin.stalder [CC BY-SA 3.0 (<http://creativecommons.org/licenses/by-sa/3.0>)], via Wikimedia Commons

Make sure you would know if an alert stopped
working

@sarahjwells

Add a unit test

```
public void shouldIncludeTriggerWordsForPublishFailureAlertInSplunk()  
{  
  
...  
  
}
```

Deliberately break things



Chaos snail



It's going to change: deal with it

@sarahjwells

Out of date information can be worse than none

@sarahjwells

Automate updates where you can

@sarahjwells

Find ways to share what's changing

@sarahjwells


In summary: to avoid alerts overload...

@sarahjwells

1

Build a system you can support





2



Concentrate on the stuff that matters





3

Cultivate your alerts



A microservice architecture lets you move fast...

@sarahjwells

But there's an associated operational cost

Make sure it's a cost you're willing to pay

@sarahjwells

Thank you

@sarahjwells