

Optimizing Mesos Utilization at Opentable

JAY CHIN

INFRASTRUCTURE ENGINEERING

1.4 Billion Online Reservations

2.3 Million Diners per Month

58 Million verified reviews

Before 2013



<https://flic.kr/p/9F6Khk>

Every 2 Months



Photo Credit : NASA



Around 2013

Opentable Codebase

Search

Reviews

Emails

Reservations

Photo
Service

Restaurant
profiles

Availability
Service

Menu API

White Label

External API

Person API

Feedback
API

Etc

Etc

Etc




Virtualization

DATACENTRE




Let's Scale !




Infrastructure Team / SRE



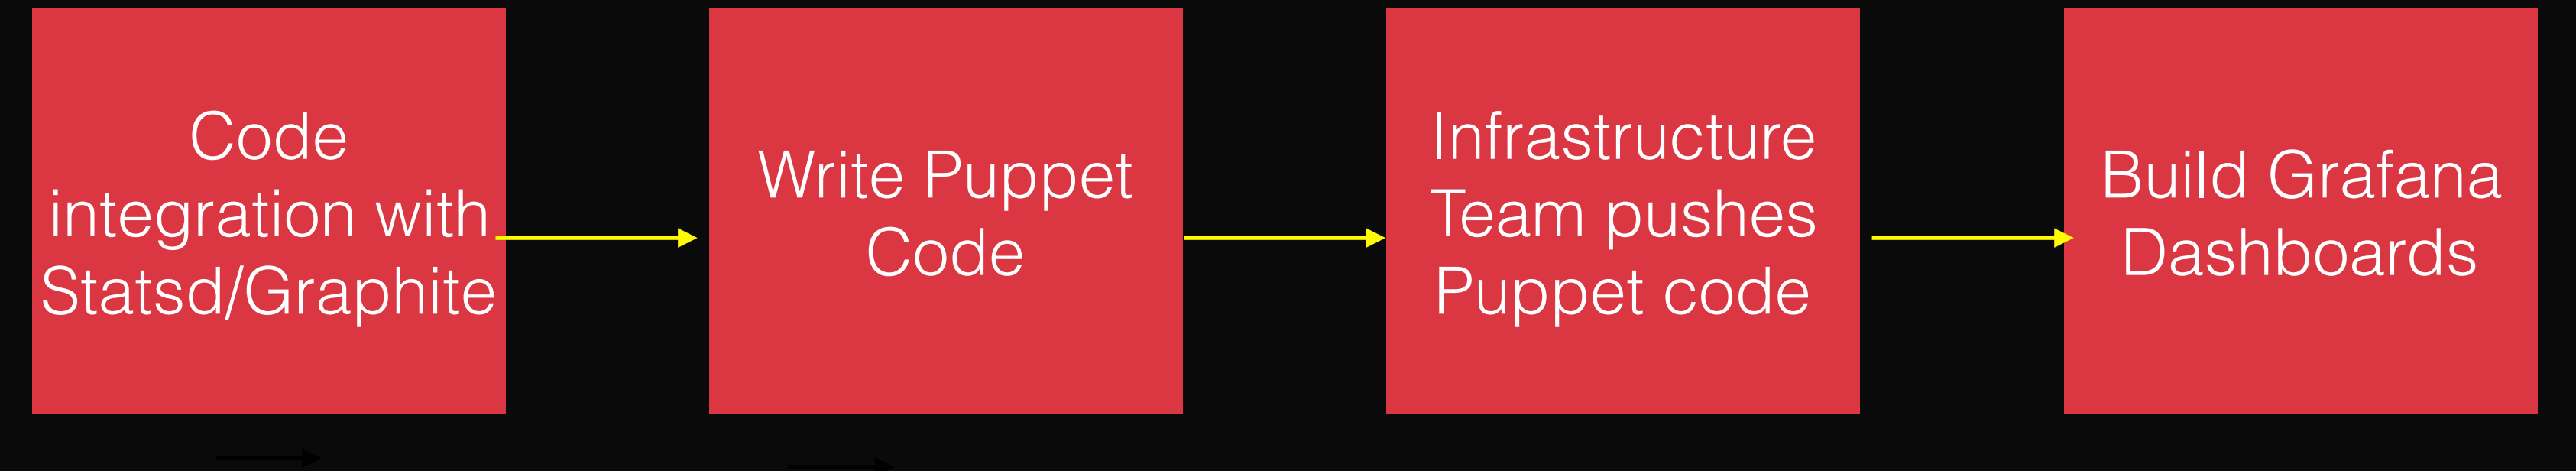
Local Build




Provision




Metrics



Monitoring





Around 2014 Explore Mesos

DATACENTRE

Mesos Cluster

Hubspot Singularity

Search

Restaurant
profiles

External API

Emails

Reviews

Reservations

Availability
Service

White Label


Photo
Service

Menu API


Feedback
API

Person API


Local Build




Provision




Metrics



Monitoring



Metrics Pipeline



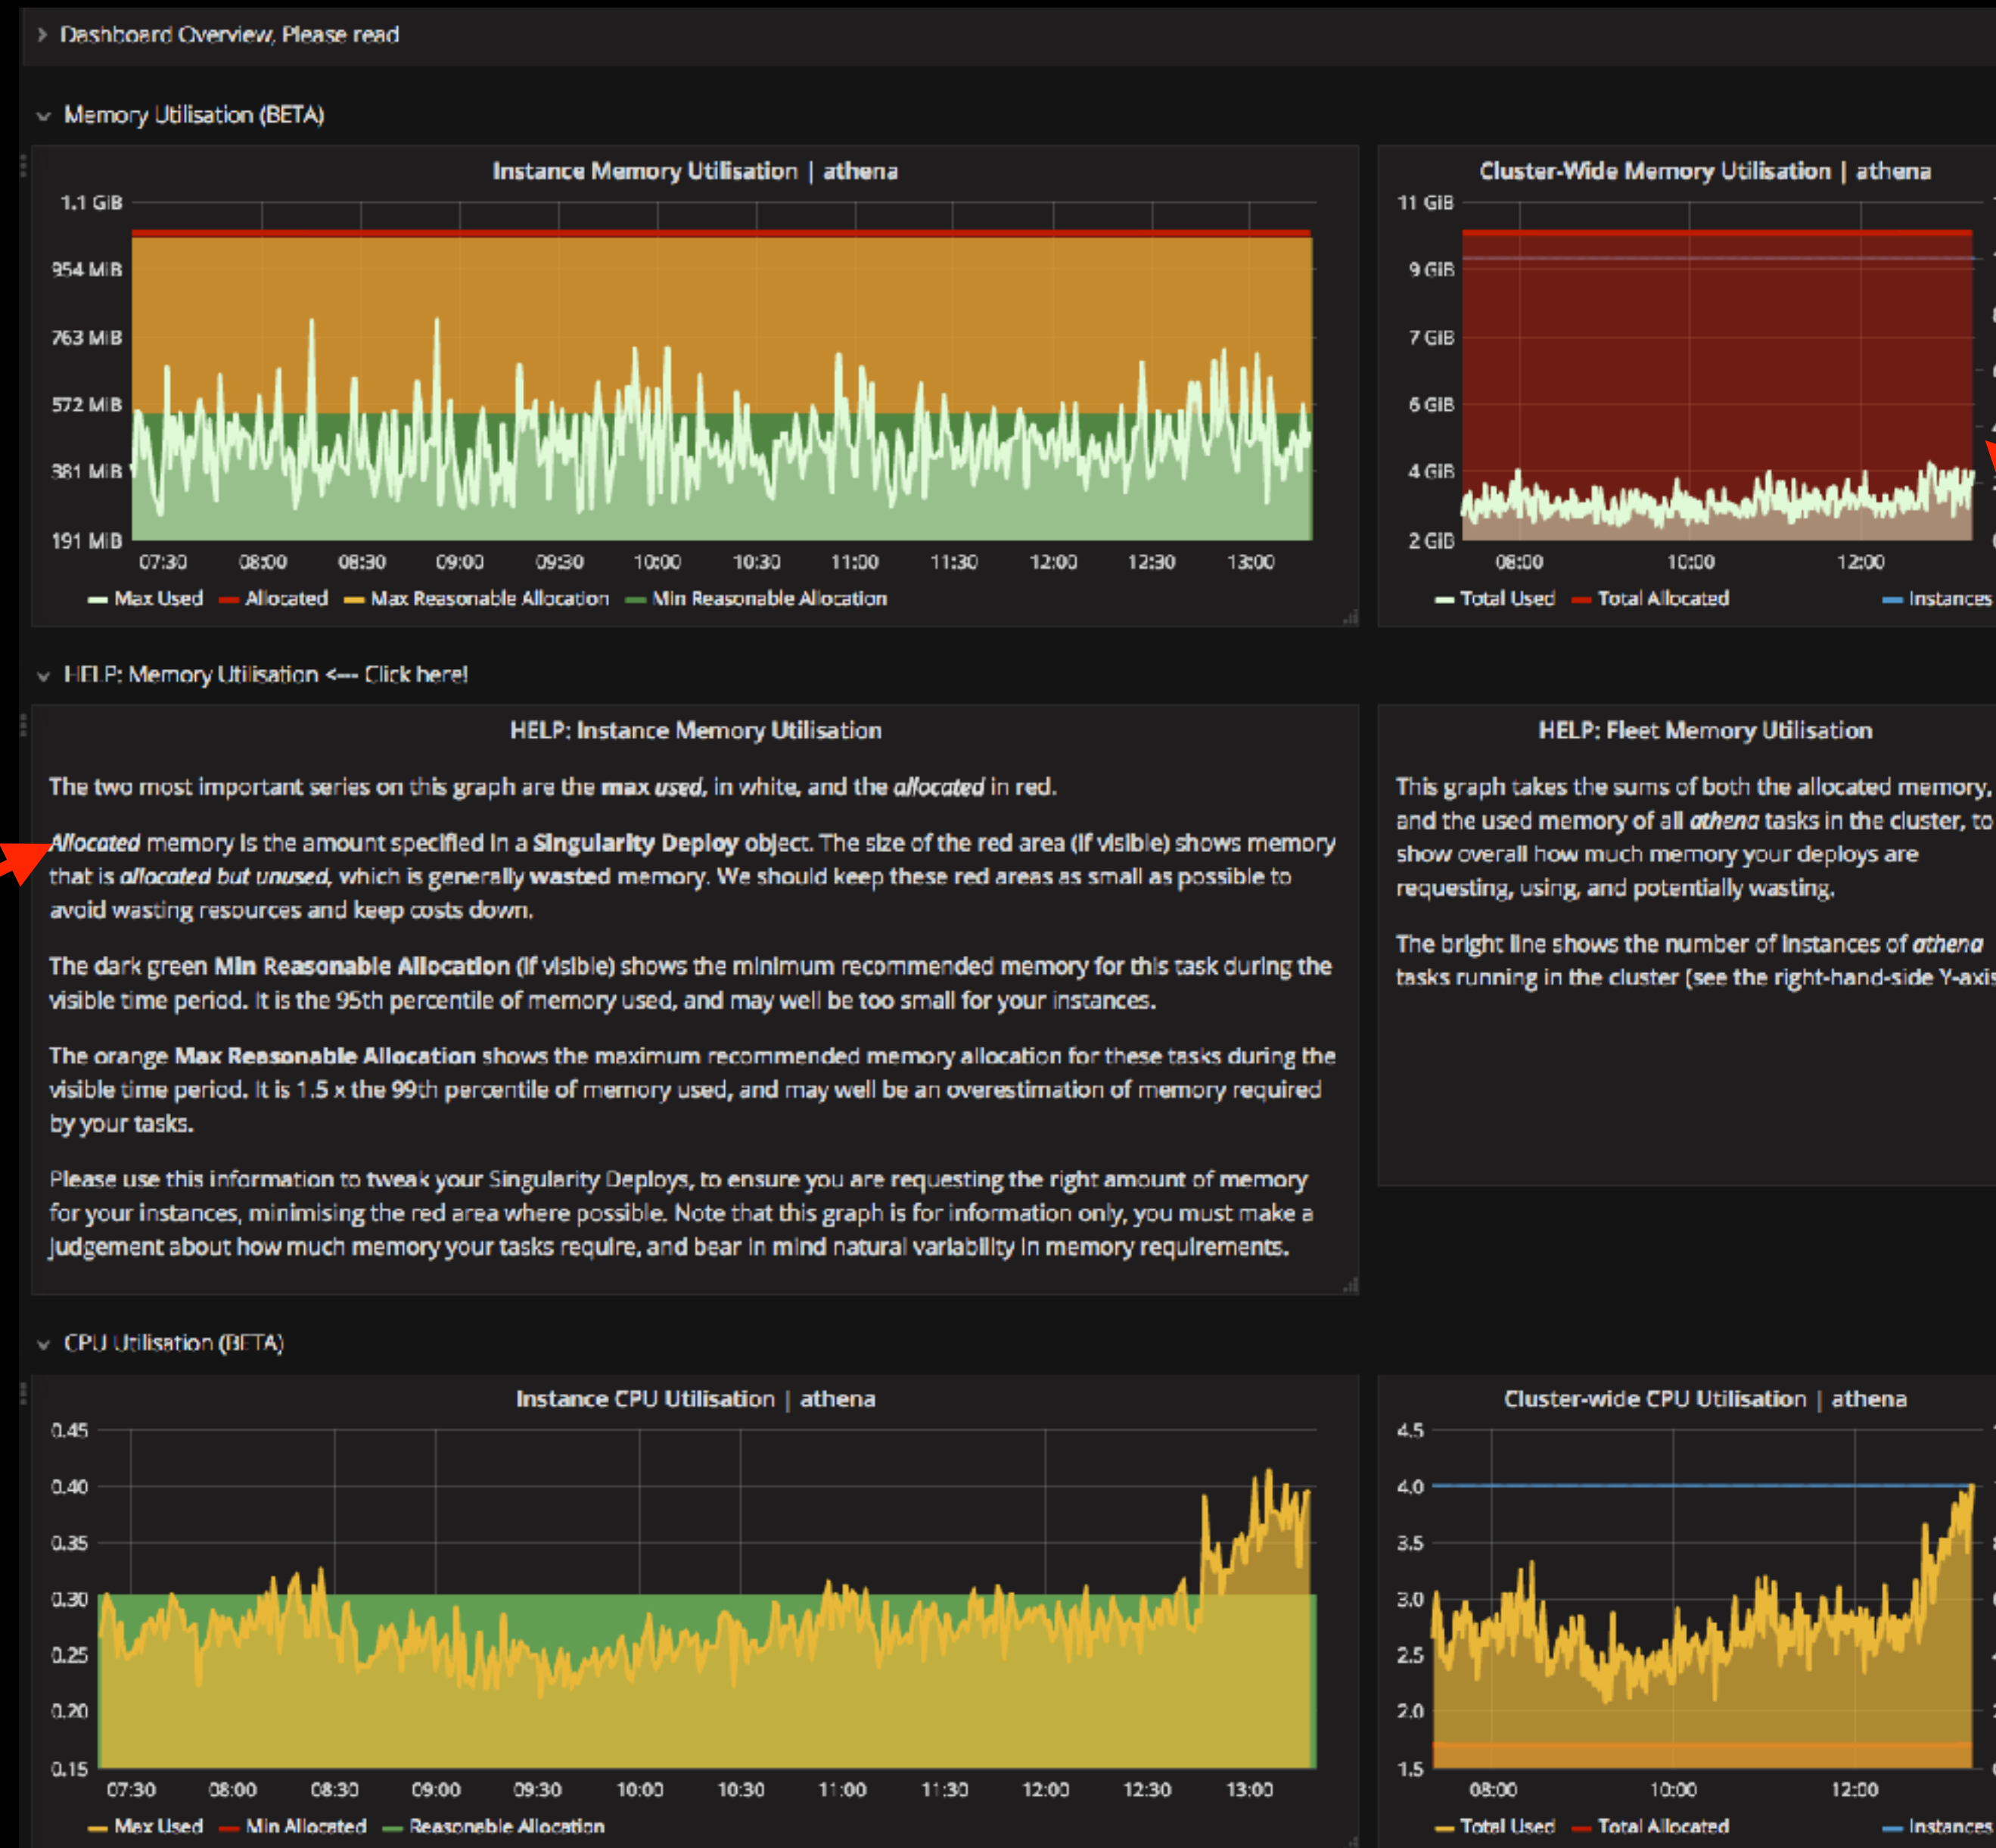
https://github.com/opentable/mesos_stats

<https://github.com/weaveworks/grafanalib>

<https://github.com/weaveworks/grafanalib>

Auto-generated Grafana Dashboard

Every Service running in Mesos will have an auto-generated dashboard



Help text explaining the graphs and what they mean

Shows cluster-wide Usage and Instance Usage

Right-sizing Resource Usage = \$\$\$ Saved

Mesos Cluster

Singularity Task


```
▼ resources: {} 4 keys
  cpus: 0.1
  memoryMb: 256
  numPorts: 1
  diskMb: 0
```

Mesos stats
and Metrics

Shows that memory is
over-provisioned
for this service

Overall Memory Utilisation (BETA)

prod-graphql-api memory utilisation (best case)




▼ HELP: Overall Memory Utilisation <--- Click here!

HELP: Task Memory Utilisation (best case)


This graph shows the maximum **used** memory per *prod-graphql-api* task (white) against the minimum **allocated** memory per *prod-graphql-api* task (red). The allocated memory is the amount specified in a Singularity Deploy object. The size of the red area shows memory that is




Local Build




Provision



Metrics




Monitoring



Sous

<https://github.com/opentable/sous>



```
1 Source: github.com/opentable/xxxx~
2 Owners:~
3 - marcopolo@opentable.com~
4 - james_bond@opentable.com~
5 Kind: http-service~
6 Deployments:~
7   Prod-London:~
8     Resources:~
9       cpus: "0.15"~
10      memory: "256"~
11      ports: "1"~
12     Env:~
13       OT_AUTOCHECK_RUNBOOK: https://wiki.otcorp.opentable.com/myfabulousservice~
14       OT_AUTOCHECK_SENSU_TEAM: the_A_team~
15       SERVICE_TYPE: super-server~
16     NumInstances: 3~
17     Volumes: []~
18     Startup:~
19       Timeout: 180~
20       CheckReadyURIPath: /servers~
21       CheckReadyURITimeout: 1~
22     Version: 0.5.45~
23   QA-uswest2:~
24     Resources:~
25       cpus: "0.15"~
26       memory: "256"~
27       ports: "1"~
28     Env:~
29       OT_AUTOCHECK_RUNBOOK: https://wiki.otcorp.opentable.com/myfabulousservice~
30       OT_AUTOCHECK_SENSU_TEAM: the_A_team~
31       SERVICE_TYPE: super-server~
32     NumInstances: 6~
33     Volumes: []~
34     Startup:~
35       Timeout: 180~
36       CheckReadyURIPath: /servers~
37       CheckReadyURITimeout: 1~
38     Version: 0.5.46~
```


Local Build




Provision



Metrics



Monitoring



Logging

**Restaurant_id == RID ==
ResID == Res_ID**

Global RequestID

https://github.com/opentable/request-timeline



Timeline Demo

Key Takeaways

- **Map out developer workflow and constantly look for opportunities to standardise, automate and enhance.**
- **Make metrics and monitoring part and parcel of the Mesos service.**
- **Engineers don't always make the best choice when deciding resource usage - help them make an informed choice.**
- **Have a common deployment pipeline across the organisation that facilitates production readiness***
- **Having a global data model for logging allows us to make more sense of logging data across the various Mesos tasks.**

Thank You

jayc@opentable.com

@jaychin

<https://www.linkedin.com/in/jayschin>