


Mesos + Singularity: PaaS automation for mortals

Gregory Chomatas @gchomatas

PaaS team
HubSpot

120 meters: My shortest travel to a Conference





Thales of Miletus - 624 BC

Those who can, do, the others philosophise...
Really?

Optionality* with Mesos

**Invest in a Mesos-powered PaaS and keep
doing what you love most; building your
product**

* Optionality is the property of asymmetric upside (preferably unlimited) with corresponding limited downside (preferably tiny)



Blogging

SEO

Social Media

CMS

Lead Management

Landing Pages

Calls-to-Action

Marketing Automation

Email

Analytics

CRM Sync

The underlying culture & structure

12-factor apps

.net monolith to microservices

Small, autonomous teams with
end-to-end ownership - no ops



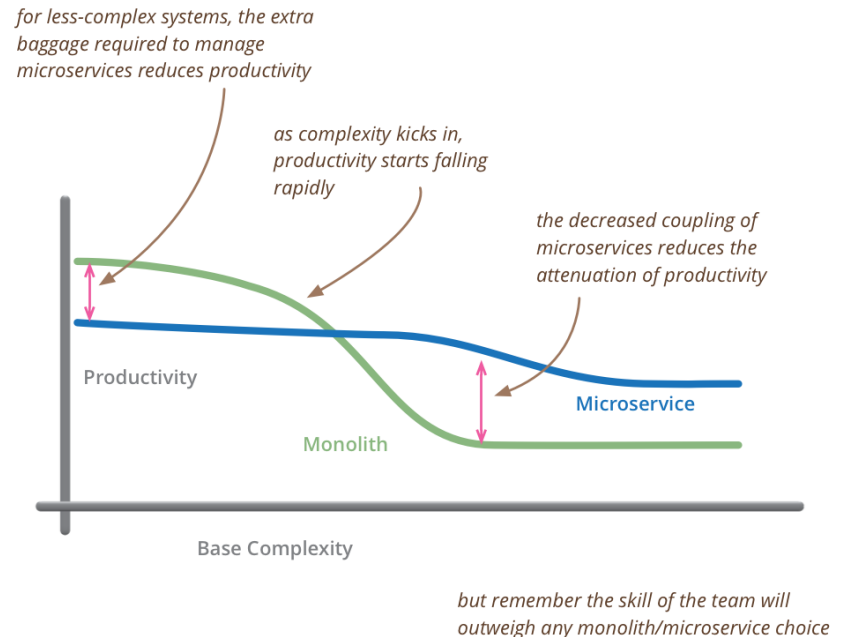
High productivity

~100 engineers

800+ components that can be updated/scaled independently

QA: ~400 small to medium AWS machines

PROD: ~750 medium to large AWS machines



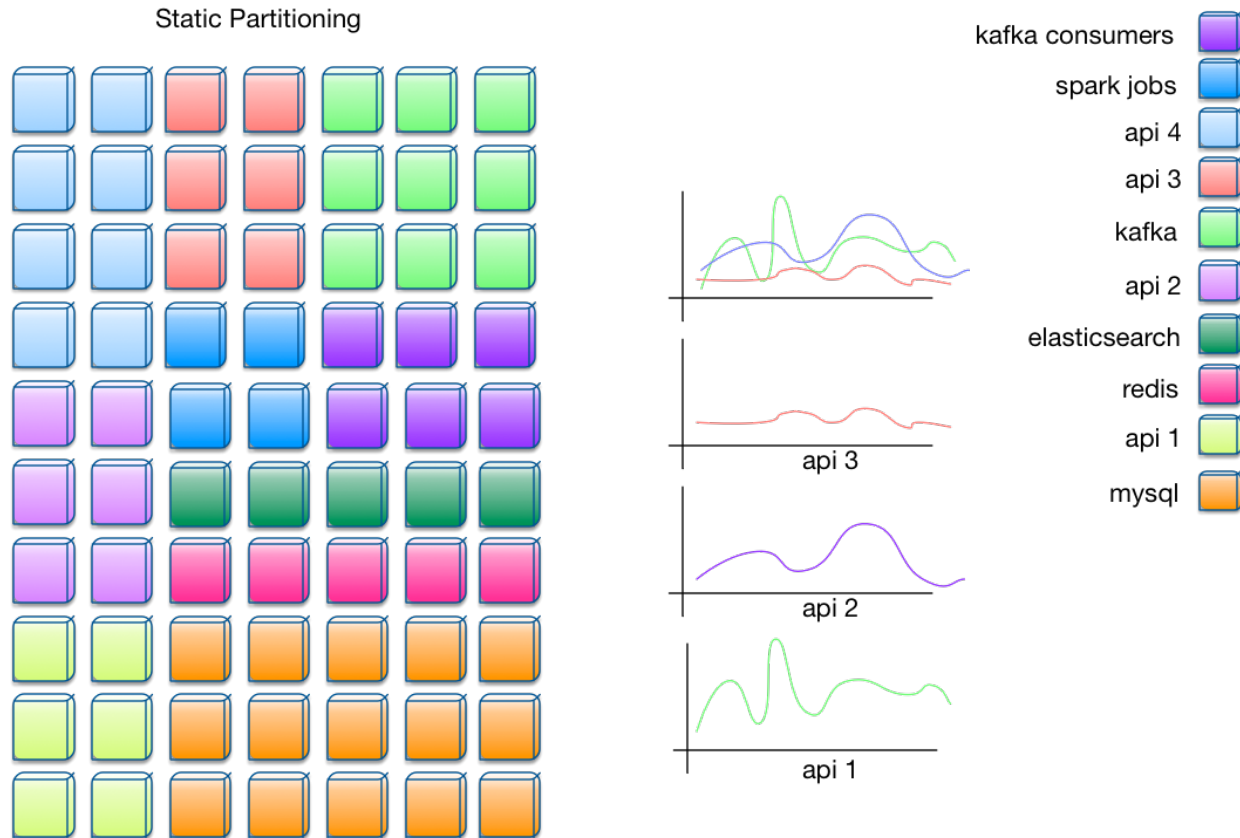
Source: Martin Fowler

<http://martinfowler.com/bliki/MicroservicePremium.html>

But for how long...

1. Develop locally
2. Provision QA aws instance
3. Deploy via local Python script
4. Provision PROD aws instance
5. Deploy via local Python script
6. Repeat 4 & 5 to scale
- ...
10. Repeat 4&5 at 4am to replace hw

Statically partitioning the cluster is inefficient



The cost of flexibility & asynchronicity

High operational overhead

Poor utilisation & elasticity

Higher rate of failures



Redress the balance with a Mesos-based PaaS

Abstract away machines

Homogenous environment

Scale out in seconds

Centralized deployables
registry



Sept 2013: Our First Mesos Cluster

To Boldly go...to Singularity



Singularity: do more with a single scheduler

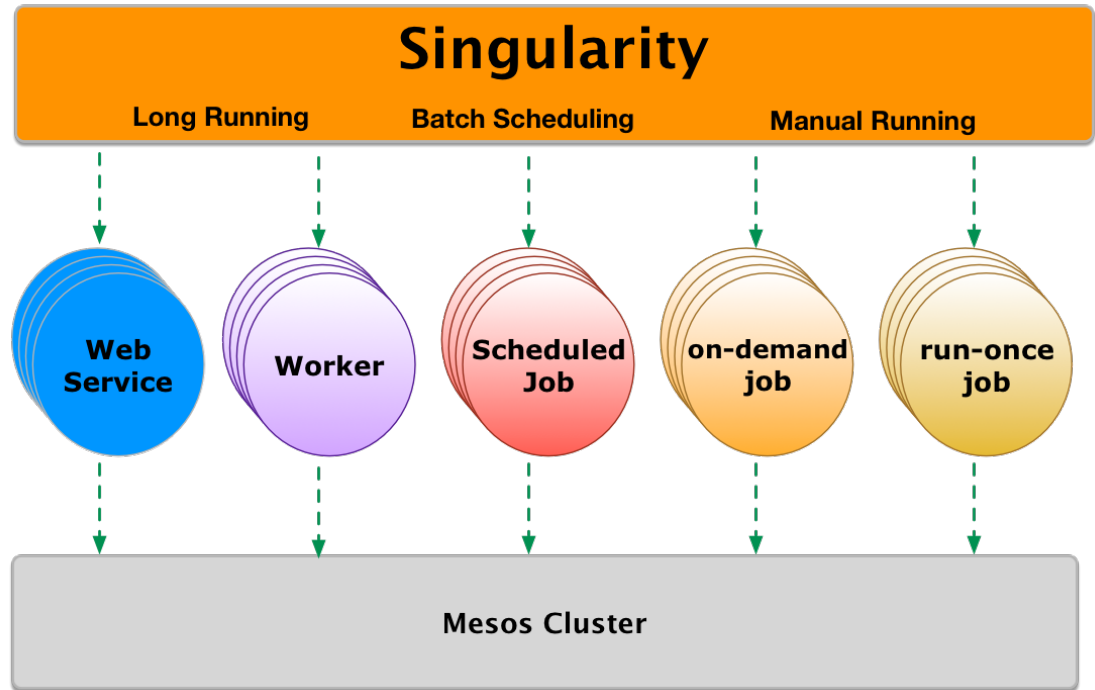
Great UI & HTTP API

Native Docker Support

Health Checks

Load Balancing API

Log Maintenance



Oct 2013: Start building Singularity

Singularity: do even more...

Security / artifact signature
verification

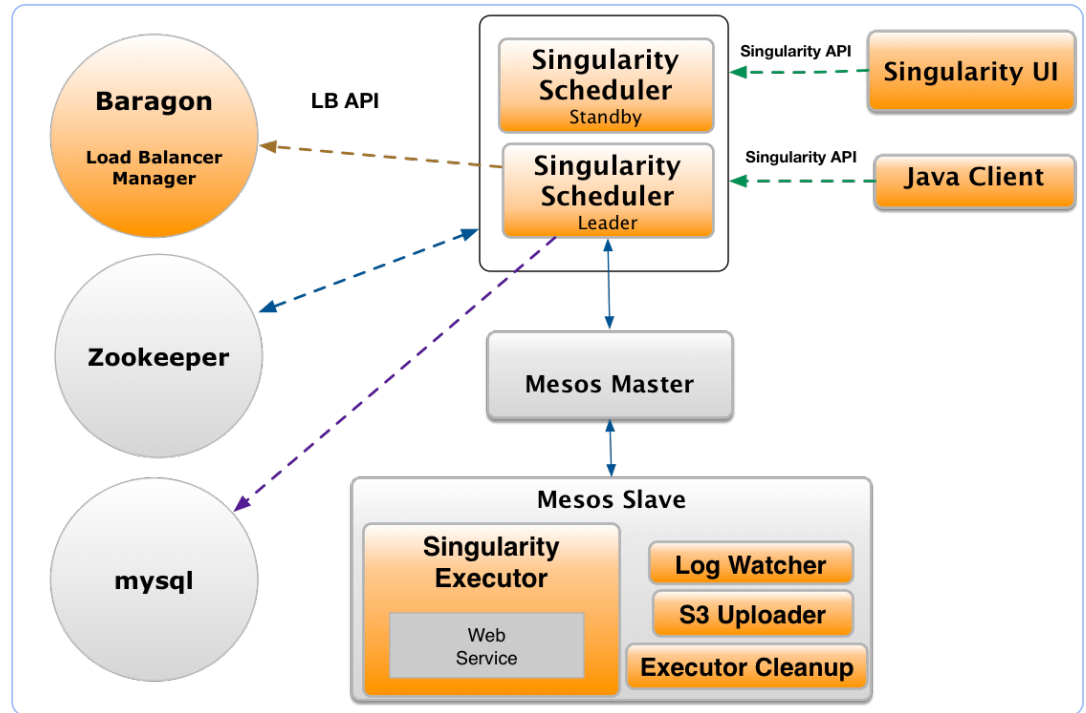
Agent & Rack maintenance

Webhooks

Auto-rollback

Email Notifications

Executor cleanup



Singularity Components

The PaaS Stack

BUILD



Jenkins

DEPLOY



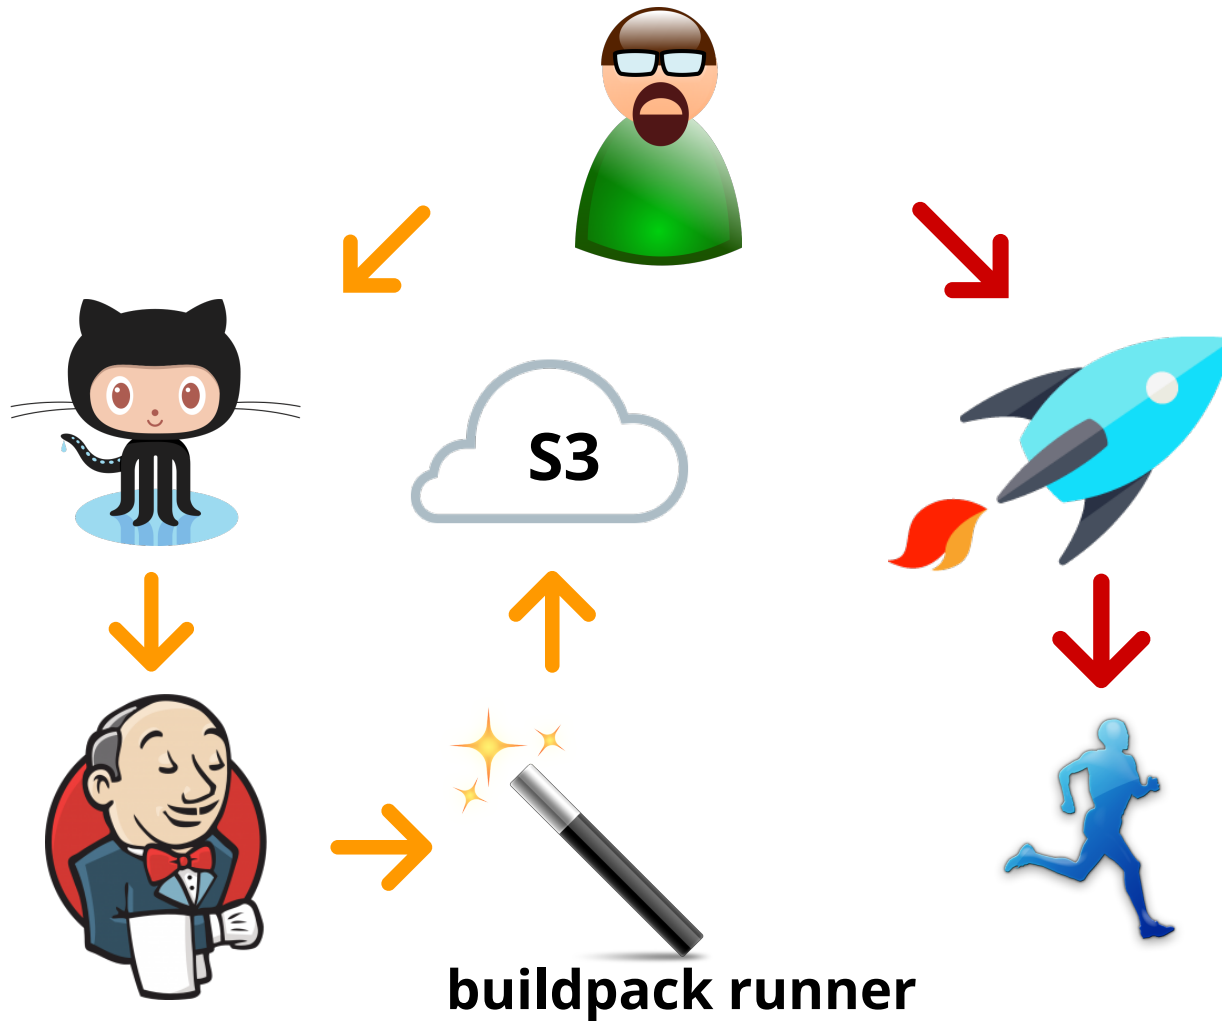
Orion

RUN



Singularity

The Build / Deploy cycle



The Deployer - Dry run

Deploys

Blockers (2)

Builds

HipChat build alerts

Deploy

New ▾

gchomatas ▾

Deploy config

ORION_NestedMetaConfig

Action

Deploy ▾

Environment

QA ▾

+ Advanced options

Deploy on QA

QA **ORION** deploy of **ORION_NestedMetaConfig**

ORION_NestedMetaConfig </>

31 (active in QA) → 33 (latest build)

worker

ORION_MetaConfig </>

31 (active in QA) → 33 (latest build)

webService

scheduledJob

worker

onDemand

run-once

ORION_TestService </>

31 (active in QA) → 33 (latest build)

web

ORION_TestJob </>

31 (active in QA) → 33 (latest build)

testjob

longcron

ORION_All_Item_Types_In_One_Config

build name	OrionTests
build number	31 → 33
repo	HubSpot/OrionTests
branch	master
committer	gchomatas@hubspot.com
commit sha	bb521846 → 1b8af26b [compare]
commit message	Changed config to get a new commit/build
committed	a minute ago

The Deployer - Deploying

Deploys Blockers (2) Builds HipChat build alerts

Deploy

New ▾

gchomatas ▾

Deploy config

ORION_NestedMetaConfig

Action

Deploy ▾

Environment

QA ▾

+ Advanced options

QA ORION deploy of **ORION_NestedMetaConfig**

FORCE Stop

Started: Thursday, October 1st 2015, 6:53:27 pm,

Deploy Time: 29 seconds

Last Update: 1 seconds ago

Deploys: WAITING DEPLOYS TO FINISH

ORION_NestedMetaConfig </> — waiting child deploys to finish (●)

ORION_MetaConfig </> — waiting child deploys to finish (●)

ORION_TestService </> — child deploys finished ✓

31 (active in QA) → 33 (latest build)

web — active ✓

ORION_All_Item_Types_In_One_Config </> — waiting child deploys to finish (●)

31 (active in QA) → 33 (latest build)

worker — will not deploy this item is paused ⚠

webService — deployed waiting activation (●)

scheduledJob — active ✓

run-once — deployed waiting activation (●)

onDemand — active ✓

ORION_TestJob </> — child deploys finished ✓

31 (active in QA) → 33 (latest build)

longcron — active ✓

testjob — active ✓

ORION_TestWorker </> — waiting child deploys to finish (●)

31 (active in QA) → 33 (latest build)

worker — deployed waiting activation (●)

Singularity

Deployable view

Singularity (rs_staging)DashboardStatusRequestsTasksAdminAPI Docs (beta)Search

activeService

ORION_TestService-web

JSONDeployScalePauseBounceEditRemove

Active deploy 33_600_350 by gchomatas -- 8 minutes ago (1 Oct 2015 18:53)

Running instances

Name	Status	Deploy ID	Started	Updated
ORION_TestService-web-33_600_350-1443722011196-1-free_dream.lad52.hubspot_networks.net-us_east_1b	Test running	33_600_350	8 minutes ago (1 Oct 2015 18:53)	8 minutes ago (1 Oct 2015 18:53) -- ()

Task history [more at once](#)

Name	Last state	Deploy	Started	Updated
ORION_TestService-web-31_599_349-1443698501118-1-round_cake.lad52.hubspot_networks.net-us_east_1e	Test killed	31_599_349	7 hours ago (1 Oct 2015 12:21)	8 minutes ago (1 Oct 2015 18:53) -- ()
ORION_TestService-web-28_598_195-1443565735607-1-mute_mode.lad52.hubspot_networks.net-us_east_1a	Test killed	28_598_195	2 days ago (29 Sep 2015 23:28)	7 hours ago (1 Oct 2015 12:22) -- ()
ORION_TestService-web-28_598_195-1443565173874-2-free_dream.lad52.hubspot_networks.net-us_east_1b	Test killed	28_598_195	21 hours ago (30 Sep 2015 21:32)	21 hours ago (30 Sep 2015 21:48) -- ()
ORION_TestService-web-28_597_194-1443552359885-1-mute_mode.lad52.hubspot_networks.net-us_east_1a	Test killed	28_597_194	2 days ago (29 Sep 2015 19:45)	2 days ago (29 Sep 2015 23:29) -- ()
ORION_TestService-web-28_596_348-1443546842336-1-mute_mode.lad52.hubspot_networks.net-us_east_1a	Test killed	28_596_348	2 days ago (29 Sep 2015 18:14)	2 days ago (29 Sep 2015 19:46) -- ()

Deploy history [more at once](#)

Deploy ID	Status	User	Timestamp
33_600_350	Succeeded	gchomatas	9 minutes ago (1 Oct 2015 18:53) ()
31_599_349	Succeeded	gchomatas	7 hours ago (1 Oct 2015 12:21) ()
28_598_195	Succeeded	blodge	2 days ago (29 Sep 2015 23:28) ()
28_597_194	Succeeded	blodge	2 days ago (29 Sep 2015 19:45) ()
28_596_348	Succeeded	gchomatas	3 days ago (28 Sep 2015 20:18) ()

Request history [more at once](#)

State	User	Created
Updated	tpetr	21 hours ago (30 Sep 2015 21:48) ()
Updated	tpetr	a day ago (30 Sep 2015 21:32) ()
Bounced	ssalinas	9 days ago (22 Sep 2015 19:16) ()
Updated	gchomatas	10 days ago (22 Sep 2015 00:33) ()
Updated	tpetr	14 days ago (17 Sep 2015 16:18) ()

S3 Logs

Log file	Size	Last modified
ORION_TestService-web-31_599_349-1443698501118-1-round_cake.lad52.hubspot_n...	1.2 KB	1 Oct 2015 18:55
ORION_TestService-web-31_599_349-1443698501118-1-round_cake.lad52.hubspot_n...	119 B	1 Oct 2015 18:55
ORION_TestService-web-28_598_195-1443565735607-1-mute_mode.lad52.hubspot_...	236 B	1 Oct 2015 12:26
ORION_TestService-web-28_598_195-1443565735607-1-mute_mode.lad52.hubspot_...	1.16 KB	1 Oct 2015 04:16
ORION_TestService-web-28_598_195-1443565735607-1-mute_mode.lad52.hubspot_...	119 B	1 Oct 2015 04:16

Singularity Task view

ORION_TestService-web-33_601_351-1444055865295-1-mute_mode.iad...

JSON

Kill task


















Task running as of 3 hours ago (5 Oct 2015 15:37) — Task running process bash - (pid: 88654)

History

Status	Message	Time
Task launched	—	3 hours ago (5 Oct 2015 15:37)
Task starting	Staging files... (executor pid: 88473)	3 hours ago (5 Oct 2015 15:37)
Task running	Task running process bash - (pid: 88654)	3 hours ago (5 Oct 2015 15:37)

Files

root / ORION_TestService-web-33_601_351-1444055865295-1-mute_mod...

Name	Size	Last modified	
 tmp/		1 Oct 2015 18:49	
 bin/		1 Oct 2015 18:49	
 conf/		1 Oct 2015 18:49	
 app/		5 Oct 2015 15:40	
 logs/		5 Oct 2015 15:37	
 ORION_TestService.yaml	1.67 KB	5 Oct 2015 15:37	
 executor.java.log	7.93 KB	5 Oct 2015 15:37	
 deploy.env	2.7 KB	5 Oct 2015 15:37	
 runner.sh	1.86 KB	5 Oct 2015 15:37	
 service.log	2.5 KB	5 Oct 2015 15:37	
 executor.bash.log	889 B	5 Oct 2015 15:37	

S3 Logs

Singularity file tailing

Singularity (hs_staging)

Dashboard

Status

Requests

Tasks

Admin ▾

API Docs (beta)

🔍 Search

[root](#) / [ORION_TestService-web-33_600_35...](#) / **service.log**

```
Running java8 with JVM arguments: -Xms2048m -Xmx2g -Xss256k -XX:+UseG1GC -XX:MaxGCPauseMillis=50 -XX:+UnlockExperimentalVMOptions
INFO [2015-10-01 17:53:37,408] com.hubspot.deploy.testservice.InjectedExceptionHandler: Added injected health checker
INFO [2015-10-01 17:53:37,447] io.dropwizard.server.ServerFactory: Starting TestService
INFO [2015-10-01 17:53:37,579] org.eclipse.jetty.setuid.SetUIDListener: Opened TestService@53e3a87a{HTTP/1.1}{0.0.0.0}
INFO [2015-10-01 17:53:37,582] org.eclipse.jetty.server.Server: jetty-9.0.z-SNAPSHOT
INFO [2015-10-01 17:53:37,709] com.sun.jersey.server.impl.application.WebApplicationImpl: Initiating Jersey application
INFO [2015-10-01 17:53:37,781] io.dropwizard.jersey.DropwizardResourceConfig: The following paths were found for the application:
GET /procfileitems/{procfileItemId} (com.hubspot.deploy.testservice.resources.ProcfileResource)
```

Singularity Health Checks & Resources

Healthchecks [collapse](#)

Beginning on **Task running**, hit </admin/healthcheck> with a 5 second timeout every 5 second(s) until **HTTP 200** is received, or 120 seconds have elapsed.

Timestamp	Duration	Status	Message	
1 Oct 2015 18:53:40	137ms	HTTP 200	<pre>{"TestServiceHealthCheck":{"healthy":true,"message":"Test service is healthy. Free memory is 1930Mbs"},"deadlocks":{"healthy":true}}</pre>	{ }


Resource Usage [collapse](#)

Memory (rss vs limit)



4.46 GB / 4.49 GB

CPU Usage



0.25 used / 12 allocated CPUs

Memory (anon) [Copy](#)

3.43 GB

Memory (file) [Copy](#)

1.02 GB

Memory (mapped file) [Copy](#)

36 KB

Singularity All Deployables view

Singularity (hs_qa) Dashboard Status **Requests** Tasks Admin ▾ API Docs (beta) 🔍 Search

All Active Cooldown Paused Pending Cleaning

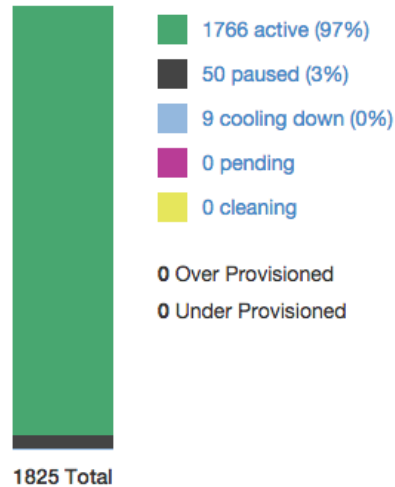
Filter requests

☒ Service ☒ Worker ☒ Scheduled ☒ On-demand ☒ Run-once Protip: press **cmd** to select multiple

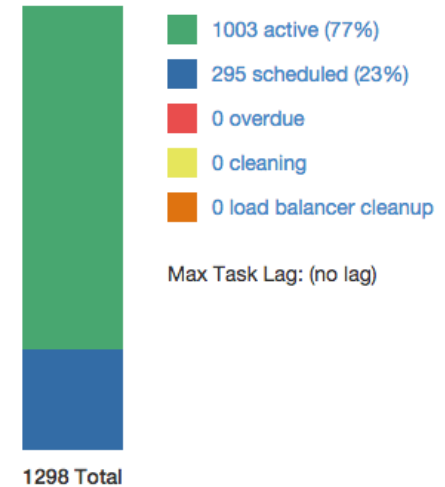
	Name	Type	State	Deploy time			
★	ORION_TestService-web	Service	Active	4 minutes ago (5 Oct 2015 19:02)	📊	🗑️	{ }
★	SignalsGatingService-web	Service	Active	4 minutes ago (5 Oct 2015 19:02)	📊	🗑️	{ }
★	LeadinFormProcessorKafka-consume	Worker	Active	8 minutes ago (5 Oct 2015 18:58)	📊	🗑️	{ }
★	LeadinJobs-sendWeeklyEmails	Scheduled	Active	8 minutes ago (5 Oct 2015 18:58)	⚡	🗑️	{ }
★	LeadinDatabaseMigration-dbmigrate	On demand	Active	8 minutes ago (5 Oct 2015 18:58)	⚡	🗑️	{ }
★	LeadinJobs-formSubmissionEmailNotifications	Worker	Active	8 minutes ago (5 Oct 2015 18:58)	📊	🗑️ Remove	{ }

Singularity Cluster Status

Requests



Tasks



Racks

3 Active Racks
0 Decommissioning Racks
0 Inactive Racks

Slaves

18 Active Slaves
0 Decommissioning Slaves
0 Inactive Slaves
0 Unknown Slaves

Deploys

0 Active Deploys

Singularity scheduler instances

Hostname	Driver status	Uptime	Time since last offer
rozel.iad02.hubspot-networks.net	Driver not started	18 minutes	—
bigpidgeon.iad02.hubspot-networks.net	Driver not started	18 minutes	—
ojofrio.iad02.hubspot-networks.net	Driver running	18 minutes	a few seconds

Singularity

Cluster Maintenance view

Slaves

Active

ID	State	Since	Rack	Host	Uptime	Activated By
20150829-130003-2967508140-0050-20229-00	Active	29 Sep 2015 21:10	us-east-1a	red-water.iad52.hubspot-networks.net	2 days	wisconsin
20150829-130003-2967508140-0050-20229-01	Active	29 Sep 2015 14:51	us-east-1a	red-sunset.iad52.hubspot-networks.net	2 days	
20150829-130003-2967508140-0050-20229-011	Active	29 Sep 2015 16:15	us-east-1a	cool-surf.iad52.hubspot-networks.net	2 days	
20150829-130003-2967508140-0050-20229-02	Active	30 Sep 2015 17:58	us-east-1b	muddy-art.iad52.hubspot-networks.net	2 days	wisconsin
20150829-130003-2967508140-0050-20229-012	Active	29 Sep 2015 16:25	us-east-1a	plain-art.iad52.hubspot-networks.net	2 days	
20150829-130003-2967508140-0050-20229-03	Active	29 Sep 2015 15:08	us-east-1b	old-water.iad52.hubspot-networks.net	2 days	
20150829-130003-2967508140-0050-20229-04	Active	29 Sep 2015 18:44	us-east-1b	sleep-band.iad52.hubspot-networks.net	2 days	wisconsin
20150829-130003-2967508140-0050-20229-05	Active	29 Sep 2015 15:16	us-east-1a	throbbing-union.iad52.hubspot-networks.net	2 days	
20150829-130003-2967508140-0050-20229-010	Active	29 Sep 2015 18:44	us-east-1b	hidden-wave.iad52.hubspot-networks.net	2 days	wisconsin
20150829-130003-2967508140-0050-20229-015	Active	29 Sep 2015 17:19	us-east-1a	crimson-cell.iad52.hubspot-networks.net	2 days	
20150829-130003-2967508140-0050-20229-016	Active	29 Sep 2015 17:35	us-east-1b	hidden-sky.iad52.hubspot-networks.net	2 days	
20150829-130003-2967508140-0050-20229-013	Active	29 Sep 2015 21:10	us-east-1a	rough-fiber.iad52.hubspot-networks.net	2 days	wisconsin
20150829-130003-2967508140-0050-20229-014	Active	30 Sep 2015 17:58	us-east-1a	late-glade.iad52.hubspot-networks.net	2 days	wisconsin
20150829-130003-2967508140-0050-20229-017	Active	29 Sep 2015 17:38	us-east-1a	restless-right.iad52.hubspot-networks.net	2 days	
20150829-130003-2967508140-0050-20229-09	Active	29 Sep 2015 16:02	us-east-1a	yellow-snow.iad52.hubspot-networks.net	2 days	
20150829-130003-2967508140-0050-20229-08	Active	29 Sep 2015 15:45	us-east-1a	gentle-fog.iad52.hubspot-networks.net	2 days	
20150829-130003-2967508140-0050-20229-014	Active	30 Sep 2015 17:58	us-east-1a	late-glade.iad52.hubspot-networks.net	2 days	wisconsin
20150829-130003-2967508140-0050-20229-017	Active	29 Sep 2015 17:38	us-east-1a	restless-right.iad52.hubspot-networks.net	2 days	

Frozen

No slaves

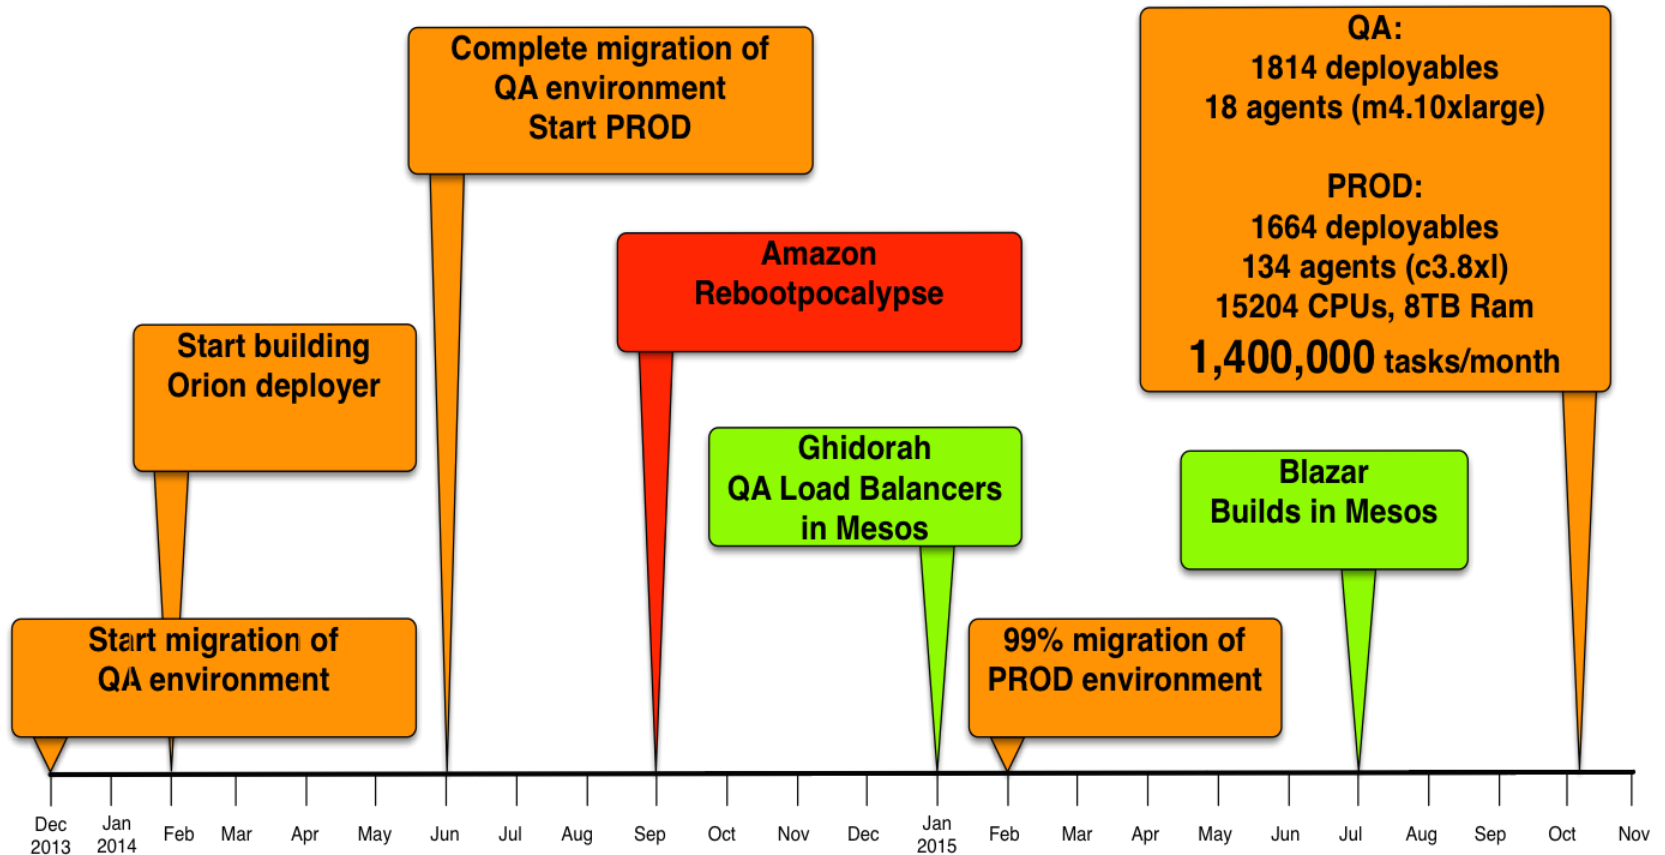
Decommissioning

No slaves

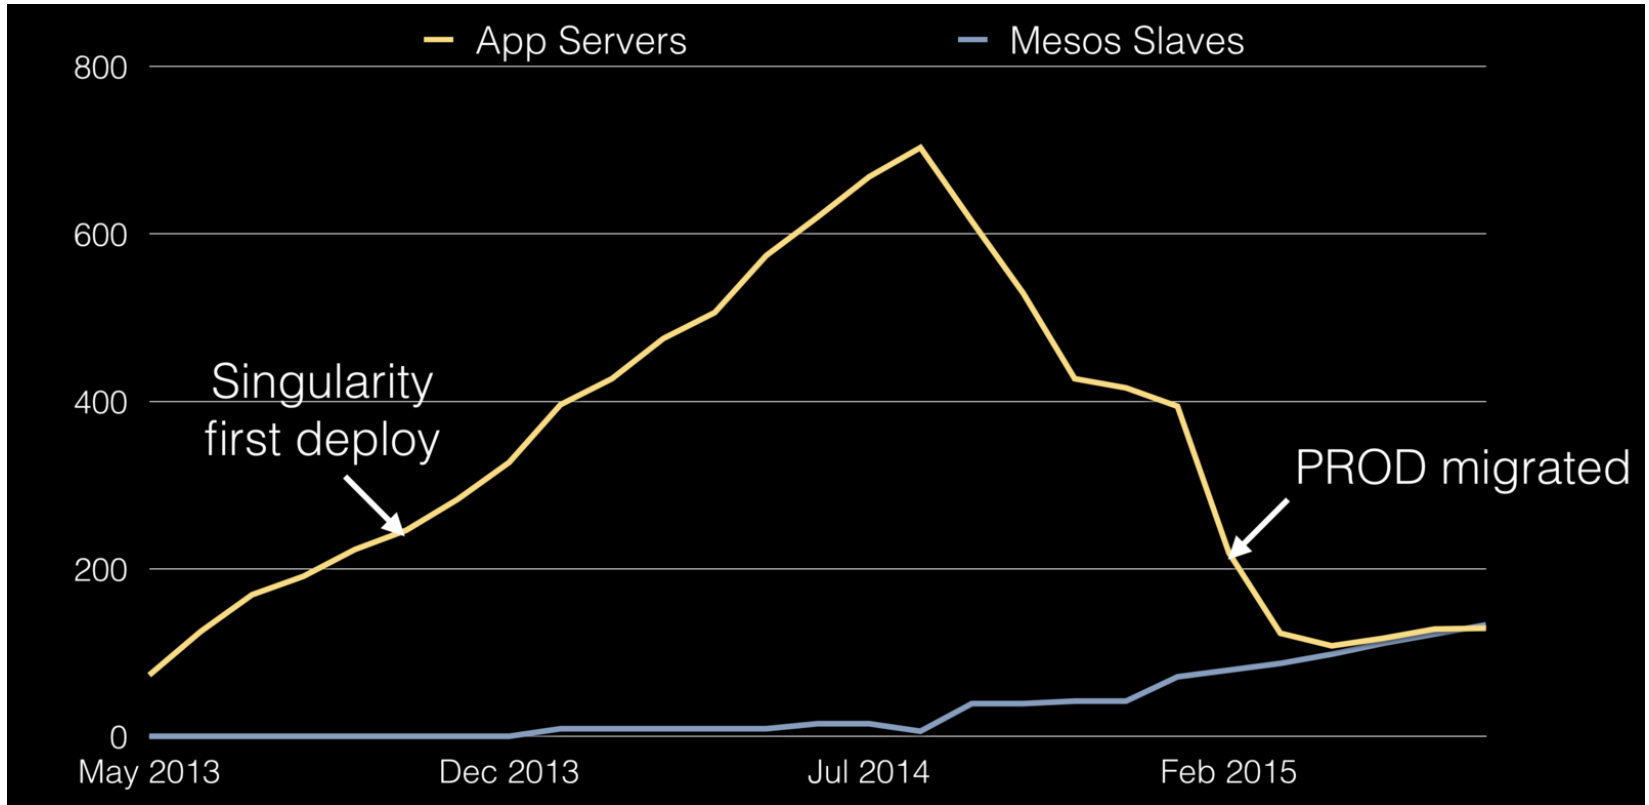
Inactive

ID	State	Since	Rack	Host	Uptime
20150326-184041-34738348-5050-8977-09	Dead	29 Sep 2015 14:44	us-east-1a	red-water.iad52.hubspot-networks.net	8 days
20150326-184041-34738348-5050-8977-08	Dead	29 Sep 2015 14:51	us-east-1a	red-sunset.iad52.hubspot-networks.net	8 days
20150326-184041-34738348-5050-8977-07	Dead	29 Sep 2015 14:56	us-east-1b	muddy-art.iad52.hubspot-networks.net	8 days
20150326-184041-34738348-5050-8977-06	Dead	29 Sep 2015 15:06	us-east-1b	old-water.iad52.hubspot-networks.net	8 days
20150326-184041-34738348-5050-8977-016	Dead	29 Sep 2015 15:11	us-east-1b	sleep-band.iad52.hubspot-networks.net	8 days
20150326-184041-34738348-5050-8977-015	Dead	29 Sep 2015 15:16	us-east-1a	throbbing-union.iad52.hubspot-networks.net	8 days
20150326-184041-34738348-5050-8977-014	Dead	29 Sep 2015 15:24	us-east-1b	super-hall.iad52.hubspot-networks.net	8 days
20150326-184041-34738348-5050-8977-013	Dead	29 Sep 2015 15:37	us-east-1a	yellow-truth.iad52.hubspot-networks.net	8 days
20150326-184041-34738348-5050-8977-017	Dead	29 Sep 2015 15:45	us-east-1a	gentle-fog.iad52.hubspot-networks.net	8 days
20150326-184041-34738348-5050-8977-012	Dead	29 Sep 2015 16:02	us-east-1a	yellow-snow.iad52.hubspot-networks.net	8 days
20150326-184041-34738348-5050-8977-011	Dead	29 Sep 2015 16:08	us-east-1b	hidden-wave.iad52.hubspot-networks.net	8 days
20150326-184041-34738348-5050-8977-010	Dead	29 Sep 2015 16:15	us-east-1a	cool-surf.iad52.hubspot-networks.net	8 days
20150326-184041-34738348-5050-8977-05	Dead	29 Sep 2015 16:25	us-east-1a	plain-art.iad52.hubspot-networks.net	8 days
20150326-184041-34738348-5050-8977-04	Dead	29 Sep 2015 16:38	us-east-1a	rough-fiber.iad52.hubspot-networks.net	8 days
20150326-184041-34738348-5050-8977-03	Dead	29 Sep 2015 16:47	us-east-1a	late-glade.iad52.hubspot-networks.net	8 days
20150326-184041-34738348-5050-8977-02	Dead	29 Sep 2015 17:19	us-east-1a	crimson-cell.iad52.hubspot-networks.net	8 days
20150326-184041-34738348-5050-8977-01	Dead	29 Sep 2015 17:38	us-east-1a	restless-right.iad52.hubspot-networks.net	8 days
20150326-184041-34738348-5050-8977-00	Dead	29 Sep 2015 17:35	us-east-1b	hidden-sky.iad52.hubspot-networks.net	8 days

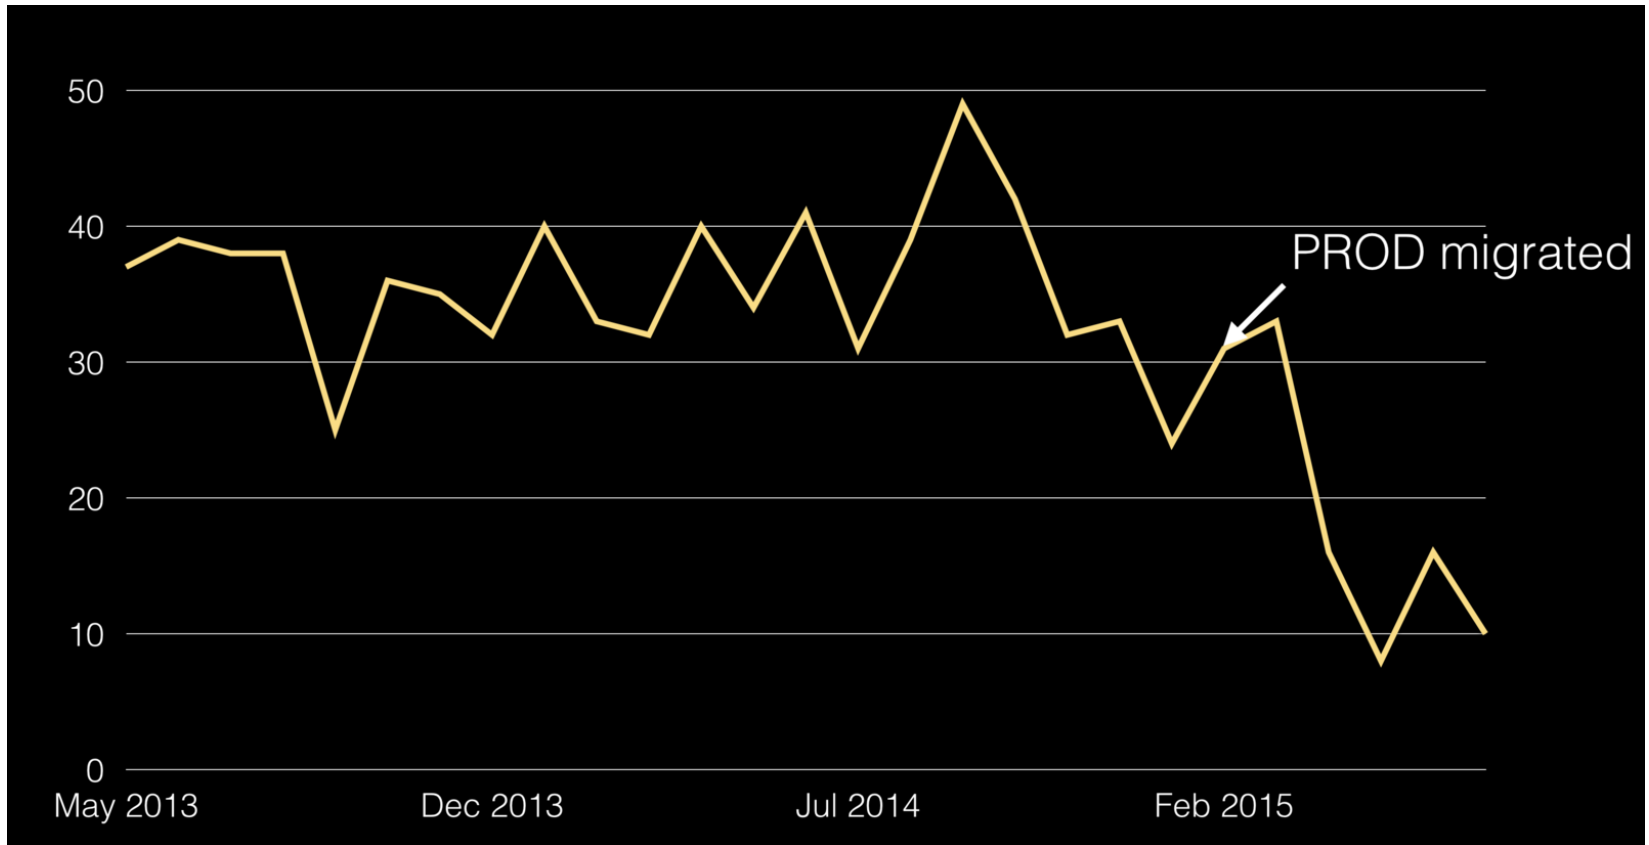
Migration to Mesos - Timeline



Manual Server Provisioning



Server provisioning UI usage



1800+ deployables

~300 deploys / day

**10 minutes from git
push to production**

I want persistent IMAP connections to 200k+
inboxes (Jul 2015)





Migration Issues

Stateful Services

Single Process services

Hard coded stationary hosts

Cgroups memory isolation

User resistance



Operational Issues

All eggs in one basket

Mesos / Framework issues (pingback port)

Failures (Zookeeper, Mesos, Singularity)

Cluster Maintenance

Missing features



Maintenance

Phased rollout of new Kernel, Instance types

Rolling upgrade of instance basic SW with puppet vars

Rolling upgrade of master/agent process with ansible

Local testing on docker cluster

Roll out at **infra-QA** then **product-QA** and last to **Production** cluster

Deploy tools deploy themselves but maintain command line alternative with fabric



Optionality with Mesos@HubSpot


Massive Builds in Mesos



Mesos Spark Cluster



Singularity

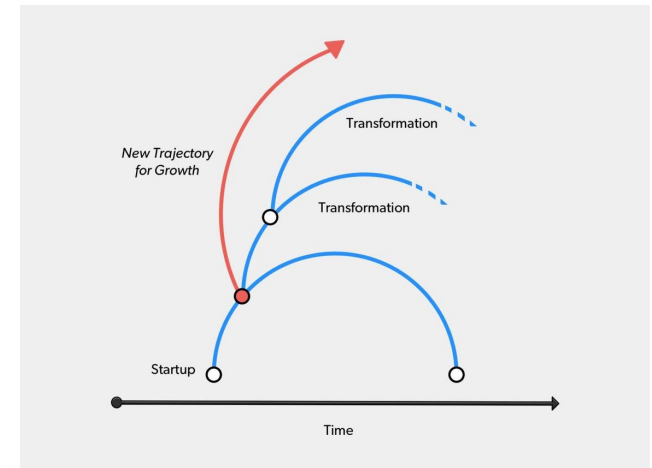
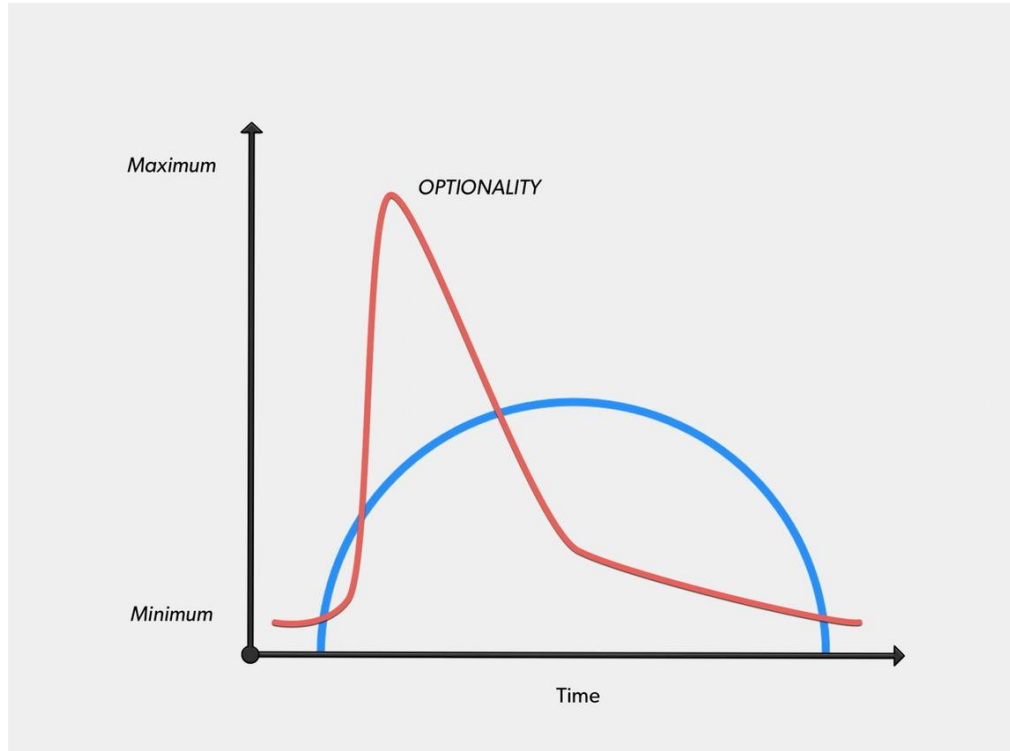


Baragon - Tasks Load Balancer Manager



Ghidorah - Load Balancers in Mesos

The sweet spot



Source: Mark Leslie (<http://firstround.com/review/The-Arc-of-Company-Life-and-How-to-Prolong-It/>)



Invest early in a deploy & build infrastructure

**Dedicate 1-2 engineers to experiment on
a Mesos powered PaaS**

Try Singularity today!

[**github.com/HubSpot/Singularity**](https://github.com/HubSpot/Singularity)

Useful links

[HubSpot Blog: How We Built Our Stack For Shipping at Scale](#)

[Blazar: An out-of-this world build system!](#)

[Baragon: Load Balancer API](#)