

# CONTINUOUS DEPLOYMENT WITH SINGULARITY

Large Scale Mission-Critical Service and Job Deployment

Gregory Chomatas  
@gchomatas



# PAAS TEAM


Implement & maintain:

- the deploy & build tools
- the PAAS platform (mesos clusters)
- load balancer tools
- logging infrastructure


Boston: Whitney Sorenson, Tom Petr, Tim Finley



Dublin: Gregory Chomatas, Kieran Manning

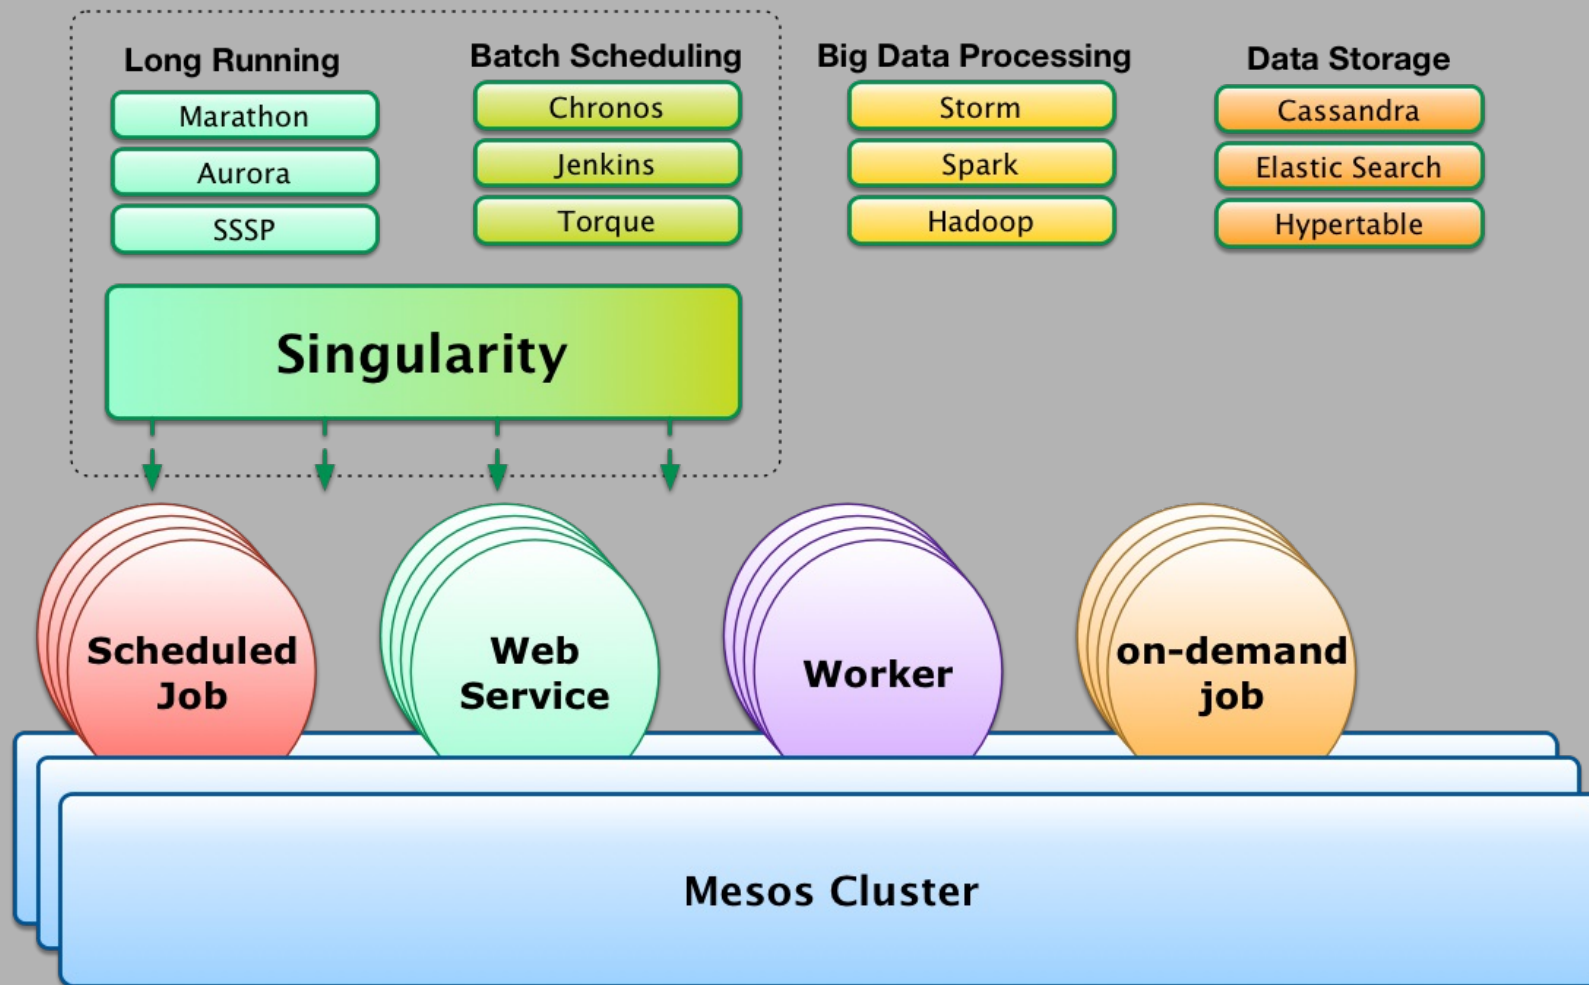


# AN ESSENTIAL SINGULARITY $\exp(1/z)$



# HUBSPOT SINGULARITY

## Mesos Frameworks



# THE WAY TO MESOS

Speed wins -> Speed Product Development

Increase change rate -> Remove Friction + Reduce size, cost, risk of change:

small teams, high trust, low process  
freedom and responsibility culture

micro services

libs & cross cutting APIs to simplify coding  
automate deployment by tooling

# SOME FACTS & NUMBERS

3-4 person teams

several micro-services & jobs per team (full operation)

1 or more services per dev

All QA in MESOS / Part of PROD with plan to move all

**400** deploys / day - **843** Deployable Items:

**219 WEBSERVICES** (long running with an API)

**246 WORKERS** (long running no API)

**205 SCHEDULED JOBS** (CRON schedule)

**173 ON-DEMAND**

# SOME FACTS & NUMBERS

## QA Environment

pre-mesos:

400 small & medium size servers (c1.xlarge)

post-mesos:

20 big servers (c3.8xlarge)

# WHY SINGULARITY

almost **no** framework 1 year ago  
get a consistent, unified API for all deployable items

mission critical / strategic tool - important to control:  
priority and delivery of bug fixes  
features and integrations  
the overall roadmap

have the resources to implement & maintain a highly complex  
piece of software



# DEPLOY CONFIGURATION

```
name: MDS_All_Item_Types_In_One_Config
buildName: MesosDeployIntegrationTestsProject
type: procfile

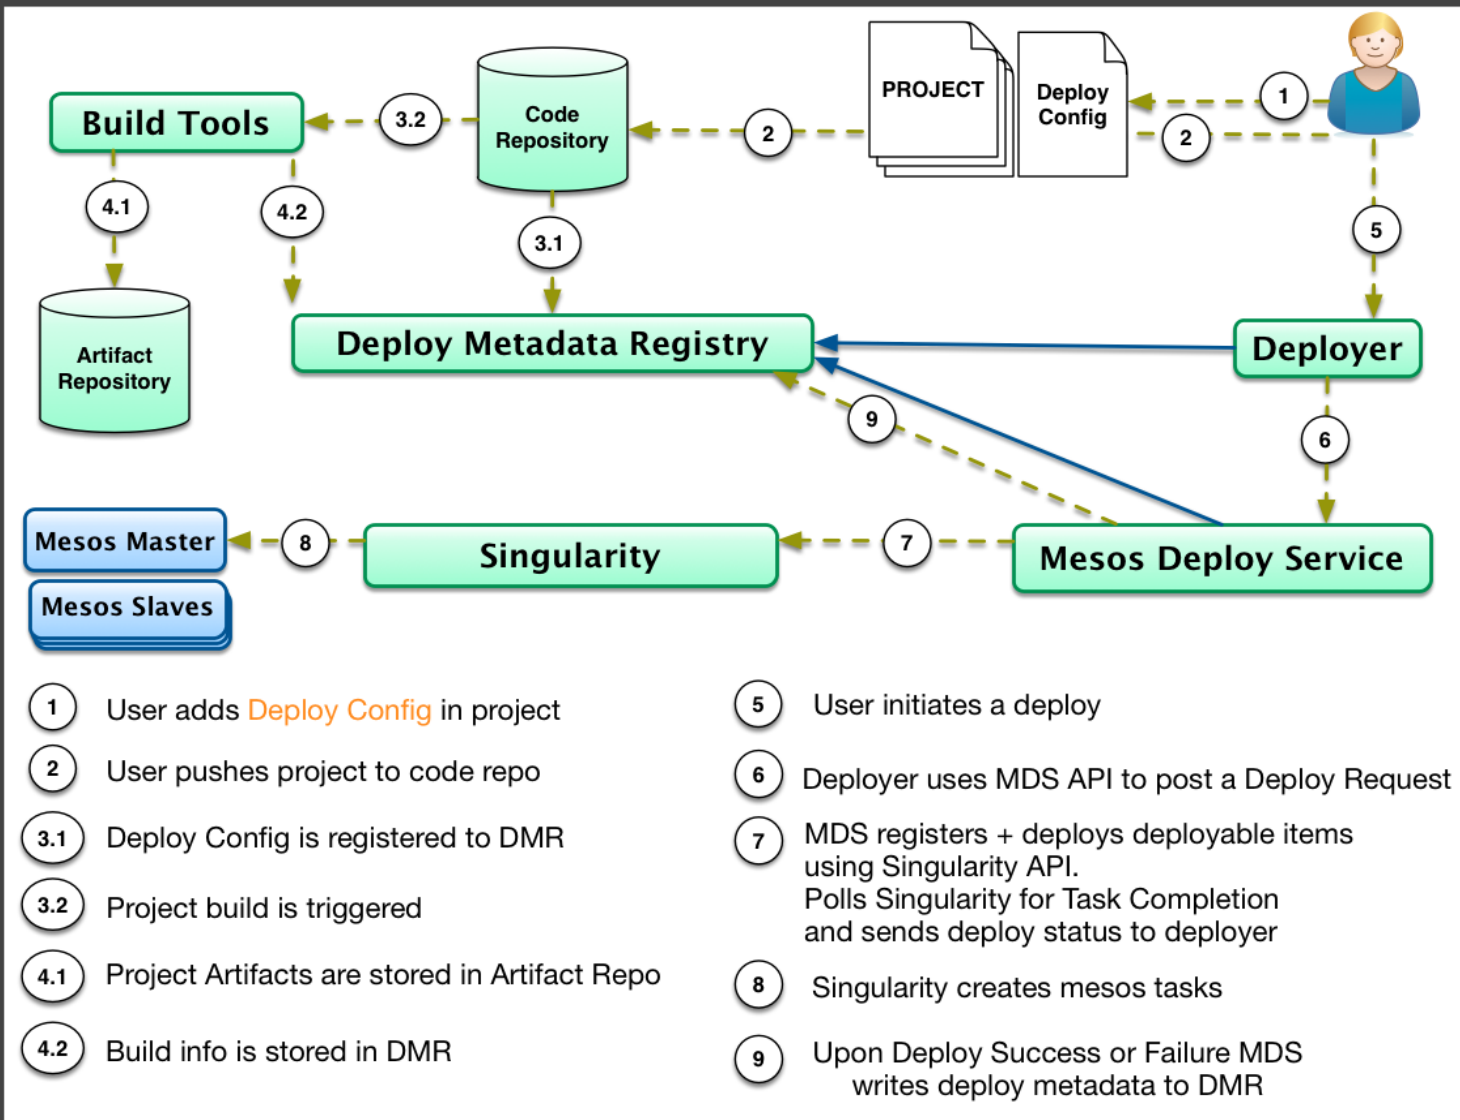
owners:
- user@hubspot.com

appRoot: /mesos-deploy-test-srv1/v1
loadBalancers:
- test

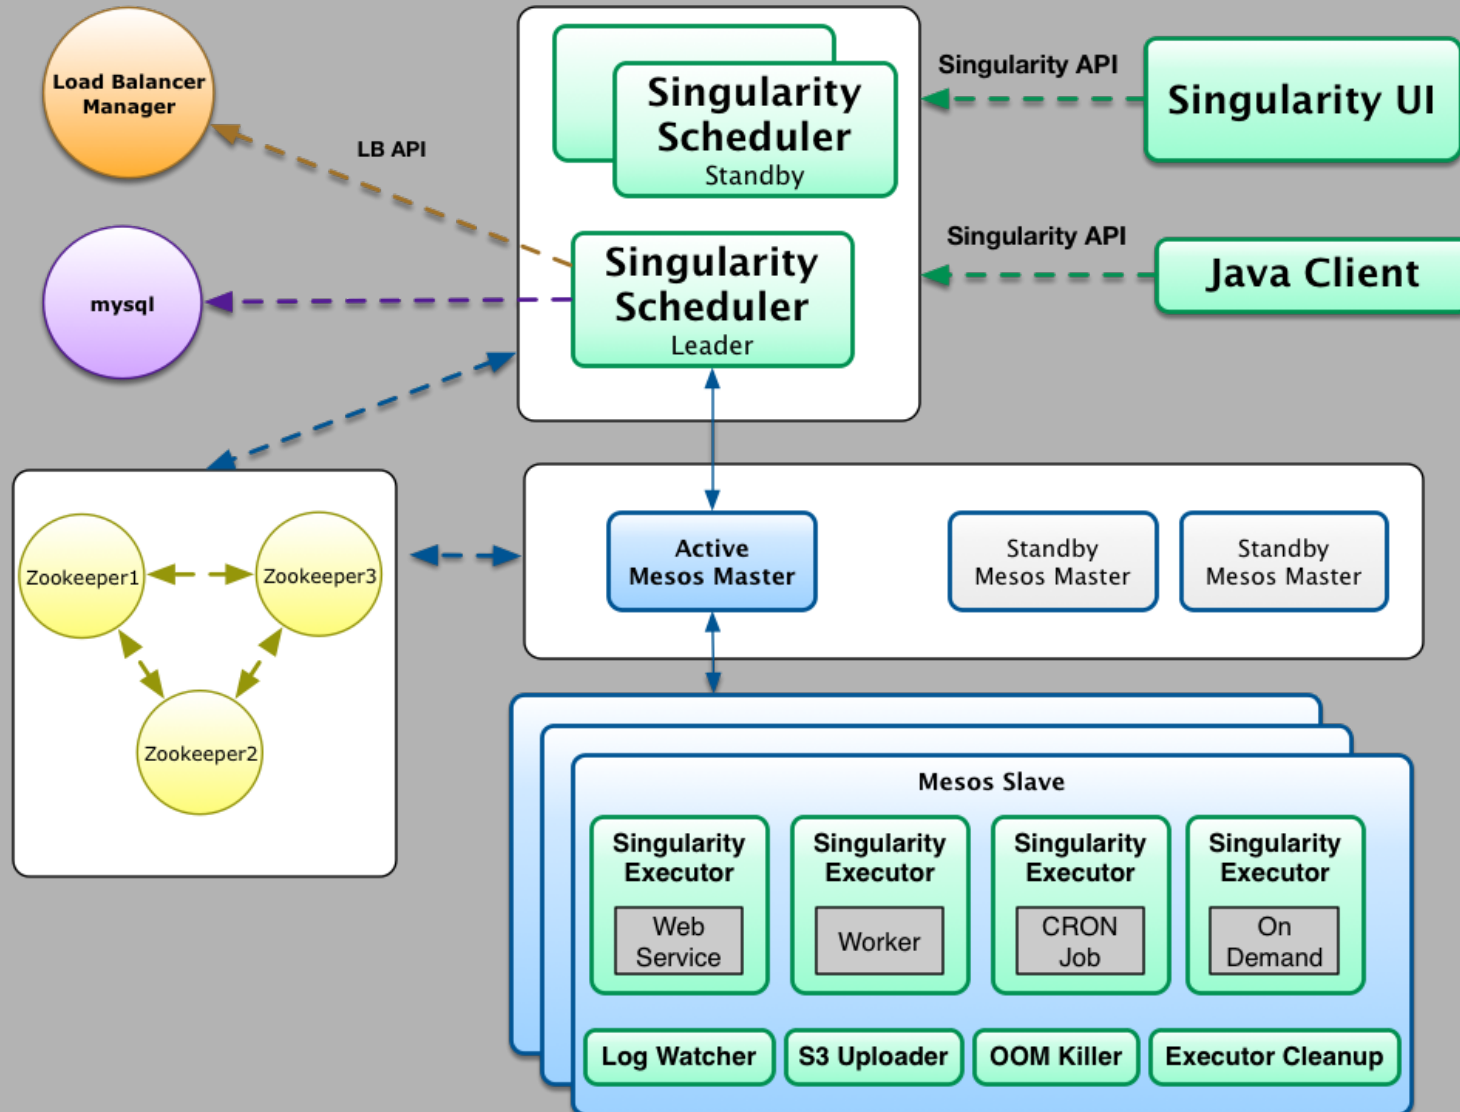
env:
  all:
    JOB_JAR: TestJob.jar

procfile:
  webService:
    cmd: java $JVM_DEFAULT_OPTS -jar TestService.jar server $CONFIG_YAML
    instances: 2
    cpus: 2
    memory: 1024
    numRetriesOnFailure: 5
  scheduledJob:
    cmd: java $JVM_DEFAULT_OPTS -jar $JOB_JAR -testjob
    schedule: '*/3 * * * *'
    numRetriesOnFailure: 5
```

# DEPLOY WITH HUBSPOT PAAS



# SINGULARITY COMPONENTS



# SINGULARITY SCHEDULER

## A DEPLOY-CENTRIC REST API TO:

- register deployable items
- execute their deploys
- view sandbox files
- get metadata / historical data

# SINGULARITY SCHEDULER

Advanced features:

Health Checking at the process and the service endpoint level

Automatic cool-down of repeatedly failing services

Load balancing of service instances (LB API)

Automatic Rollback of failed deploys

Reconciliation of LOST tasks

Decommissioning of Slaves & Racks

# SINGULARITY EXECUTOR

Log Rotation

Task Sandbox Cleanup

Graceful Task Killing with configurable timeout

Environment Setup

Task Runner Script

# ADVANCED SLAVE SERVICES

**Log Watcher** : Tail & Stream Logs

**S3 uploader** : Archive logs with AWS S3 Service

**Executor Cleanup** : Clean failed executor tasks

**OOM Killer** : replace the default memory limit checking supported by Linux Kernel CGROUPS

# KEY SINGULARITY ABSTRACTIONS

## SINGULARITY REQUEST OBJECT

```
{  
  "id": "TestService",  
  "owners": [  
    "feature_x_team@mycompany.com",  
    "developer@mycompany.com"  
  ],  
  "daemon": true,  
  "instances": 3,  
  "rackSensitive": true,  
  "loadBalanced": true  
}
```



# KEY SINGULARITY ABSTRACTIONS

## SINGULARITY DEPLOY OBJECT

**RESOURCES:** Memory, CPUs, network ports

**HEALTH CHECKS:** Timeouts and URLs

**LOAD BALANCING** of web service instances (LB groups, api base path)

**EXECUTOR INFORMATION:** execution environment, executable artifacts, configuration files, command to execute, executor to use, etc.

```
{  
  "requestId": "MDS_TestService",  
  "id": "71_7",  
  "customExecutorCmd": ".../singularity-executor",  
  "resources": {  
    "cpus": 1,  
    "memoryMb": 896,  
    "numPorts": 3  
  },  
  "env": {  
    "DEPLOY_MEM": "768",  
    "JVM_MAX_HEAP": "384m",  
  },  
}
```

```
"executorData": {
  "cmd": "java -Xmx$JVM_MAX_HEAP -jar .../TestService.jar server $CONFIG_YAML",
  "embeddedArtifacts": [
    {
      "name": "rawDeployConfig",
      "filename": "TestService.yaml",
      "content": "bmFtZT..."
    }
  ],
  "externalArtifacts": [],
  "s3Artifacts": [
    {
      "name": "executableSlug",
      "filename": "TestService.tar.gz",
      "md5sum": "313be85c5979a1c652ec93e305eb25e9",
      "filesize": 81055833,
      "s3Bucket": "hubspot.com",
      "s3ObjectKey": "build_artifacts/.../TestService.tar.gz"
    }
  ],
}
```

# SINGULARITY API

## MANAGE DEPLOYABLE ITEMS

ENDPOINT: `/requests`

`register / update / unregister` an item

get info about an item

list items in `active | paused | cool-down` state

`run / restart / pause / un-pause` an item

# SINGULARITY API

## DEPLOY THE DEPLOYABLE ITEMS

ENDPOINT: `/deploys`

`deploy` an already registered item

`cancel` a pending deploy

# SINGULARITY API

## MANAGE DEPLOYABLE ITEM INSTANCES (TASKS )

ENDPOINT: `/tasks`

get the list of all scheduled tasks (not yet active)

get scheduled tasks for a specific item

list tasks in `active | cleaning | lbcleanup` state

info about a specific task

active tasks in a slave

Kill a task

# SINGULARITY API

Historical Information about deployable items & their tasks

ENDPOINT: [/history](#)

a single task history

tasks that have run in the past

all previous item updates

search for historical items by item id

all item deploys

a specific item deploy

# SINGULARITY API

## LIST & DOWNLOAD FILES IN ACTIVE TASK SANDBOX

ENDPOINT: `/sandbox`

list all task files

read file chunks

download a file



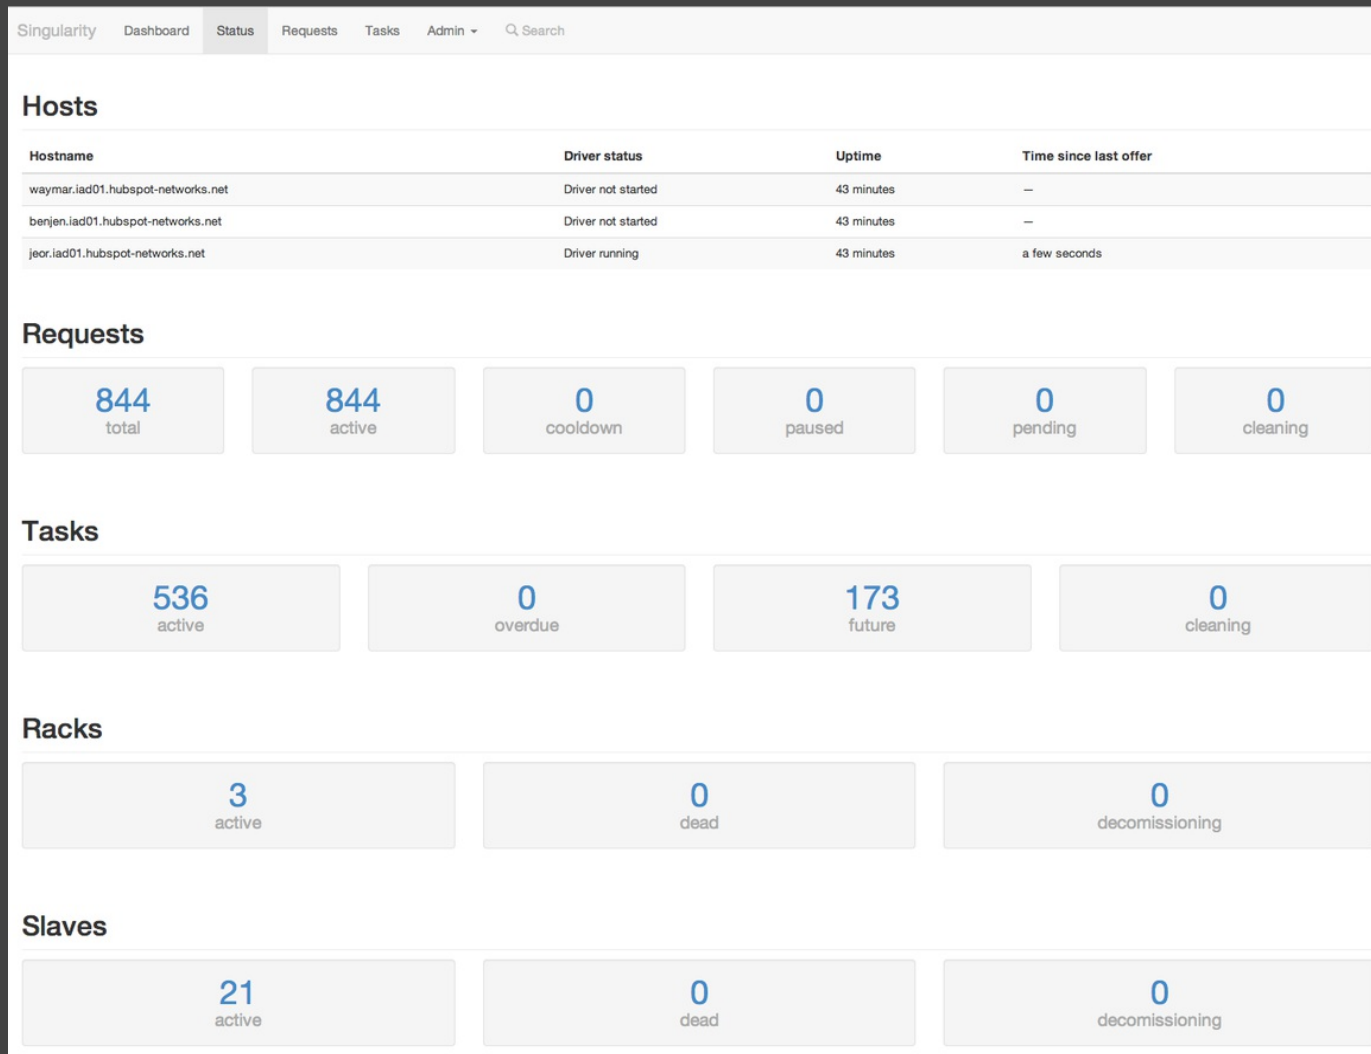
# SINGULARITY API

## Cluster STATE Information

ENDPOINT: `/state`

```
{
  activeTasks: 567,
  activeRequests: 843,
  cooldownRequests: 1,
  scheduledTasks: 142,
  pendingRequests: 0,
  lbCleanupTasks: 1,
  activeSlaves: 21,
  deadSlaves: 0,
  decommissioningSlaves: 0,
  activeRacks: 3,
  deadRacks: 0,
  futureTasks: 142,
  maxTaskLag: 0,
  overProvisionedRequests: 0,
  underProvisionedRequests: 0,
  allRequests: 844
}
```

# SINGULARITY UI - GLOBAL CLUSTER STATUS



# SINGULARITY UI - DASHBOARD

Search **Dashboard** Status Requests Tasks Admin ▾

Singularity

**gchomatas** [change](#)

## Requests

6  
total

0  
daemon

4  
on-demand

2  
scheduled

## Starred requests

	Request	Requested	Deploy user	Instances
★	<a href="#">MDS_TestService_TestJob_TestWorker_TestOnDemand-on-demand</a>	a minute ago	gchomatas	
★	<a href="#">MDS_TestService_Traveview_Enabled-web</a>	6 days ago (8/1/2014 4:55pm)	gchomatas	2

# SINGULARITY UI - DEPLOYABLE ITEM LIST

Singularity

Dashboard

Status

Requests

Tasks

Admin ▾

🔍 Search

All

Active

Cooldown

Paused

Pending

Cleaning
















MDS

✓ Daemon

✓ On-demand

✓ Scheduled

Protip: press **cmd** to select multiple

Name	State	Instances	Deploy ID	Deploy user	Deploy time	Schedule	
★ <a href="#">MDS_All_Item_Types_In_One_Config-onDemand</a>	Active	1	78_1	gchomatas	2 minutes ago		  
★ <a href="#">MDS_All_Item_Types_In_One_Config-worker</a>	Active	2	78_1	gchomatas	2 minutes ago		 
★ <a href="#">MDS_All_Item_Types_In_One_Config-scheduledJob</a>	Active	1	78_1	gchomatas	2 minutes ago	0 */3 * * * ?	  
★ <a href="#">MDS_All_Item_Types_In_One_Config-webService</a>	Active	2	78_1	gchomatas	2 minutes ago		 
★ <a href="#">MDS_TestService_Traveview_Enabled-web</a>	Active	2	71_7	gchomatas	6 days ago		 
★ <a href="#">MDS_TestJob-testjob</a>	Active	1	63_1m	gchomatas	10 days ago	0 */20 * * * ?	  

# SINGULARITY UI - DEPLOYABLE ITEM

Singularity

Dashboard

Status

Requests

Tasks

Admin ▾

🔍 Search

active

daemon

JSON

Pause





Bounce

Remove





## MDS\_All\_Item\_Types\_In\_One\_Config-webService

Active deploy 78\_2 by gchomatas — 11 minutes ago (7 Aug 2014 23:00) [Deploy history](#)



### Running instances

Name	Status	Deploy ID	Started	Updated	
<a href="#">MDS_All_Item_Types_In_One_Config-webService-78_2-1407448842109-1-darry-us_east_1b</a>	Task running	78_2	11 minutes ago	11 minutes ago	 
<a href="#">MDS_All_Item_Types_In_One_Config-webService-78_2-1407448843093-2-courtenay-us_east_1a</a>	Task running	78_2	11 minutes ago	11 minutes ago	 


### Task history

Name	Last state	Deploy	Started	Updated	
<a href="#">MDS_All_Item_Types_In_One_Config-webService-78_1-1407448460420-2-arianne-us_east_1a</a>	Task killed	78_1	18 minutes ago	11 minutes ago	 
<a href="#">MDS_All_Item_Types_In_One_Config-webService-78_1-1407448459429-1-aegon-us_east_1b</a>	Task killed	78_1	18 minutes ago	11 minutes ago	 

### Deploy history

Deploy ID	Status	User	Timestamp	
78_2	Succeeded	gchomatas	11 minutes ago (7 Aug 2014 23:00)	
78_1	Succeeded	gchomatas	18 minutes ago (7 Aug 2014 22:54)	

### Request history

State	User	Created	
Created	gchomatas	18 minutes ago	

# SINGULARITY UI - DEPLOYABLE ITEM TASK

[Singularity](#) [Dashboard](#) [Status](#) [Requests](#) [Tasks](#) [Admin](#) [Search](#)

Request [MDS\\_All\\_Item\\_Types\\_In\\_One\\_Config-webService](#)

[JSON](#) [Kill task](#)

## MDS\_All\_Item\_Types\_In\_One\_Config-webService-78\_2-140778235...

















Task running as of 15 minutes ago (11 Aug 2014 19:39) — Task running process bash - (pid: 473)

## History

Status	Message	Time
Task starting	Staging files... (executor pid: 385)	15 minutes ago (11 Aug 2014 19:39)
Task running	Task running process bash - (pid: 473)	15 minutes ago (11 Aug 2014 19:39)

## Files

MDS\_All\_Item\_Types\_In\_One\_Config-webService-78\_2-140778235556...

Name	Size	Last modified
 app/		11 Aug 2014 19:39
 bin/		7 Aug 2014 22:50
 conf/		7 Aug 2014 22:52
 logs/		11 Aug 2014 19:39
 MDS_All_Item_Types_In_One_Config.yaml	1 KB	11 Aug 2014 19:39 
 deploy.env	2 KB	11 Aug 2014 19:39 
 executor.bash.log	433 B	11 Aug 2014 19:39 
 executor.java.log	5 KB	11 Aug 2014 19:39 
 runner.sh	1 KB	11 Aug 2014 19:39 
 service.log	112 KB	11 Aug 2014 19:53 

# SINGULARITY UI - HISTORICAL TASK

SingularityDashboardStatusRequestsTasksAdminSearch

Request MDS\_All\_Item\_Types\_In\_One\_Config-webServiceJSON

MDS\_All\_Item\_Types\_In\_One\_Config-webService-78\_2-140778234...

Task killed as of 22 minutes ago (11 Aug 2014 19:39) — Task killed. Process exited gracefully with code 143

History

Status	Message	Time
Task starting	Staging files... (executor pid: 3807)	22 minutes ago (11 Aug 2014 19:39)
Task running	Task running process bash - (pid: 3908)	22 minutes ago (11 Aug 2014 19:39)
Task cleaning	SCALING_DOWN	22 minutes ago (11 Aug 2014 19:39)
Task killed	Task killed. Process exited gracefully with code 143	22 minutes ago (11 Aug 2014 19:39)

Files

MDS\_All\_Item\_Types\_In\_One\_Config-webService-78\_2-140778234160...

Name	Size	Last modified
conf/		7 Aug 2014 22:52
bin/		7 Aug 2014 22:50
logs/		11 Aug 2014 19:41
MDS_All_Item_Types_In_One_Config.yaml	1 KB	11 Aug 2014 19:39
executor.bash.log	433 B	11 Aug 2014 19:39
tail_of_finished_service.log	31 KB	11 Aug 2014 19:39
logrotate.status	786 B	11 Aug 2014 19:39
deploy.env	2 KB	11 Aug 2014 19:39
runner.sh	1 KB	11 Aug 2014 19:39
service.log	0 B	11 Aug 2014 19:39
executor.java.log	10 KB	11 Aug 2014 19:39

S3 Logs

Log file	Size	Last modified
MDS_All_Item_Types_In_One_Config-webService-78_2-1407782341600-2-elia-us_east...	23 KB	11 Aug 2014 19:41

# SINGULARITY UI - RACKS & SLAVES

Search Dashboard Status Requests Tasks Admin			Singularity
Racks			
Active racks			
ID	Uptime		
us_east_1b	19 hours (8/6/2014 9:08pm)	Decommission	
us_east_1a	19 hours (8/6/2014 9:08pm)	Decommission	
us_east_1e	19 hours (8/6/2014 9:08pm)	Decommission	

Search Dashboard Status Requests Tasks Admin				Singularity
Slaves				
Active slaves				
ID	Rack	Host	Uptime	
20140724-180527-3634914058-5050-7632-29	us_east_1b	darklyn	19 hours (8/6/2014 9:08pm)	Decommission
20140724-180527-3634914058-5050-7632-27	us_east_1e	lornel	19 hours (8/6/2014 9:08pm)	Decommission
20140724-180527-3634914058-5050-7632-37	us_east_1b	darry	19 hours (8/6/2014 9:08pm)	Decommission
20140724-180527-3634914058-5050-7632-28	us_east_1a	blackhair	19 hours (8/6/2014 9:08pm)	Decommission
20140724-180527-3634914058-5050-7632-25	us_east_1b	zhoe	19 hours (8/6/2014 9:08pm)	Decommission



# DEVELOP WITH SINGULARITY

java 7  
guice

dropwizard  
(jersey, jackson, liquibase)

maven

backbone  
nodejs  
brunch

# ROADMAP / NEW FEATURES

Enhance Job Scheduler

Support deploy of Docker containers

Add advanced slave affinity algorithms to support data locality for  
Big Data Analysis tasks

Open source Deployer

(A simplified version of Deploy Metadata Registry + Mesos  
Deploy Service + Deployer UI)

# USEFUL LINKS

<http://getsingularity.com/>

<https://github.com/HubSpot/Singularity>

[https://github.com/HubSpot/Singularity/blob/master/Docs/Singularity\\_API\\_Reference.md](https://github.com/HubSpot/Singularity/blob/master/Docs/Singularity_API_Reference.md)

[https://github.com/HubSpot/Singularity/blob/master/Docs/Singularity\\_Local\\_Setup\\_For\\_Testing.md](https://github.com/HubSpot/Singularity/blob/master/Docs/Singularity_Local_Setup_For_Testing.md)

<https://mesosphere.io/resources/mesos-case-study-hubspot/>