

Cracking the Container Scale Problem with Apache Mesos

Connor Doyle
connor@mesosphere.io

Sunil Shah
sunil@mesosphere.io

We are Mesosphere



**240 million monthly active
users**

500 million tweets per day

Up to 150k tweets per second

**More than 100TB per day of
compressed data**

“Mesos is the cornerstone of our elastic compute infrastructure – it's how we build all our new services and is critical for Twitter's continued success at scale. It's one of the primary keys to our data center efficiency.” – Chris Fry, SVP of Engineering at Twitter



airbnb

ebay

vimeo



HubSpot



xogito
...radical thinking...

PayPal

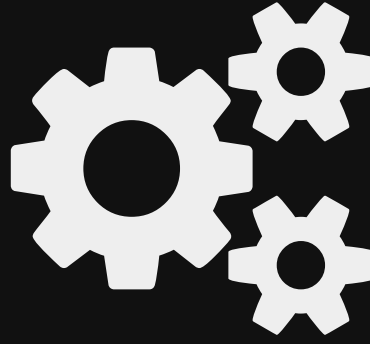


Atlassian





Mesoswhat?



Marathon

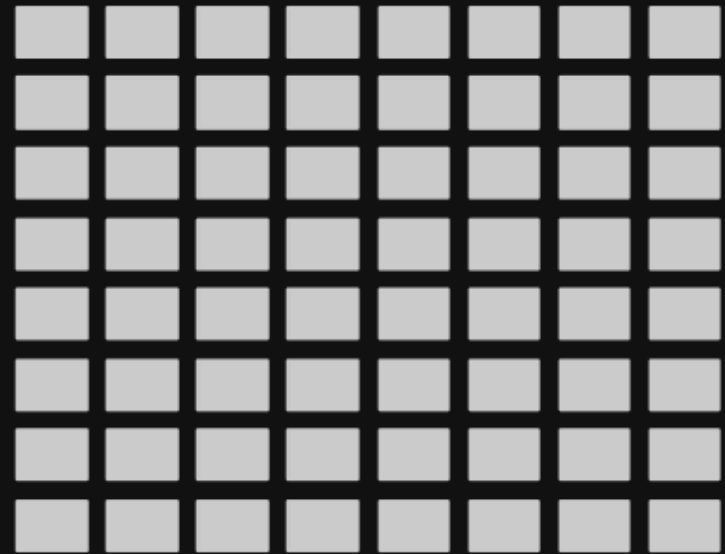
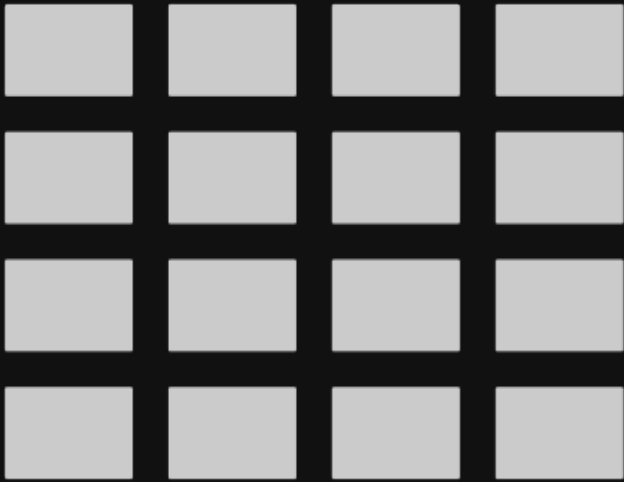


Chronos

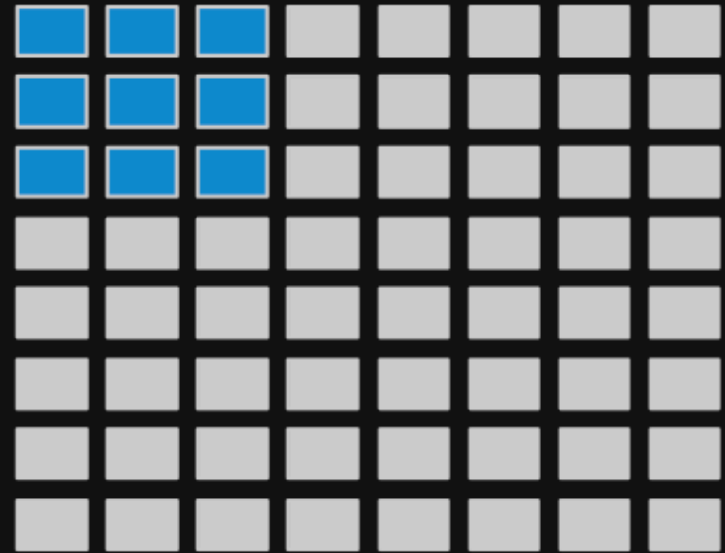
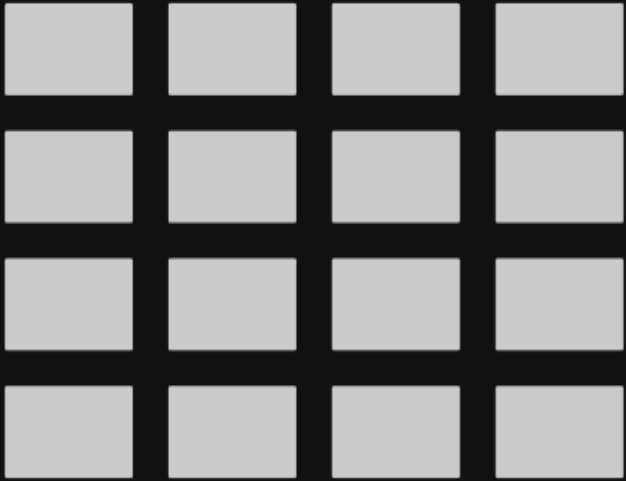


Demo`s!

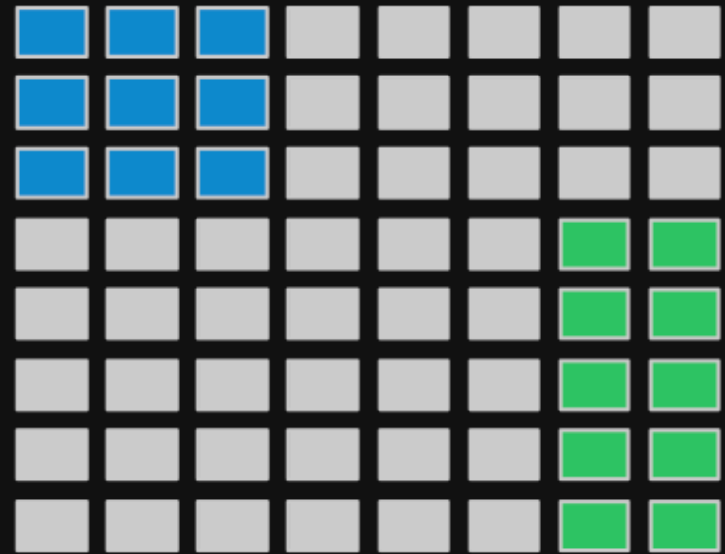
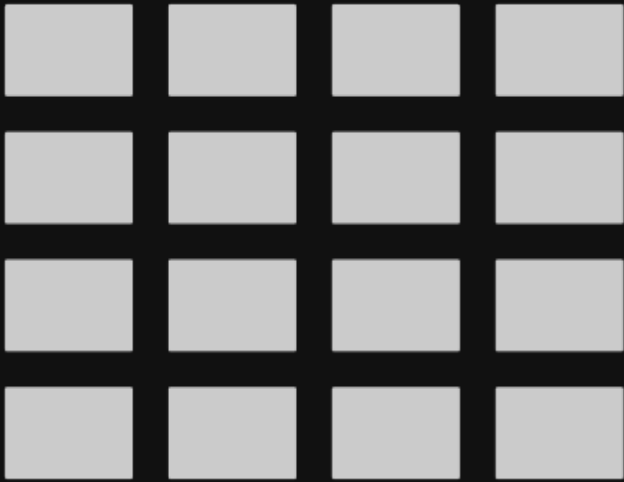
**Status quo is static partitioning
and use of virtual machines**



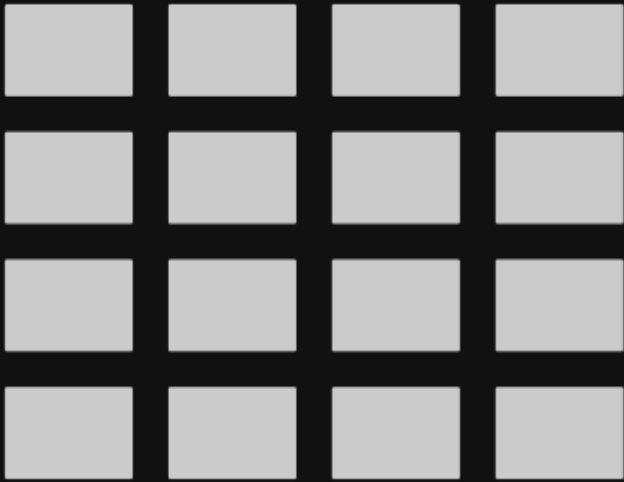
Add some virtual machines



Provision Hadoop



Provision a web service

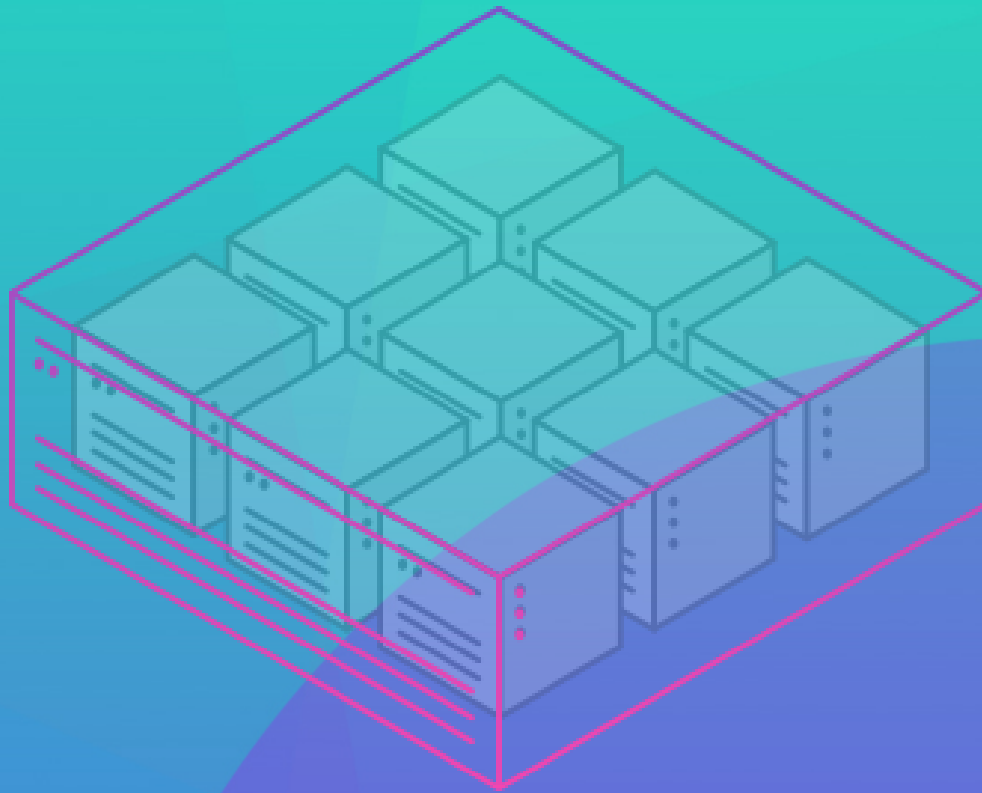


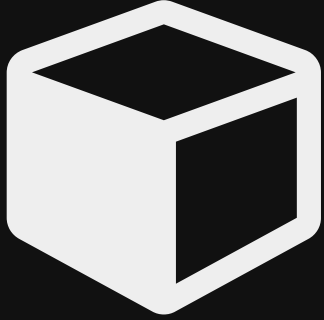
Moar data, moar Hadoop

Mesos let us treat a cluster of nodes...

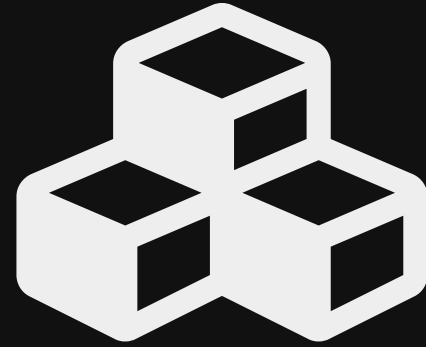


As one big computer





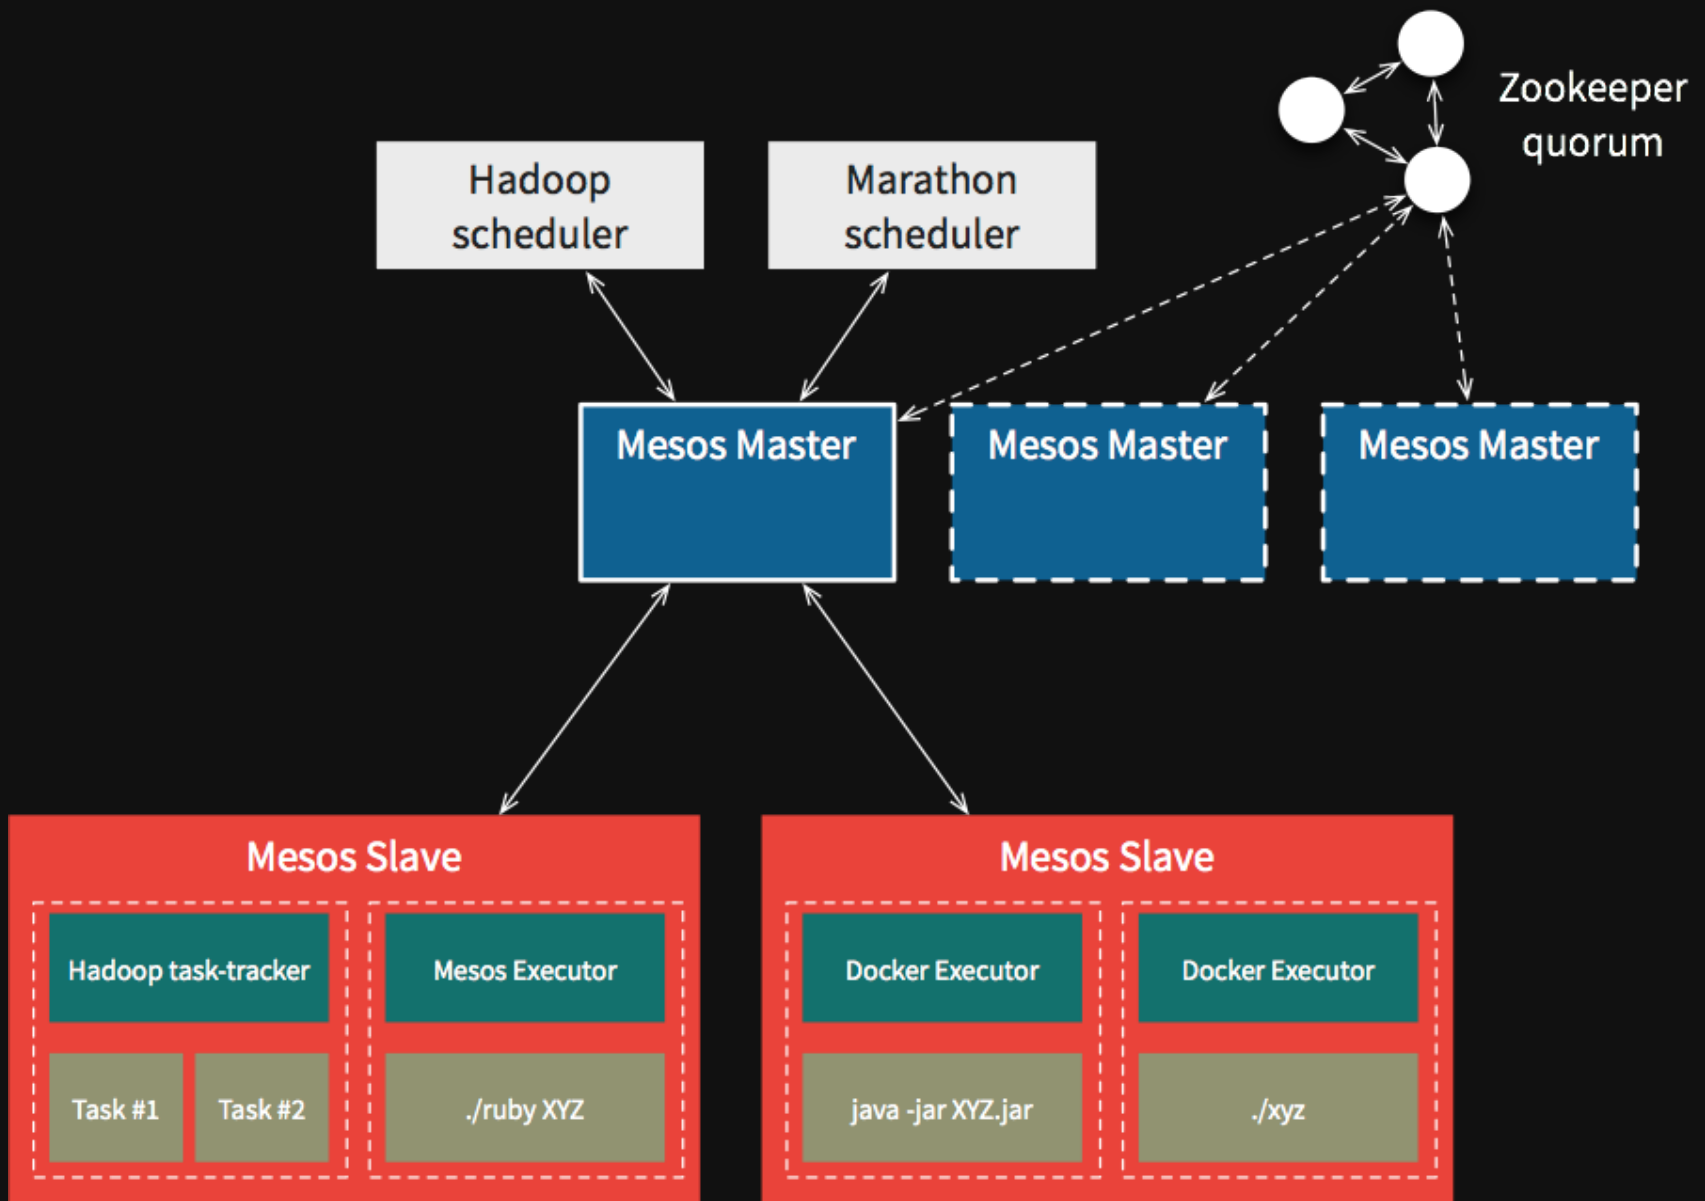
Not as individual machines

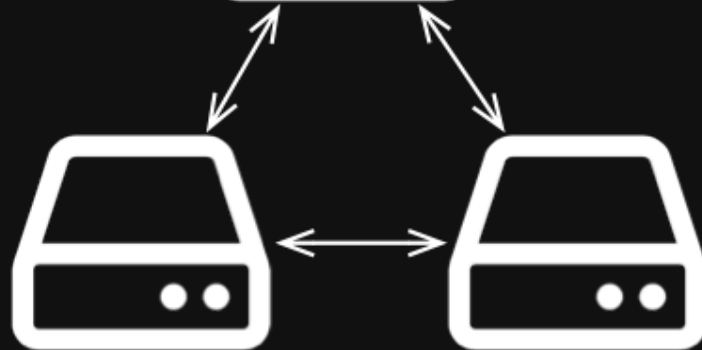


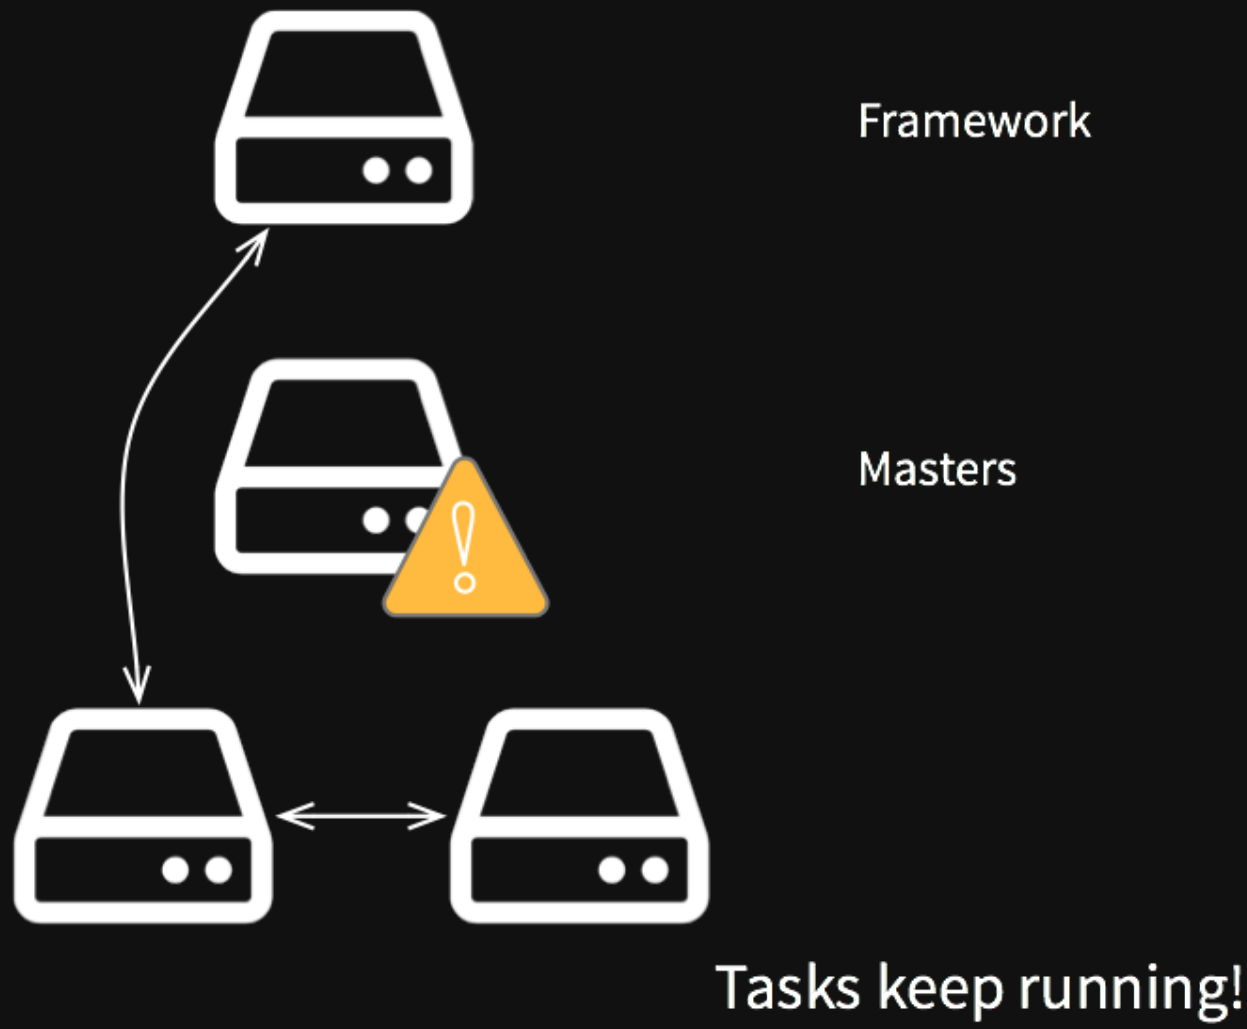
Not as VMs

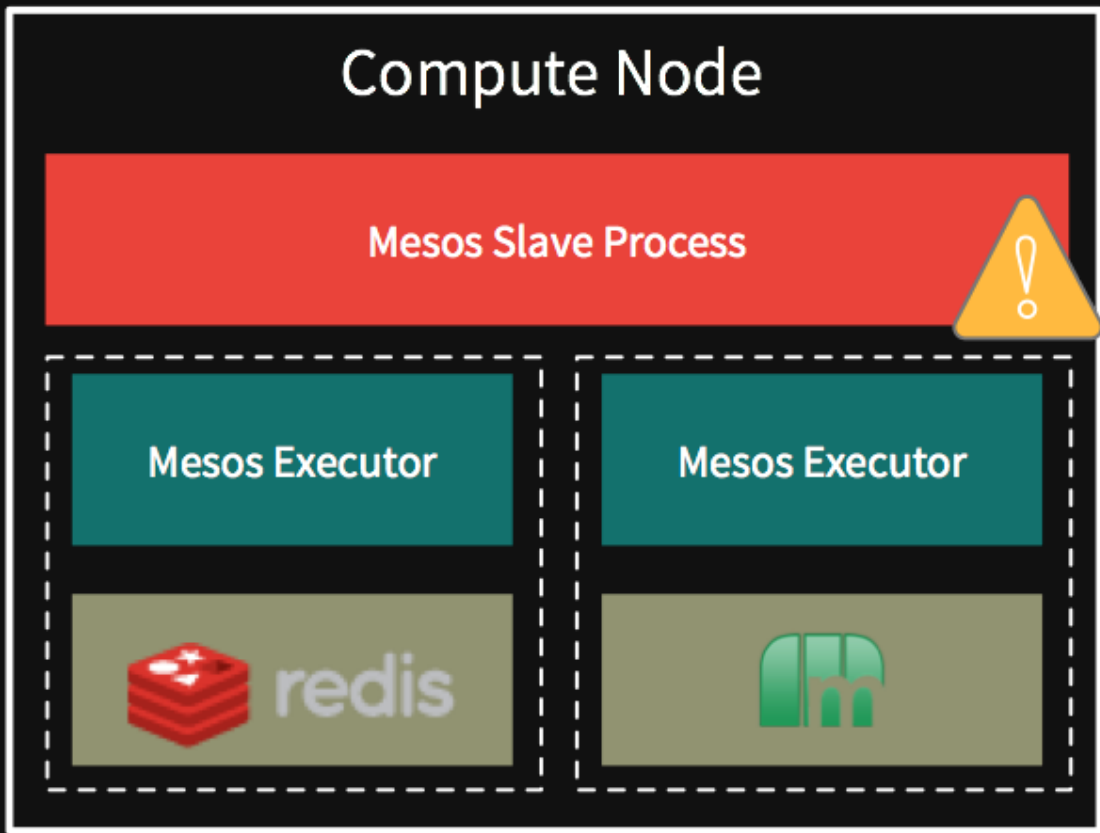


**But as computational resources
like cores, memory, disks, etc.**









Tasks keep running!

Mesos for all the things



Marathon

Kubernetes

Cassandra

Spark

Hadoop

Mesosphere Mesos Distribution

Linux

Linux

Linux

Linux

Linux

Linux






Linux

Linux

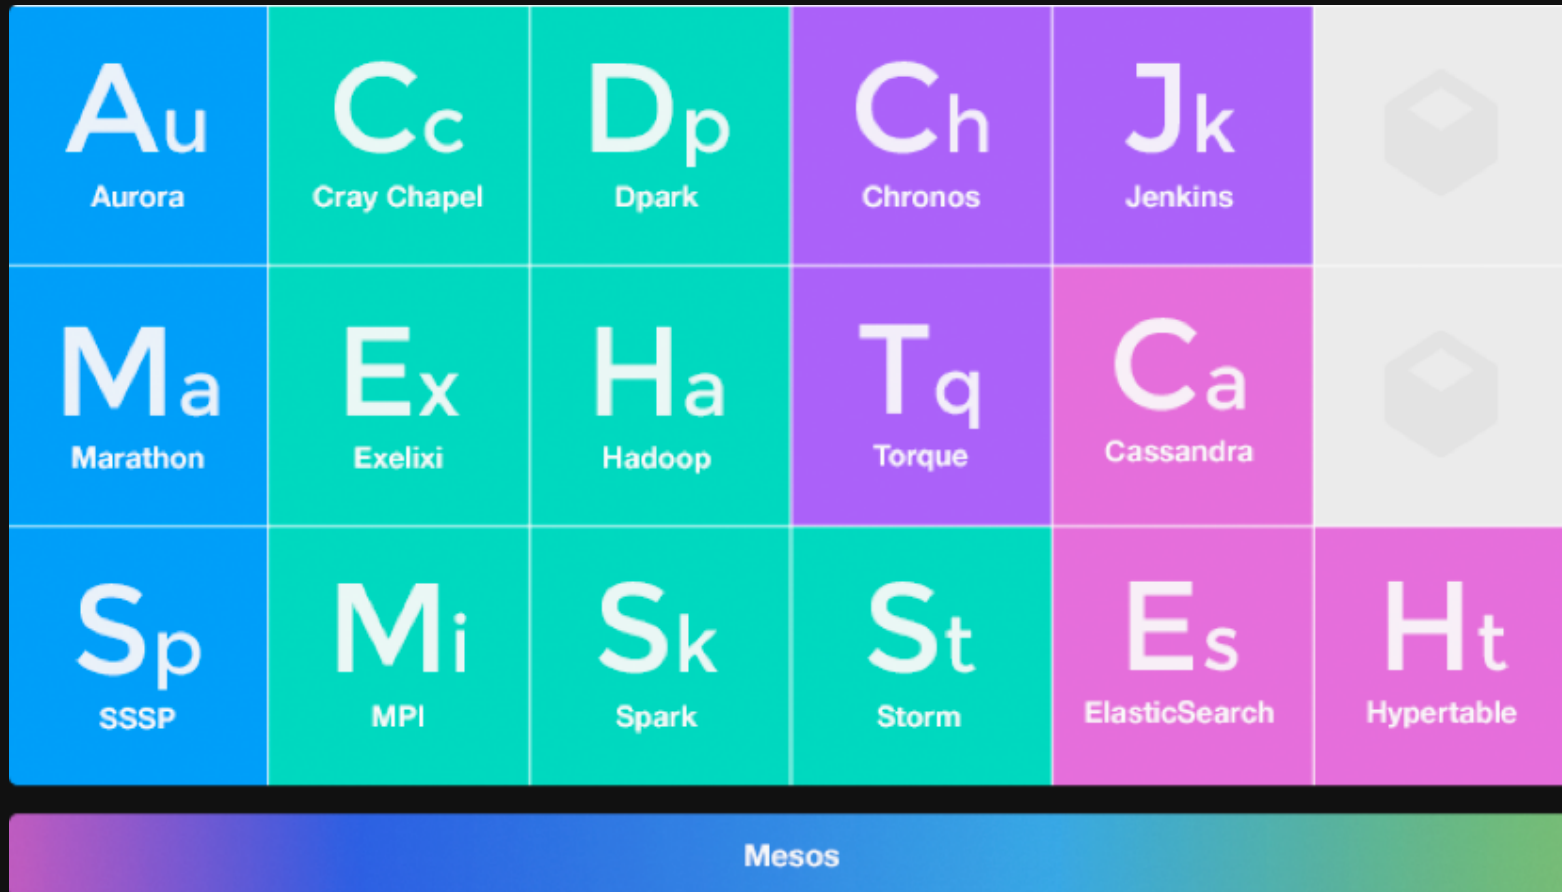
Linux

Linux

Mesos is...

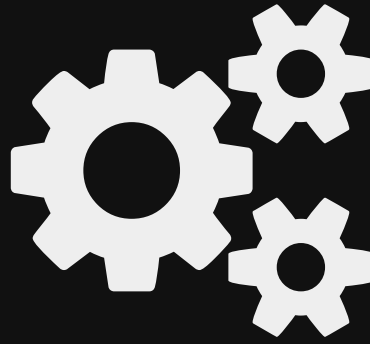
-  **Open Source Apache project**
-  **Cluster Resource Manager**
-  **Scalable to 10,000s of nodes**
-  **Fault-tolerant, battle-tested**
-  **An SDK for distributed apps**

The Mesos ecosystem is growing





Mesoswhat?



Marathon



Chronos



Demo!

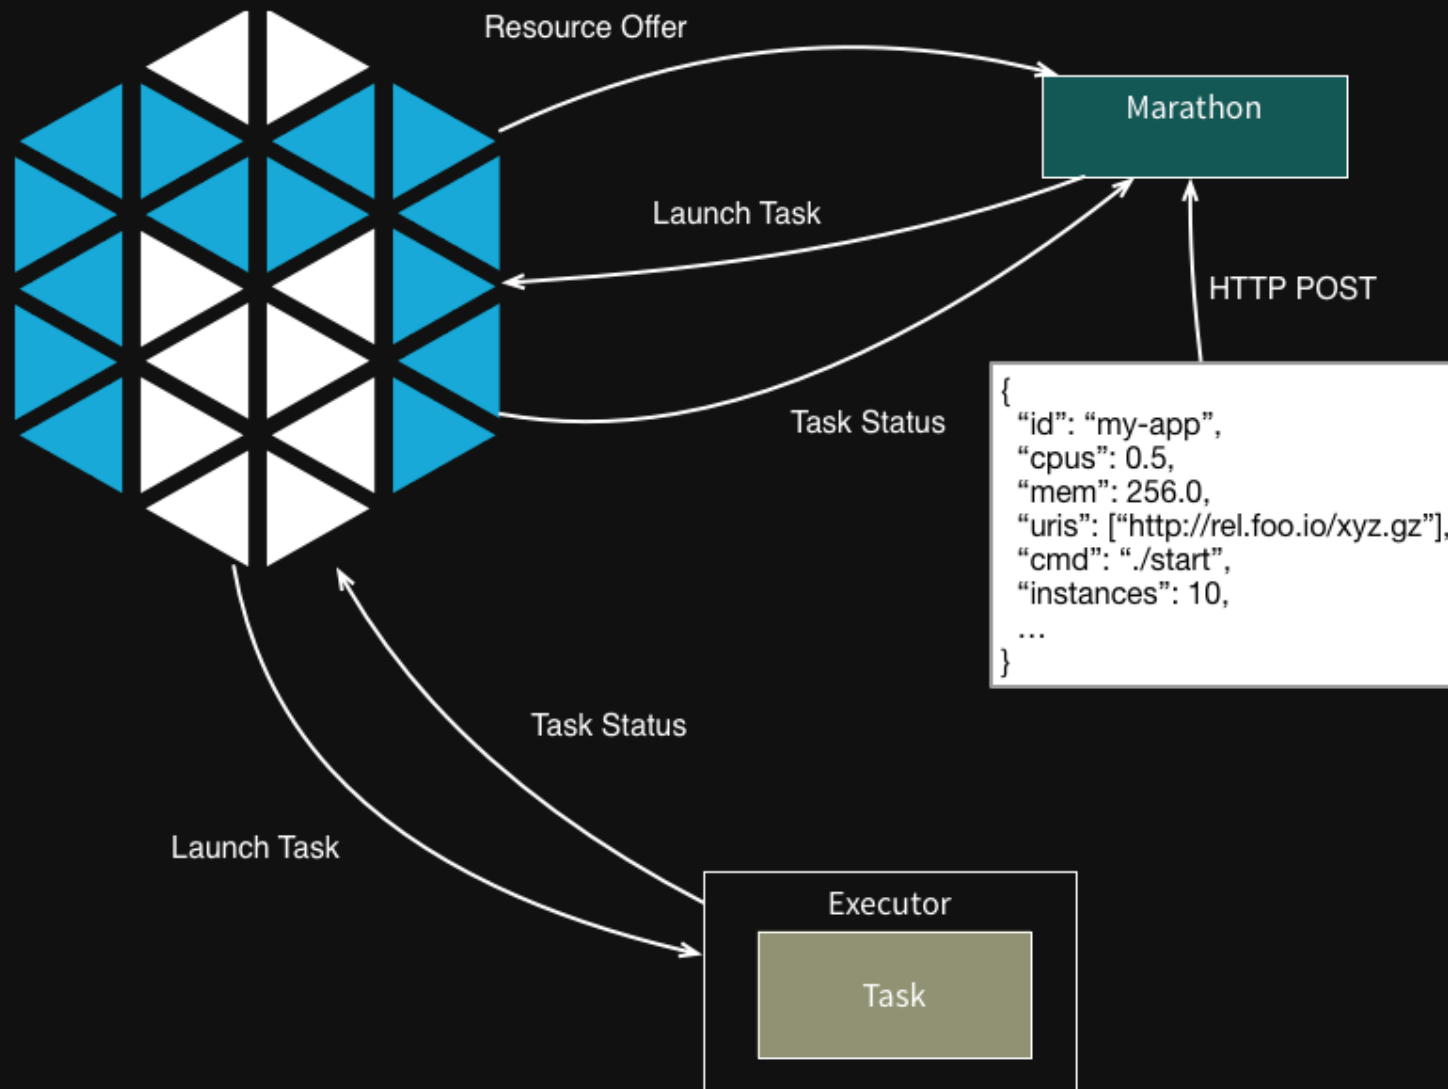
Say hi to Marathon



MARATHON

a self-serve interface to your cluster

distributed "init" for long-running services



a private fault-tolerant PaaS

a container orchestration platform



Marathon does it!

- **Start, stop, scale, update apps**
- **Nice web interface, API**
- **Highly available, no SPoF**
- **Native Docker support**
- **Pluggable event bus**
- **Rolling deploy / restart**
- **Application health checks**
- **Artifact staging**

Service Discovery

Set environment variables

Read config from device (rsync'ed to fs)

Read from K-V store

Use DNS

HAProxy works pretty well

Marathon ♥ REST

POST /v2/apps

GET /v2/apps

PUT /v2/apps/{appId}

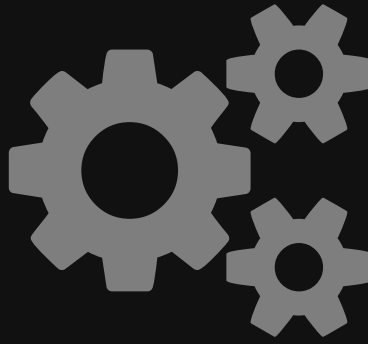
GET /v2/apps/{appId}/tasks

DELETE /v2/apps/{appId}/tasks/{taskId}

...



Mesoswhat?



Marathon



Chronos



Demo!

Introducing Chronos

The screenshot displays the Chronos web interface. On the left, a sidebar shows the 'CHRONOS' logo, a search bar with 'hostings', and statistics: 256 TOTAL JOBS and 16 FAILED JOBS. Below these are buttons for 'Dependency Graph' and 'New Job'. The main area is a table of jobs, all with a 'success' status. A modal window on the right shows the configuration for a job named 'Steve_Jobs', including a command to echo 'FOO' to a file, a schedule of 'T 11:44:03 T24H', and an owner 'ateam@airbnb.com'. A calendar widget is overlaid on the modal, showing the date 2013-03-15 selected.

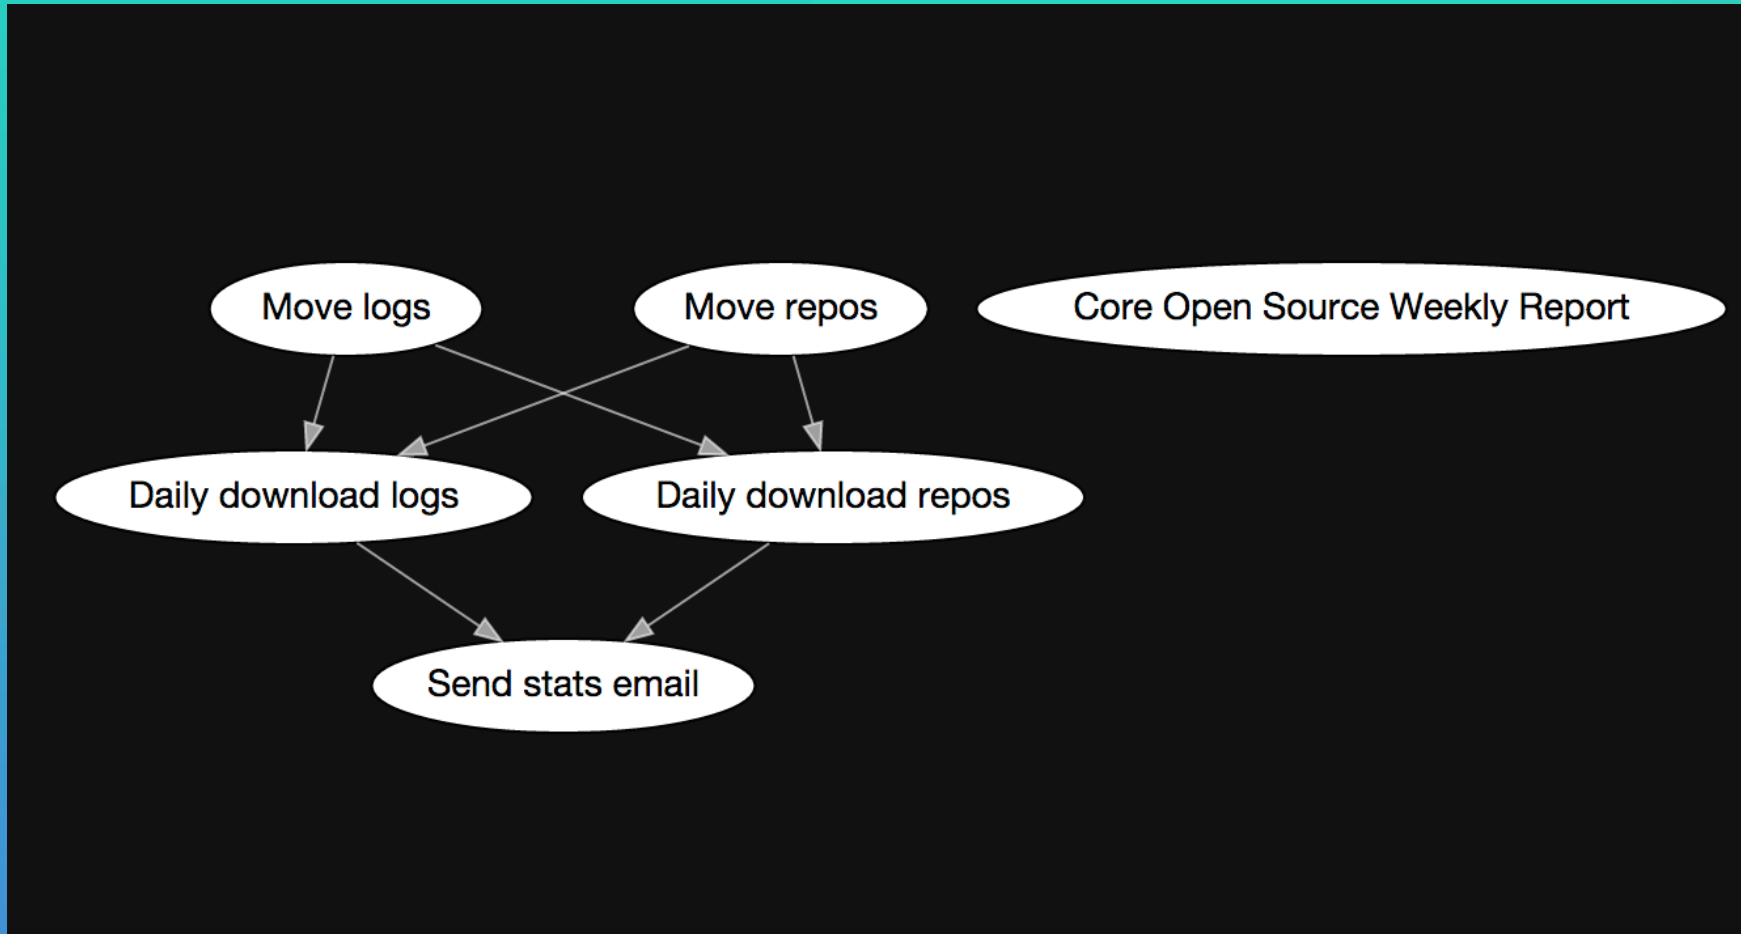
NAME	GRAPH	LAST
create_airbed_dump_table_hostings		success
create_airbed3_dump_table-hostings_first...		success
create_airbed_dump_table_hostings_with...		success
create_airbed_dump_table_collection_host...		success
create_omg_table-affiliate_events_hostings		success
hostings_summary		success
daily_gibson-import_airbed3_hostings		success
db_export-airbed_hostings		success
hostings_summary_2_quality_score		success
hostings_summary_1_pre		success
hostings_impressions_normalize		success
hostings_impressions_normalize_prepare		success
daily-update_hostings_summary_history		success
hostings_earnings_summary#async		success
daily-create_hostings_history		success
daily-update_hostings_history		success
daily-create hostings summary history		success

Modal Configuration:

- NAME: Steve_Jobs
- COMMAND: `echo 'FOO' >> /tmp/steve.txt`
- PARENTS: Choose parents...
- OWNER: ateam@airbnb.com
- SCHEDULE: R / 2013-03-15 T 11:44:03 Z/ P T24H
- EPSILON: PT15M
- EXECUTOR: /custom/execut...
- NAME: daily_gibson-default_data_cooked_table_import
- COMMAND:

a scheduler for batch and one-off jobs

Distribute a graph of jobs



Dependency graph for execution

Features

- **Distributed job scheduler**
- **Web interface, API**
- **Highly available, no SPoF**
- **Native Docker support**
- **Easy scheduling with repeating intervals**

Chronos ♥ REST

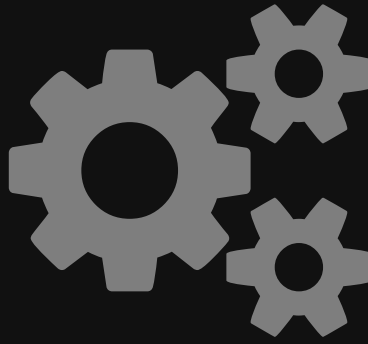
```
PUT chronos-node:8080/scheduler/job/job1
```

```
GET chronos-node:8080/scheduler/jobs
```

```
DELETE chronos-node:8080/scheduler/task/kill/job2
```




Mesoswhat?



Marathon

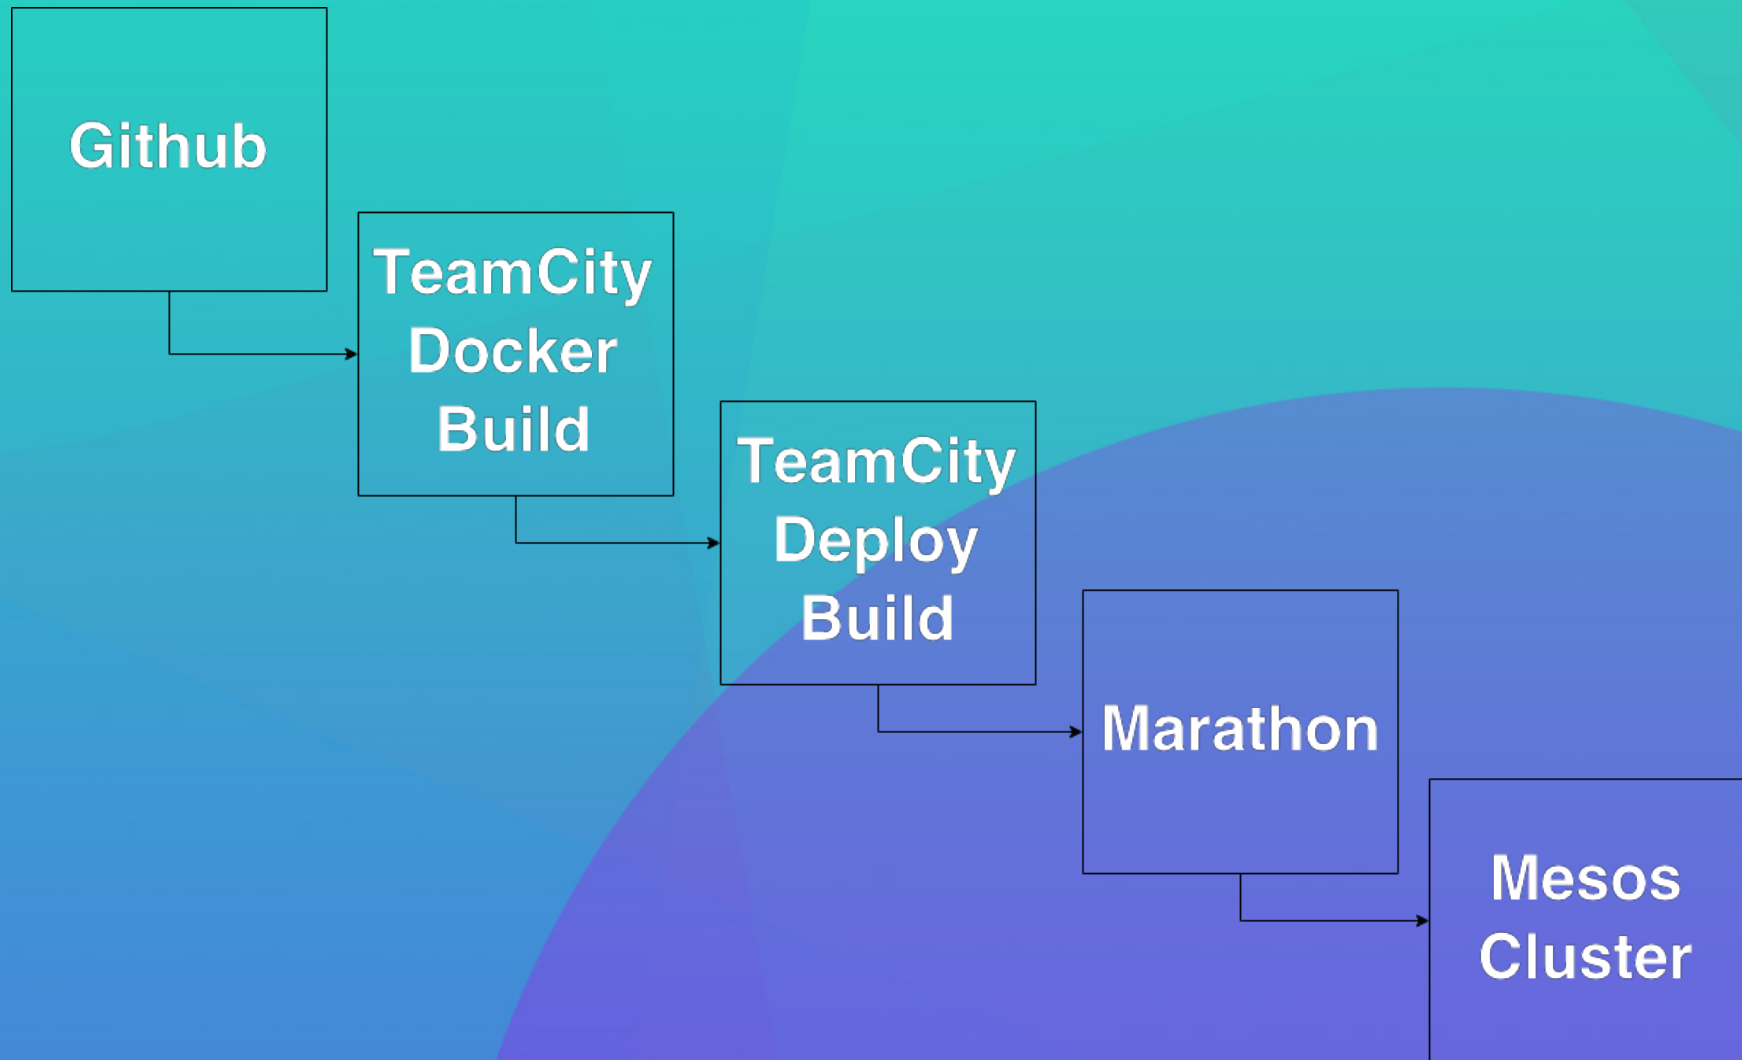


Chronos



Demo!

Continuous Delivery Pipeline



Code Base

This screenshot shows a GitHub repository page for `mesosphere / presentations`. The user is viewing the `ba-infracoders-2015` branch, which is 3 commits ahead of master. The latest commit by BenWhitehead, titled "Add Dockerfile and marathon.json", is shown with a commit hash of `8e1c2e8489`. Below this, a commit history table lists three recent changes:

File	Commit Message	Time
app	Move files into app directory	13 minutes ago
Dockerfile	Add Dockerfile and marathon.json	12 minutes ago
marathon.json	Add Dockerfile and marathon.json	12 minutes ago

The footer of the page includes copyright information for GitHub, Inc. (© 2015), navigation links (Terms, Privacy, Security, Contact), the GitHub logo, and additional links (Status, API, Training, Shop, Blog, About).

Docker Build

Steps

1. `docker login`
2. `docker build`
3. `docker push`
4. `report tag`

Artifacts

1. `docker-tag`
2. `target/marathon.json`

Dockerfile

```
FROM nginx
MAINTAINER Mesosphere support@mesosphere.io

EXPOSE 80

ADD app/ /usr/share/nginx/html
```

marathon.json

```
{
  "id": "/mesosphere/cd-demo-app",
  "instances": 1,
  "cpus": 1,
  "mem": 512,
  "container": {
    "type": "DOCKER",
    "docker": {
      "image": "mesosphere/cd-demo-app:$tag",
      "network": "BRIDGE",
      "portMappings": [
        { "servicePort": 28080, "containerPort": 80, "hostPort": 0, "protocol": "tcp" },
        { "servicePort": 28443, "containerPort": 443, "hostPort": 0, "protocol": "tcp" }
      ]
    }
  },
  "healthChecks": [
    {
      "gracePeriodSeconds": 120,
      "intervalSeconds": 30,
      "maxConsecutiveFailures": 3,
      "path": "/"
    }
  ]
}
```

Report Tag

```
mkdir -p target  
  
echo %TAG% > target/docker-tag  
  
cat marathon.json | \  
jq '.container.docker.image |= "%DOCKER_IMAGE%:%TAG%"' > target/marathon.json
```


Artifacts

Demo > Bay Area Infracoders Presentation > #4 (25 Mar 15 21:39) | ▾

Overview Changes **5** Build Log Parameters Dependencies Artifacts

- docker-tag (63 B)
- marathon.json (884 B)

Total size: 947 B

Deploy

1. `http PUT http://marathon/v2/apps < marathon.json`
2. Send slack message

Slack Message

```
echo '{
  "username"    :"%SLACK_USERNAME%",
  "channel"     :"%SLACK_CHANNEL%",
  "text"        :"%SLACK_MESSAGE%",
  "icon_emoji" :"%SLACK_EMOJI%",
  "mrkdwn"     :%SLACK_MARKDOWN%
}' | http --print=HhBb --json POST %SLACK_WEBHOOK_URL%
```

Build Parameters

Parameters (11) [edit »](#)

Configuration parameters:

Name	Value
MARATHON_HOST	[REDACTED]
MARATHON_JSON_FILE	marathon.json
MARATHON_JSON_TYPE	apps
MARATHON_PORT	8080
SCHEME	http
SLACK_CHANNEL	#demo
SLACK_EMOJI	:teamcity:
SLACK_MARKDOWN	true
SLACK_MESSAGE	Heads Up! %teamcity.build.triggeredBy% just deployed <

Roll Out

Marathon will now begin rolling out¹ the updated version of the application and your team has just been notified of the deployment.

[1] <https://mesosphere.github.io/marathon/docs/deployments.html#rolling-restarts>

Run a Docker container

```
http -v POST http://10.53.20.188:8080/v2/apps @app-ruby.json
```

```
{
  "container": {
    "type": "DOCKER",
    "docker": {
      "image": "superguenter/demo-app"
    }
  },
  "cmd": "rails server -p $PORT",
  "id": "rails-demo",
  "instances": 1,
  "cpus": 0.01,
  "mem": 256,
  "ports": [3000]
}
```

Scale!

```
http -v PUT http://10.53.20.188:8080/v2/apps/demo @scale-app.json
```

```
{  
  "instances": 3  
}
```


Deploy a new Python version

```
http -v PUT http://10.53.20.188:8080/v2/apps/demo @app-python.json
```

```
{
  "container": {
    "type": "DOCKER",
    "docker": {
      "image": "superguenter/demo-app"
    }
  },
  "cmd": "python -m SimpleHTTPServer $PORT",
  "id": "demo",
  "instances": 3,
  "cpus": 0.01,
  "mem": 256,
  "ports": [3000]
}
```

Use HTTP healthchecks

```
http -v POST http://10.53.20.188:8080/v2/apps @app-with-healthcheck.json
```

```
{
  "container": {
    "type": "DOCKER",
    "docker": {
      "image": "superguenter/demo-app"
    }
  },
  "cmd": "rails server -p $PORT",
  "id": "healthcheck-demo",
  "instances": 1,
  "cpus": 0.01,
  "mem": 256,
  "ports": [3000],
  "healthChecks": [
    {
      "path": "/",
      "portIndex": 0,
      "protocol": "HTTP",
      "gracePeriodSeconds": 30,
      "intervalSeconds": 30,
      "timeoutSeconds": 30,
      "maxConsecutiveFailures": 0
    }
  ]
}
```

Add a UNIQUE hostname constraint

```
http -v POST http://10.53.20.188:8080/v2/apps @app-with-constraint.json
```

```
{
  "container": {
    "type": "DOCKER",
    "docker": {
      "image": "superguenter/demo-app"
    }
  },
  "cmd": "python -m SimpleHTTPServer $PORT",
  "id": "constraint-demo",
  "instances": 6,
  "cpus": 0.01,
  "mem": 256,
  "ports": [3000],
  "healthChecks": [
    {
      "path": "/",
      "portIndex": 0,
      "protocol": "HTTP"
    }
  ],
  "constraints": [
    { "hostname": "UNIQUE" }
  ]
}
```



Thanks!

Come and talk to us

P.S., we're hiring!

Get Mesosphere packages: mesosphere.io/downloads

Read about Marathon: github.com/mesosphere/marathon

Read about Chronos: github.com/mesos/chronos

Try out Mesosphere on GCE: google.mesosphere.io

Come work with us: mesosphere.io/jobs