



Is DC/OS a better way to run Docker on AWS?

Chien Huey | DevOps Engineer

xo group/ the knot [the nest] the bump

XO loves Docker

Goals

XO loves Docker

Pick 2

Goals

Goals

XO loves Docker

Pick 2

DC/OS 👍👎

Goals

XO loves Docker

Pick 2

DC/OS 👍👎

The way forward

Disclaimer

About XO Group

theKnot.com

xo group/

© 2016 XO GROUP INC. ALL RIGHTS RESERVED.



About XO Group

theKnot.com

Rails, nodeJS



xo group/

© 2016 XO GROUP INC. ALL RIGHTS RESERVED.

About XO Group

theKnot.com

Rails, nodeJS

~11M visitors/month

xo group/

© 2016 XO GROUP INC. ALL RIGHTS RESERVED.



About XO Group

theKnot.com

Rails, nodeJS

~11M visitors/month

recruiting@xogrp.com



xo group/

© 2016 XO GROUP INC. ALL RIGHTS RESERVED.

Squad

Squad

Squad

Smaller, autonomous squads

xo group/

© 2016 XO GROUP INC. ALL RIGHTS RESERVED.

DevOps 1.0: XO Toolbox

Engineering Productivity

Squad

Squad

Squad

With some DevOps on top

The Good

Docker on Beanstalk

Consistent, repeatable deployment

Docker on Beanstalk

Consistent, repeatable deployment
End to end solution

Docker on Beanstalk

Consistent, repeatable deployment

End to end solution

HA across AZs

Docker on Beanstalk

Consistent, repeatable deployment

End to end solution

HA across AZs

User-friendly interface

The Bad

Docker on Beanstalk

One container per instance

Docker on Beanstalk

One container per instance

Resizing instances not easy

ServiceRole: aws-elasticbeanstalk-service-role
aws:elasticbeanstalk:healthreporting:system:
 AWSEBHealthdGroupId: REDACTED
 SystemType: enhanced
aws:autoscaling:launchconfiguration:
 IamInstanceProfile: aws-elasticbeanstalk-ec2-role
 InstanceType: m3.xlarge
 EC2KeyName: REDACTED
aws:autoscaling:updatepolicy:rollingupdate:
 RollingUpdateType: Health
 RollingUpdateEnabled: **true**
EnvironmentTier:
 Type: Standard
 Name: WebServer
ConfigurationTemplateVersion: **1.1.0.0**

Docker on Beanstalk 🙄

One container per instance

Resizing instances not easy

No cross-region HA

Docker on Beanstalk 🙄

One container per instance

Resizing instances not easy

No cross-region HA

AWS-specific

Docker on Beanstalk 🙄

One container per instance

Resizing instances not easy

No cross-region HA

AWS-specific

Expensive

The Ugly

Underutilization

CPU > 10%

Underutilization

CPU > 10% - 25%

Underutilization

CPU > 10% - 25%

50%

Underutilization

CPU > 10% - 25%

50%

60%

Underutilization

CPU > 10% - 25%

50%

60%

72%

Not cool












EB replacement criteria









User friendly

EB replacement criteria

User friendly

Many containers, one instance





EB replacement criteria

User friendly

Many containers, one instance

Vendor-agnostic

EB replacement criteria

User friendly

Many containers, one instance

Vendor-agnostic

HA across failure zones

Why DC/OS?



Docker and Mesos: Like peanut butter and jelly

You want to run Docker containers, but how do you do so at hyper scale? Apache Mesos may be the answer. Matt Asay explains.

By Matt Asay  | February 28, 2015, 4:23 AM PST

Installing

DC/OS Installation

CF template

DC/OS Installation

CF template

CLI/Advanced installer

```
1 ---
2 agent_list:
3 - 10.88.3.26
4 - 10.88.3.27
5 bootstrap_url: http://10.88.3.28:8080
6 cluster_name: 'dcos-testing'
7 exhibitor_storage_backend: static
8 ip_detect_filename: /genconf/ip-detect
9 log_directory: /genconf/logs
10 master_list:
11 - 10.88.3.29
12 - 10.88.3.30
13 - 10.88.3.31
14 resolvers:
15 - 8.8.4.4
16 - 8.8.8.8
17 ssh_key_path: /genconf/ssh-key
18 ssh_port: '22'
19 ssh_username: core
```

DC/OS Installation

CF template

CLI/Advanced installer

Demo



Dashboard



Services



Nodes




Universe



System


75%

3 of 4 Shares




15%

2 GiB of 11 GiB



0%


Current Failure Rate



Services Health

jenkins	Healthy
swarm	Healthy
marathon	Healthy
kafka	N/A

Tasks



3
Tasks running

0
Tasks staging

Component Health

Admin Router	Healthy
Admin Router Reloader	Healthy
Admin Router Reload...	Healthy
Cluster ID	Healthy
Diagnostics	Healthy

The Good

DC/OS Evaluation

Easy cluster creation

DC/OS Evaluation 🍌

Easy cluster creation

One-button package installs



arangodb

0.3.0

Install Package



cassandra

1.0.5-2.2.5

Install Package



chronos

2.4.0

Install Package



jenkins

0.2.3

Install Package



kafka

1.0.9-0.10.0.0

Install Package



marathon

1.1.1

Install Package



spark

1.0.0-1.6.1-2

Install Package

DC/OS Evaluation 😊

Easy cluster creation

One-button package installs

Vertically scaling beyond instance size

DC/OS Evaluation 😊

Easy cluster creation

One-button package installs

Vertically scaling beyond instance size

Turn key for homogenous workload

The Bad

DC/OS Evaluation 🙄

Private Docker registries

Demo

Renewing private registry authentication

DC/OS Evaluation 🙄

Private Docker registries

Unable to autoscale cluster

DC/OS Evaluation 🙄

Private Docker registries

Unable to autoscale cluster

<https://dcos.io/docs/1.7/usage/faq/>



Apache
MESOS

[Getting Started](#)

[Blog](#)

[Documentation](#)

[Downloads](#)

[Community](#)

If you're new to Mesos

See the [getting started](#) page for more information about downloading, building, and deploying Mesos.

If you'd like to get involved or you're looking for support

See our [community](#) page for more details.

Mesos Observability Metrics

This document describes the observability metrics provided by Mesos master and agent nodes. This document also provides some initial guidance on which metrics you should monitor to detect abnormal situations in your cluster.



```
{"containerizer\/mesos\/container_destroy_errors":0.0,"containerizer\/mesos\/provisioner\/bind\/remove_rootfs_errors":0.0,"containerizer\/mesos\/provisioner\/remove_container_errors":0.0,"slave\/container_launch_errors":0.0,"slave\/cpus_percent":1.1,"slave\/cpus_revocable_percent":0.0,"slave\/cpus_revocable_total":0.0,"slave\/cpus_revocable_used":0.0,"slave\/cpus_total":1.0,"slave\/cpus_used":1.1,"slave\/disk_percent":0.0,"slave\/disk_revocable_percent":0.0,"slave\/disk_revocable_total":0.0,"slave\/disk_revocable_used":0.0,"slave\/disk_total":3628.0,"slave\/disk_used":0.0,"slave\/executor_directory_max_allowed_age_secs":136901.818755422,"slave\/executors_preempted":0.0,"slave\/executors_registering":0.0,"slave\/executors_running":1.0,"slave\/executors_terminated":0.0,"slave\/executors_terminating":0.0,"slave\/frameworks_active":1.0,"slave\/invalid_framework_messages":0.0,"slave\/invalid_status_updates":0.0,"slave\/mem_percent":0.198323004010208,"slave\/mem_revocable_percent":0.0,"slave\/mem_revocable_total":0.0,"slave\/mem_revocable_used":0.0,"slave\/mem_total":2743.0,"slave\/mem_used":544.0,"slave\/recovery_errors":0.0,"slave\/registered":1.0,"slave\/tasks_failed":0.0,"slave\/tasks_finished":0.0,"slave\/tasks_killed":0.0,"slave\/tasks_killing":0.0,"slave\/tasks_lost":0.0,"slave\/tasks_running":1.0,"slave\/tasks_staging":0.0,"slave\/tasks_starting":0.0,"slave\/uptime_secs":149218.787748608,"slave\/valid_framework_messages":0.0,"slave\/valid_status_updates":1.0,"system\/cpus_total":1.0,"system\/load_15min":0.05,"system\/load_1min":0.0,"system\/load_5min":0.02,"system\/mem_free_bytes":1275125760.0,"system\/mem_total_bytes":3950710784.0}
```

```
2 "containerizer\mesos\container_destroy_er
3 "containerizer\mesos\provisioner\bind\r
4 "containerizer\mesos\provisioner\remove_
5 "slave\container_launch_errors":0.0,
6 "slave\cpus_percent":1.1,
7 "slave\cpus_revocable_percent":0.0,
8 "slave\cpus_revocable_total":0.0,
9 "slave\cpus_revocable_used":0.0,
10 "slave\cpus_total":1.0,
11 "slave\cpus_used":1.1,
12 "slave\disk_percent":0.0,
13 "slave\disk_revocable_percent":0.0,
14 "slave\disk_revocable_total":0.0,
15 "slave\disk_revocable_used":0.0,
16 "slave\disk_total":3628.0
```

DC/OS Evaluation 🙄

Private Docker registries

Unable to autoscale cluster

<https://dcos.io/docs/1.7/usage/faq/>

Expensive

Cost analysis

8:1 replacement ratio

Cost analysis

8:1 replacement ratio

Cost of masters & LB 1.5 agents

Cost analysis

8:1 replacement ratio

Cost of masters & LB 1.5 agents

OSS version only

DC/OS Evaluation 😞

Private Docker registries

Unable to autoscale cluster

[http://mesos.apache.org/documentation/
latest/monitoring/](http://mesos.apache.org/documentation/latest/monitoring/)

Expensive

Incomplete

The Verdict

Next steps for XO

Mesos 👍

Next steps for XO

Mesos 👍

Swarm (on Mesos?)

Next steps for XO

Mesos 👍

Swarm (on Mesos?)

Staying aware of DC/OS

References

<http://sched.co/6jtW> - Huawei

<http://sched.co/6lLy> - Allegro

<http://bit.ly/1Xg0T6R> - Tendril marathon
autoscaler



Yours truly

Chien Huey
chuey@xogrp.com

[@chspinning](#)

xogroupinc.com
recruiting@xogrp.com