
Container Orchestrators and DCOS

Timothy Chen

About me:


- Previously Mesosphere Lead Engineer, (ex-Microsoft/VMWare)
- Apache Mesos, Drill PMC
- Help maintain Apache Spark on Mesos
- Co-founder on new Container startup (still stealth)

Outline

- The Need of Container Orchestration
- What is DCOS?
- Breaking down DCOS
- Demo



docker



The Open Container Initiative

is an open governance structure for the express purpose of creating open industry standards around container formats and runtime.



INTERNET

Web Application

Operating System

Hardware

SCALE & HIGH AVAILABILITY

Web App

Web App

Web App

Operating
System

Operating
System

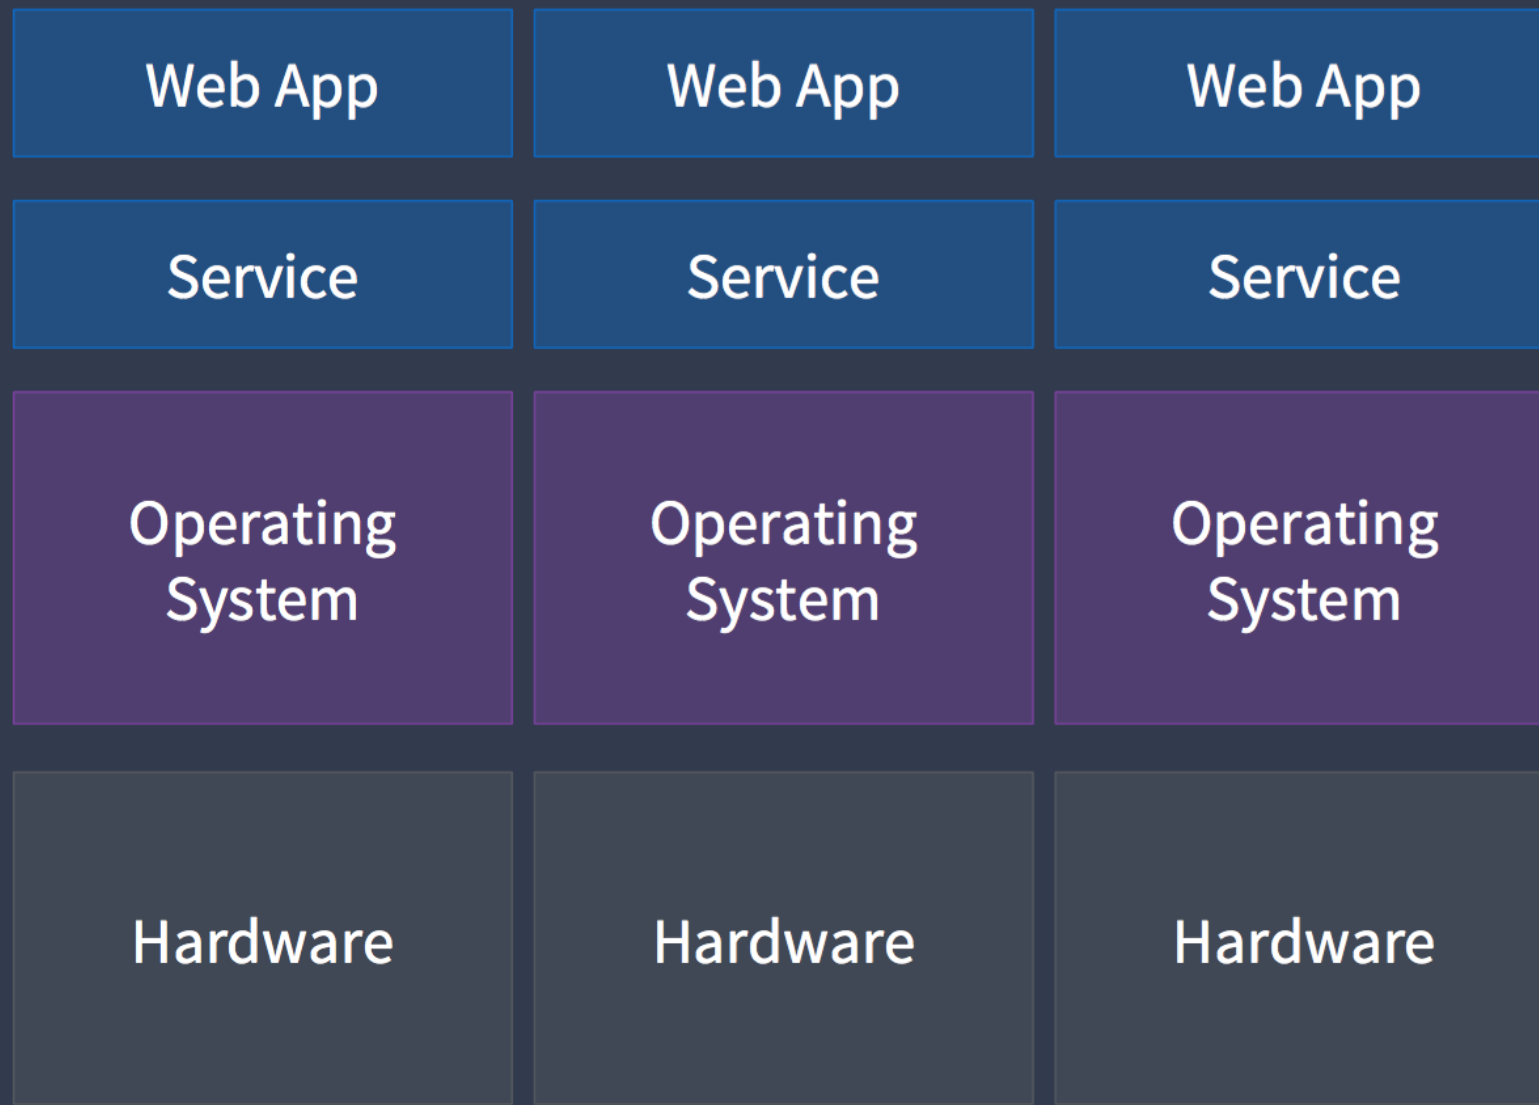
Operating
System

Hardware

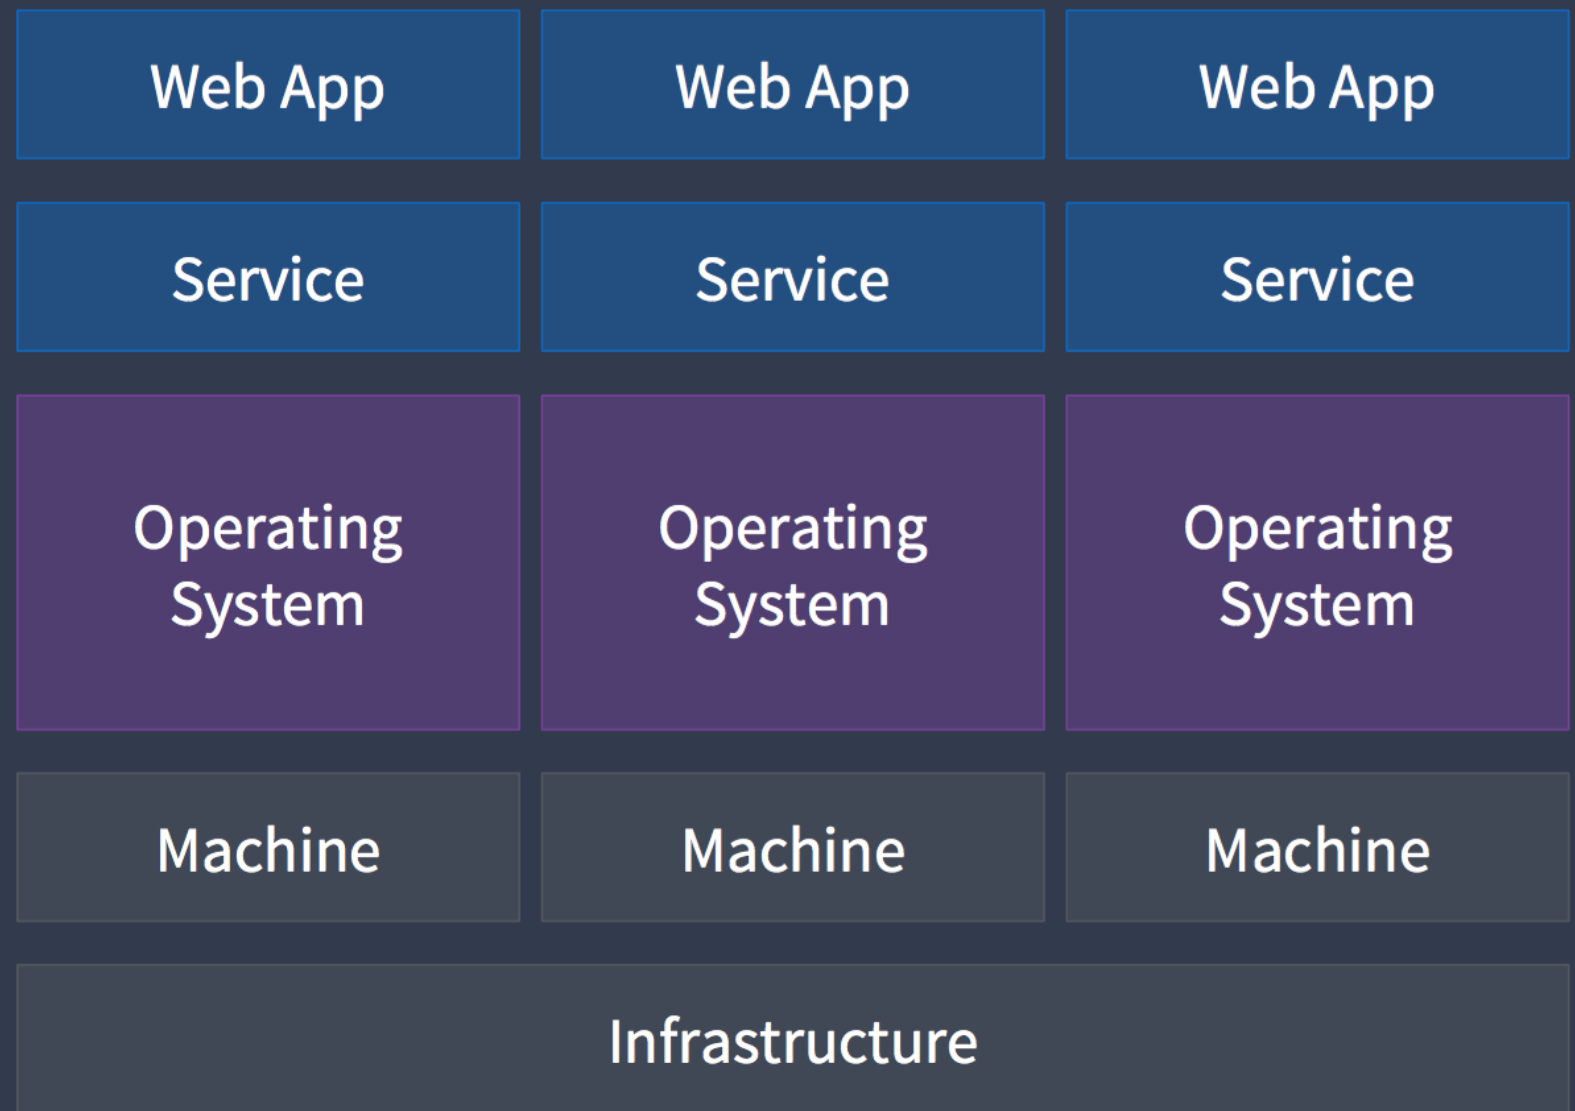
Hardware

Hardware

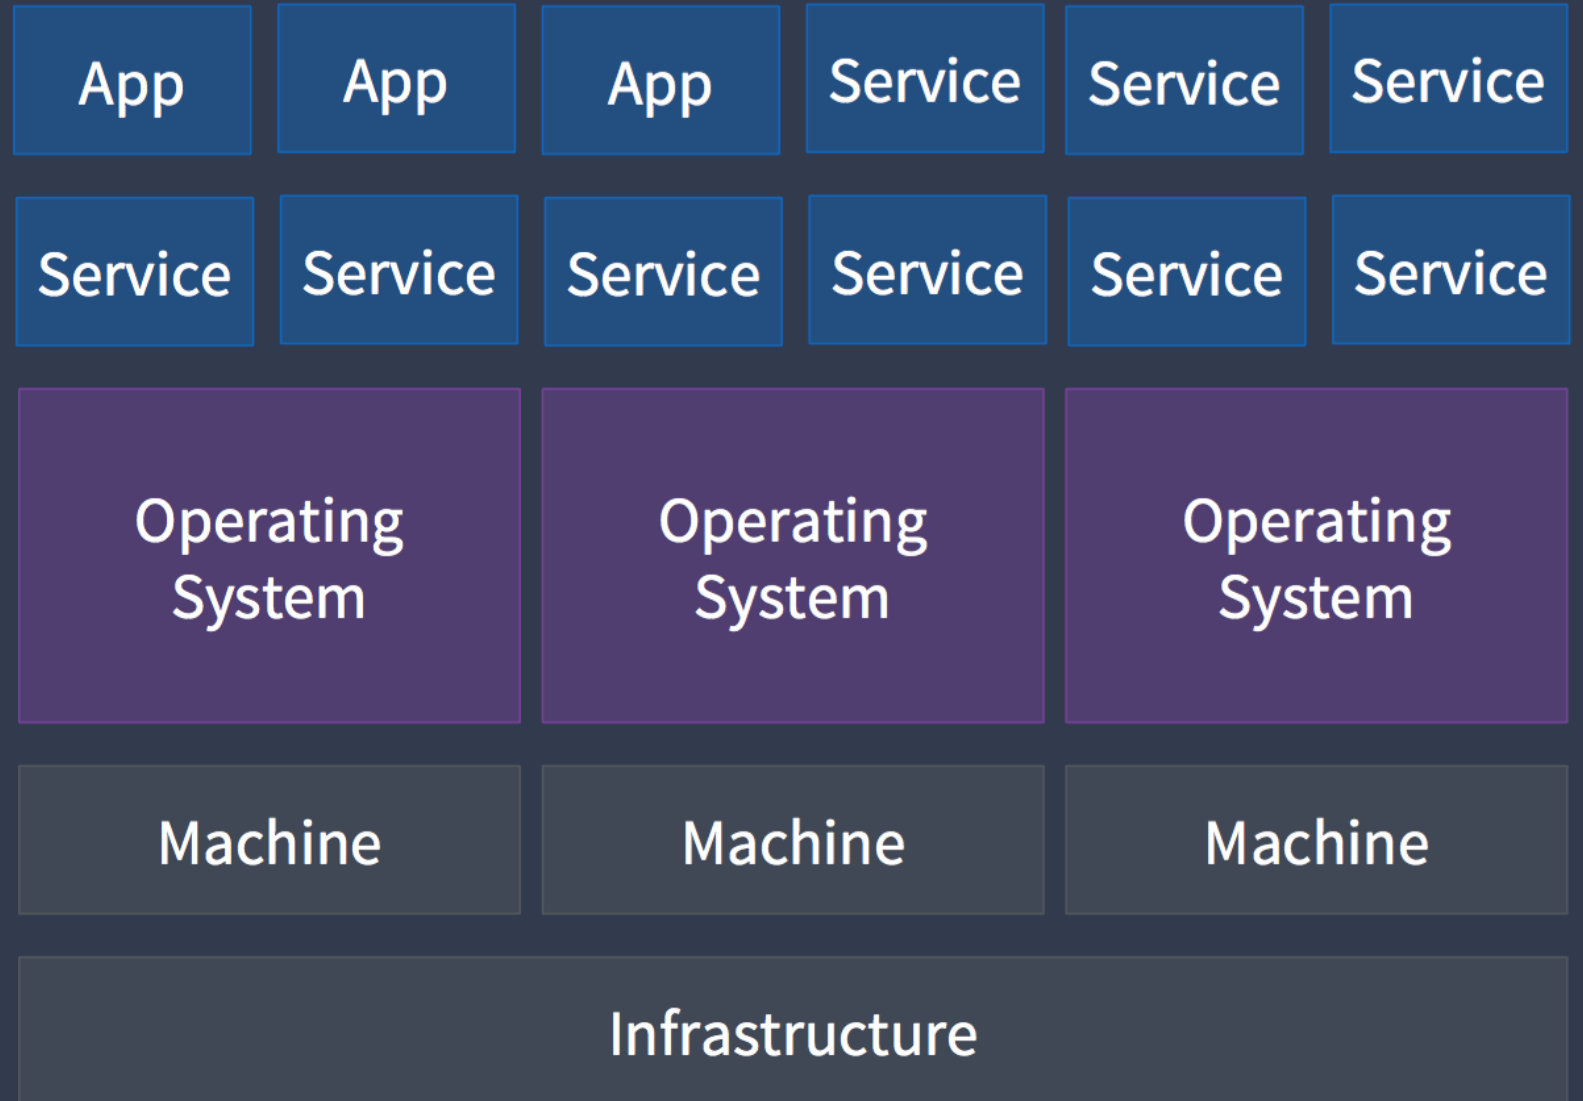
SERVICE-ORIENTED ARCHITECTURE



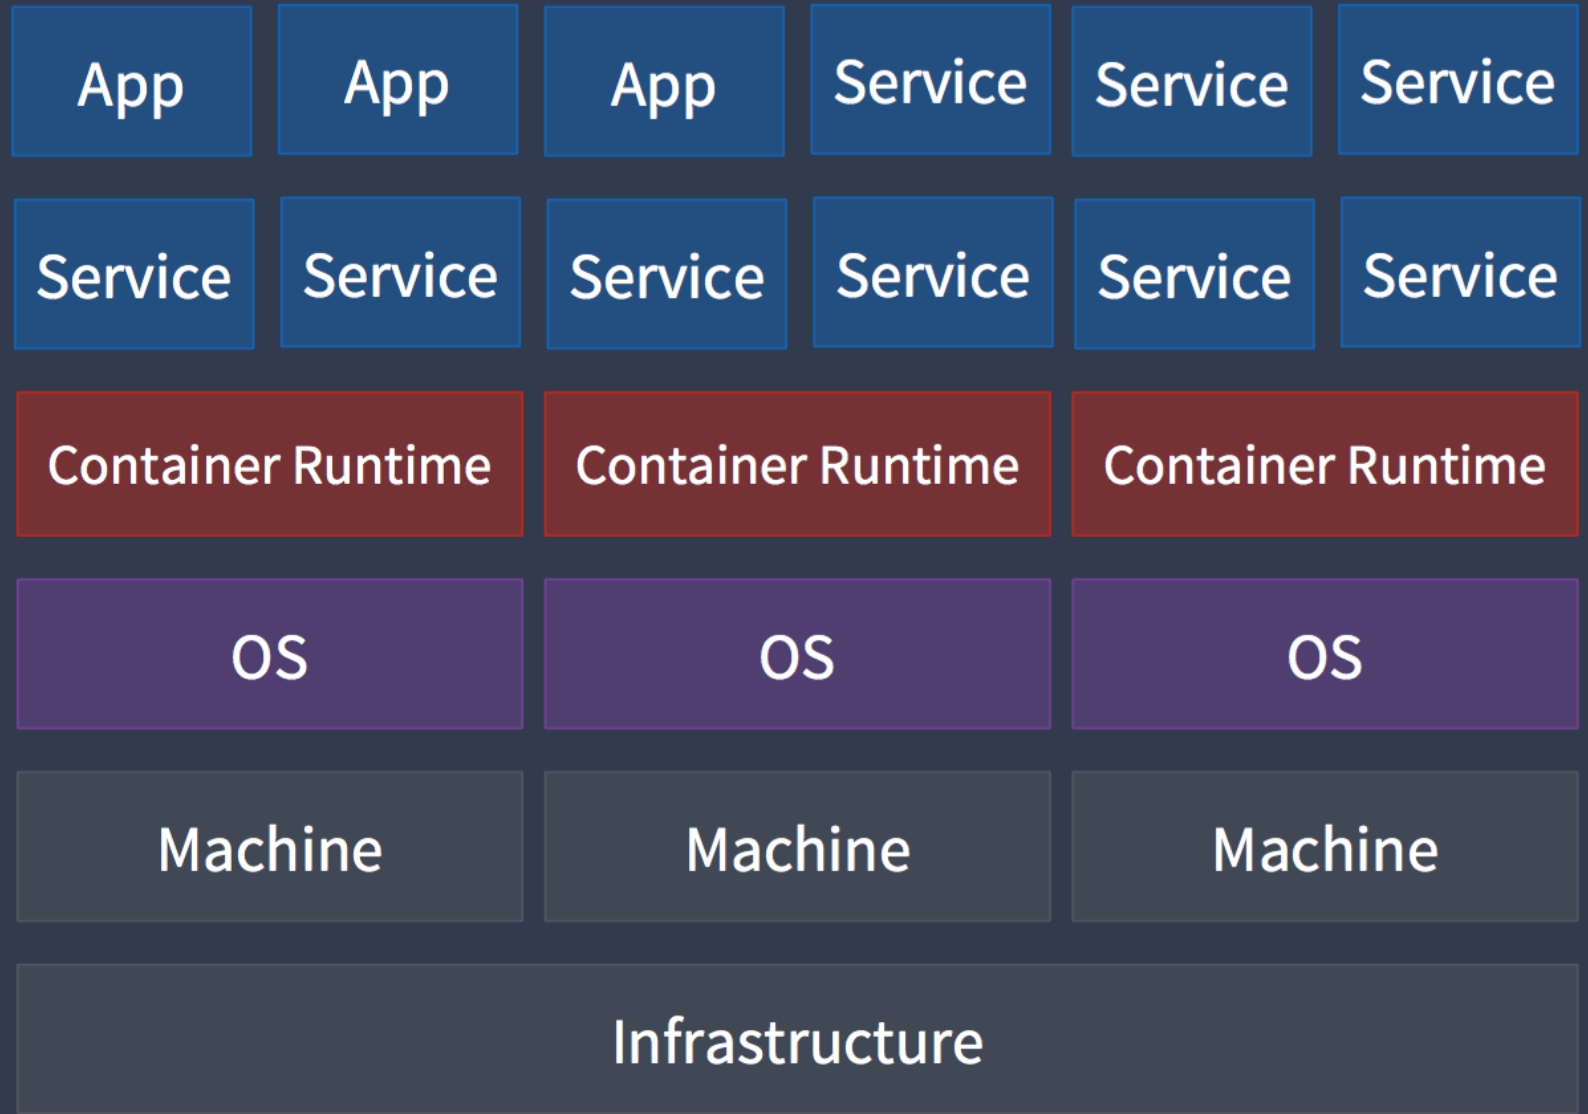
HARDWARE VIRTUALIZATION



MICROSERVICES



CONTAINERIZED MICROSERVICES

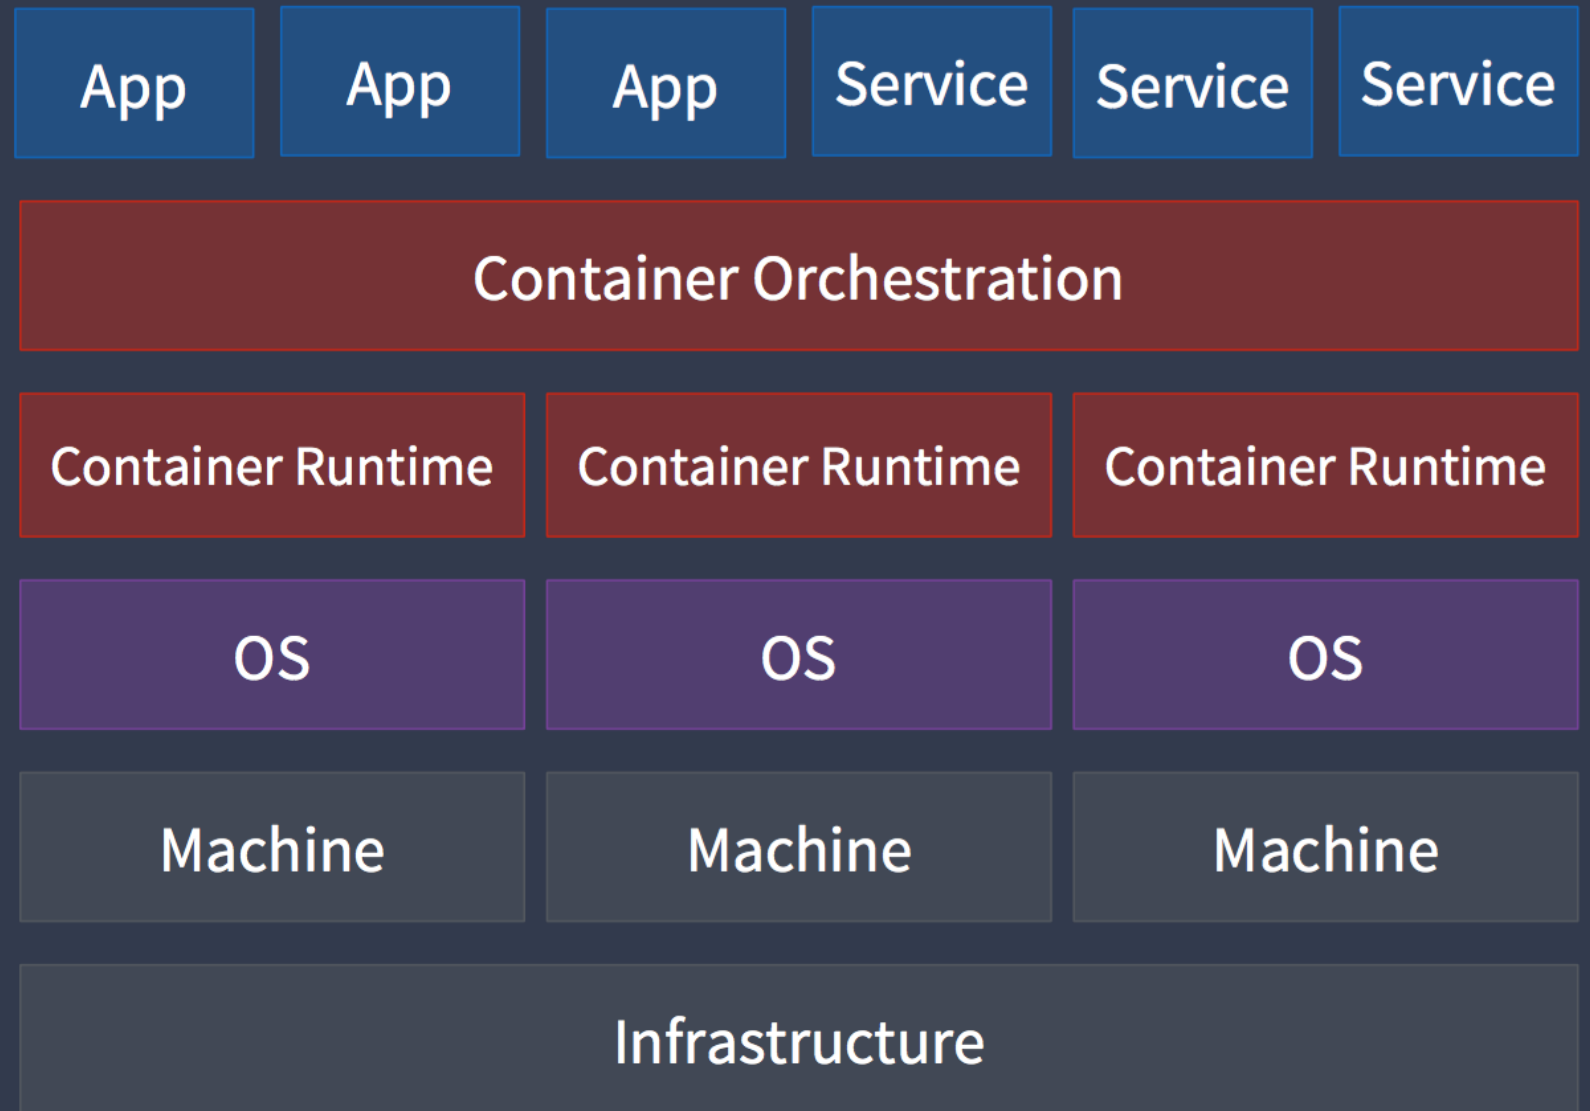




**Time spent on getting app
to run in prod**

**Time spent on
Application Logic**

CONTAINER ORCHESTRATION



Facebook -> Tupperware

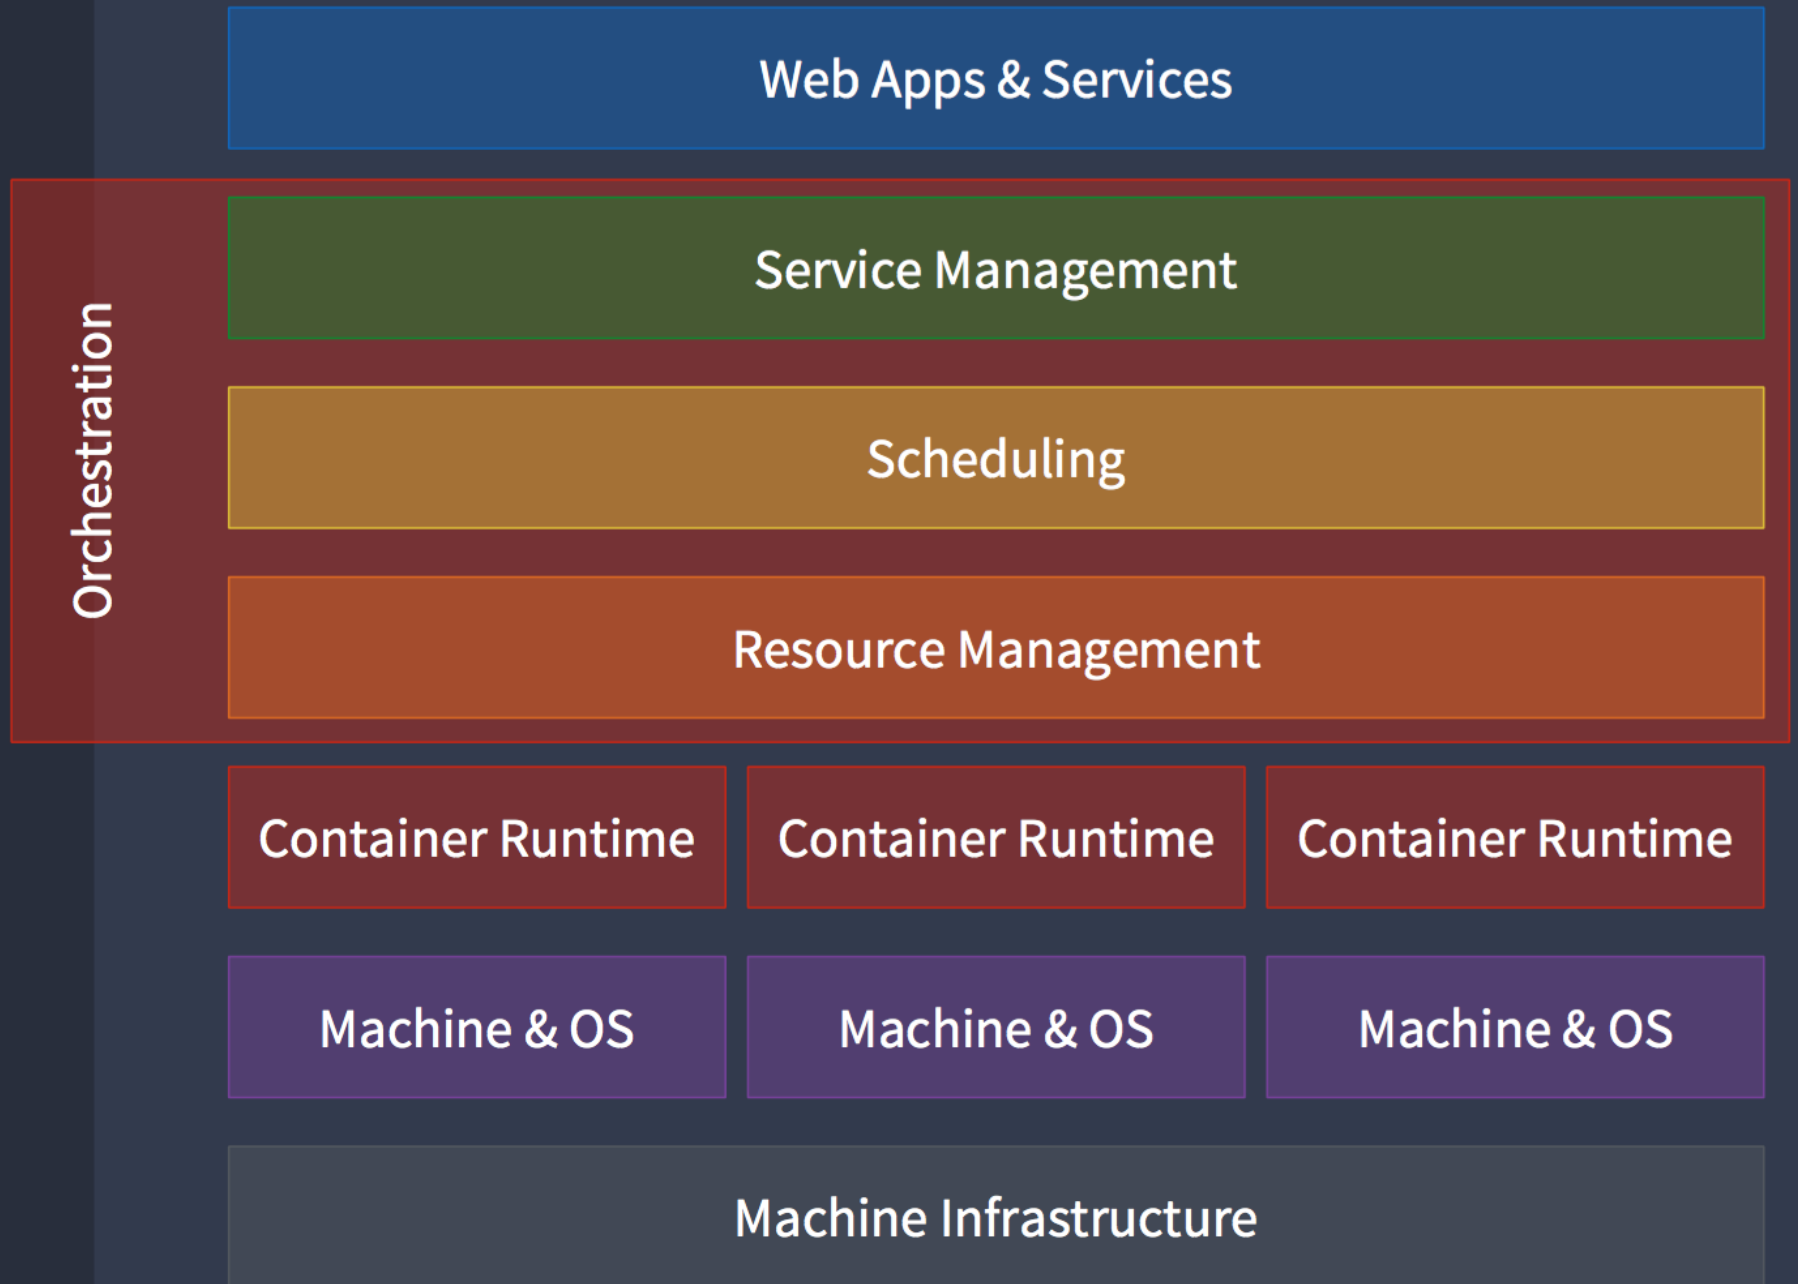
Google -> Borg / Omega

Yahoo -> YARN (Hadoop)

Twitter -> Mesos / Aurora

CONTAINER ORCHESTRATION

Revised



FUNCTIONAL CAPABILITIES

SCHEDULING

- Placement
- Replication/Scaling
- Resurrection
- Rescheduling
- Rolling Deployment
- Upgrades
- Downgrades
- Collocation

RESOURCE MANAGEMENT

- Memory
- CPU
- GPU
- Volumes
- Ports
- IPs

SERVICE MANAGEMENT

- Labels
- Groups/Namespaces
- Dependencies
- Load Balancing
- Readiness Checking

NON-FUNCTIONAL QUALITIES

SCALABILITY

Performance, Responsiveness, Efficiency

AVAILABILITY

Fault Tolerance, Robustness, Reliability, Resilience,
Disaster Recovery

FLEXIBILITY

Format Support, Portability, Interoperability,
Extensibility

USABILITY

Familiarity, Maintainability, Compatibility,
Debuggability

PORTABILITY

Container Runtimes, Host OS, Hosted, Cloud, Bare-
Metal

SECURITY

Auditability, Secrets Management, Encryption,
Isolation

BRIDGING REALITIES: ORCHESTRATION AND PROGRAMMABLE INFRASTRUCTURE

22 Jun 2016 9:10am, by [Benjamin Ball](#) and [Alex Williams](#)



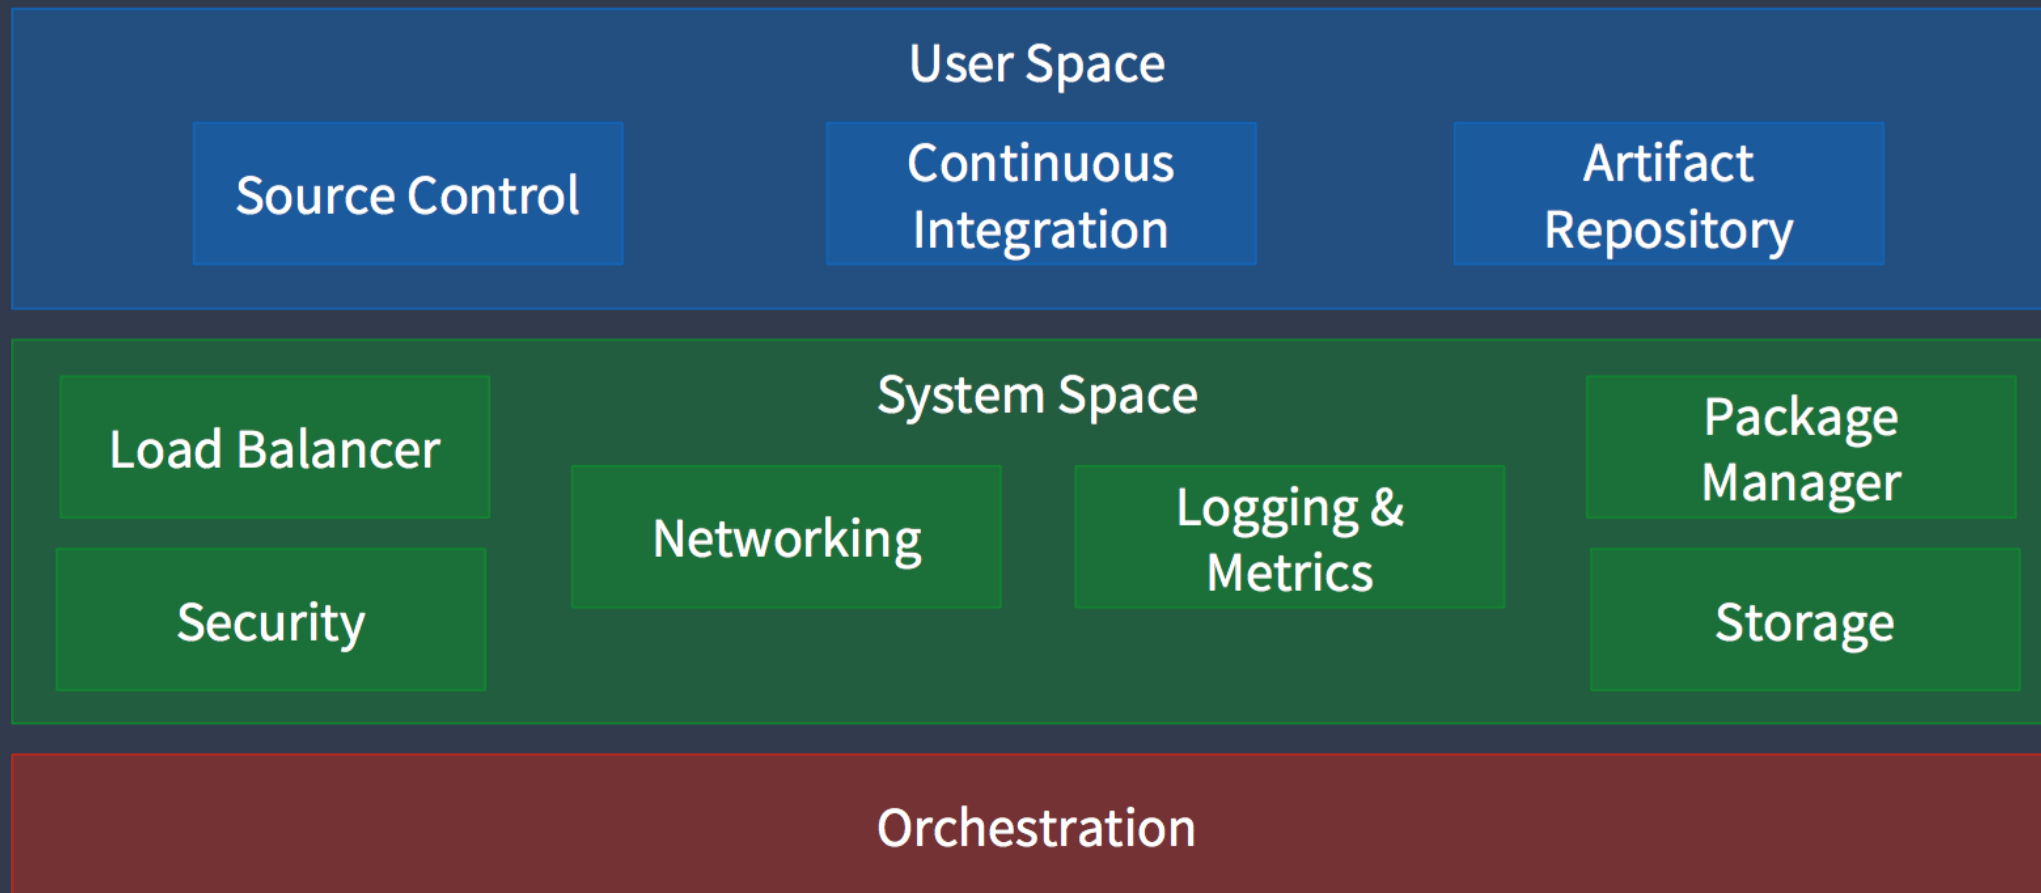
PAAS?

“Platform as a service (PaaS) is a category of cloud computing services that provides a platform allowing customers to develop, run, and manage applications without the complexity of building and maintaining the infrastructure typically associated with developing and launching an app.”

- Wikipedia



CONTAINER OPERATIONS



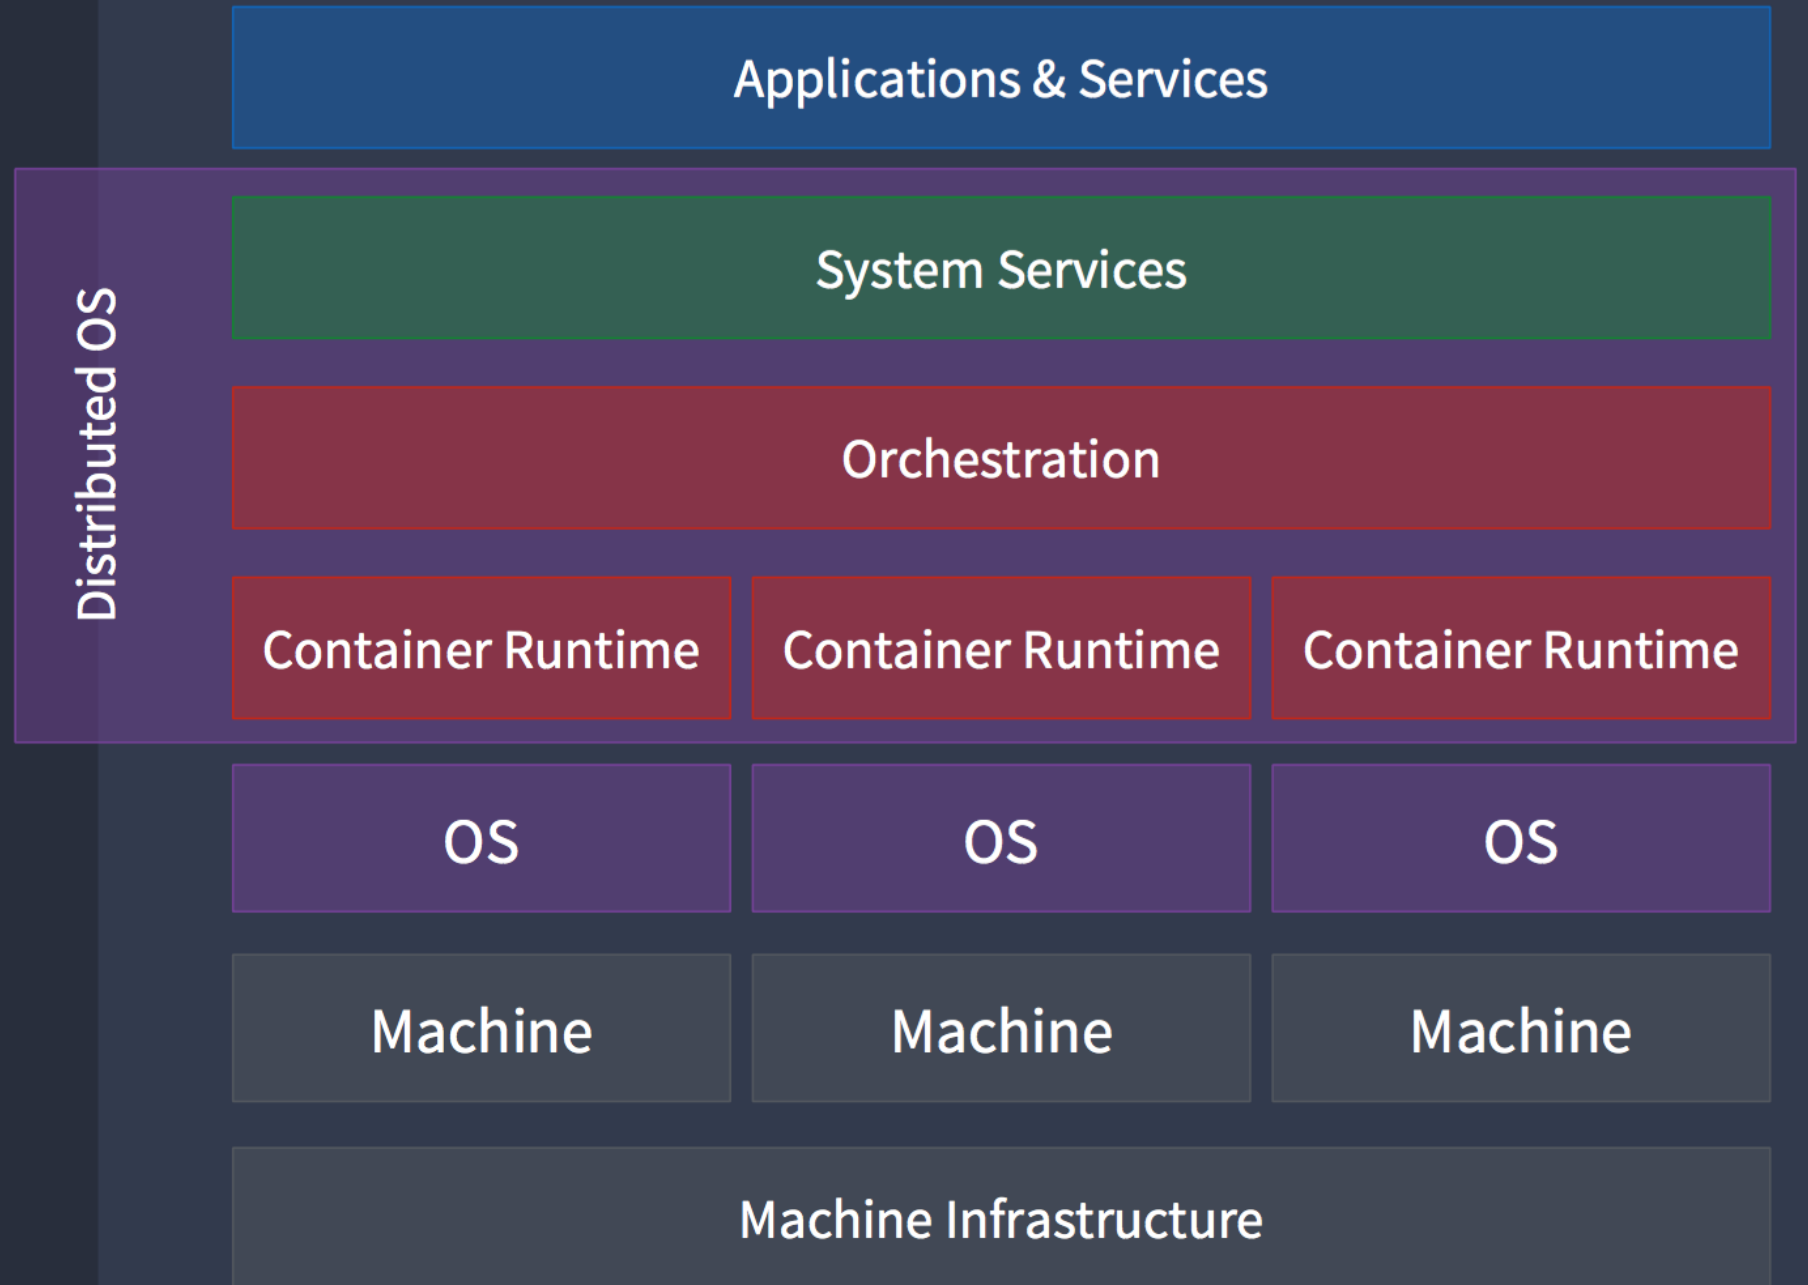
MEET THE
DATACENTER
OPERATING
SYSTEM



DC/OS

DISTRIBUTED OPERATING SYSTEM

Revised



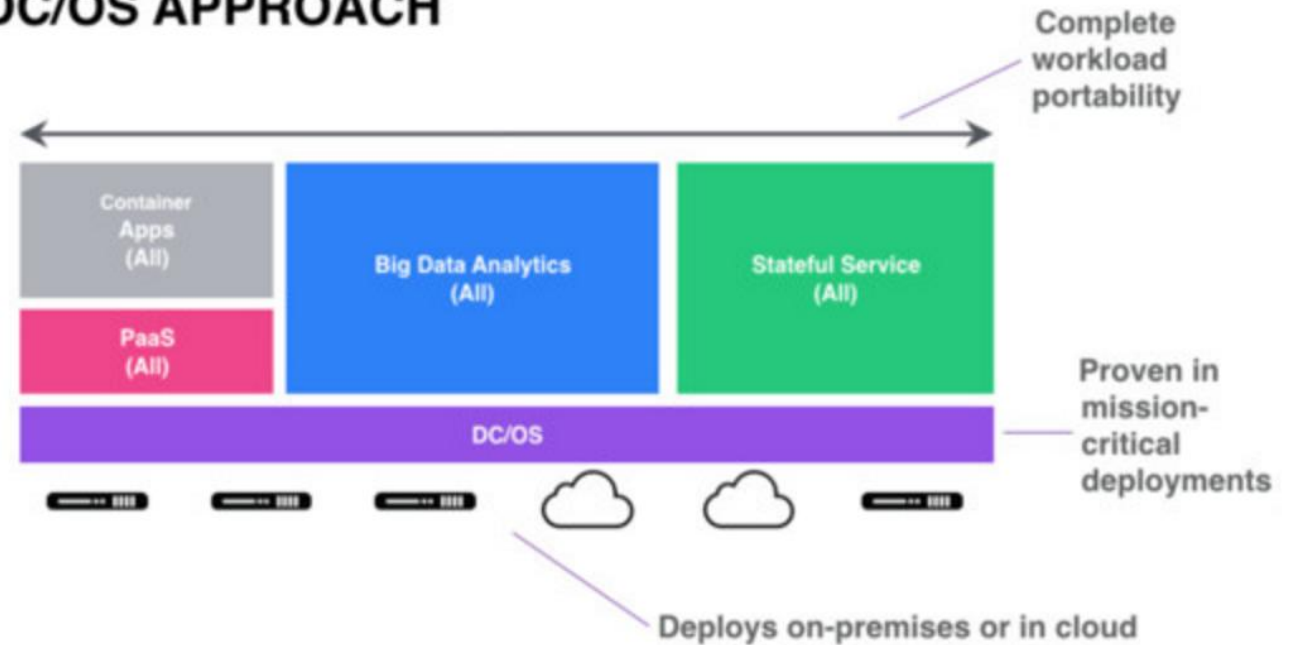
DC/OS IS A DISTRIBUTED OS

TRADITIONAL APPROACH

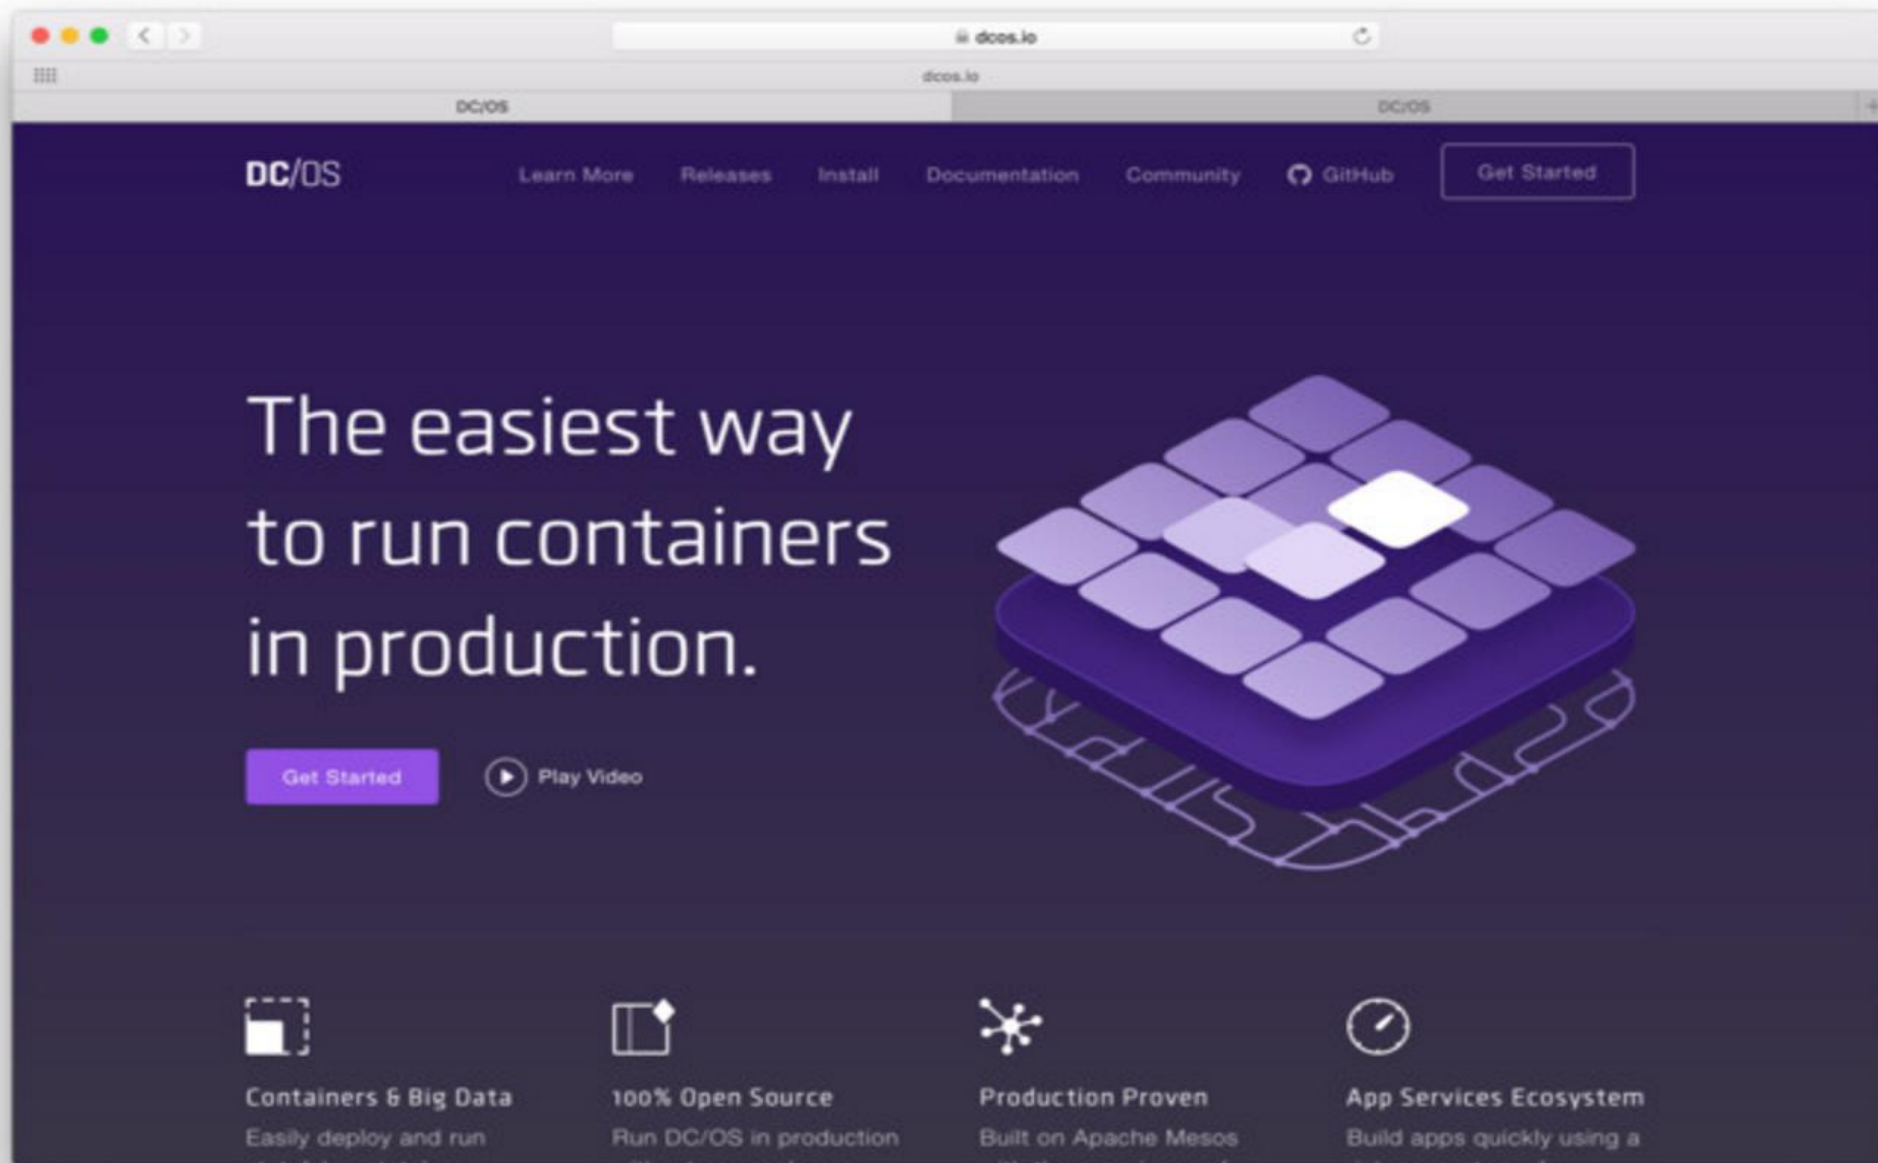


- Many "snowflakes"
- Management nightmare
- Lengthy cycles to deploy code
- Low utilization

DC/OS APPROACH



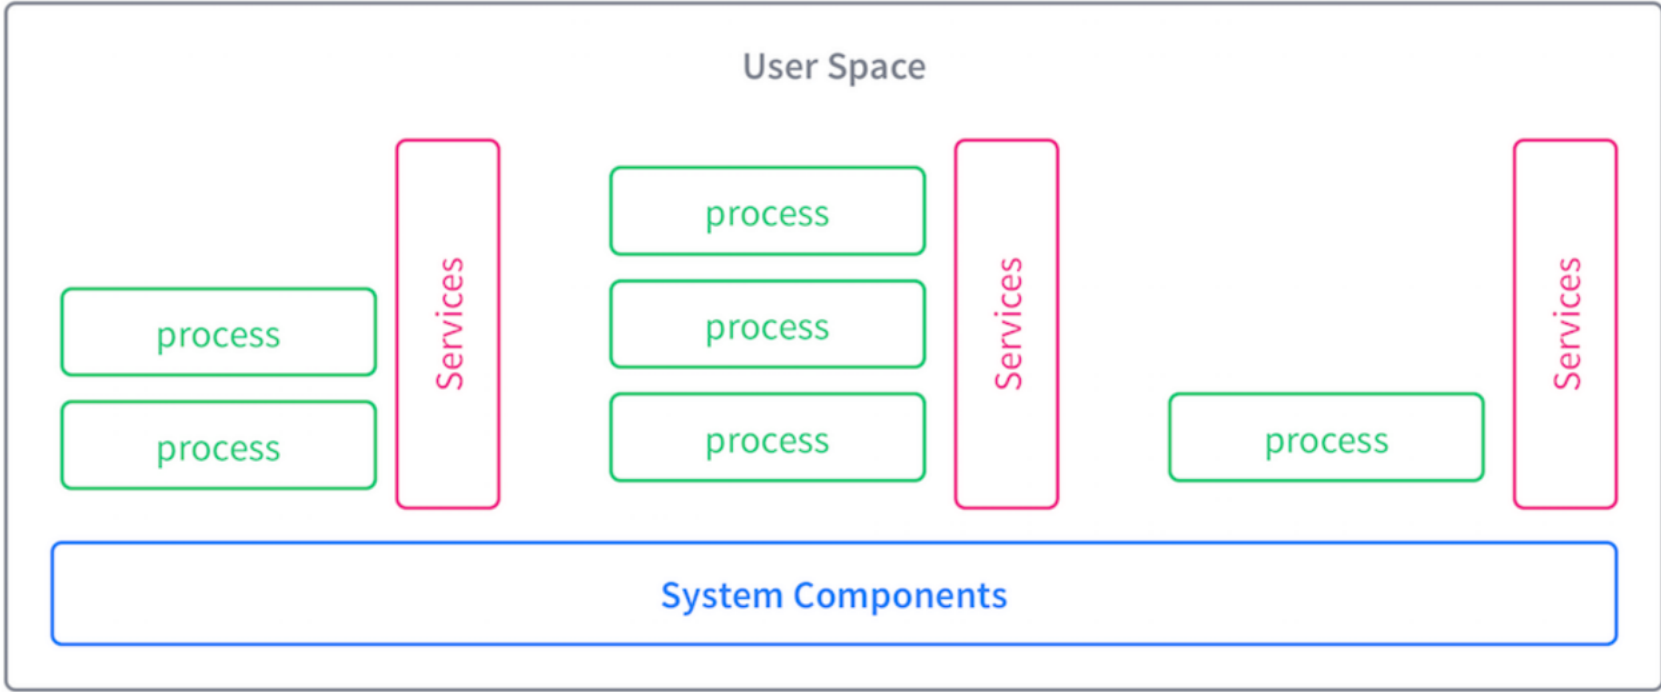
- High performance and resource isolation
- Easy scalability and multi-tenancy
- Fault tolerant and highly available
- Highly efficient with highest utilization



<https://dcos.io>

Why is DCOS awesome?

- Containers! (High resource utilization, etc)
- Extensible Isolation
- Public and Private Service repositories
- Cloud Agnostic Installer
- Web and Command Line Interfaces
- Service Discovery and Load Balancing



APACHE MESOS



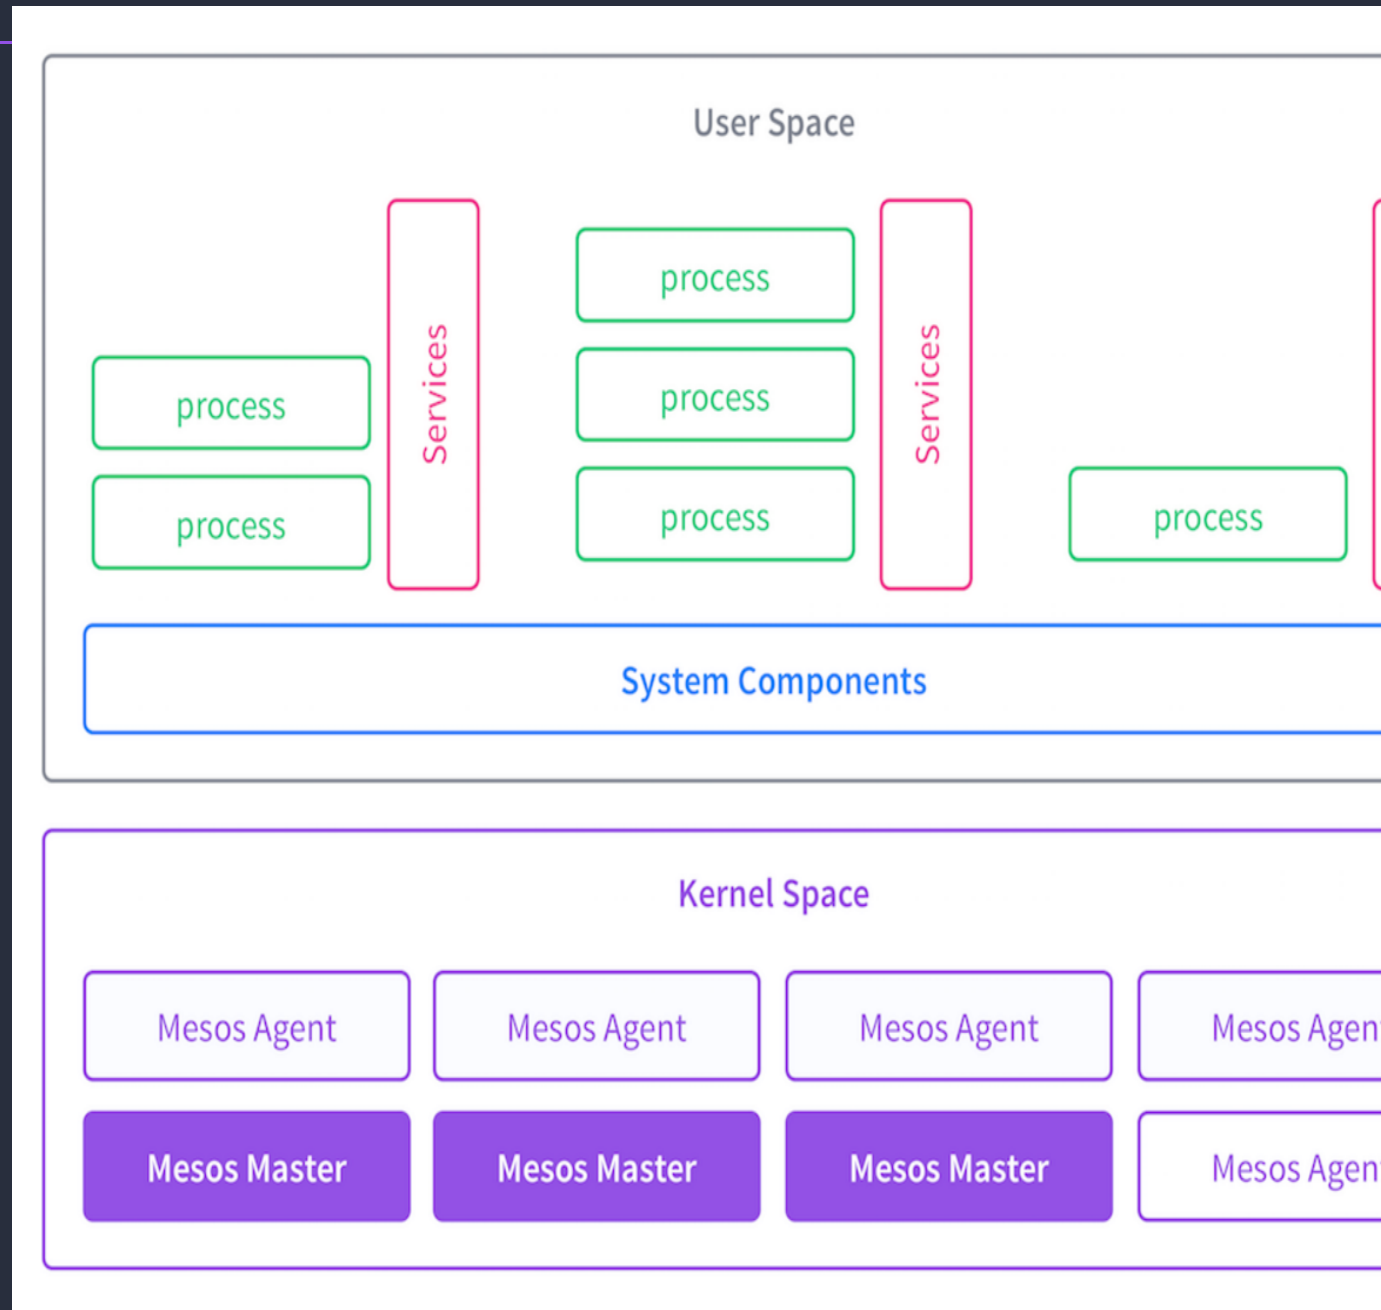
- A top-level ASF project
- A cluster resource negotiator
- Scalable to 10,000s of nodes but also useful for a handful of nodes
- Fault-tolerant, battle-tested
- An SDK for distributed apps
- Native Docker support

mesos.apache.org

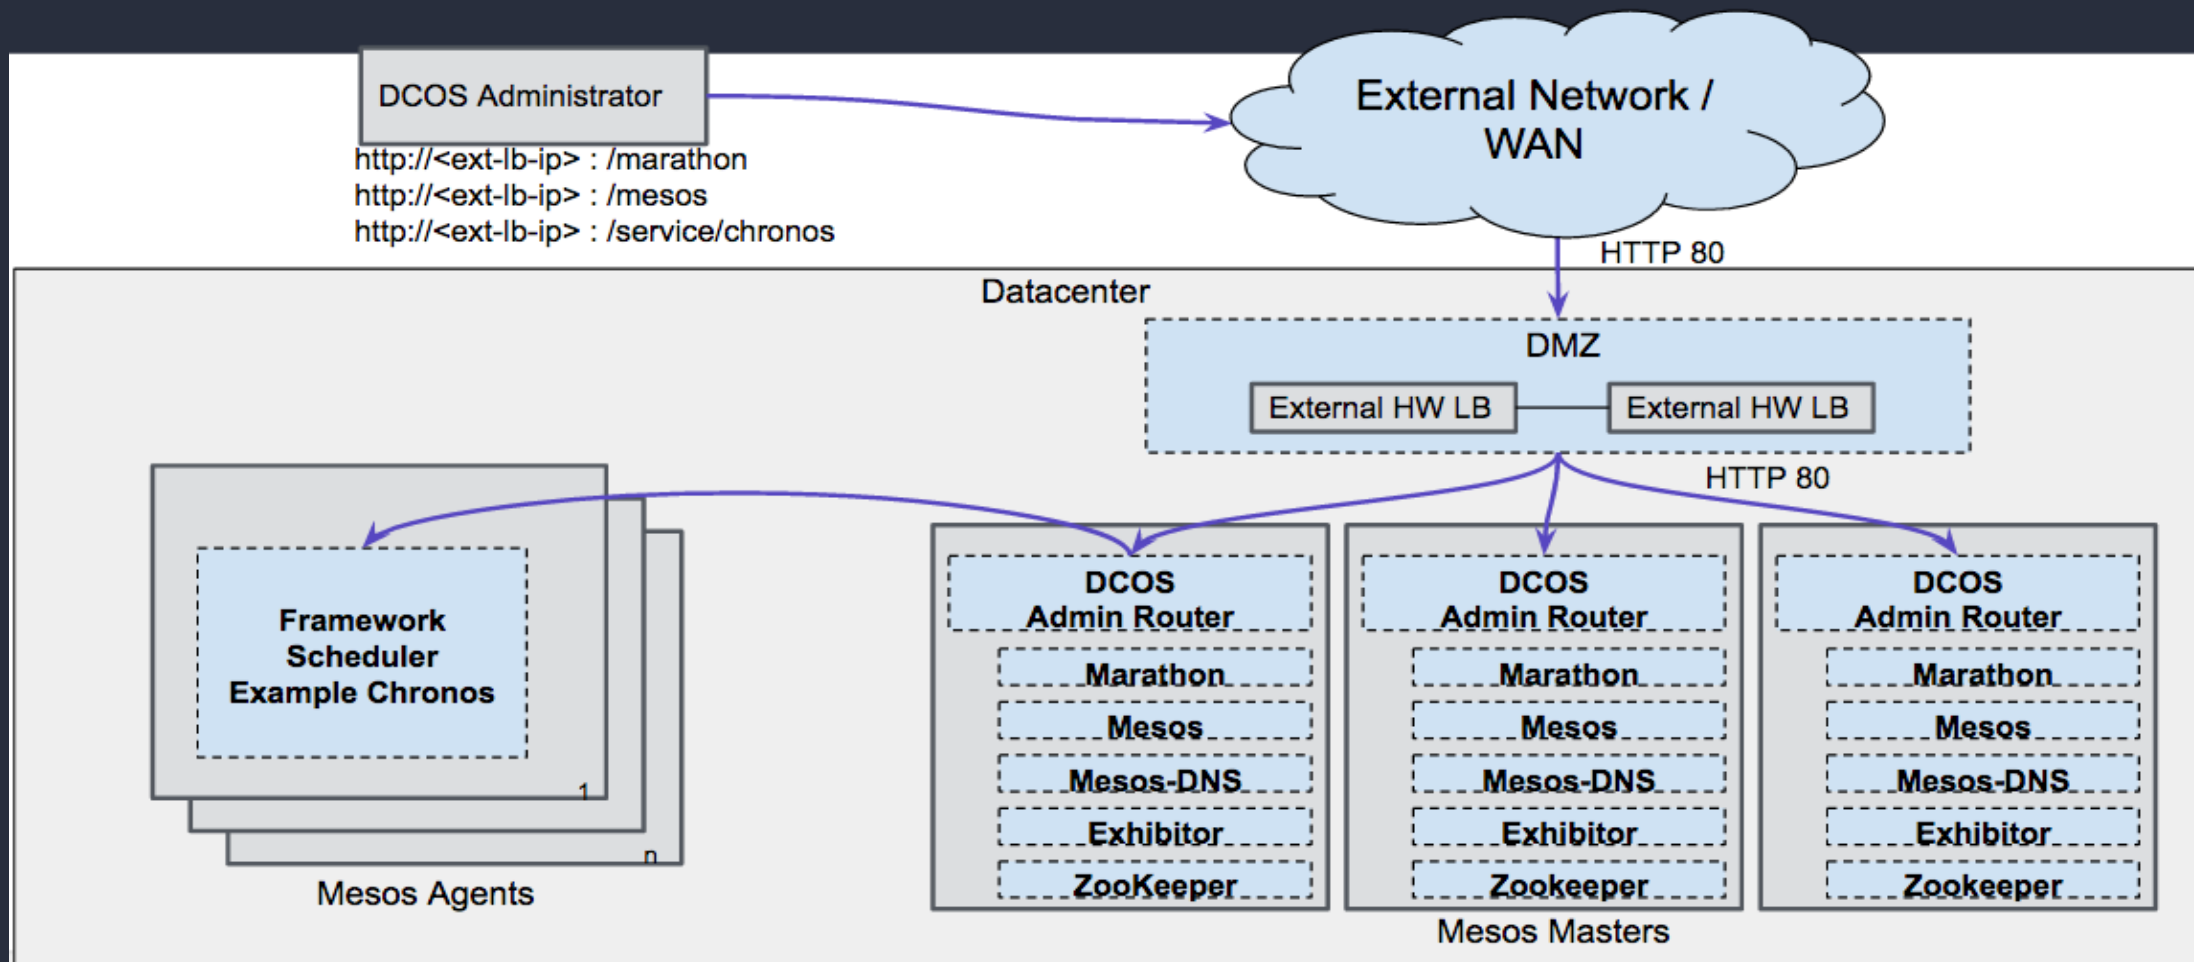


System Components

- Admin Router
- Mesos DNS
- Cosmos Service
- Marathon
- Diagnostic + Signal Service
- Minuteman Service
- DNS Proxy
- History Service
- Auth Service

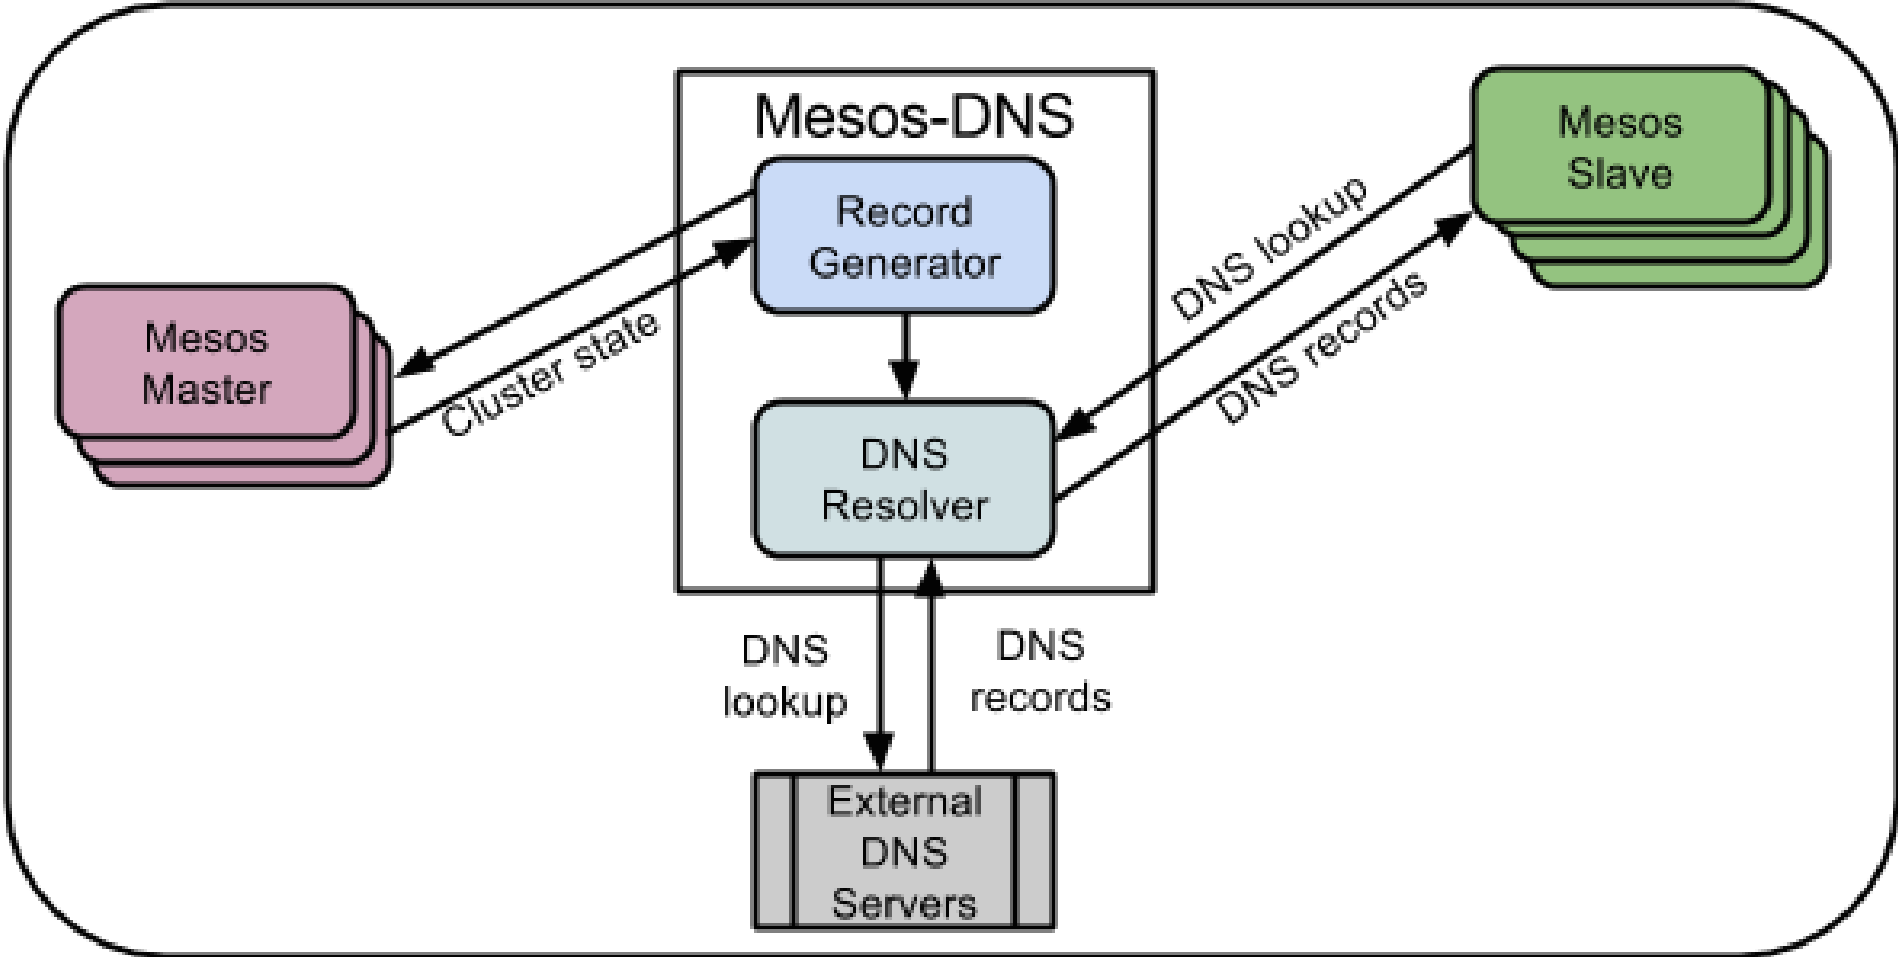


DCOS Admin Router



Component	Open Source Mesos Host : Port	With DCOS & Admin Router via HTTP
Mesos Master Web UI / API	<mesos-master> : 5050	<mesos-master> /mesos/
Mesos Agent UI / API	<mesos-agent> : 5051	<mesos-master> /slave/<slave-id>
Exhibitor UI / API	<mesos-master> : 8181	<mesos-master> /exhibitor/
Marathon UI / API	<mesos-master> : 8080	<mesos-master> /marathon/
Mesos-DNS API	<mesos-master> : 8123	<mesos-master> /mesos_dns/v1/...
Frameworks UI	<framework-scheduler> : xyz	<mesos-master> /service/<service-id>
DCOS History Svc API	N/A	<mesos-master> /dcos-history-service/history/...

Mesos DNS



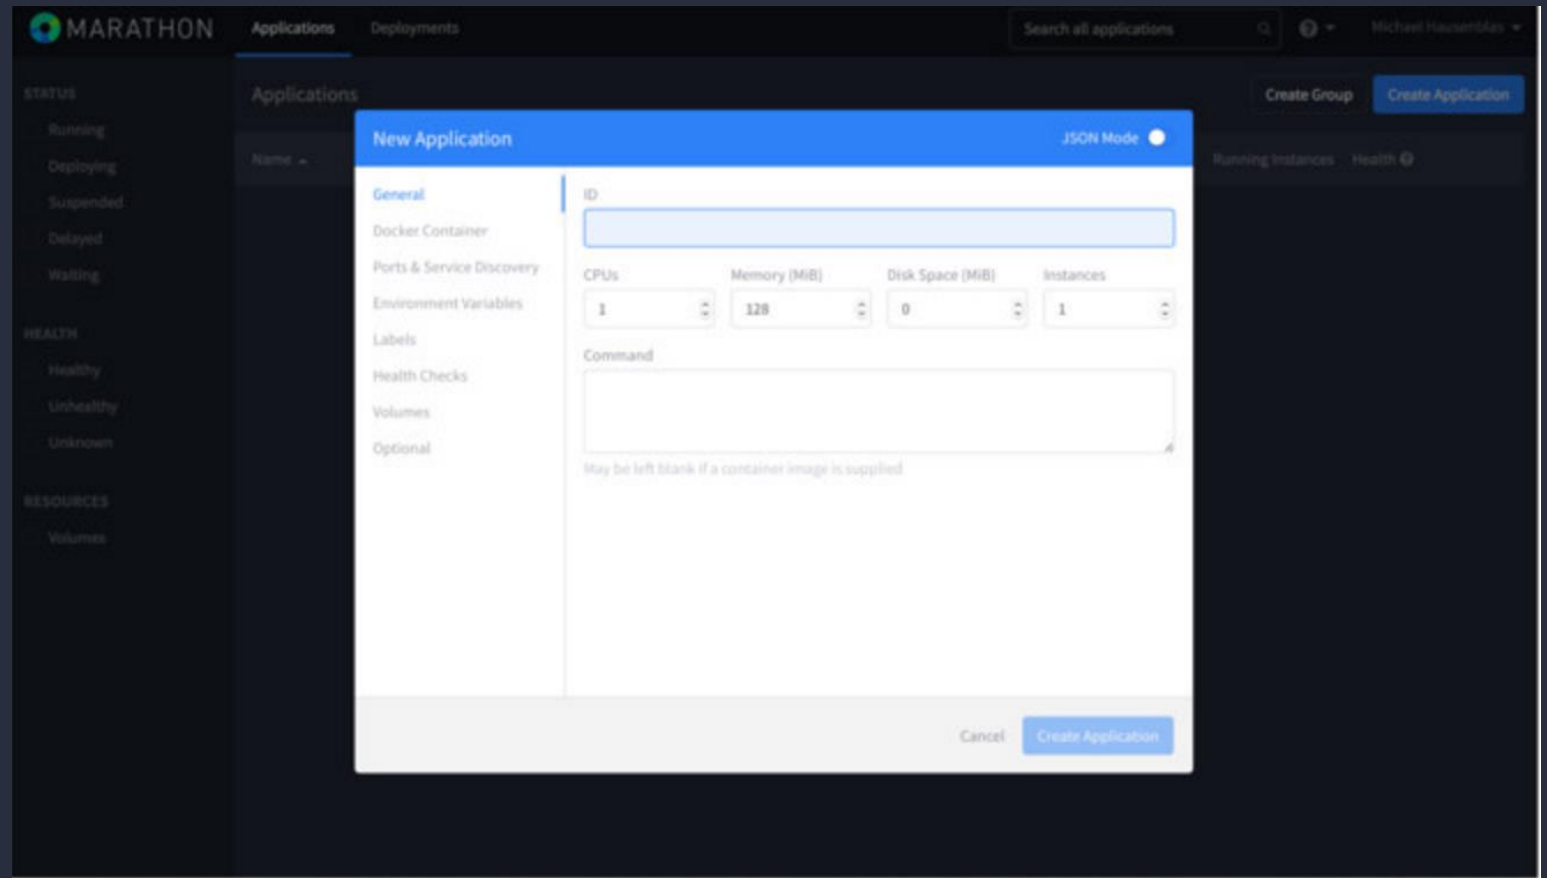
Cosmos

- **Package management server**
- **Handles packaging discovery, serving, install**

<https://github.com/dcos/cosmos>

Marathon

- DCOS Init system
- Deployment
- HA
- Health Checks



The screenshot displays the Marathon web interface. The main header shows 'MARATHON' and navigation tabs for 'Applications' and 'Deployments'. A search bar and user profile 'Michael Hausenblas' are visible. The left sidebar lists application statuses: Running, Deploying, Suspended, Delayed, and Waiting. Below this, the 'HEALTH' section shows Healthy, Unhealthy, and Unknown, and the 'RESOURCES' section shows Volumes. The main content area is titled 'Applications' and features a 'New Application' modal form. The form has a blue header with 'JSON Mode' and a 'Name' field. The 'General' tab is active, showing fields for 'ID', 'CPU's (set to 1), 'Memory (MiB)' (set to 128), 'Disk Space (MiB)' (set to 0), and 'Instances' (set to 1). There are also sections for 'Ports & Service Discovery', 'Environment Variables', 'Labels', 'Health Checks', 'Volumes', and 'Optional'. A 'Command' text area is present with a note: 'May be left blank if a container image is supplied'. At the bottom of the modal are 'Cancel' and 'Create Application' buttons.



MARATHON

Diagnostics / Signal Service

- **3DT monitoring agent + REST API, collects node health and DC/OS specific services health**
- **Signal service forwards telemetry to consumer (SegmentIO, etc)**

<https://github.com/dcos/3dt>

<https://github.com/dcos/dcos-signal>

Minuteman

- **DC/OS Load balancing + Service discovery solution**
- **Distributed L4 load balancer with VIPs**

Install DC/OS

- Local Vagrant
- Cloud Provisioned (AWS, Azure)
- Custom installation (GUI, scriptable)

joel-enterprise
52.39.80.140

Dashboard

Services

Nodes

Network

Universe

System

CPU Allocation

0%
0 of 8 Shares



Memory Allocation

0%
0 B of 27 GiB



Task Failure Rate

0%
Current Failure Rate



Services Health

marathon Idle

View all Services

Tasks

0
Total Tasks

0
Tasks running

0
Tasks staging

Component Health

- Admin Router Healthy
- Admin Router Reloader Healthy
- Admin Router Reloader Timer Healthy
- Cluster ID Healthy
- Diagnostics Healthy

View all 29 Components

Nodes

2
Connected Nodes

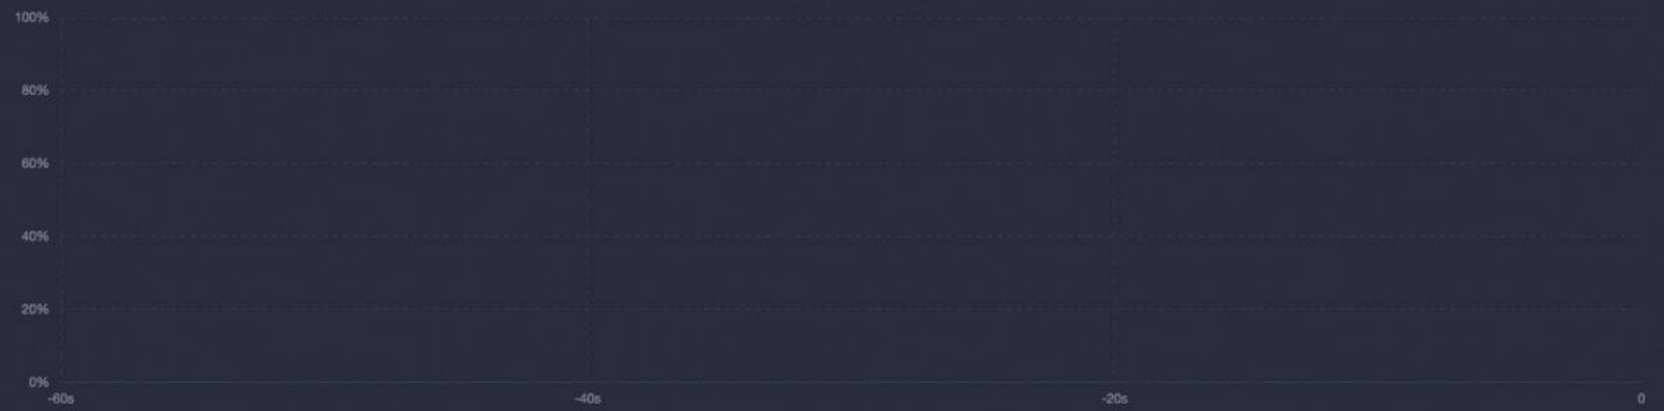


joel-enterprise
52.39.80.140

- Dashboard
- Services**
- Nodes
- Network
- Universe
- System

CPU Memory Disk

CPU Allocation Rate
1 Total Services



1 Services

All 1
Healthy 0
Unhealthy 0

SERVICE NAME ▲	HEALTH	TASKS	CPU	MEM	DISK
marathon	Idle	0	0	0 B	0 B



Enterprise DC/OS v.1.7.0

Dashboard

Services

Nodes

Network

Universe

System

1 Virtual IPs

Filter

VIRTUAL IP ▾

SUCCESSES

1.1.1.1:30000

3425

Back

1.1.1.1:30000

Virtual IP

Backends

Successes and Failures per Minute ▾

3 Total Backends



3 Backends

Filter

BACKEND NAME ▾	SUCCESSES	FAILURES	P99 LATENCY
10.0.2.164:24745	1823	3	2.4ms
10.0.2.167:2293	1180	0	4.11ms
10.0.2.168:4707	1602	0	2.55ms

DC/OS Packages

- A package:
 - Mesos framework (Marathon JSON)
 - CLI
 - Metadata (Description, Icon)
- Universe stores all packages index and versions
- (<http://github.com/mesosphere/universe>)
- DCOS CLI contacts Cosmos to un/install packages

joel-0hbj7g7
52.39.40.226








- Dashboard
- Services
- Nodes
- Network
- Universe**
- System

Enterprise DC/OS v1.7.0

Bootstrap superuser

Packages Installed

Selected Packages

 arangodb 0.3.0 Install Package	 cassandra 1.0.0-2.2.5 Install Package	 chronos 2.4.0 Install Package	 jenkins 0.2.3 Install Package
 kafka 1.0.0-0.9.0.1 Install Package	 marathon 0.15.3 Install Package	 spark 1.6.1-6 Install Package	

Community Packages

Demo!

<http://github.com/mesosphere/tweeter>

Next big question moving forward...

How do we run containers much more efficiently at scale?

Thanks!

@tnachen

tnachen@gmail.com