
BALUG, 19 June 2018



Introduction to DC/OS

Elizabeth K. Joseph
@pleia2



Elizabeth K. Joseph, Developer Advocate

- ❑ Developer Advocate at Mesosphere
- ❑ 15+ years working in open source communities
- ❑ 10+ years in Linux systems administration and engineering roles
- ❑ Founder of OpenSourceInfra.org
- ❑ Author of The Official Ubuntu Book and Common OpenStack Deployments



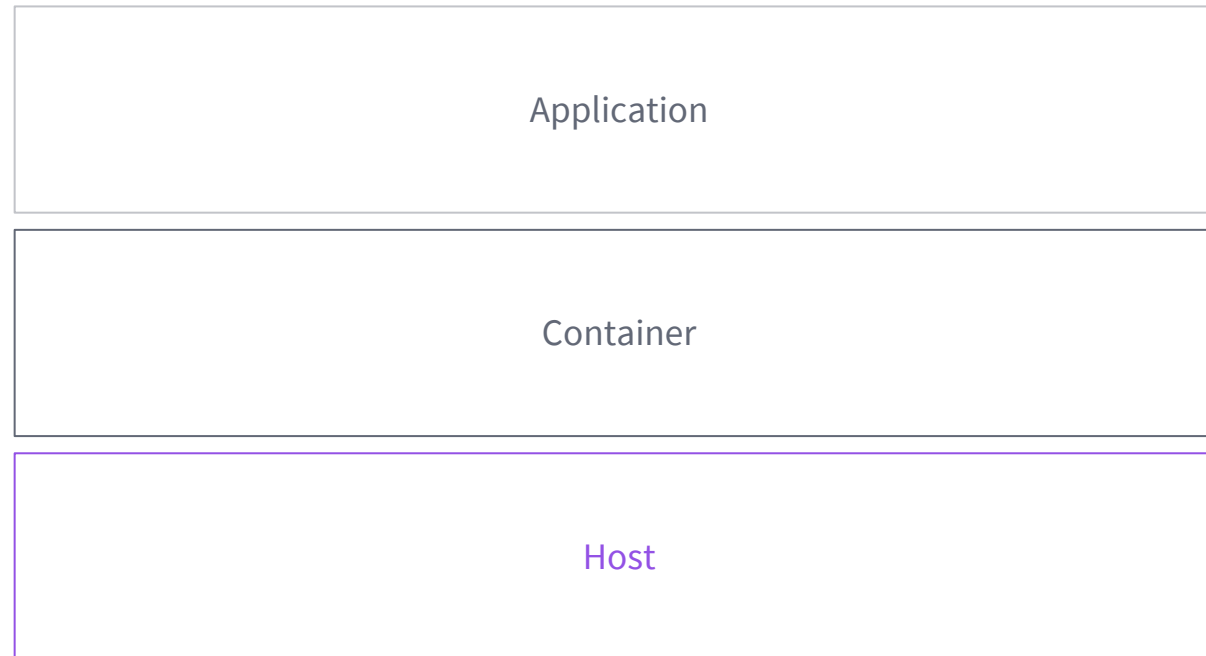
What are Cloud-Native Systems?

You no longer have a single server with everything running on it.

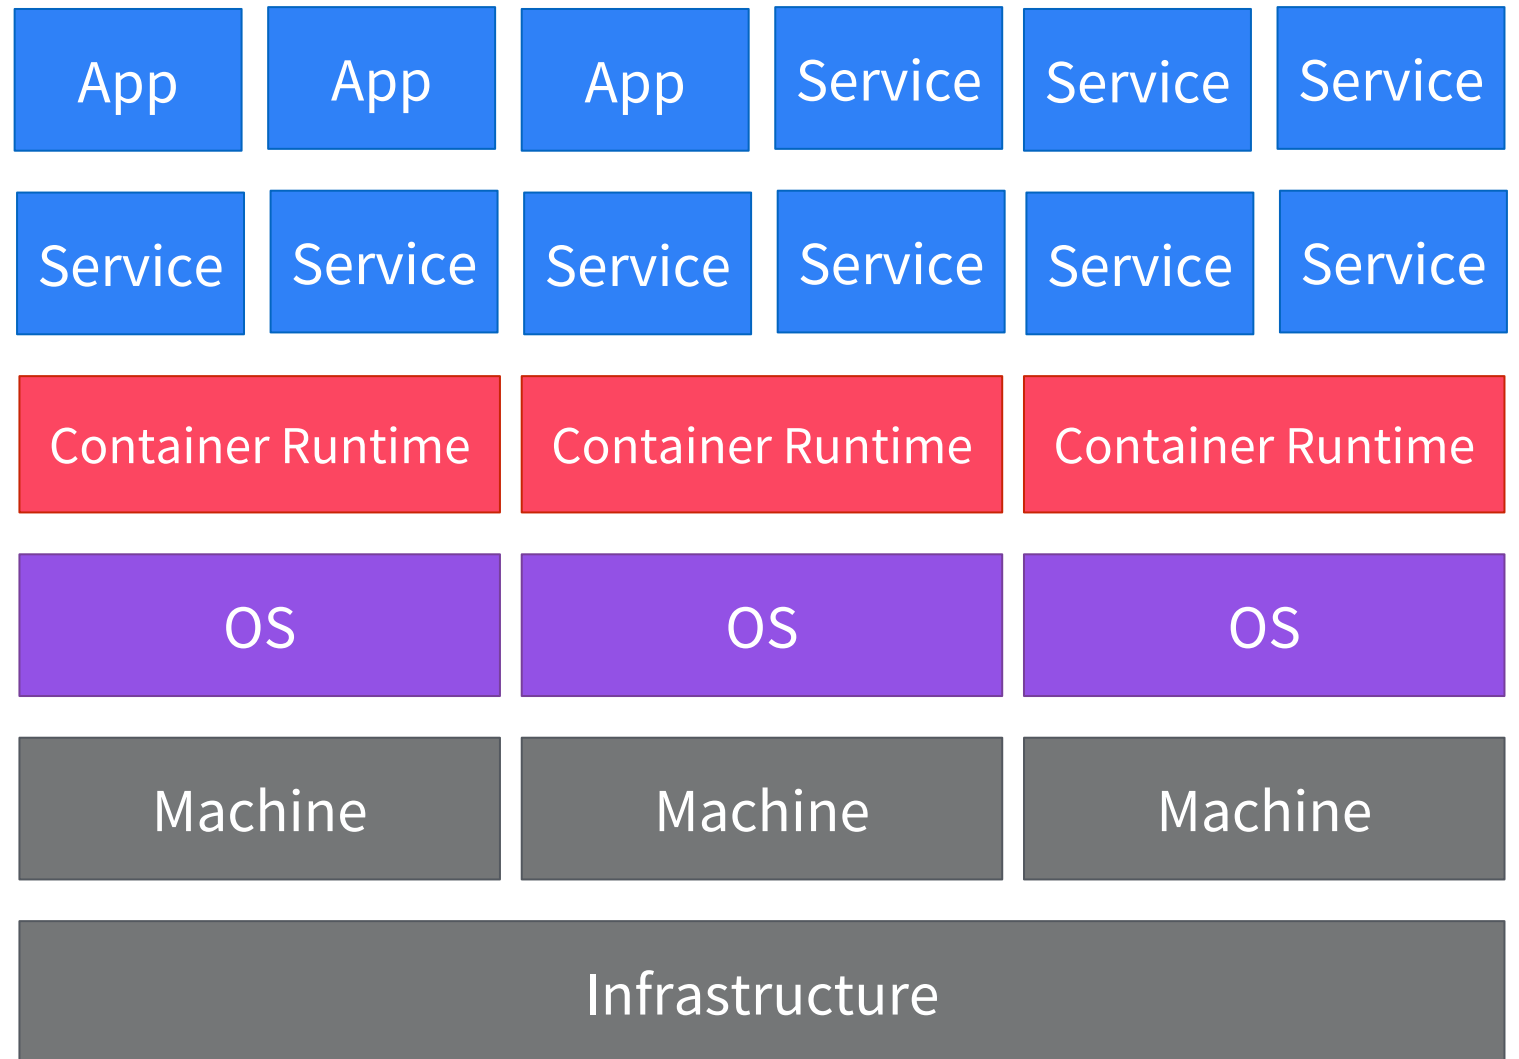
You have a multi-tier system with various layers and owners down the stack:


- ❑ Hardware
- ❑ Network
- ❑ Resource abstraction
- ❑ Scheduler
- ❑ Containers
- ❑ Virtual network
- ❑ Application
- ❑ ...

Cloud-native scopes



Cloud-Native with Containers






Apache Mesos: The datacenter kernel

<http://mesos.apache.org/>

Building block of the modern internet

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



verizon^v



<http://mesos.apache.org/documentation/latest/powered-by-mesos/>

THE BIRTH OF MESOS

Spring 2009



CS262B

Ben Hindman, Andy Konwinski and Matei Zaharia create “Nexus” as their CS262B class project.

TWITTER TECH TALK

The grad students working on Mesos give a tech talk at Twitter.

March 2010



September 2010



MESOS PUBLISHED

Mesos: A Platform for Fine-Grained Resource Sharing in the Data Center is published as a technical report.

December 2010



APACHE INCUBATION


Mesos enters the Apache Incubator.

April 2016




DC/OS

MULTIPLEXING OF DATA, SERVICES, USERS, ENVIRONMENTS



Typical Datacenter
siloed, over-provisioned servers,
low utilization



Apache Mesos
automated schedulers, workload multiplexing onto the
same machines



Master 62dff48e-dfaa-4309-94f0-73d5e94ab01e

Cluster: ejoseph-te4msh6

Leader: 10.0.5.237:5050

Version: 1.4.0

Built: 5 days ago by

Started: 53 minutes ago

Elected: 53 minutes ago

LOG

Agents

Activated 5

Deactivated 0

Unreachable 0

Tasks

Staging 0

Starting 0

Running 11

Unreachable 0

Killing 0

Finished 1

Killed 0


Active Tasks

Find...

Framework ID	Task ID	Task Name	Role	State	Started ▼	Host	
62dff48e-dfaa-4309-94f0-73d5e94ab01e-0001	bus-demo_dashboard.37943816-8677-11e7-b432-425ffc45b8	dashboard.bus-demo	slave_public	RUNNING	a minute ago	10.0.5.101	Sandbox
62dff48e-dfaa-4309-94f0-73d5e94ab01e-0001	bus-demo_ingest.0999da65-8676-11e7-b432-425ffc45b8	ingest.bus-demo	slave_public	RUNNING	9 minutes ago	10.0.1.204	Sandbox
62dff48e-dfaa-4309-94f0-73d5e94ab01e-0004	broker-2__581647a0-6953-4cfe-af96-356d04535c38	broker-2	kafka-role	RUNNING	12 minutes ago	10.0.3.240	Sandbox
62dff48e-dfaa-4309-94f0-73d5e94ab01e-0004	broker-1__d24b1885-860b-4ae9-9feb-502ffcdded5fe	broker-1	kafka-role	RUNNING	13 minutes ago	10.0.3.7	Sandbox
62dff48e-dfaa-4309-94f0-73d5e94ab01e-0004	broker-0__eb077cd0-f416-4918-9cbd-1f5b1ea8c10d	broker-0	kafka-role	RUNNING	13 minutes ago	10.0.1.204	Sandbox
62dff48e-dfaa-4309-94f0-73d5e94ab01e-0001	kafka.8a668774-8675-11e7-b432-425ffc45b8	kafka	slave_public	RUNNING	13 minutes ago	10.0.0.68	Sandbox
62dff48e-dfaa-4309-94f0-73d5e94ab01e-0003	node-2__a9c29921-d7c1-4a32-8eb5-4fd37b25665d	node-2	cassandra-role	RUNNING	14 minutes ago	10.0.3.7	Sandbox 10

Marathon

- Mesos can't run applications on its own.
- A Mesos framework is a distributed system that has a scheduler.
- Schedulers like Marathon start and keep your applications running. A bit like a distributed init system.
- Learn more at <https://mesosphere.github.io/marathon/>



STATUS

- Running 5
- Deploying
- Suspended
- Delayed
- Waiting

HEALTH

- Healthy 4
- Unhealthy
- Unknown 1

LABEL

RESOURCES

- Volumes


Applications

[Create Group](#)
[Create Application](#)

Name ▲	CPU	Memory	Status ?	Running Instances	Health ?
bus-demo	0.2	4 GiB		2 of 2	<div style="width: 100%; height: 10px; background-color: green;"></div> ...
cassandra DCOS_MIGRATION_API_PATH:/v1/plan ...	0.5	2 GiB	Running	1 of 1	<div style="width: 100%; height: 10px; background-color: green;"></div> ...
kafka DCOS_MIGRATION_API_PATH:/v1/plan ...	1.0	1 GiB	Running	1 of 1	<div style="width: 100%; height: 10px; background-color: green;"></div> ...
spark ...	1.0	1 GiB	Running	1 of 1	<div style="width: 100%; height: 10px; background-color: green;"></div> ...

Containers

- Rapid deployment
- Some service isolation
- Dependency handling
- Container image repositories



CONTAINER ORCHESTRATION

CONTAINER SCHEDULING



RESOURCE MANAGEMENT




SERVICE MANAGEMENT



DC/OS brings it all together

- Resource management
- Task scheduling
- Container orchestration
- Logging and metrics
- Network management
- “Universe” catalog of pre-configured apps (including Apache Spark, Apache Kafka...), browse at <http://universe.dcos.io/>
- And much more <https://dcos.io/>



DC / OS

DC/OS is ...

- 100% open source (ASL2.0)
- An umbrella for ~30 OSS projects
 - + Roadmap and designs
 - + Documentation and tutorials
- Not limited in any way
- Familiar, with more features
 - + Networking, Security, CLI, UI, Service Discovery, Load Balancing, Packages, ...

DC/OS Architecture Overview


Services & Containers



DC/OS



ANY INFRASTRUCTURE



Interact with DC/OS (1/2)

Web-based GUI

<https://dcos.io/docs/latest/usage/webinterface/>

The screenshot displays the DC/OS Services web interface. The left sidebar shows the navigation menu for 'dcos-cluster-2' with the 'Services' option highlighted. The main content area shows a table of services with columns for NAME, STATUS, CPU, MEM, and DISK. The services listed are cassandra, kafka, marathon-lb, and zeppelin, all in a 'Running' state.

NAME	STATUS	CPU	MEM	DISK
cassandra	Running (4 Instances)	2	8 GiB	0 B
kafka	Running (4 Instances)	4	4.8 GiB	0 B
marathon-lb	Running (1 Instance)	2	1 GiB	0 B
zeppelin	Running (1 Instance)	1	2 GiB	0 B

Interact with DC/OS (2/2)

CLI tool

<https://dcos.io/docs/latest/usage/cli/>

API

<https://dcos.io/docs/latest/api/>

Catalog of Applications (Universe)

ejoseph-te4msh6
ejoseph@mesosphere.io

- Dashboard
- Services
- Jobs
- Catalog**

RESOURCES

- Nodes
- Networking












SYSTEM

- Overview
- Components
- Settings
- Organization


Catalog

Certified

Certified packages are verified by Mesosphere for interoperability with DC/OS.

 arangodb3 3.2.x CERTIFIED	 artifactory 5.1.4 CERTIFIED	 cassandra 1.0.25-3.0.10 CERTIFIED	 chronos 2.5.0 CERTIFIED
 confluent-kafka 1.1.19.1-3.2.2 CERTIFIED	 dcos-enterprise-cli 1.2.0 CERTIFIED	 elastic 1.0.8-5.2.2 CERTIFIED	 gitlab 1.0-9.1.0 CERTIFIED
			

Install an Application

 **gitlab**
1.0-9.1.0

service

routing

email

high-availability

single-node

enterprise

service

GitLab service properties

NAME * ?

CPUS * ?

MEM * ?

ROLE ?

HOST-VOLUME * ?

HOST-SHARED-VOLUME * ?

CANCEL

REVIEW AND DEPLOY

Application JSON


```
{
  "service": {
    "name": "gitlab",
    "cpus": 1,
    "mem": 2048,
    "role": "*",
    "host-volume": "/srv/gitlab",
    "host-shared-volume": "/srv/gitlab-data"
  },
  "routing": {
    "https-redirect": false,
    "ssh-port": 22222,
    "registry-port": 50000
  },
  "email": {
    "enabled": false,
    "port": 25,
    "authentication": "login",
    "enable-starttls-auto": true,
    "openssl-verify-mode": "peer",
    "tls": false
  },
  "high-availability": {
    "enabled": false,
    "postgres": {},
    "redis": {}
  },
  "single-node": {
    "local-volumes": {},
    "external-volumes": {
      "enabled": false
    }
  },
  "enterprise": {
    "enterprise-edition": false
  }
}
```

1,1

All

Building a Real-World Pipeline

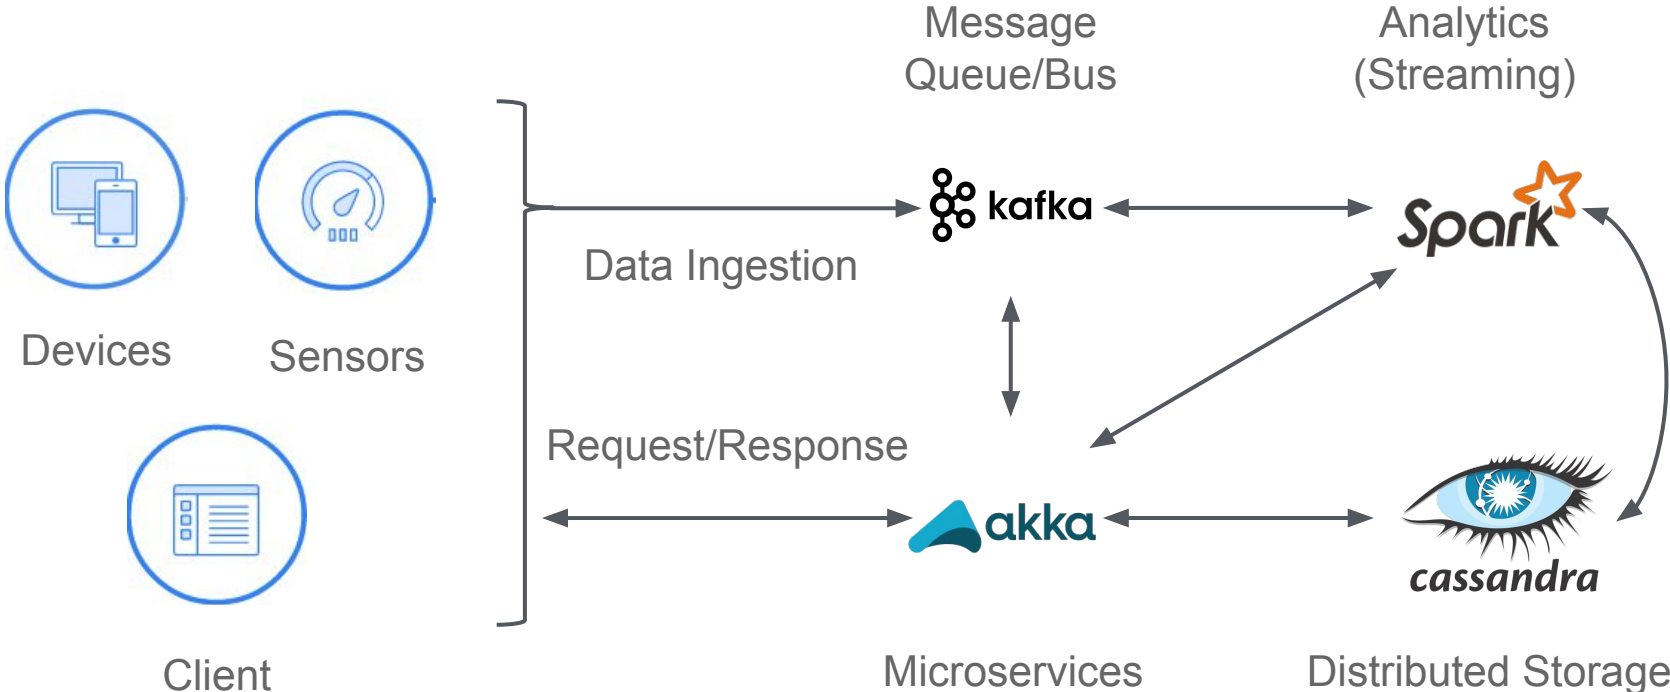
MODERN APPLICATION -> FAST DATA BUILT-IN



Use Cases:

- Anomaly detection
- Personalization
- IoT Applications
- Predictive Analytics
- Machine Learning

The SMACK Stack



Use Cases:

- Anomaly detection
- Personalization
- IoT Applications
- Predictive Analytics
- Machine Learning

Keeping things running: Day 2 Operations

Metrics & Monitoring

- Collecting metrics
- Routing events
- Downstream processing
 - Alerting
 - Dashboards
 - Storage (long-term retention)

Logging

- Scopes
- Local vs. Central
- Security considerations

Day 2 Operations con't

Maintenance

- Cluster Upgrades
- Cluster Resizing
- Capacity Planning
- User & Package Management
- Networking Policies
- Auditing
- Backups & Disaster Recovery

Troubleshooting

- Debugging
 - Services
 - System
 - Access
- Tracing
- Chaos Engineering

Demo

<https://dcos.io/demos/>

Questions? Feedback?

Elizabeth K. Joseph
Twitter: @pleia2
Email: lyz@princessleia.com



@dcos



chat.dcos.io



users@dcos.io



/dcos

/dcos/examples

/dcos/demos