

MesosCon
EUROPE

DataStax Enterprise on DC/OS

Yes, It's Possible; Customer Case Studies

Chris Splinter - Product Manager, DataStax

Agenda


- Production Quality Stateless Services
- Evolution of the Framework
- Current State of the Project
- Customer Deployments
- Demos



What is DataStax?

(10 second overview)

DataStax Overview



- Peer-to-Peer
- NoSQL Database
- Based on Cassandra
 - + Search
 - + Analytics
 - + Graph
 - + Security



PERSONALIZATION



MESSAGING



INTERNET OF THINGS



FRAUD DETECTION




PLAYLISTS



Production Quality Stateful Services


Terminology

- Node: a single instance
- Data Center: a logical grouping of racks or nodes
- Cluster: a logical grouping of datacenters

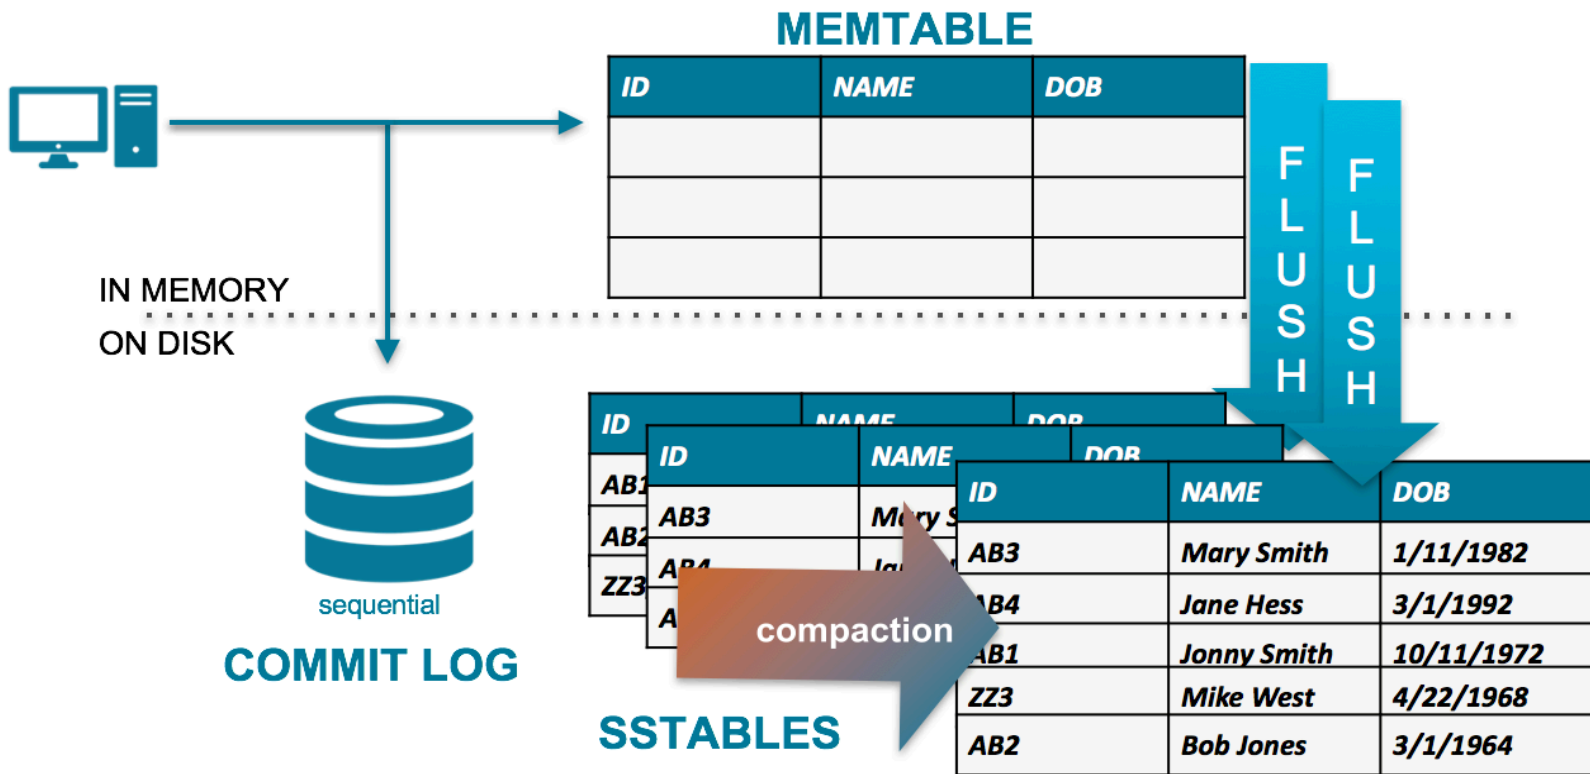


Gossip Protocol

- Gossip makes it all work
- Nodes interact every second
- Share information about themselves and others
- Slow or down nodes are marked as such
- Requests are routed based on this peer knowledge



Write Path



Configuration Management

- 30+ configuration files
 - cassandra.yaml = ~1k settings
 - dse.yaml
- Production Changes Likely
- Workload Dependent

```
./resources/cassandra/tools/bin/cassandra.in.sh
./resources/graph/gremlin-console/bin/gremlin.sh
./resources/hadoop2-client/conf/hadoop-env.sh
./resources/hadoop2-client/conf/httpfs-env.sh
./resources/hadoop2-client/conf/kms-env.sh
./resources/hadoop2-client/conf/mapred-env.sh
./resources/hadoop2-client/conf/yarn-env.sh
./resources/hadoop2-client/libexec/hadoop-config.sh
./resources/hadoop2-client/libexec/hdfs-config.sh
./resources/hadoop2-client/libexec/httpfs-config.sh
./resources/hadoop2-client/libexec/kms-config.sh
./resources/hadoop2-client/libexec/mapred-config.sh
./resources/hadoop2-client/libexec/yarn-config.sh
./resources/spark/bin/load-spark-env.sh
resources/spark/conf/dse-spark-env.sh
resources/spark/conf/spark-env.sh
resources/spark/sbin/spark-config.sh
resources/spark/sbin/spark-daemon.sh
resources/spark/sbin/start-history-server.sh
resources/spark/sbin/start-thriftserver.sh
resources/spark/sbin/stop-history-server.sh
resources/spark/sbin/stop-thriftserver.sh
resources/spark/spark-jobserver/kill-process-tree.sh
resources/spark/spark-jobserver/manager_start.sh
resources/spark/spark-jobserver/server_start.sh
resources/spark/spark-jobserver/server_stop.sh
resources/spark/spark-jobserver/setenv.sh
resources/spark/spark-jobserver/settings.sh
resources/cassandra/conf/cassandra-topology.yaml
resources/cassandra/conf/cassandra.yaml
resources/conf/metrics-reporter-config-sample.yaml
resources/tools/cqlstress-counter-example.yaml
resources/tools/cqlstress-example.yaml
resources/tools/cqlstress-insanity-example.yaml
```

Stateful Services are Hard

Deployment

Maintenance



Operations

Resources

Backup +
Restore

Configuration

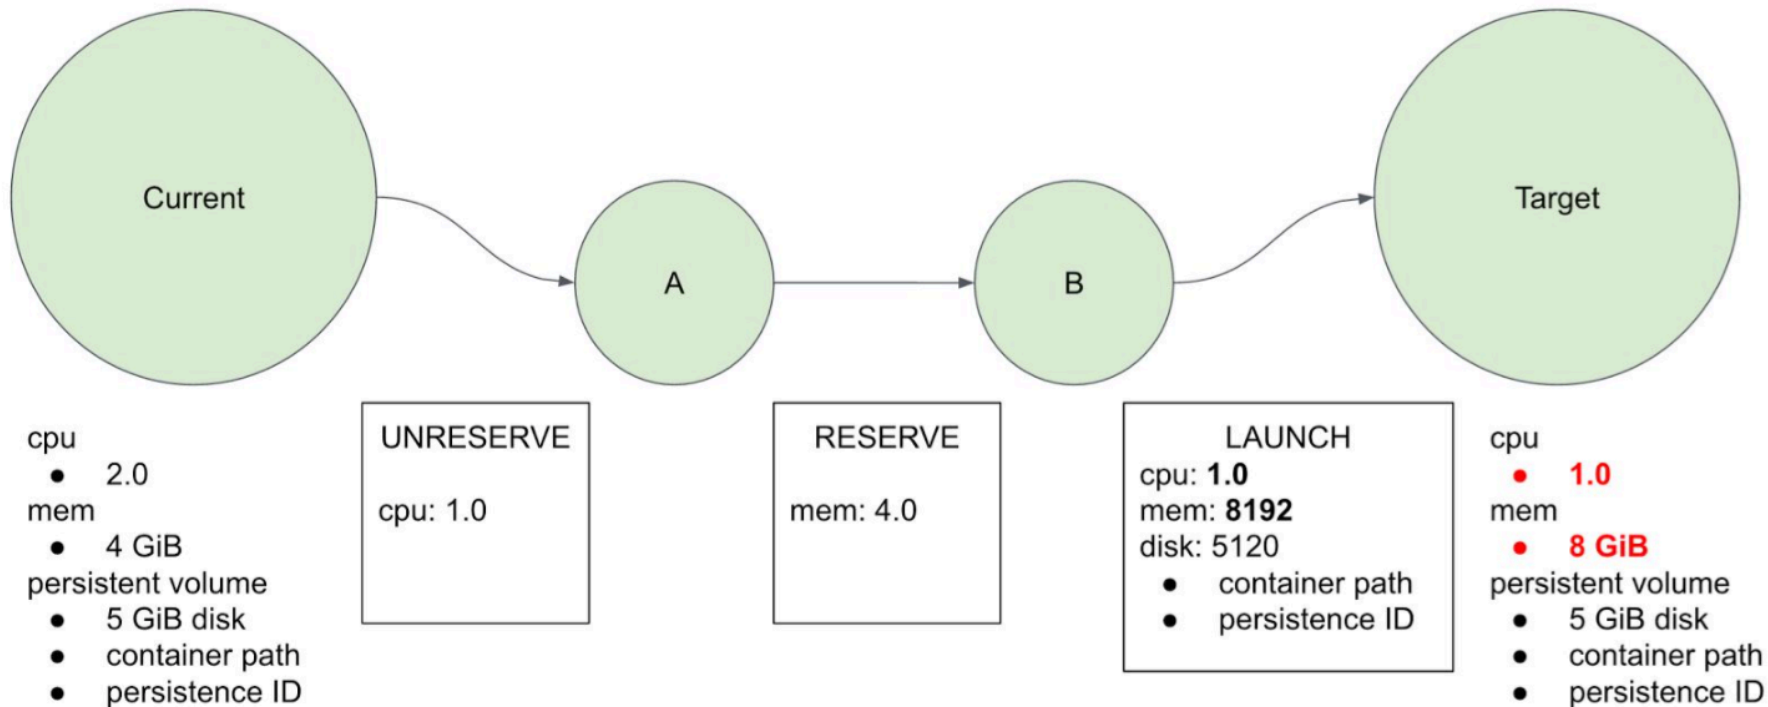
Resilience



DC/OS
SOFTWARE DEVELOPMENT KIT

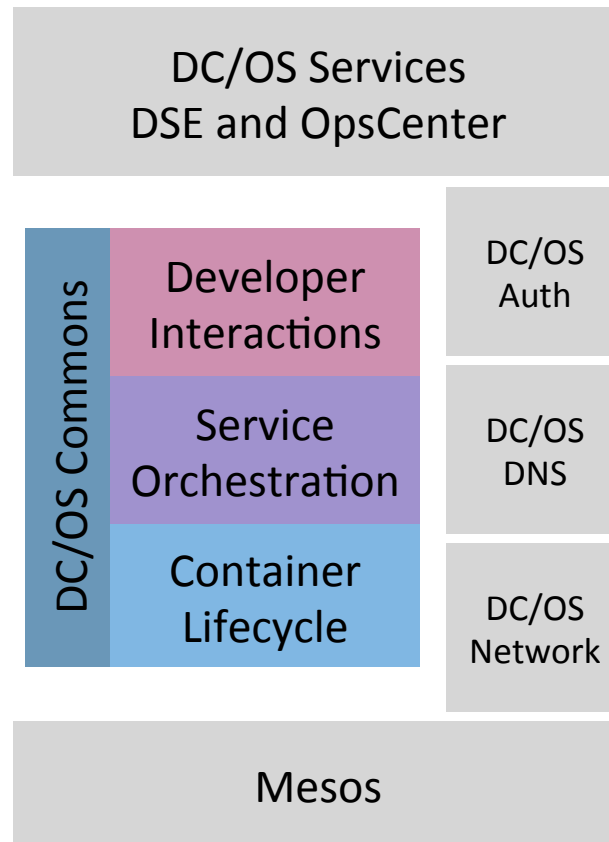
Configuration Management

Update



DCOS-Commons

- Open Source library to achieve these goals
- Robust deployment
 - install
 - configuration
 - software updates
- Maintenance of Services
- Accomplishes these tasks through a declarative or “goal oriented” approach
- <https://github.com/mesosphere/dcos-commons>





Evolution of the Framework

Version One

- Design Principles
 - Easy to Develop and Maintain
 - Easy to Extend and Customize
 - Single Universe Package
- Outcome
 - Customers wanted to fly a plane, not build one









Version Two – Give them a plane

- Major Changes
 - Dedicated DataStax engineering team
 - Hand-in-Hand collaboration with Mesosphere
 - Beta the heck out of it
 - Joint Support Agreements
- Better Together
 - Configuration Management – Rollback Changes
 - Automated vertical (resources) and horizontal (node count) scaling
 - Uniform deployment of Enterprise Applications

Current State

- mesosphere-dse... ▾
- Dashboard
- Services
- Jobs
- Catalog**
- RESOURCES
- Nodes
- Networking ▸
- SYSTEM
- Overview
- Components

Catalog

 cassandra CERTIFIED	 chronos CERTIFIED	 confluent-kafka CERTIFIED	 datastax-dse CERTIFIED
 datastax-ops CERTIFIED	 dcos-enterprise-cli CERTIFIED	 elastic CERTIFIED	 hdfs CERTIFIED

Version Two – Current State

- Full Platform Support
- Node Placement – Node / Task Failure Recovery
- Strict Mode Support
- Multi-Tenancy
- Pod Replace with local storage
- Networking Management / CNI
- Monitoring



Customer Deployments

Customer Use Cases

- Distributing Real-Time Microservices to a large travel company
 - Repeatable, Consistent deployments of an application stack
 - DSE + Kafka + API Management

- Platform as a Service by Financial Institution
 - DSE, Kafka, Others as a Service

Demos





Install DSE

(the database)

- mesosphere-dse...
- Dashboard
- Services
- Jobs
- Catalog
- SOURCES
- Nodes
- Networking
- STEM
- Overview
- Components

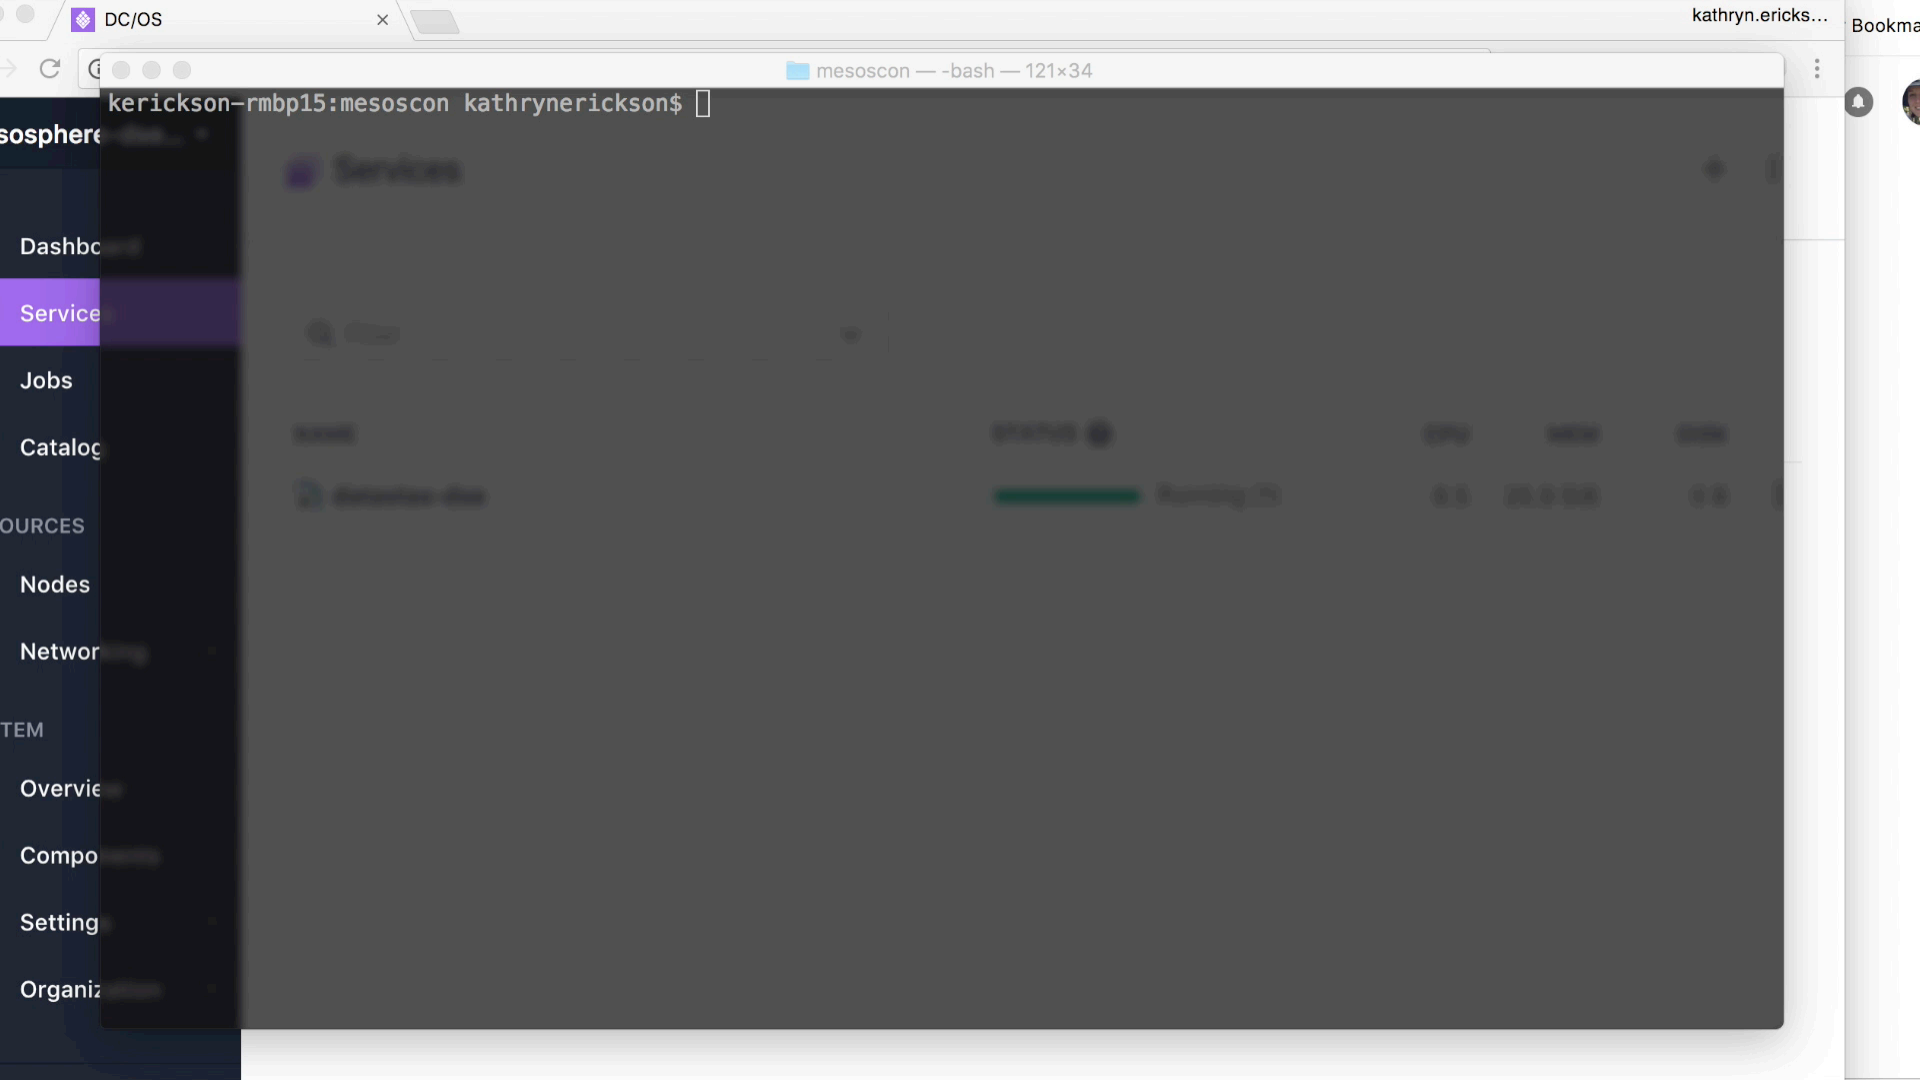
Services

NAME	STATUS ?	CPU	MEM	DISK
 datastax-ops	 Running (1)	3	4.9 GiB	10 GiB




Install OpsCenter

(Management GUI)



DC/OS 1.10 CLI

con-dse superuser








Services >  datastax-dse Running (1)

Instances Configuration Debug

Showing 7 of 9 tasks [\(Clear\)](#)

Filter

ALL 9 ACTIVE 7 COMPLETED 2





<input type="checkbox"/>	ID	NAME	HOST	STATUS	HEALTH		CPU
<input type="checkbox"/>	dse-2-node__00c837f1-f161-41ec-982e-20b232983a6d	dse-2-node	10.0.3.173	Running	●		2
	dse-2-agent__210942d0-5940-4a25-ba67-1890595f6959	dse-2-agent	10.0.3.173	Running	●		0.5
	dse-1-node__51ea31cd-a2e1-4af2-b286-84da155aecb8	dse-1-node	10.0.3.51	Running	●		2
	dse-1-agent__c5f2c82b-8b69-4ec3-88e0-85594cdd6128	dse-1-agent	10.0.3.51	Running	●		0.5
	dse-0-node__cd370287-9150-4a4a-8453-196d5ddac9b2	dse-0-node	10.0.3.22	Running	●		2
	dse-0-agent__1f24a760-b6d1-440f-bc87-a1500ba3aba5	dse-0-agent	10.0.3.22	Running	●		0.5
<input type="checkbox"/>	datastax-dse.da751c80-972d-11e7-abea-f2ea4739ecaa	datastax-dse	10.0.2.128	Running	●		1

Add a Node

- Dashboard
- Services
- Jobs
- Catalog
- SOURCES
- Nodes
- Networking
- ITEM
- Overview
- Components
- Settings
- Organization

Services

Filter

NAME	STATUS ?	CPU	MEM	DISK
 datastax-dse	 Running (1)	8.5	25.9 GiB	0 B
 datastax-ops	 Running (1)	3	4.9 GiB	10 GiB



Take it for a spin

<https://github.com/012345/MesosCon>

MesosCon

EUROPE