CSDN

DC/OS 1.9 DEEP DIVE

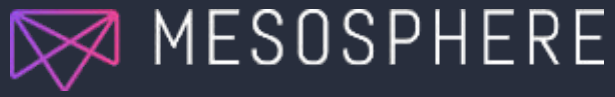

Mesosphere DC/OS 1.9

AGENDA

- Data Services Ecosystem
- Operations
- Workloads
- Compliance Reports

MESOSPHERE DC/OS: RUN YOUR DATACENTER AS A GIANT COMPUTER

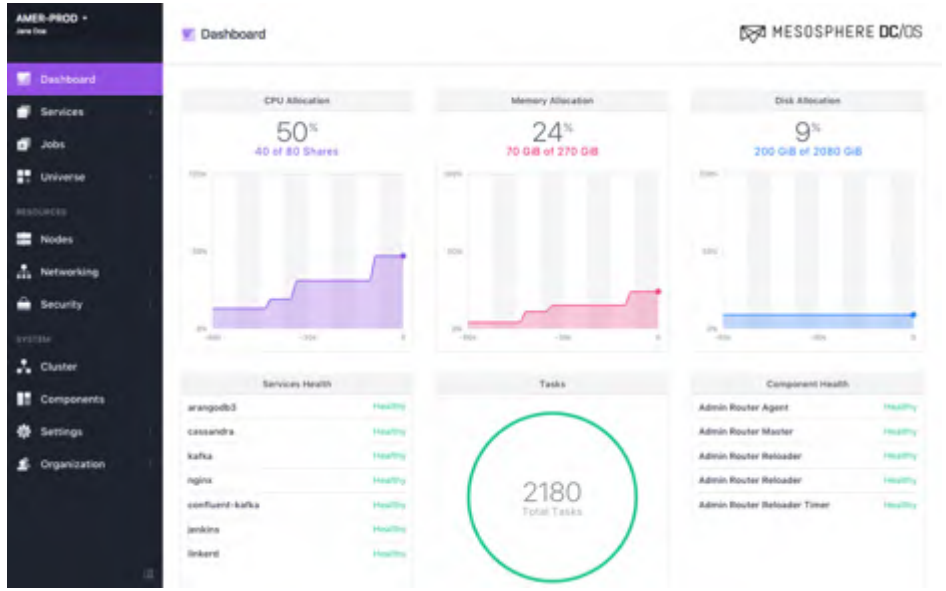

Manage your infrastructure as a single giant computing pool

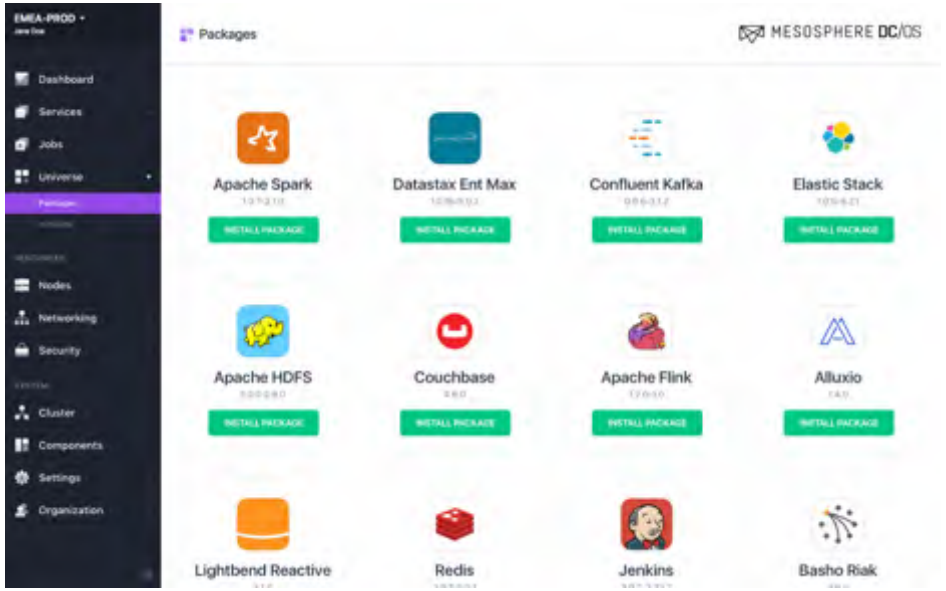

Deploy & run complex services with an app store like experience

MESOSPHERE DC/OS: - GOING A LITTLE DEEPER

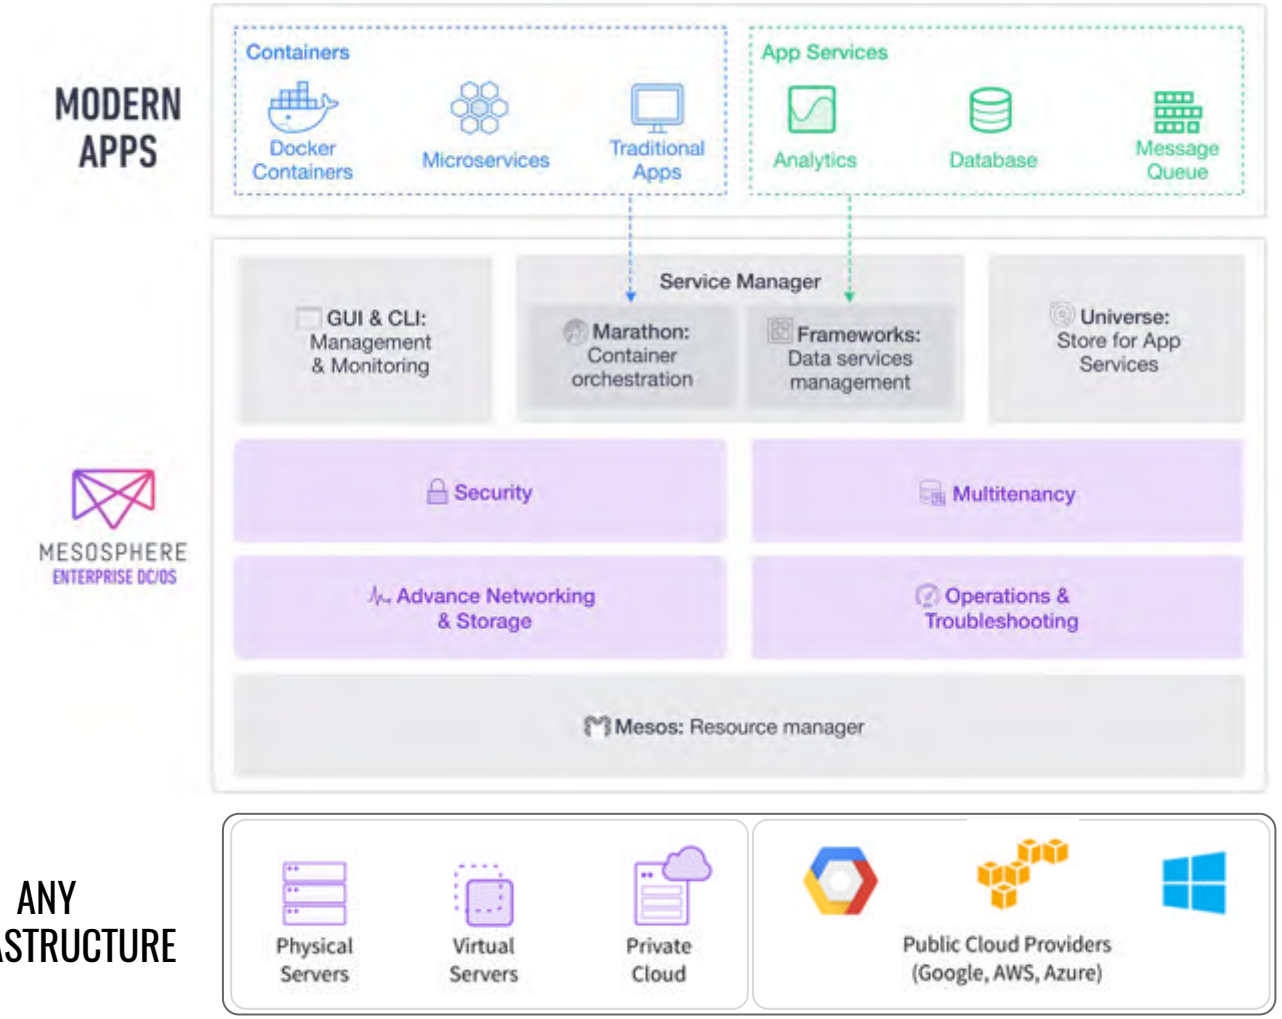

DC/OS

NEW UI

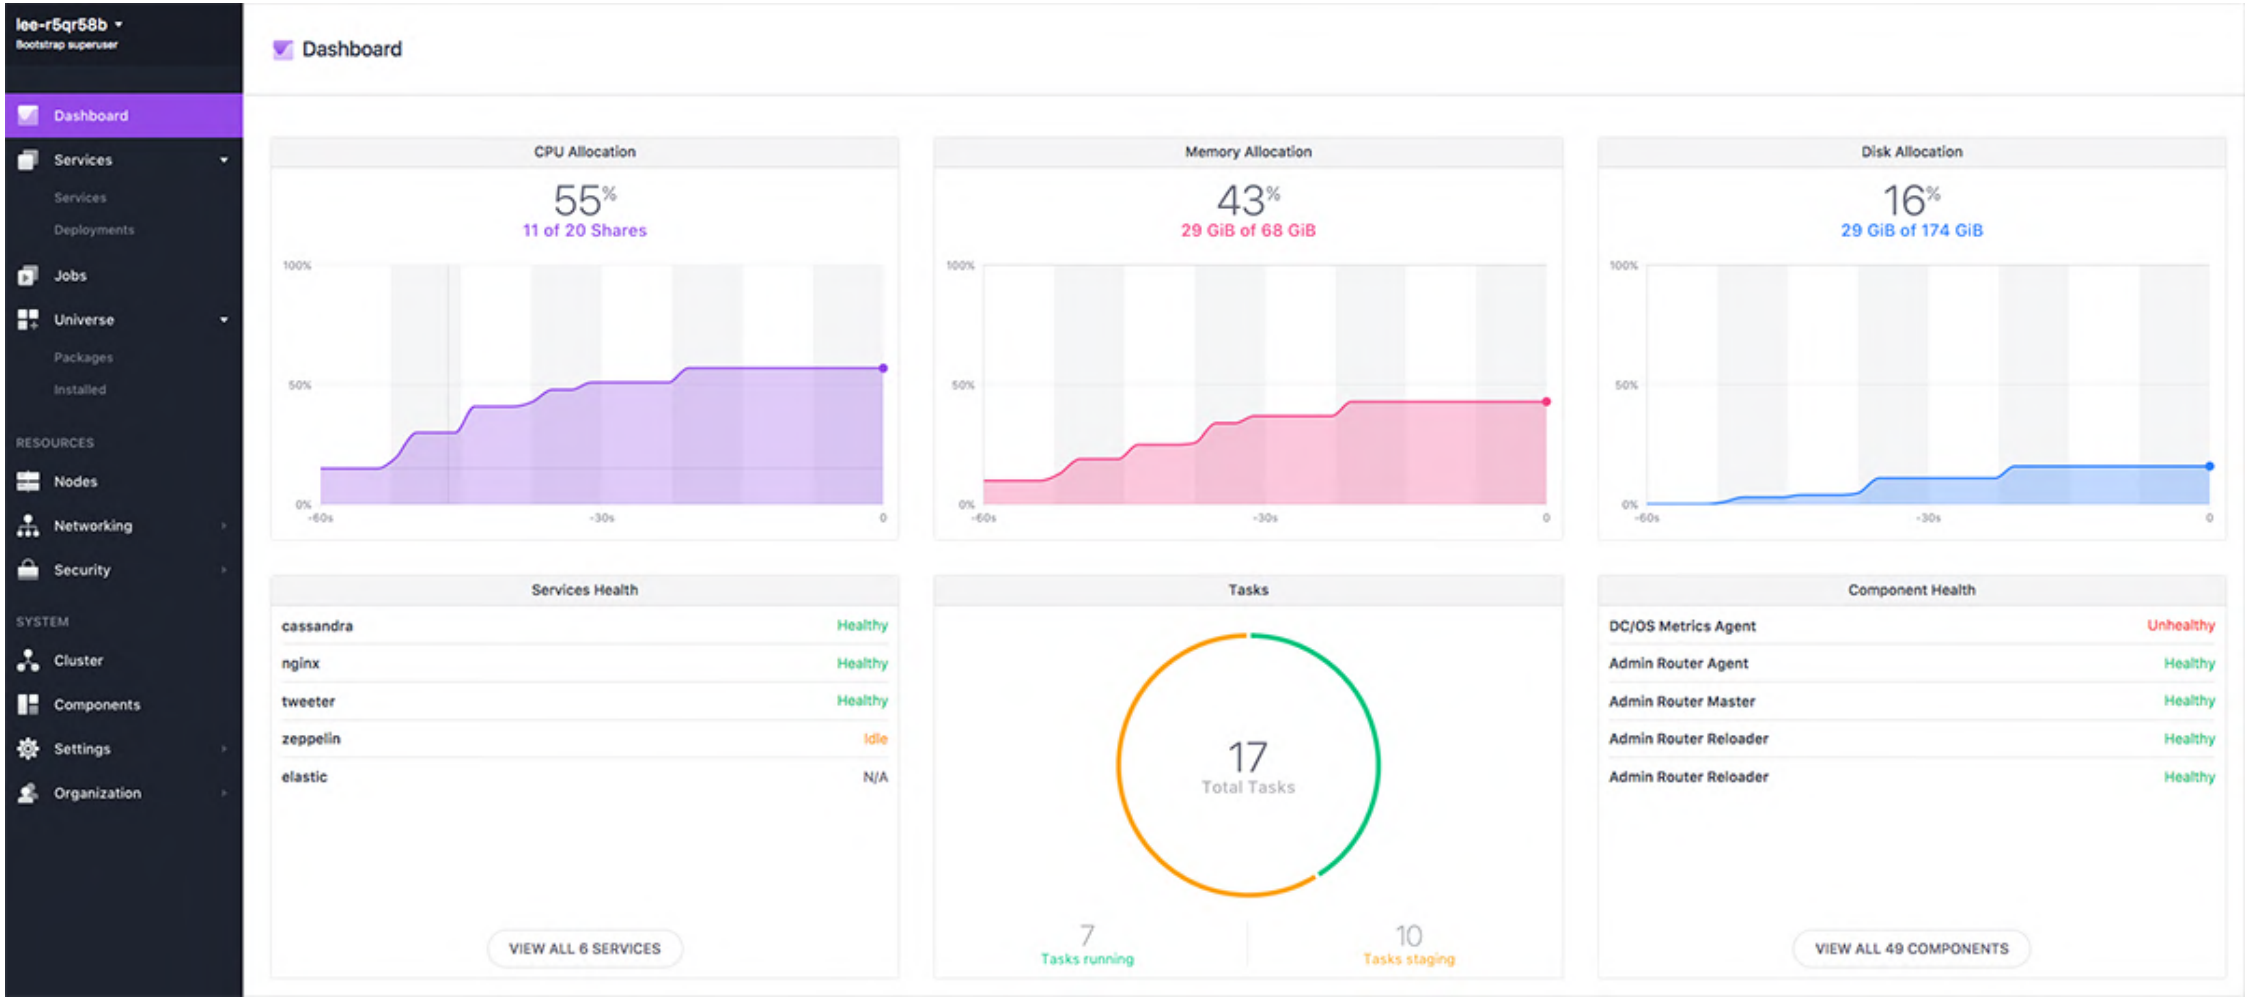

DC/OS

NEW CREATE SERVICE FLOW

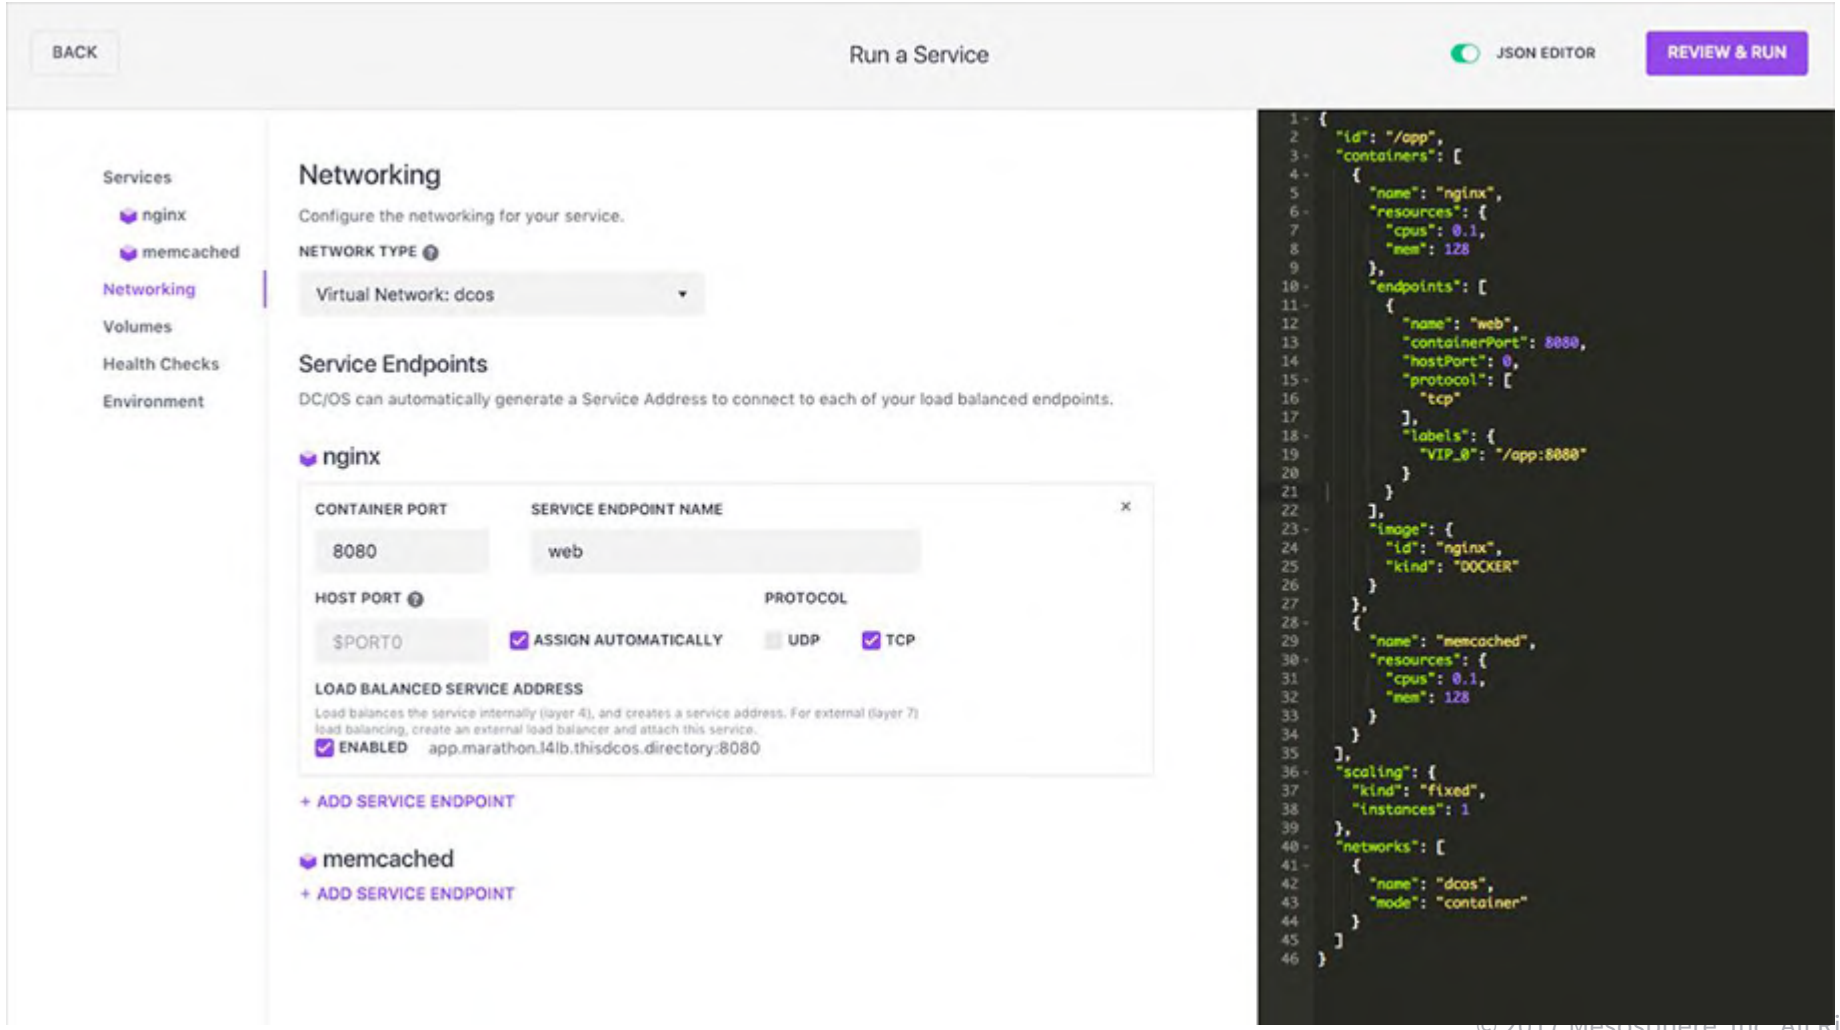

KEY THEMES - DC/OS 1.9

DATA SERVICES ECOSYSTEM

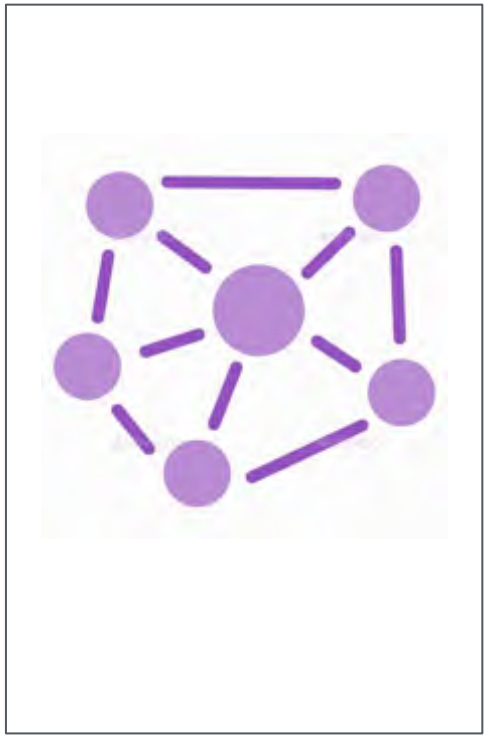

OPERATIONS

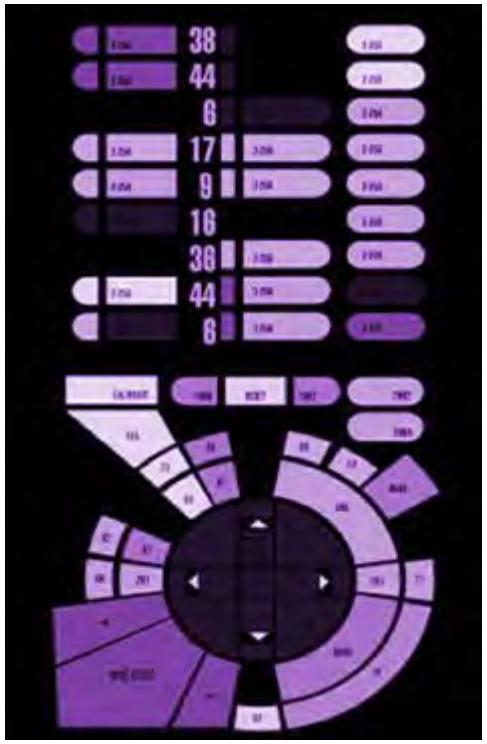

WORKLOADS

COMPLIANCE REPORTS

KEY THEMES - DC/OS 1.9

DATA SERVICES ECOSYSTEM

- *Alluxio*
- *Couchbase*
- *Datastax DSE*
- *Elastic (ELK)*
- *Redis*
- *Apache Flink*

OPERATIONS

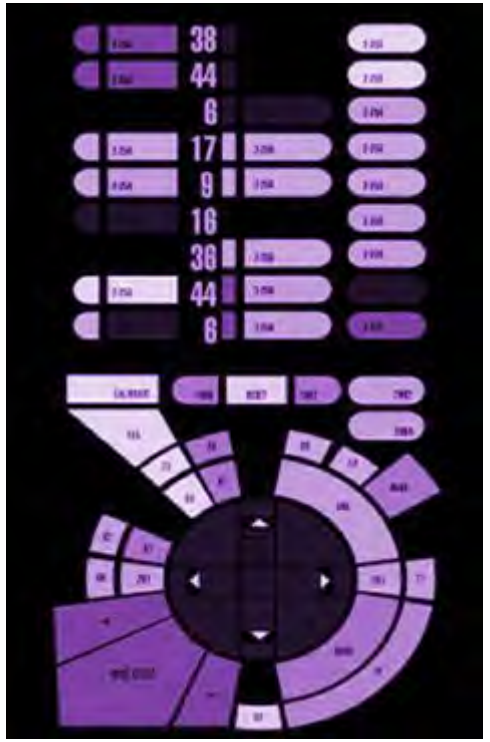

WORKLOADS

COMPLIANCE REPORTS

DATA INFRASTRUCTURE CHALLENGES

Deploying each data service is time consuming

Operating data services is manual and error-prone

Infrastructure silos with low utilization

DC/OS: DATA SERVICES ECOSYSTEM

NEW SERVICES PARTNERSHIPS IN DC/OS 1.9

And more than +100 Services available in our service catalog “DC/OS Universe”

KEY THEMES - DC/OS 1.9

DATA SERVICES ECOSYSTEM

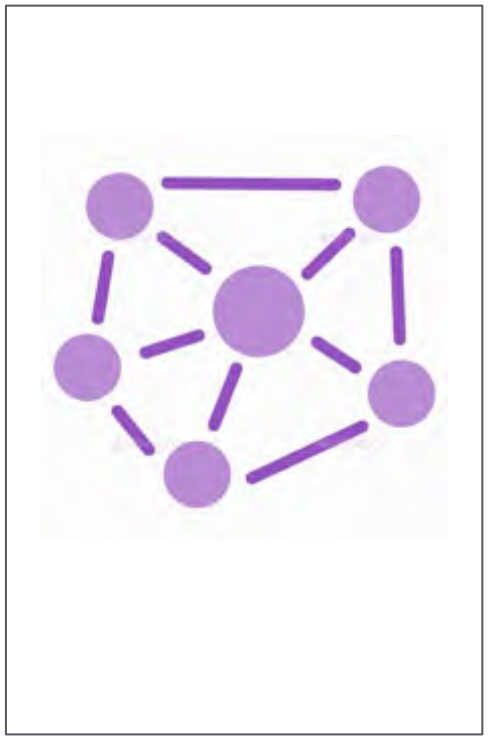

OPERATIONS

- Remote Container Shell
- Unified Metrics
- Unified Logging
- Deployment Failure Debugging
- Upgrades & Configuration updates

WORKLOADS

COMPLIANCE REPORTS

DC/OS: OPERATIONS

REMOTE CONTAINER SHELL

- Open encrypted, interactive, remote session to your containers
- Remotely execute commands for real time app troubleshooting
- Provide developers access to their own applications, not the entire host or cluster
- Restrict access to specific applications with Access Control Lists [Enterprise DC/OS only]

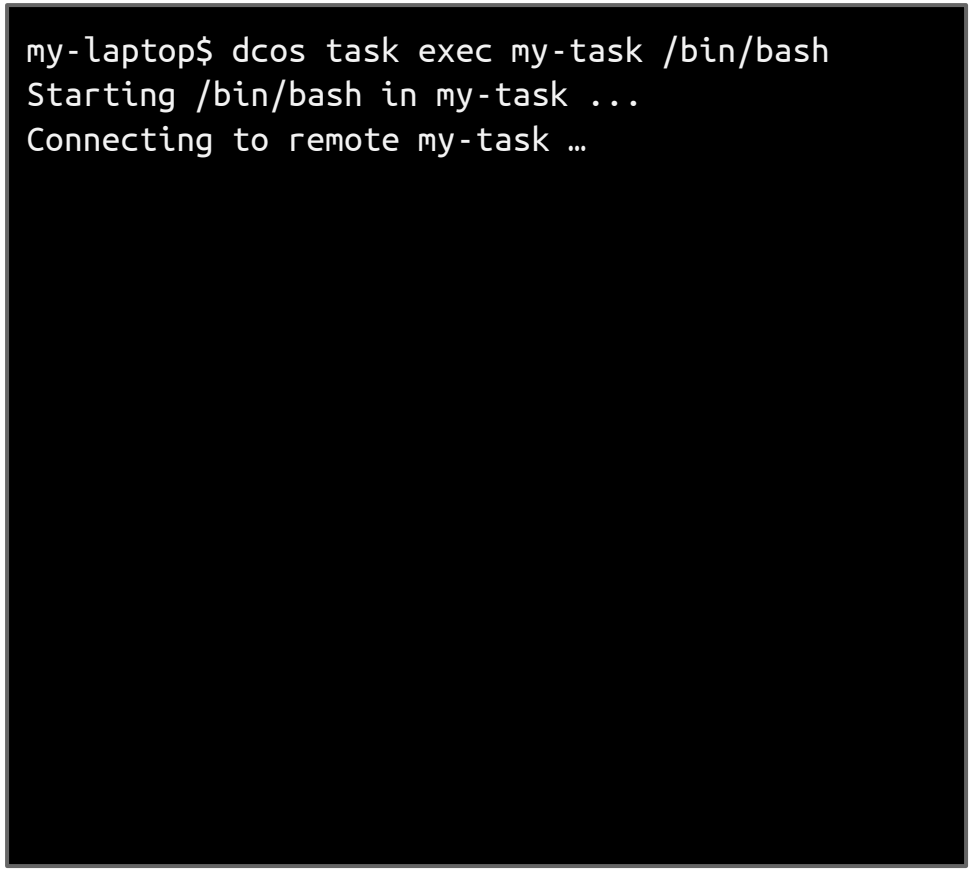

DC/OS: OPERATIONS

UNIFIED LOGGING

- Access application, DC/OS and OS logs
- Easily troubleshoot applications with critical metadata such as container id and app id
- Integrate easily with existing logging systems
- Restrict access to specific app logs with Access Control Lists [Enterprise DC/OS only]

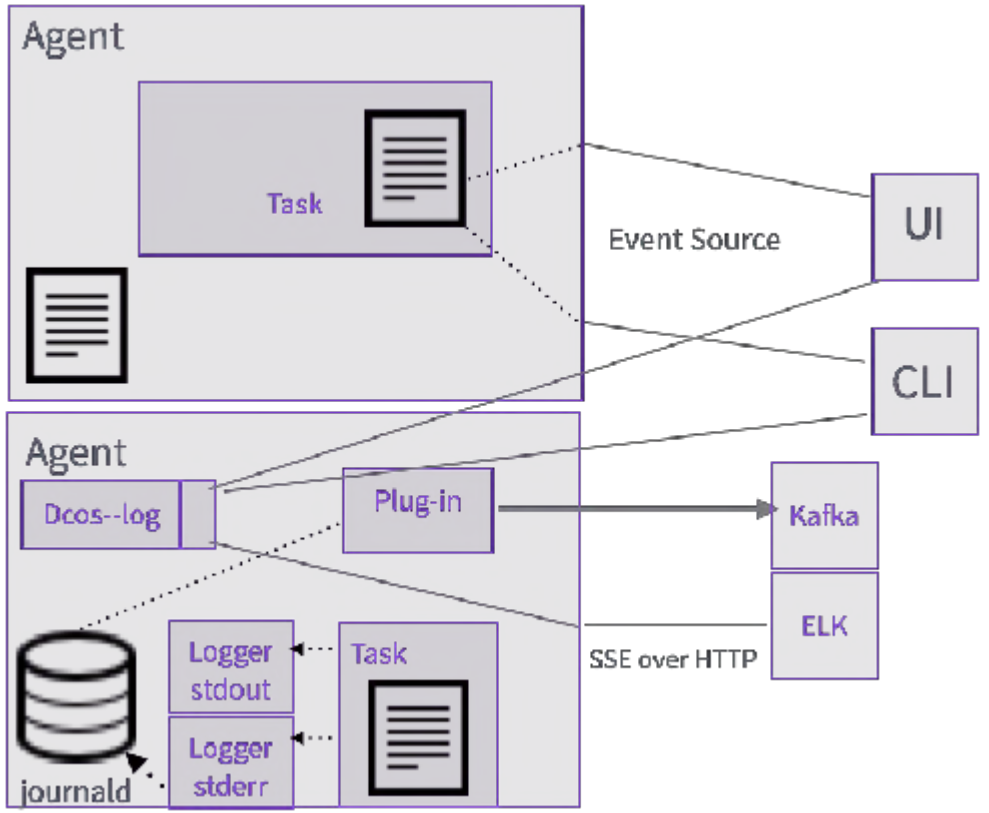

DC/OS: OPERATIONS

UNIFIED METRICS

- Single API for system, container and application metrics
- Metadata such as host id and container id are automatically added to assist in debugging
- Integrate easily with existing metrics systems
- Restrict access to specific app metrics with Access Control Lists [Enterprise DC/OS only]

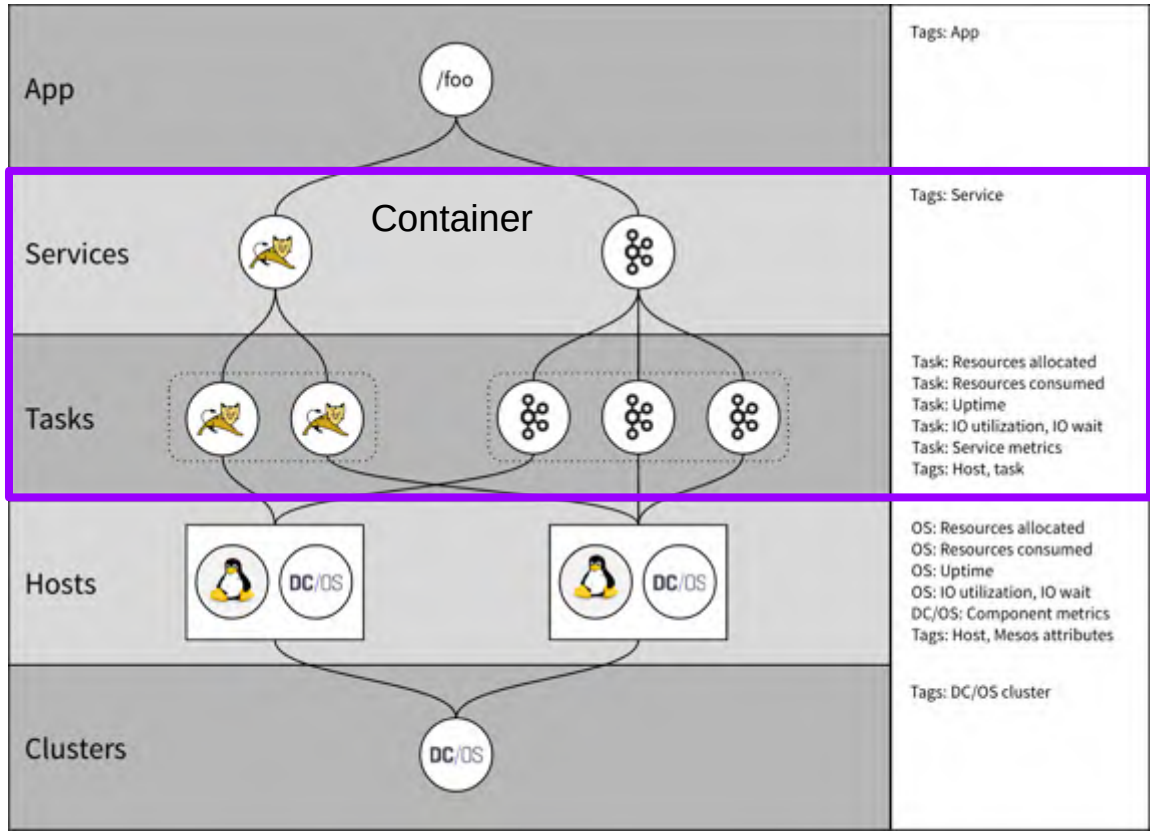

DC/OS: OPERATIONS

DEPLOYMENT FAILURE DEBUGGING

- Understand why your application is not deploying
- Understand which nodes in the cluster can accommodate the role, constraints, cpu, mem, disk and port requirements for your app

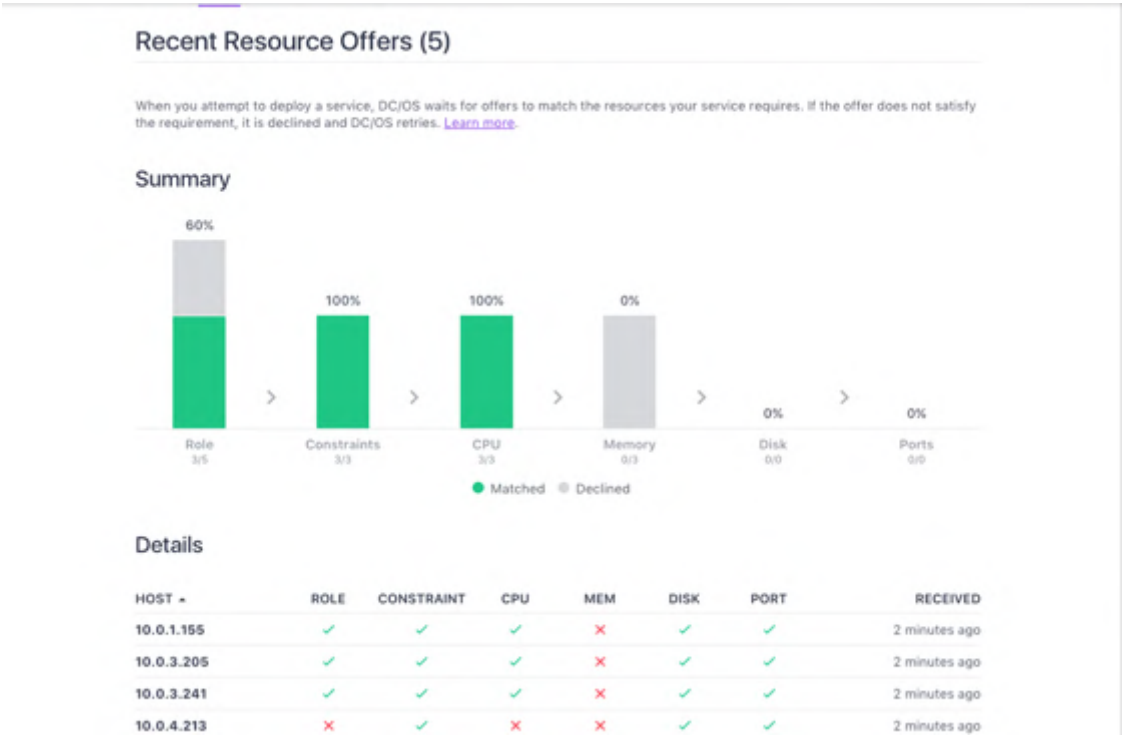

DC/OS: OPERATIONS

UPGRADES AND CONFIG UPDATES

- Generate new config for cluster nodes

```
$ dcos_generate_config.sh --generate-node-upgrade-script
<installed_cluster_version>
```

- Single command upgrade script for individual nodes

```
$ curl -O <Node upgrade script URL>
$ sudo bash ./dcos_node_upgrade.sh
```

KEY THEMES - DC/OS 1.9

DATA SERVICES ECOSYSTEM

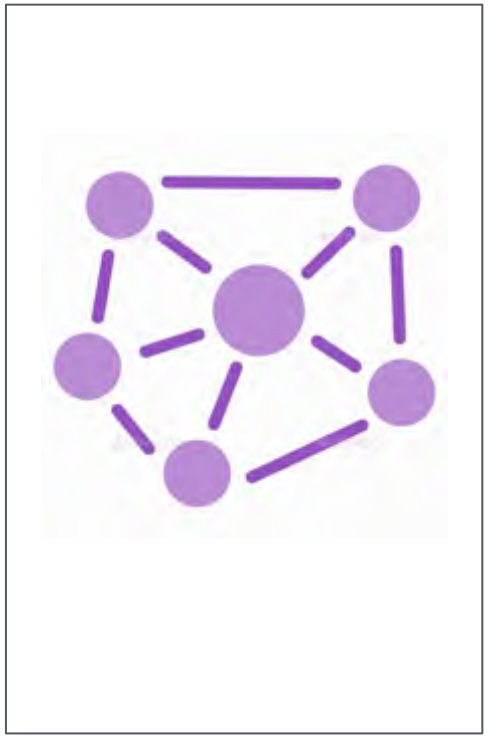

OPERATIONS

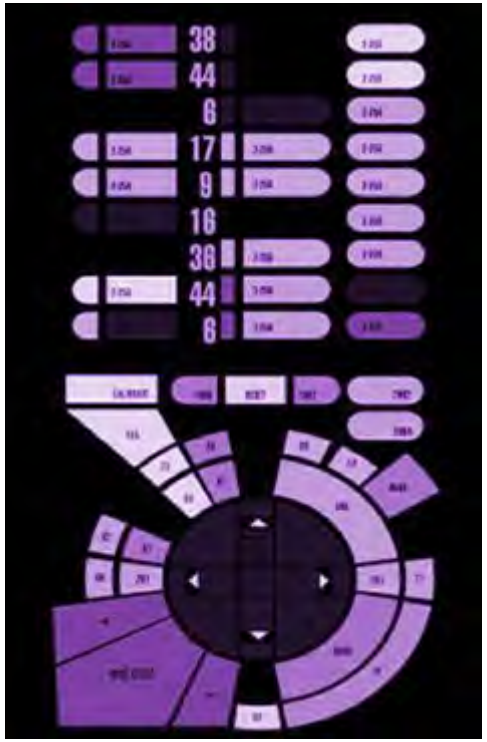

WORKLOADS

- Pods
- GPU based scheduling

COMPLIANCE REPORTS

DC/OS: WORKLOADS

PODS

- Schedule, deploy and scale multiple containers on the same host(s) while sharing IP address and storage volumes
- All containers in a pod instance run as if they are running on a single host in pre-container world
- Useful for migrating legacy applications or building advanced micro services (side car containers)

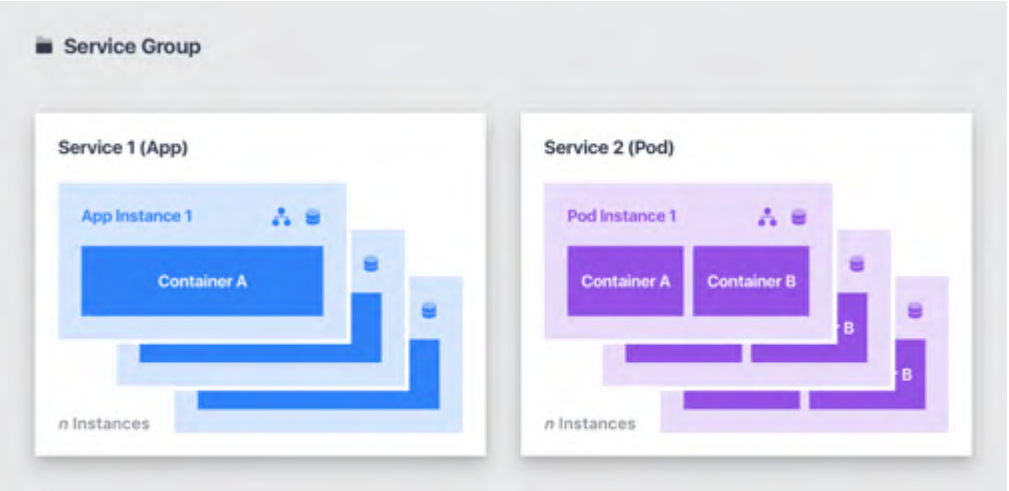

DC/OS: WORKLOADS

PODS: MIGRATING LEGACY APPS TO CONTAINERS

- Traditional monolithic apps on VMs usually have support services such as log shipper, message queuing clients
- Many support services assume col-location on same host, and local-host access to networking and storage
- Pods simplify moving legacy monolithic apps to containers, reducing risk and accelerating migrations

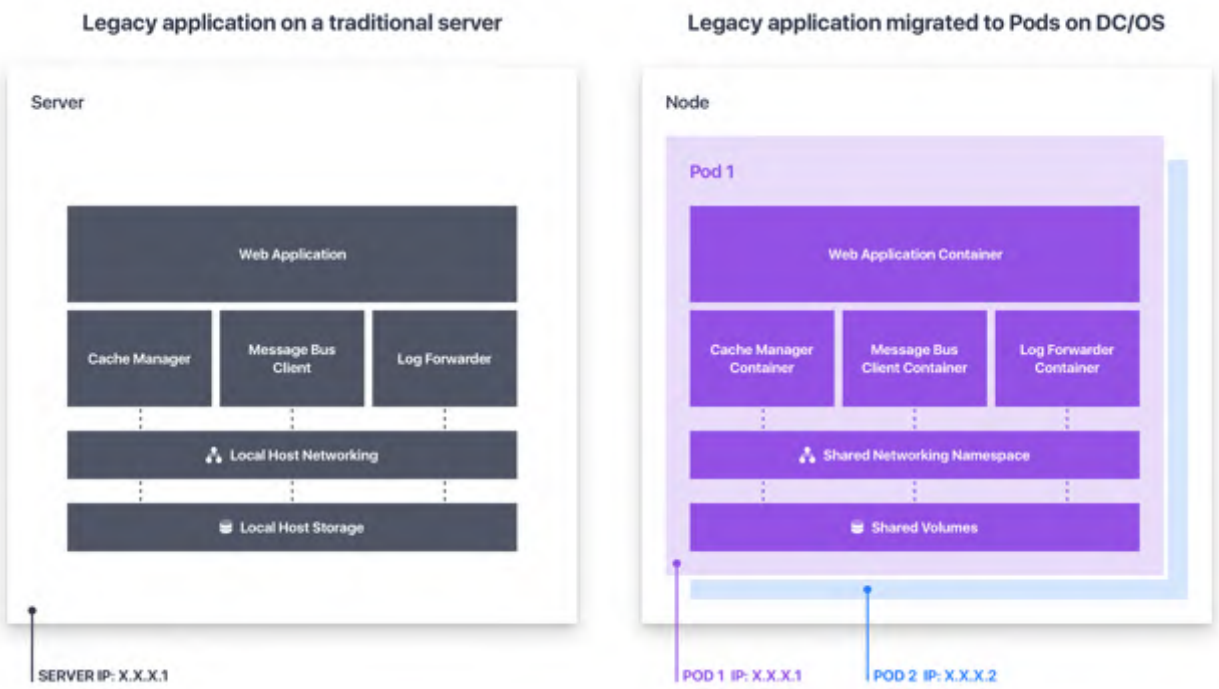

PODS: SUPPORT SERVICES (SIDE-CAR CONTAINERS)

- Advanced Micro Services patterns require colocating containers together
- Support services include for example:
  - Logging or monitoring agents,
  - Backup tooling & Proxies
  - Data change watchers & Event publishers
- Pods simplify the building and maintenance of complex such microservices

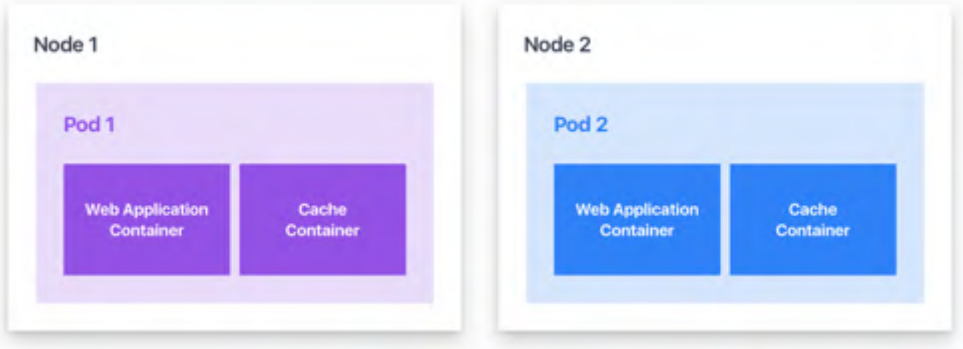

DC/OS: WORKLOADS

GPU: WHY GPU?

- GPUs are needed for many machine learning and deep learning applications
- GPUs are essential for real-time or near real time machine learning models
- GPUs deliver from 10X to 100X performance for some applications, resulting lower $$$/IOPS and more productivity to data science teams
- GPU applications include real time fraud detection, genome sequencing, cohort analysis and many others

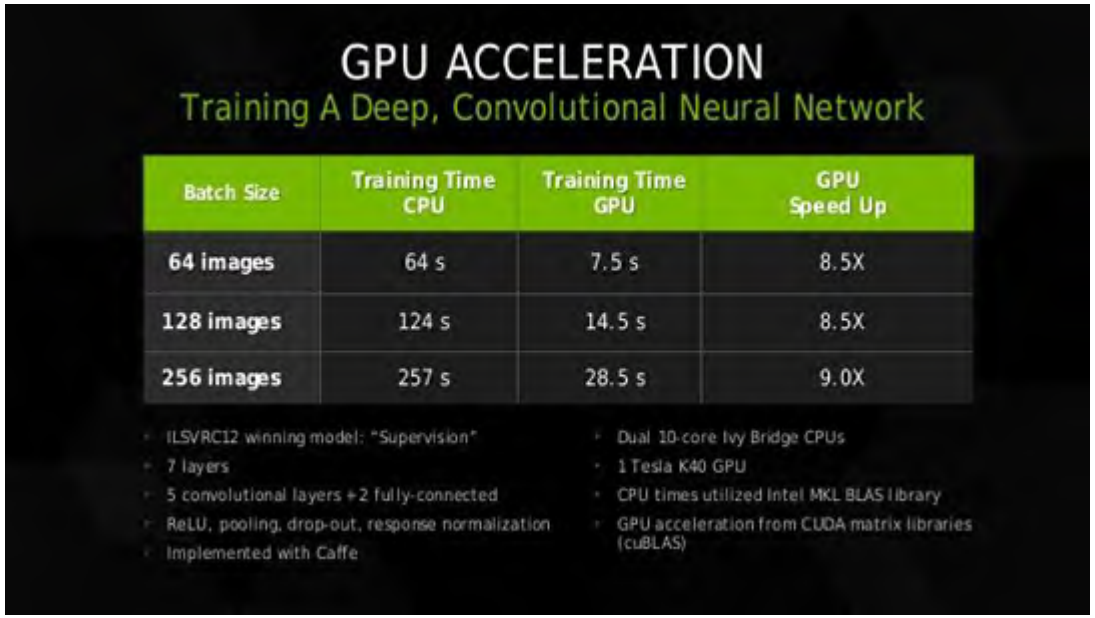

DC/OS: WORKLOADS

CHALLENGES WITH GPU CLUSTERS

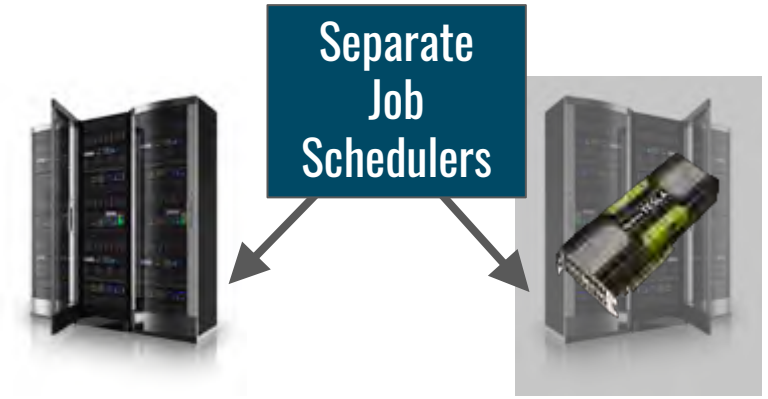

SILOS

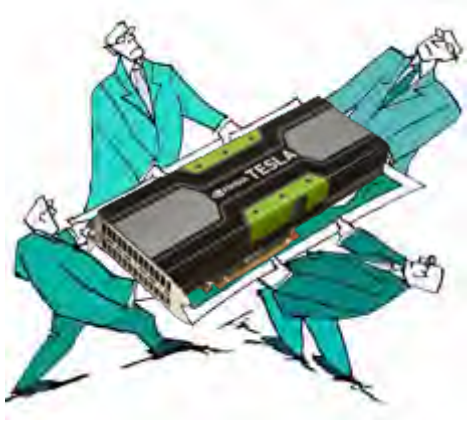

NO ISOLATION

NO PORTABILITY

DC/OS: WORKLOADS

GPU BASED SCHEDULING

- Simplify migrating machine learning models across from dev to production, and across clouds
- Test locally with Nvidia-Docker, deploy to production with DC/OS
- Isolate GPU instances and schedule workloads just like CPU and memory, guaranteeing performance
- Efficiently share GPU resources across data science team

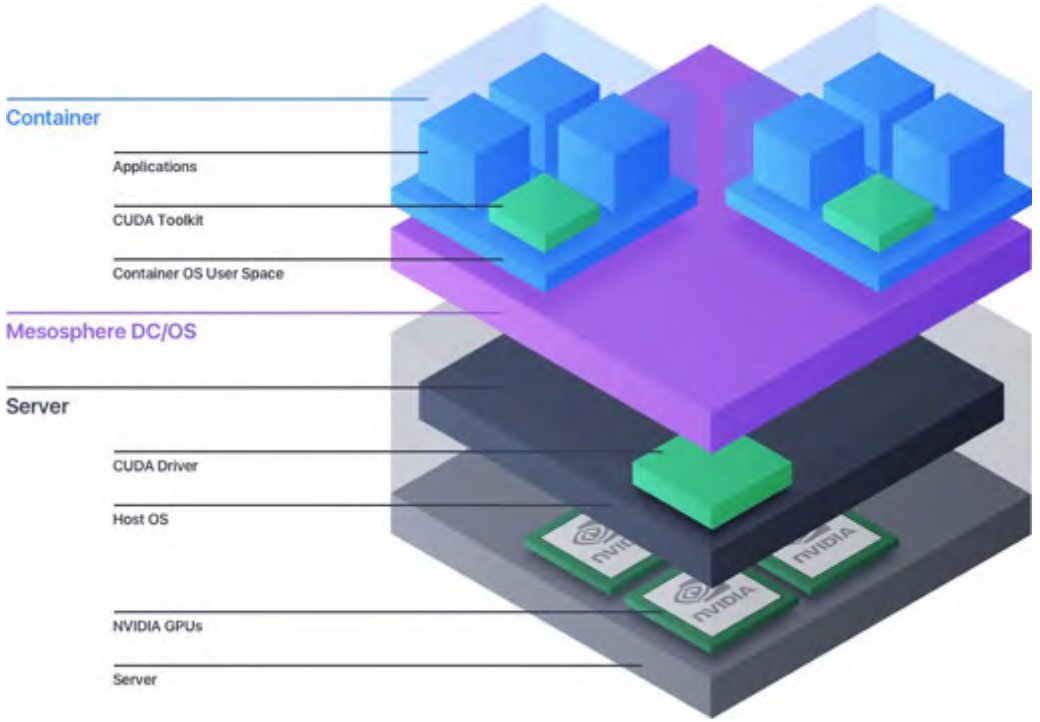

KEY THEMES - DC/OS 1.9

DATA SERVICES ECOSYSTEM

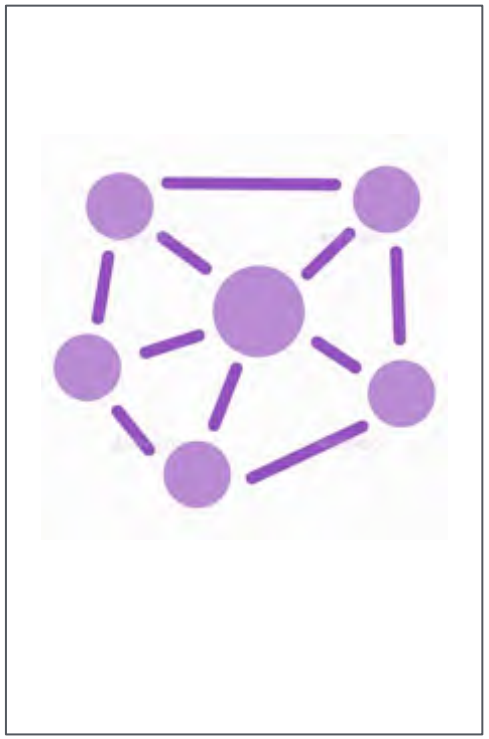

OPERATIONS

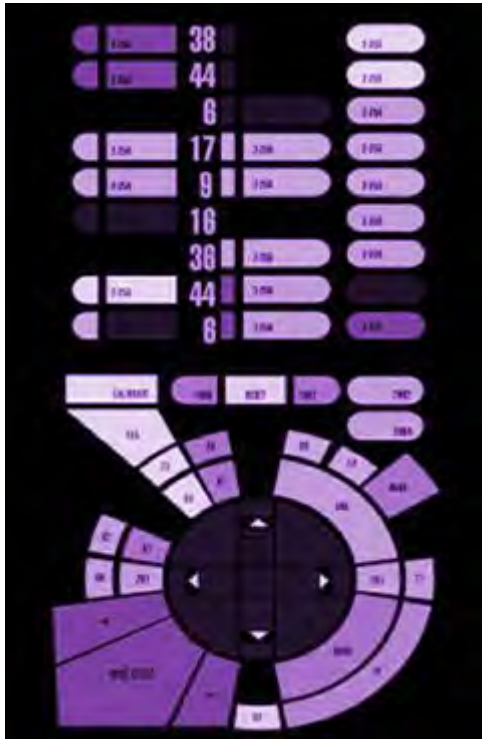

WORKLOADS

COMPLIANCE REPORTS

- NIST-800/53
- ISO 27001
- FISMA-Mod.
- SOC-2
- HIPAA

DC/OS: COMPLIANCE

WHICH CERTIFICATION FOR WHAT?

FISMA-Moderate

- ***Top Level Federal/DoD Controls report***
- ***Superset of***
  - *NIST-800*
  - *Specific STIGs*
  - *FIPS 140-2 controls*

ISO-27001

- *International Standard around Information Security Risk procedures.*
- *Widely accepted by non-US GEOs*

SOC-2

- *Required by anyone who stores or processes any kind of business financial data*

HIPPA

- ***Required by any organization that processes or stores healthcare data.***

DC/OS: COMPLIANCE

ACCELERATE CERTIFICATION WITH DC/OS COMPLIANCE REPORTS

- Jump Start Complete installation compliance audit process.
- Understand how DC/OS meets various compliance controls.
- Understand where to put workaround or human controls
- Identify any high value gaps and possible remediation plans
- Accelerate time to on-boarding business critical applications.

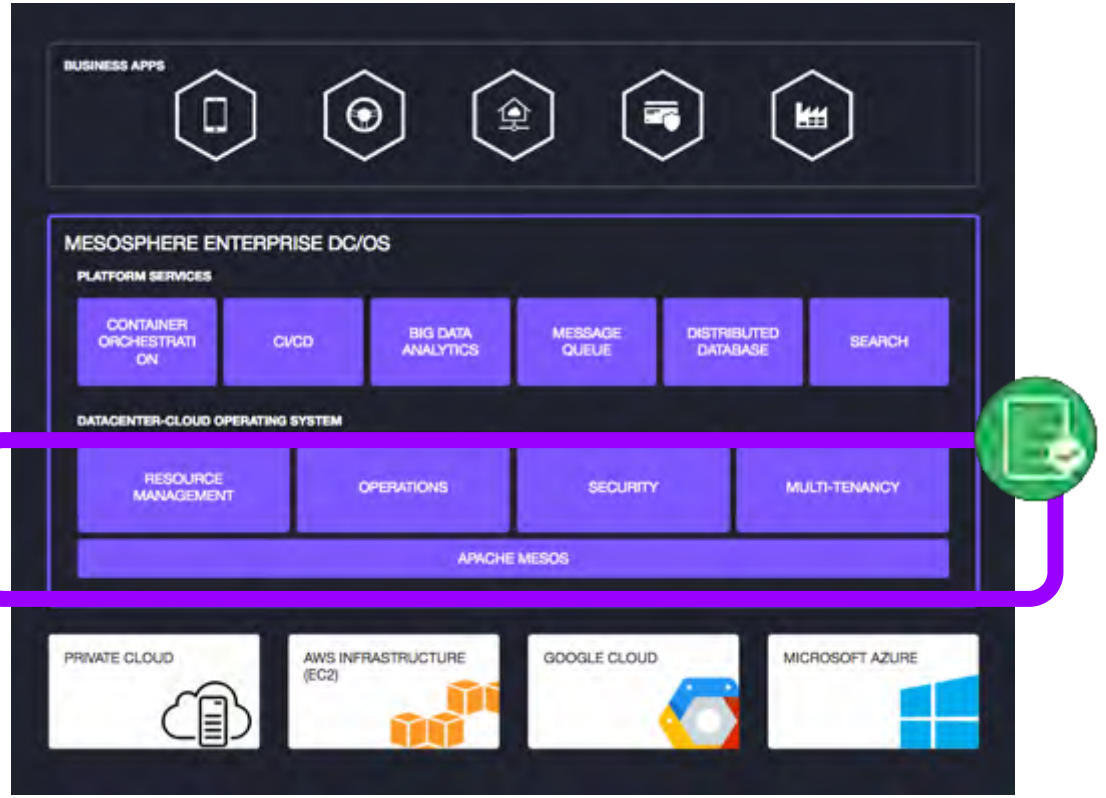

DC/OS

OTHER IMPROVEMENTS

Mesos 1.2

Marathon 1.4

Docker 1.12 and 1.13 (17.03-ce) support

Centos 7.3 and CoreOS 1235.12.0 support

Performance improvements across all networking features.

CNI support for 3rd party CNI plugins.

100s of additional bugfixes and tests

MESOSPHERE DC/OS: THE PLATFORM FOR MODERN DATA-RICH APPS

Mesosphere DC/OS provides the power of AWS-like platform services and the freedom to run on any infrastructure

New with Mesosphere DC/OS 1.9:

- **Fast-growing data services ecosystem with:** Alluxio, Couchbase, Datastax DSE, Elastic (ELK), Apache Flink, Redis; (Now 100+ services available)
- Simplified operations with **remote container shell, unified metrics and logging**
- Migrate legacy workloads with **Pods** and accelerate machine learning applications with **GPU based scheduling**
- Accelerate mission critical apps with **compliance reports**

MESOSPHERE

MESOSCON ASIA ‘17
BEIJING 6/20-22

Learn more by visiting
http://events.linuxfoundation.org/events/mesoscon-asia