

MesosCon  
NORTH AMERICA

# Day 2 Operations Best Practices

**Janet Yu**, Software Engineer, SignalFx  
**Ben Lin**, APAC Tech Lead, Mesosphere

# Agenda

- Overview
- Architecture
- Metrics API
- Demo

# Continuously Connected World

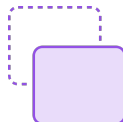
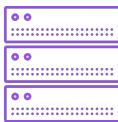


Mobile 4.4B



Internet of Things (IoT) 6B

Modern Enterprise Architecture



# App Transformation

## Traditional Enterprise Apps

### App



Monolithic packaged software (in VMs)

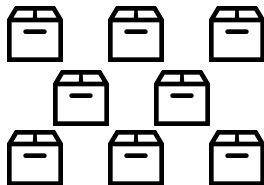
### Data



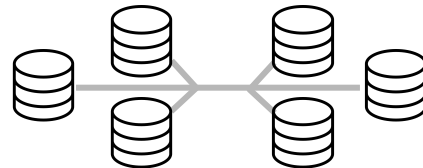
Big databases (e.g., Oracle, SQL Server)



## Modern Enterprise Apps



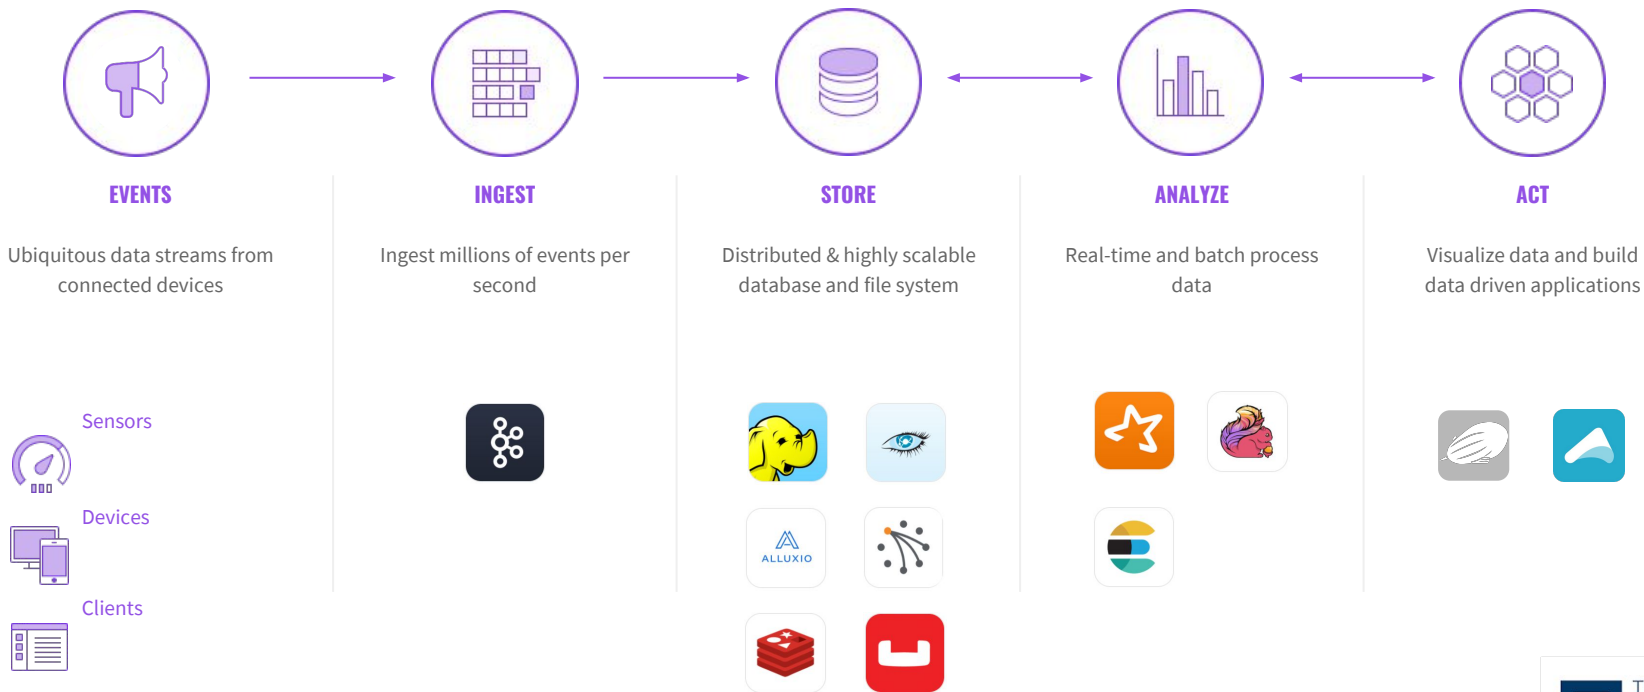
Microservices (in containers)



Cloud native data services (e.g., Spark, Kafka, Cassandra)



# Data Intensive



# Key Challenges

- Scalable Capacity
- Dynamic Architecture
- Load Balancing

# Scalable Capacity

Benefit:

Nodes added or removed, based on load

Concern:

When does it need to occur

# Dynamic Architecture

## Benefit:

One piece can be easily swapped out with another

## Concern:

Obtaining meaningful view of application as a whole when pieces can change



# Load Balancing

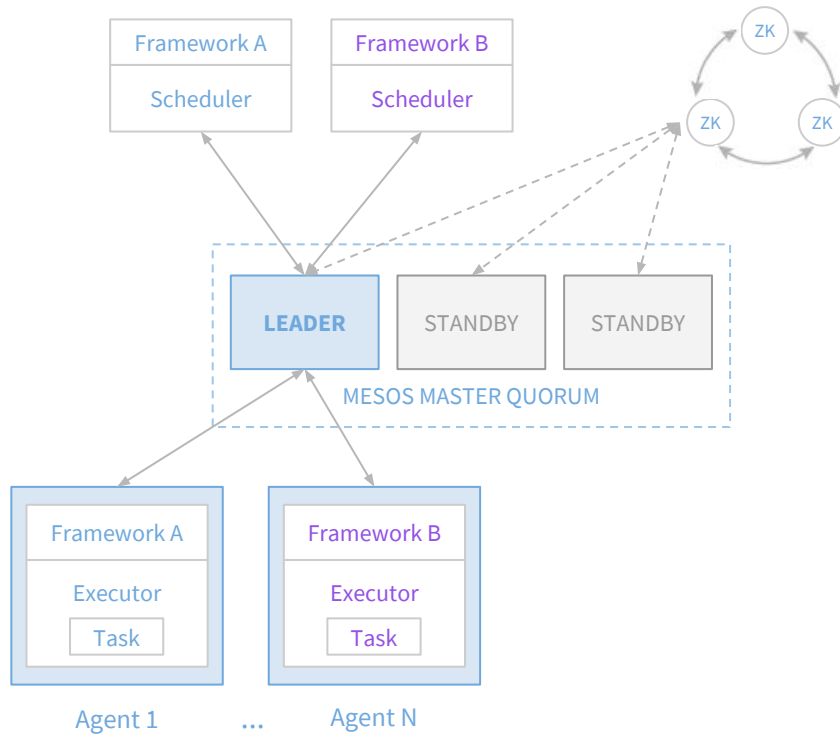
Benefit:

Work is fairly shared among resources

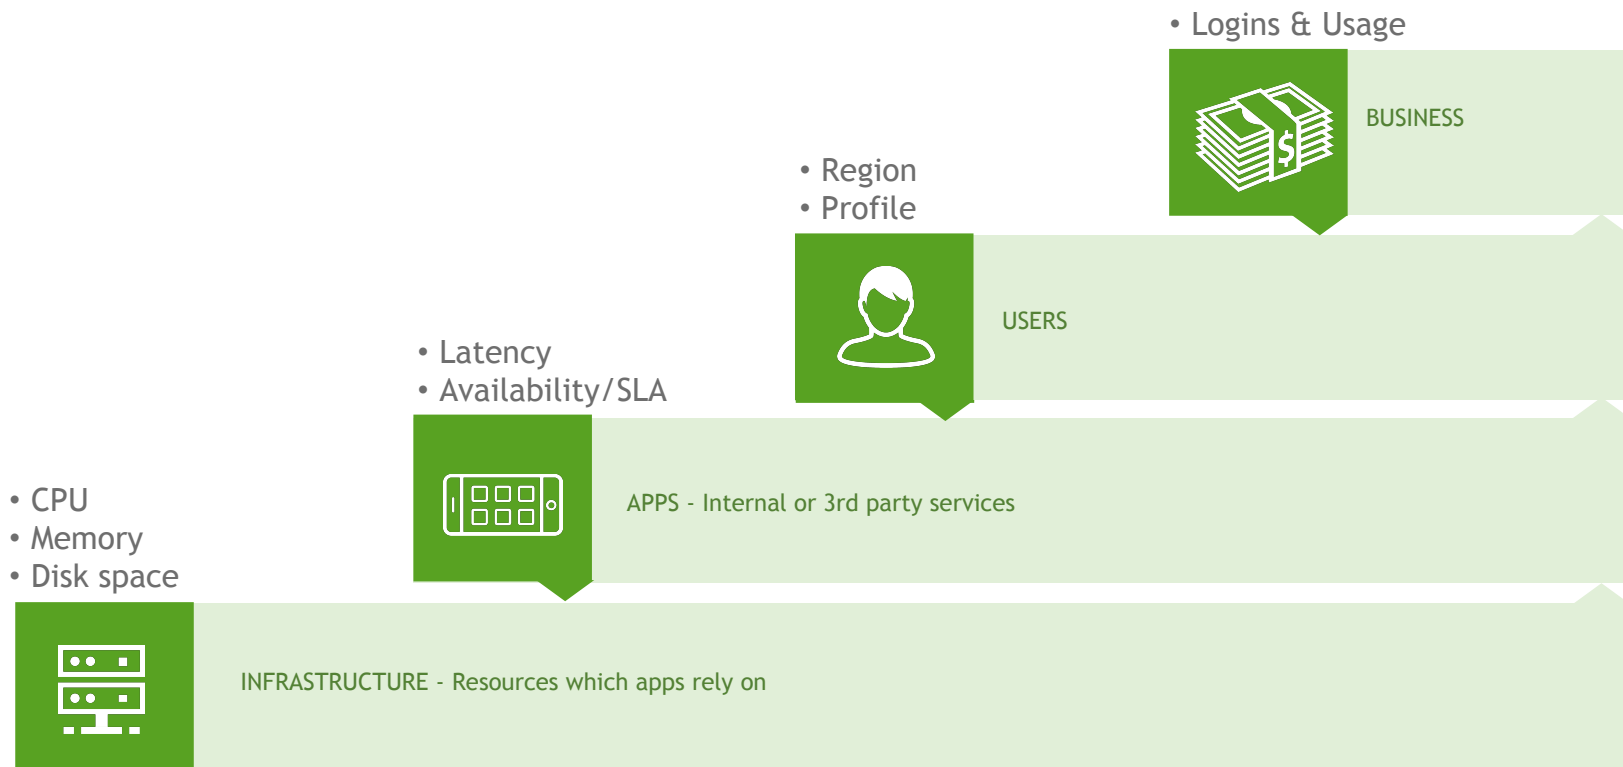
Concern:

How effective is the algorithm

# Mesos Architecture



# Metric Categories



# Metrics

Metric: Anything that is measurable and variable

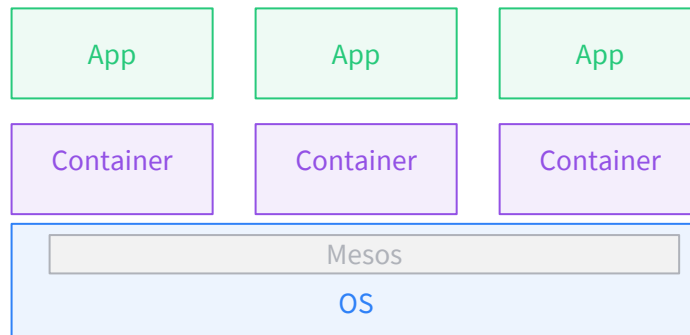
Measurements captured to determine health and performance of cluster:

- How utilized is the cluster?
- Are resources being optimally used?
- Is the system performing better or worse over time?
- Are there bottlenecks in the system?
- What is the response time of applications?



# Mesos Metric Sources

- Mesos metrics
  - Resource, frameworks, masters, agents, tasks, system, events
- Container Metrics
  - CPU, mem, disk, network
- Application Metrics
  - QPS, latency, response time, hits, active users, errors



# Master Metrics

- Metrics for the master node are available at the following URL:
  - `http://<mesos-master-ip>/mesos/master/metrics/snapshot`
  - The response is a JSON object that contains metrics names and values as key-value pairs.

- Metric Groups:

- Resources
- Master
- System
- Slaves
- Frameworks
- Tasks
- Messages
- Event Queue
- Registrar

```
1 {  
2   "allocator/event_queue_dispatches": 0,  
3   "master/cpus_percent": 0.35625,  
4   "master/cpus_revocable_percent": 0,  
5   "master/cpus_revocable_total": 0,  
6   "master/cpus_revocable_used": 0,  
7   "master/cpus_total": 16,  
8   "master/cpus_used": 5.7,  
9   "master/disk_percent": 0,  
10  "master/disk_revocable_percent": 0,  
11  "master/disk_revocable_total": 0,  
12  "master/disk_revocable_used": 0,  
13  "master/disk_total": 130164,  
14  "master/disk_used": 0,  
15  "master/dropped_messages": 2,  
16  "master/elected": 1,  
17  "master/event_queue_dispatches": 4,  
18  "master/event_queue_http_requests": 0,  
19  "master/event_queue_messages": 0,
```

# Master Basic Alerts

Metric Value	Inference
master/uptime_secs is low	The master has restarted
master/uptime_secs < 60 for sustained periods of time	The cluster has a flapping master node
master/tasks_lost is increasing rapidly	Tasks in the cluster are disappearing. Possible causes include hardware failures, bugs in one of the frameworks or bugs in Mesos
master/slaves_active is low	Slaves are having trouble connecting to the master
master/cpus_percent > 0.9 for sustained periods of time	DCOS Cluster CPU utilization is close to capacity
master/mem_percent > 0.9 for sustained periods of time	DCOS Cluster Memory utilization is close to capacity
master/disk_used & master/disk_percent	DCOS Disk space consumed by Reservations
master/elected is 0 for sustained periods of time	No Master is currently elected

# Agent Metrics

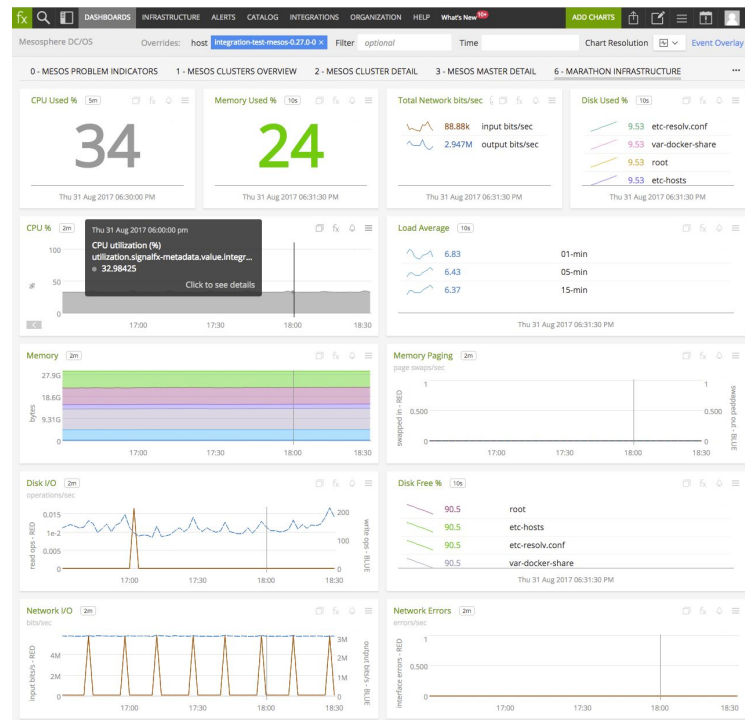
- Metrics for the agent node are available at the following URL:  
`http://<mesos-agent-ip>:5051/metrics/snapshot`
  - The response is a JSON object that contains metrics names and values as key-value pairs.
- Metric groups:
  - Resources
  - Slave
  - System
  - Executors
  - Tasks
  - Messages

```
1 {
2   "containerizer/mesos/container_destroy_errors": 0,
3   "containerizer/mesos/provisioner/bind/remove_rootfs_errors": 0,
4   "containerizer/mesos/provisioner/remove_container_errors": 0,
5   "slave/container_launch_errors": 0,
6   "slave/cpus_percent": 0.7,
7   "slave/cpus_revocable_percent": 0,
8   "slave/cpus_revocable_total": 0,
9   "slave/cpus_revocable_used": 0,
10  "slave/cpus_total": 4,
11  "slave/cpus_used": 2.8,
12  "slave/disk_percent": 0.281119982008321,
13  "slave/disk_revocable_percent": 0,
14  "slave/disk_revocable_total": 0,
15  "slave/disk_revocable_used": 0,
16  "slave/disk_total": 35572,
17  "slave/disk_used": 10000,
18  "slave/executor_directory_max_allowed_age_secs": 151040.386469261,
19  "slave/executors_preempted": 0,
20  "slave/executors_registering": 0,
```



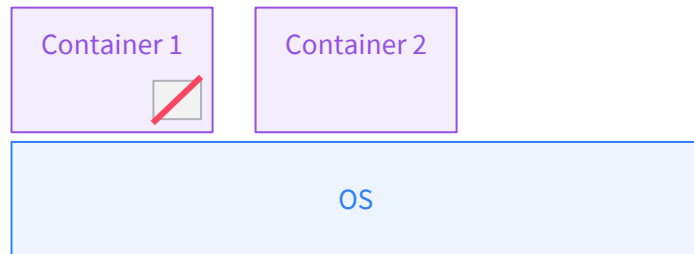
# Marathon Metrics

- Metrics for Marathon are available at the following URL:
  - `http://<marathon-ip>:8080/metrics`
  - for DC/OS `http://<master-ip>:/marathon/metrics`
- Redirect metrics to graphite when you start the Marathon process by adding the following flag: `--reporter_graphite tcp://<graphite-server>:2003?prefix=marathon-test&interval=10`



# Container Level Metrics

- Monitoring agent per container?
  - Not scalable
  - Increased footprint



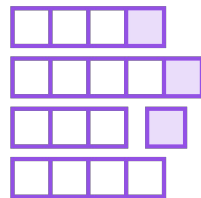
# Mesos Metrics Module

## Simplified config

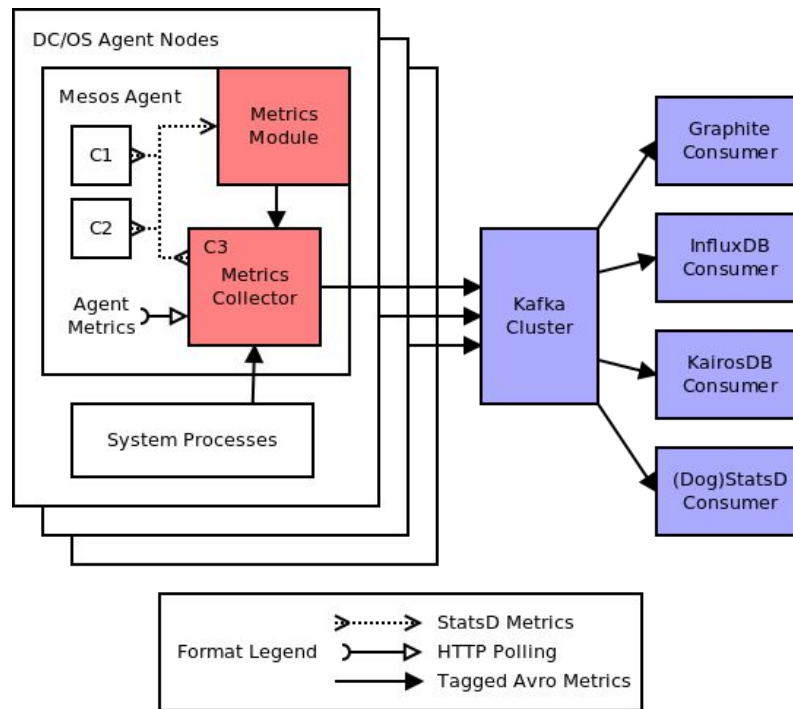
- Container metrics (automated)
- Application metrics (statsd env vars)

## Context injection

- Automated source tagging (container, agents, ...)



# Metrics API Architecture



# Metrics API

Poll for data about cluster, hosts, containers, applications

```
GET http://<cluster>/system/v1/agent  
/<agent_id>/metrics/v0/<resource_path>  
Accept: application/json  
Authorization: token=<token_string>
```

# Metrics API Response

```
"datapoints": [  
  {  
    "name": "processes",  
    "value": 209,  
    "unit": "",  
    "timestamp": "2017-08-31T01:00:19Z"  
  },  
  ...  
],  
"dimensions": {  
  "mesos_id": "a29070cd-2583-4c1a-969a-3e07d77ee665-S0",  
  "hostname": "10.0.2.255"  
}
```

# Metrics API Tips

- Get authentication token

```
POST http://<cluster>/acs/api/v1/auth/login  
{ "username": "<user>", "password": "<pw>" }
```

- Datapoint timestamp format may vary

```
2017-09-01T00:25:23.502867353Z,  
2017-09-01T06:25Z
```

- Error check datapoint value type

```
{u'timestamp': u'2017-09-06T21:07:03Z',  
u'unit': u'', u'name':  
u'org.apache.cassandra.metrics.Table.ReadLatency  
.system.peer_events.mean', u'value': u'NaN'}
```

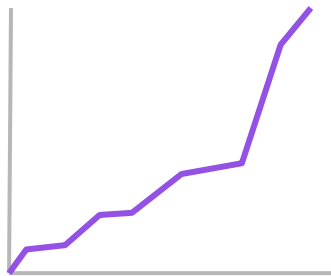
# Datapoint

Single reported value of a metric from a particular source at a particular time

- Metric name
- Value
- Timestamp
- Metric type
- Dimensions



# Metric Types



**Counters**

Discrete events that are monotonically increasing.

- # of failed tasks
- # of agent registrations



**Gauges**

An instantaneous sample of some magnitude.

- % of used memory in cluster
- # of connected slaves

# Dimensions

- Key/value pairs
- Set of dimensions represents the source of a datapoint
- Correlates related datapoints, patterns
- Enables classification, aggregation, filtering

# Metrics vs. Dimensions

## Metrics

- *cpu.idle*
- *disk\_ops.write*
- *load*
- *memory.free*
- *memory.used*
- *lf\_octets.rx*

## Dimensions

- *Division*
- *Datacenter*
- *Region*
- *Browser Type*
- *Machine type*

# Metric + Dimensions = Time Series



Figure 1

cpu.idle

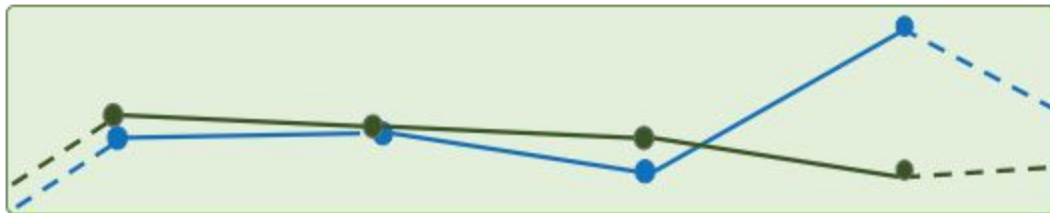


Figure 2

cpu.idle



datacenter: eastdc  
datacenter: westdc



# Tips for Sending Metrics

- Structure names hierarchically
- Use a single, consistent delimiter for wildcard searches
- Separate dimensions from metric names
- Don't use dimensions with high cardinality
  - Timestamps, task ids
- Don't send metric type as a dimension
  - Gauges average, counters summed

# Monitoring

Send data to monitoring app for analysis

POST <https://ingest.signalfx.com>

Content-Type: application/json

X-SF-TOKEN: <token\_string>

```
{ "gauges": [{  
    "metric": "processes",  
    "dimensions": { "host": "10.0.2.255", ...},  
    "value": 209}, ...]], ...}
```

# DEMO

# Key Takeaways

- Scalable Capacity
  - Collect system and custom **metrics**, find outliers that might be bottlenecks
- Dynamic Architecture
  - Use **dimensions** common across all related pieces vs. tracking per-instance identifier
- Load Balancing
  - Compare **time series**, calculate ratios



# Resources

Visit the SignalFx and Mesosphere booths :)

- <http://mesos.apache.org/documentation/latest/monitoring/>
- <https://mesosphere.github.io/marathon/docs/metrics.html>
- <https://dcos.io/docs/1.9/metrics/metrics-api/>
- <https://developers.signalfx.com/docs/signalfx-api-overview>
- <https://github.com/signalfx/collectd-mesos>

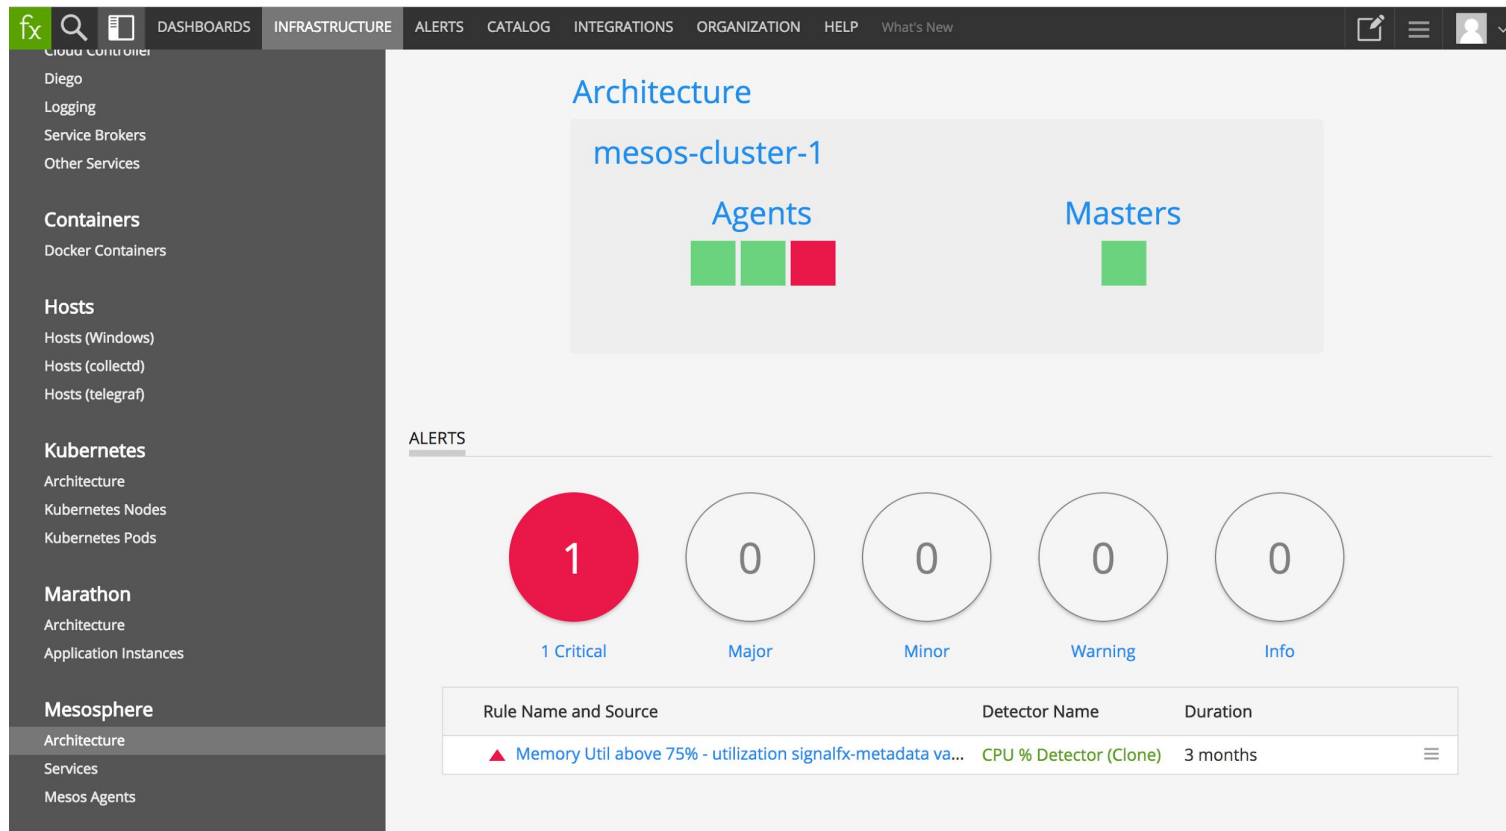


# BACKUP SLIDES

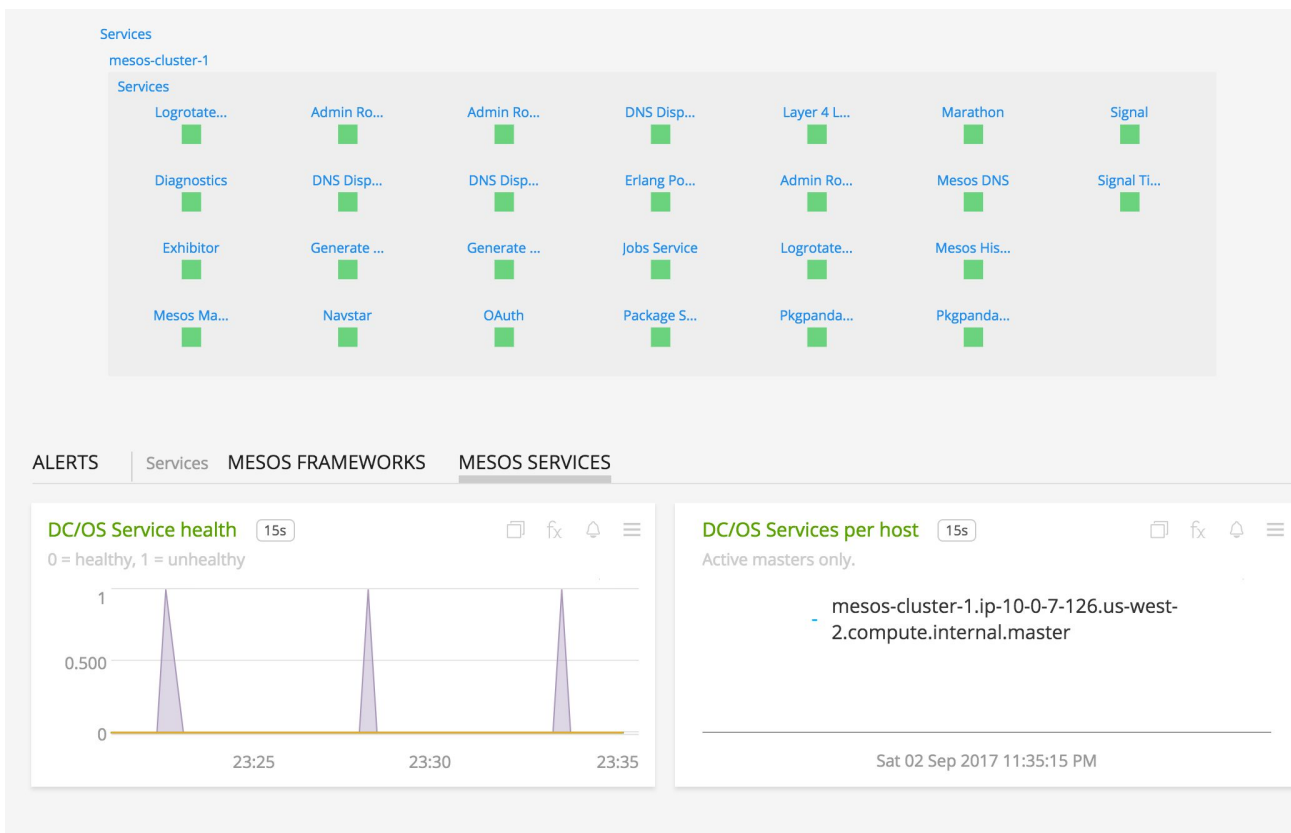
# Logging

# Troubleshooting

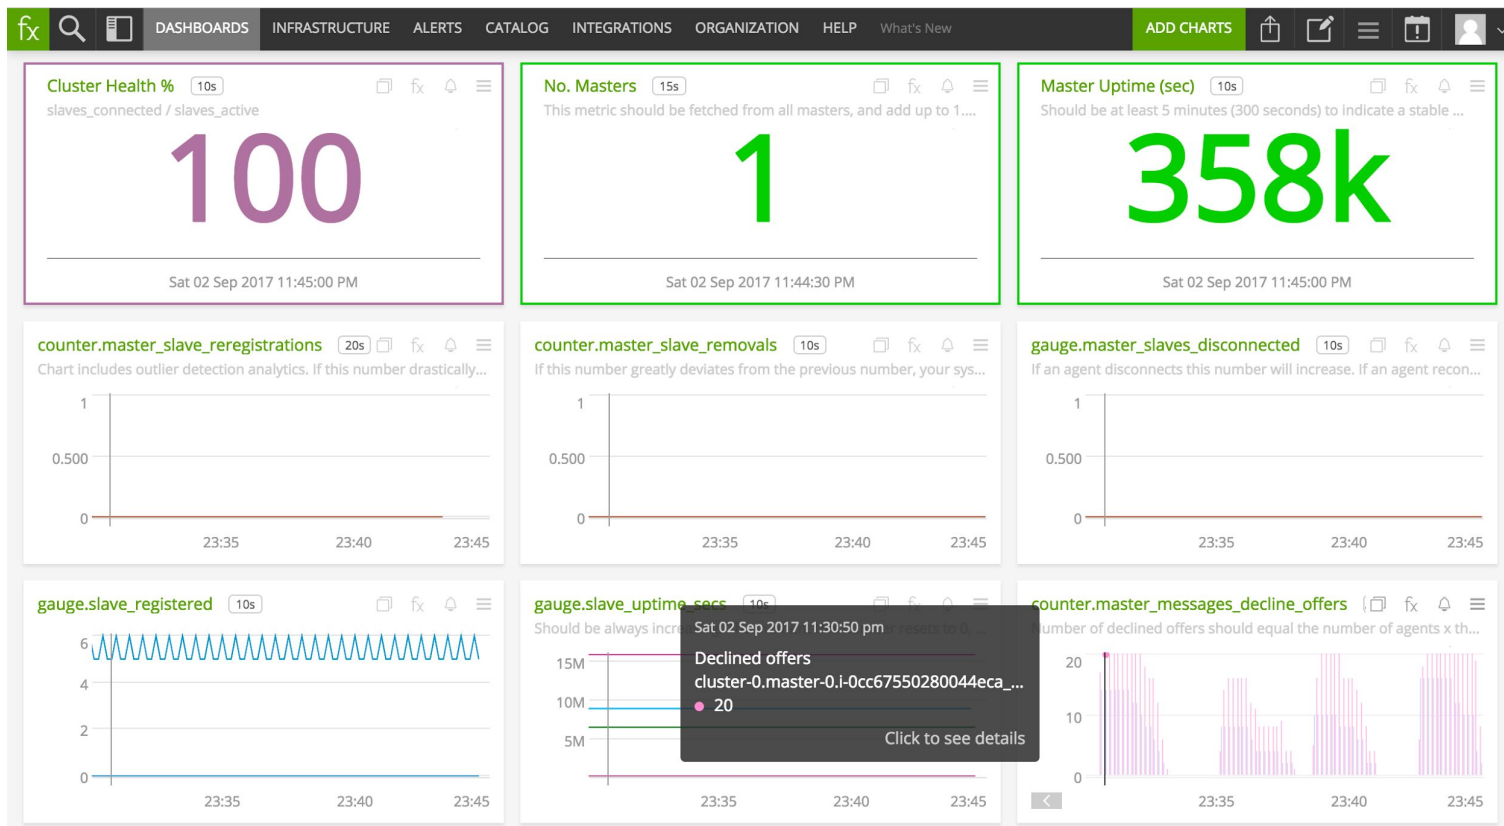
# Infrastructure Outliers



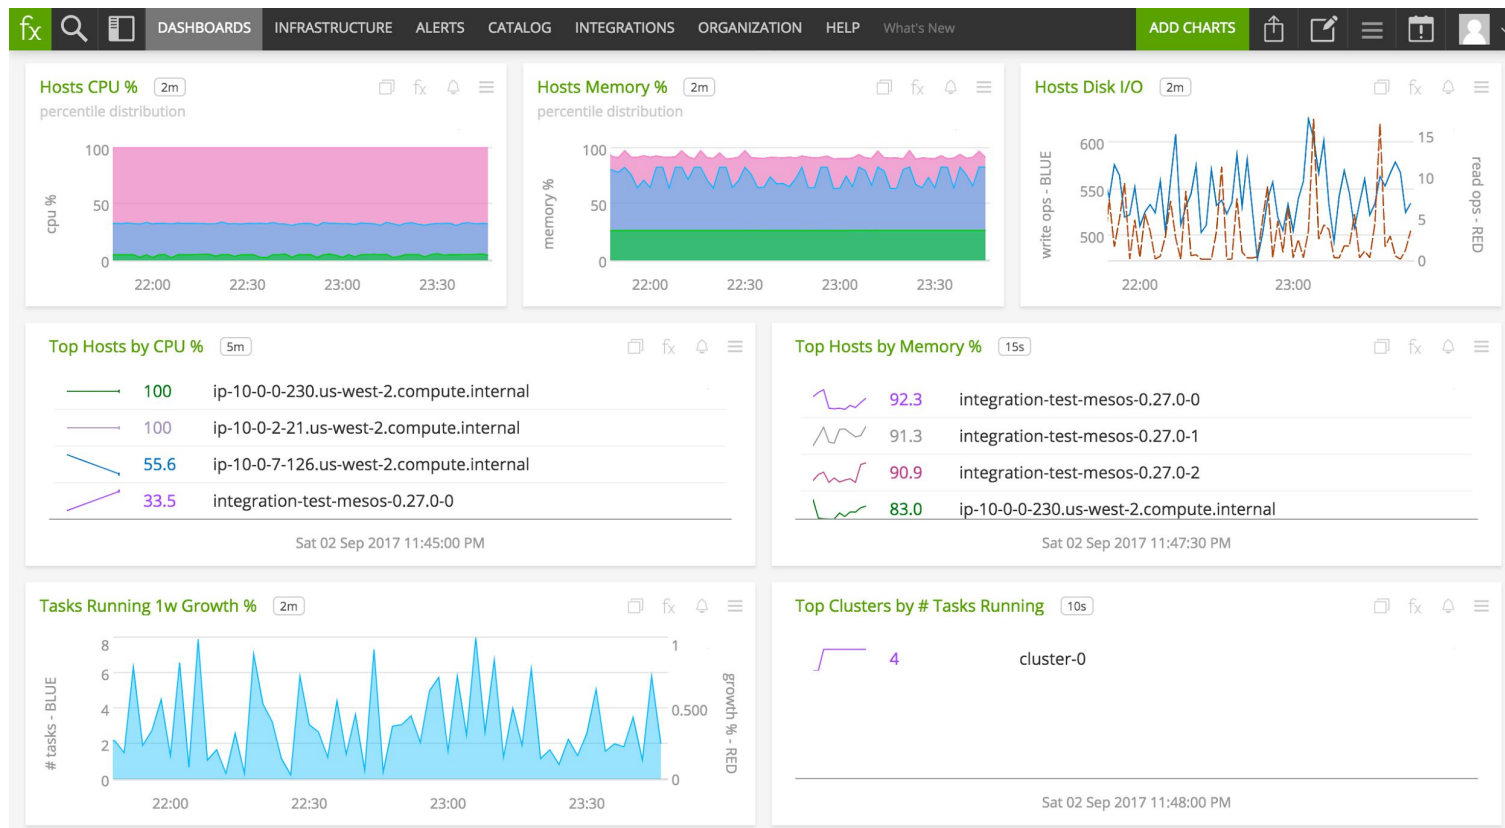
# Service Health



# Problem Indicators

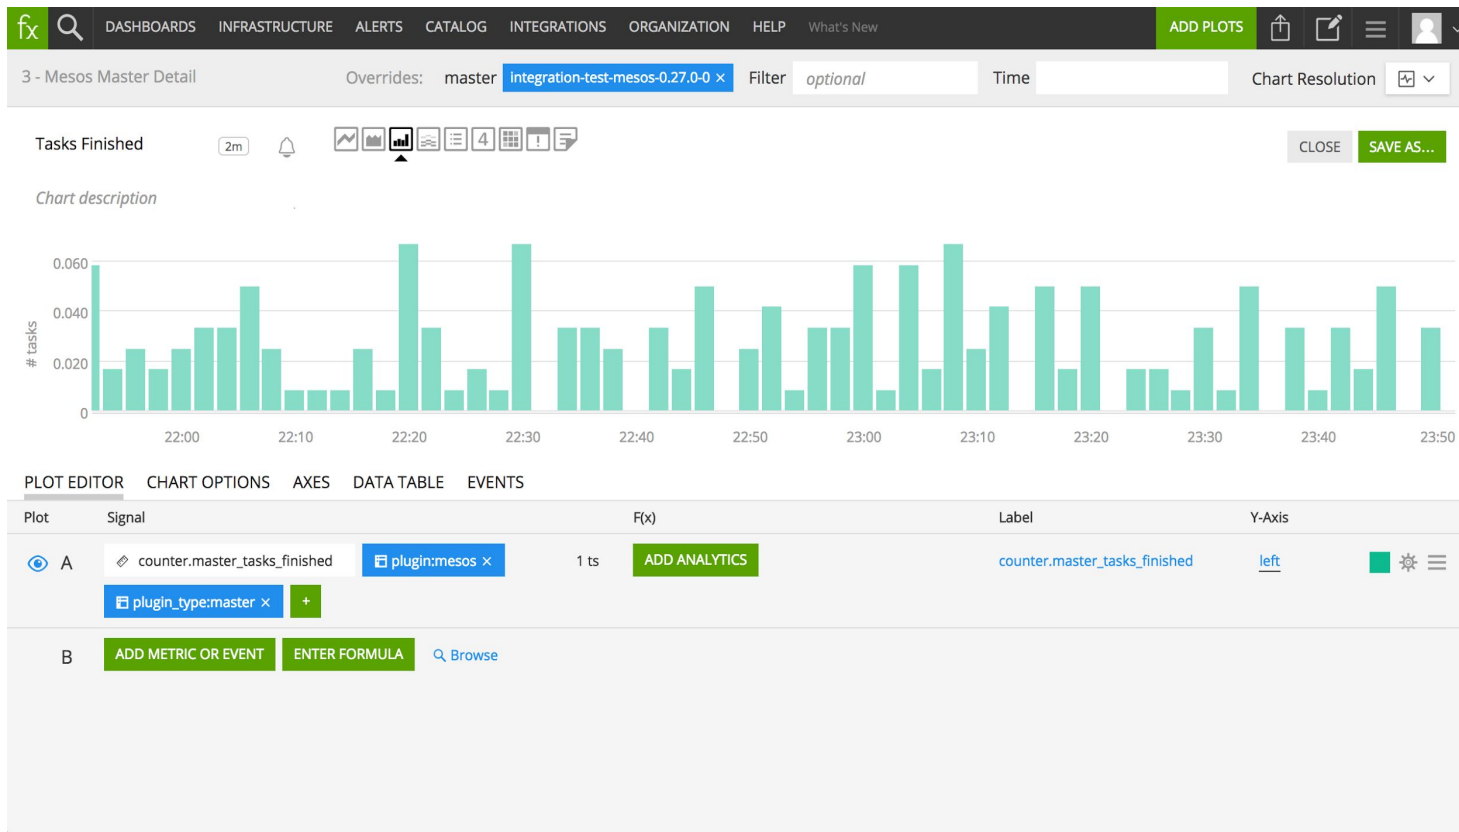


# Cluster Trends





# Filtering by Dimension



# Inputs / Outputs

Input: **StatsD**

- Text records: either one-per-packet or newline separated.
- Optional tagging

**memory.usage\_mb:5**|g

**frontend.query.latency\_ms:46**|g|#shard\_id:6,section:frontpage

Pseudocode:

```
if (env["STATSD_UDP_HOST"] and env["STATSD_UDP_PORT"]) {  
    // 1. Open UDP socket to the endpoint  
    // 2. Send StatsD-formatted metrics  
}
```

Output: **Apache Avro**

# Marathon App Performance

\$ curl <leader.mesos>/marathon/v2/apps/sleep | jq .

- Find the appld (sleep), "host", and "id" (task ID) fields

```
"tasks": [  
  {  
    "id": "sleep.cb536c16-c6cf-11e5-a84d-0a43d276f399",  
    "host": "10.0.3.226",  
    "ports": [  
      10466  
    ],  
    "startedAt": "2016-01-29T21:32:28.443Z",  
    "stagedAt": "2016-01-29T21:32:27.644Z",  
    "version": "2016-01-29T21:32:27.599Z",  
    "slaveId": "caa0847c-3751-456f-a2fd-30feb7a1fda5-S1",  
    "appId": "/sleep"  
  }  
]
```

# Marathon App Performance

Curl the Agent host and look for the Marathon Task ID from previous step

```
$ curl http://<agent-internal-IP>:5051/monitor/statistics | jq .
```

```
{
  "executor_id": "sleep.cb536c16-c6cf-11e5-a84d-0a43d276f399",
  "executor_name": "Command Executor (Task:
sleep.cb536c16-c6cf-11e5-a84d-0a43d276f399) (Command: sh -c 'env && sleep...')",
  "framework_id": "caa0847c-3751-456f-a2fd-30feb7a1fda5-0000",
  "source": "sleep.cb536c16-c6cf-11e5-a84d-0a43d276f399",
  "statistics": {
    "cpus_limit": 0.2,
    "cpus_system_time_secs": 0,
    "cpus_user_time_secs": 0.01,
    "mem_limit_bytes": 50331648,
    "mem_rss_bytes": 200704
  }
}
```