
Dublin Apache Kafka Meetup, 30 August 2017

The SMACK Stack:

Spark*, Mesos*, Akka, Cassandra*, Kafka*



Elizabeth K. Joseph
@pleia2

* ASF projects

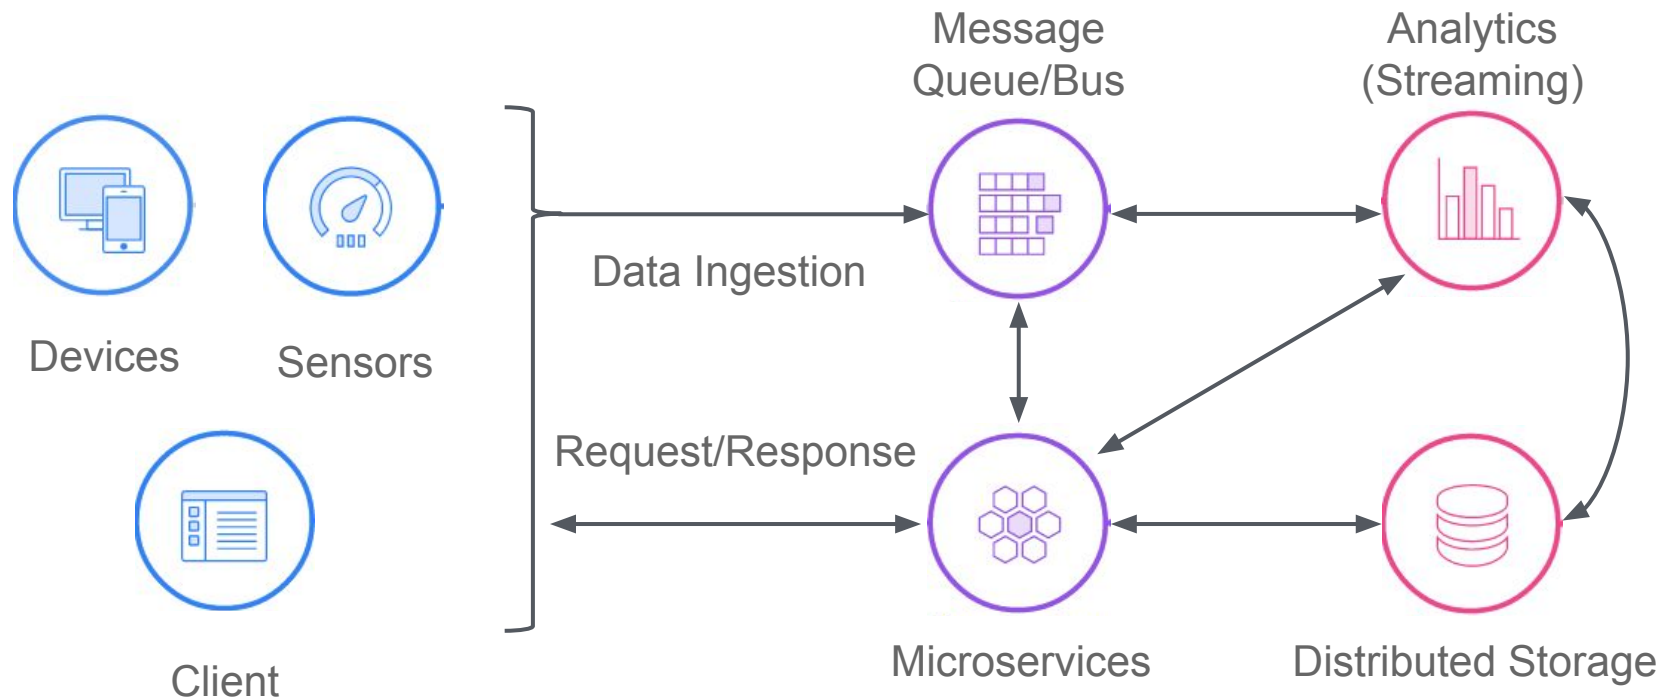


Elizabeth K. Joseph, Developer Advocate

- ❑ 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



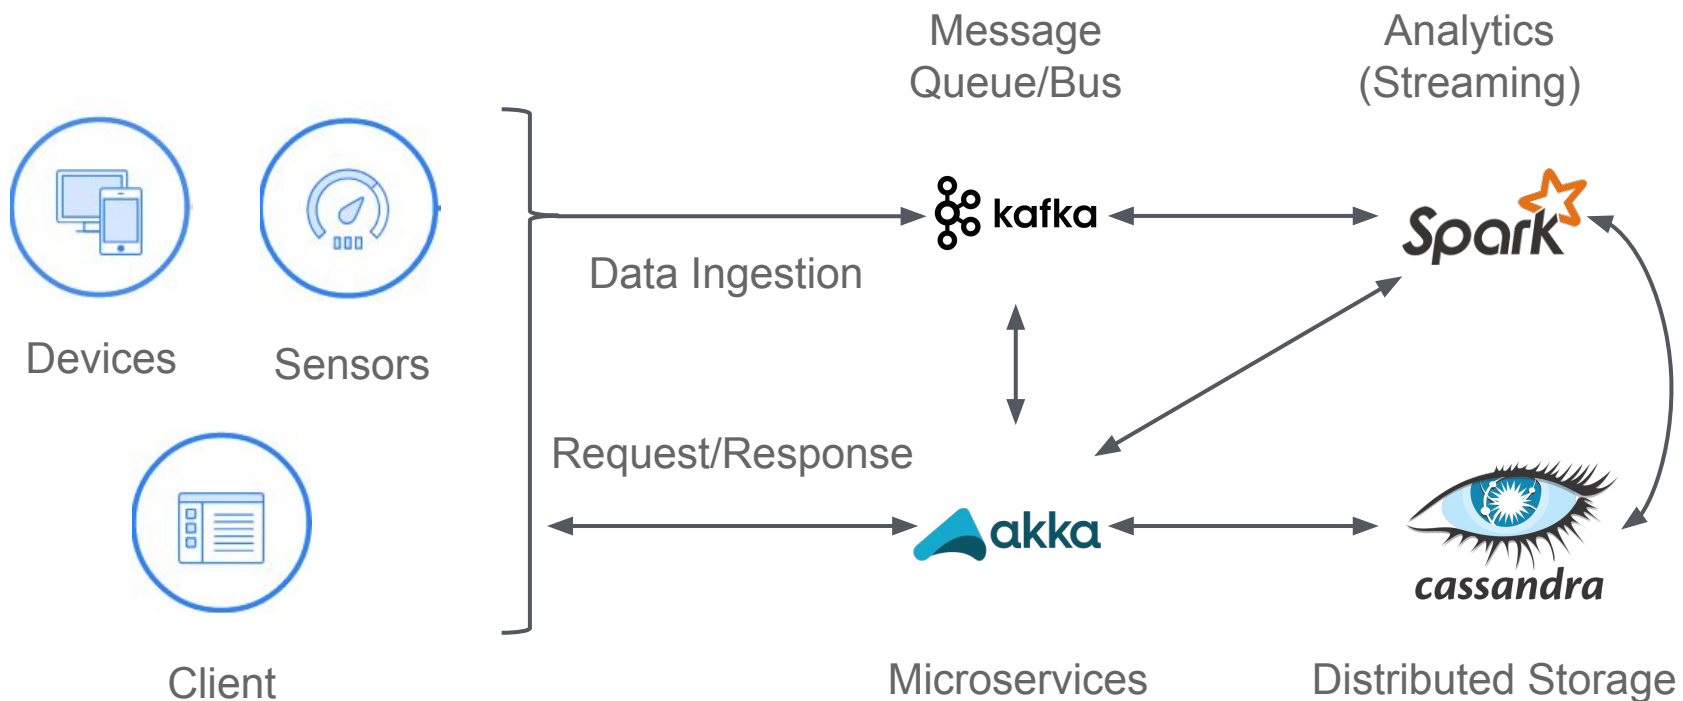
MODERN APPLICATION -> FAST DATA BUILT-IN



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

SMACK Stack

P

A

R

K

STREAMING ANALYTICS

Microbatching


- Apache Spark (Streaming)

Native Streaming

- Apache Flink
- Apache Storm/Heron
- Apache Apex
- Apache Samza



APACHE SPARK (STREAMING)




Typical Use: distributed, large-scale data processing; micro-batching

Why Spark Streaming?

- Micro-batching creates very low latency, which can be faster
- Well defined role means it fits in well with other pieces of the pipeline

SMACK Stack

E
S
O
S



Apache Mesos: The datacenter kernel

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

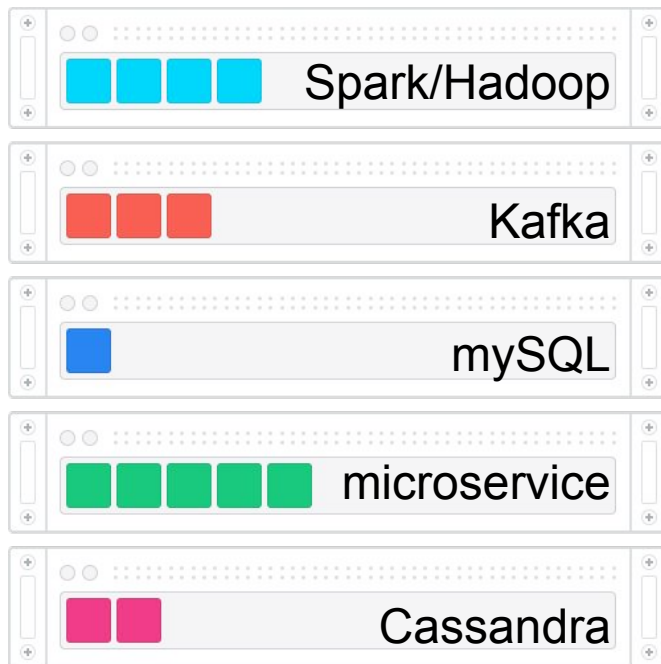
Building block of the modern internet

- 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




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

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 12

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.
- Mesos mechanics are fair and HA
- 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



DC/OS Architecture Overview


Services & Containers



DC/OS



ANY INFRASTRUCTURE



DC/OS brings it all together

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



DC/OS

DC/OS is ...



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

Interact with DC/OS (1/2)

Web-based GUI

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

The screenshot displays the DC/OS web-based GUI. The left sidebar shows the navigation menu for 'dcos-cluster-2' (user: joel@mesosphere.io). The 'Services' menu item is highlighted. The main content area shows the 'Services' page with a search filter and a table of running services.

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

 kafka
1.1.19.1-0.10.1.0

service
brokers
executor
kafka

CPUS * ?
1

MEM * ?
2304

heap
The Kafka process JVM heap configuration object

SIZE * ?
2048

DISK * ?
5000

DISK_TYPE ?
ROOT

COUNT * ?
3

BACK

REVIEW AND INSTALL

Application JSON

```
{
  "service": {
    "name": "kafka",
    "user": "root",
    "principal": "kafka-principal",
    "placement_strategy": "NODE",
    "phase_strategy": "INSTALL",
    "enable_replacement": false,
    "recover_in_place_grace_period_secs": 1200,
    "min_delay_between_recoverys_secs": 600,
    "enable_health_check": false,
    "health_check_delay_sec": 15,
    "health_check_interval_sec": 10,
    "health_check_timeout_sec": 20,
    "health_check_grace_period_sec": 10,
    "health_check_max_consecutive_failures": 3
  },
  "brokers": {
    "cpus": 1,
    "mem": 2304,
    "disk": 5000,
    "disk_type": "ROOT",
    "count": 3,
    "port": 0,
    "heap": {
      "size": 2048
    }
  },
  "jmx": {
    "enable": false,
    "remote": false,
    "remote_port": 9999,
    "remote_registry_ssl": false,
    "remote_ssl": false,
    "remote_authenticate": false,
    "remote_ssl_need_client_auth": false
  },
  "statsd": {
    "port": 0
  }
},
"executor": {
  "cpus": 0.5,
  "mem": 256,
  "disk": 0
},
"kafka": {
  "kafka_zookeeper_uri": "master.mesos:2181",
  "kafka_advertise_host_ip": true,

```

SMACK Stack

K

K

A

Akka-driven applications

Akka is a toolkit for building highly concurrent, distributed, and resilient message-driven applications for Java and Scala.

- Simple
- Highly Performant
- Elastic
- Reactive

SMACK Stack

A
S
S
A
N
D
R
A

DISTRIBUTED STORAGE

NoSQL

- ArangoDB
- mongoDB
- Apache Cassandra
- Apache HBase

SQL

- MemSQL

Filesystems

- Quobyte
- HDFS

Time-Series Datastores

- InfluxDB
- OpenTSDB
- KairosDB
- Prometheus

see also iot-a.info



APACHE CASSANDRA



Typical Use: No-dependency, time series database

Why Cassandra?

- A top level Apache project born at Facebook and built on Amazon's Dynamo and Google's BigTable
- Offers continuous availability, linear scale performance, operational simplicity and easy data distribution

SMACK Stack

A

F

K

A

MESSAGE QUEUES



Message Brokers

- Apache Kafka
- ØMQ, RabbitMQ, Disque

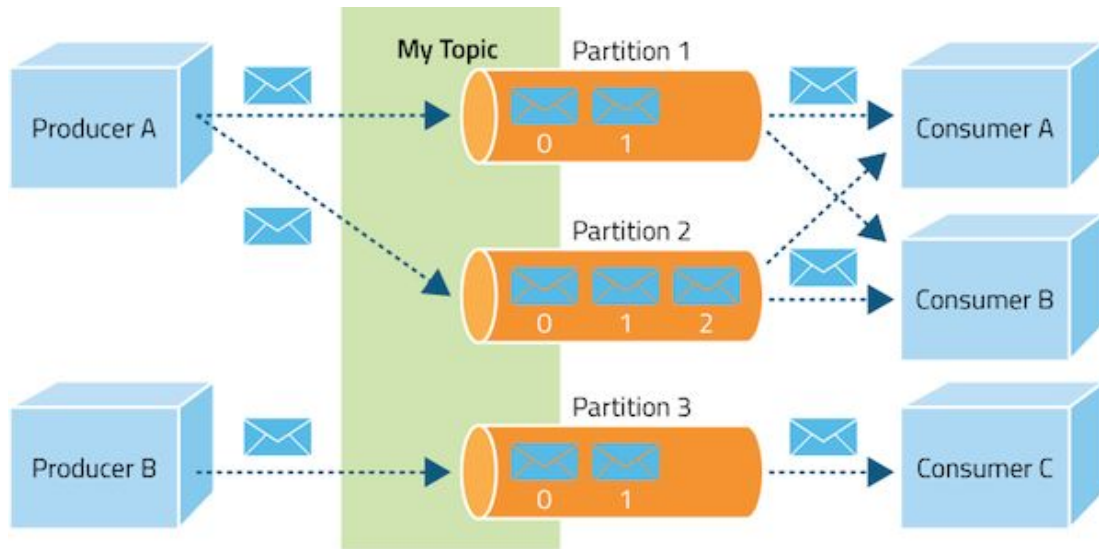
Log-based Queues

- fluentd, Logstash, Flume

see also queues.io



APACHE KAFKA



Typical Use: A reliable buffer for stream processing

Why Kafka?

- High-throughput, distributed, persistent publish-subscribe messaging system
- Created by LinkedIn; used in production by 100+ web-scale companies [1]

[1] <https://cwiki.apache.org/confluence/display/KAFKA/Powered+By>

DELIVERY GUARANTEES

- **At most once**—Messages may be lost but are never re-delivered
- **At least once**—Messages are never lost but may be redelivered (Kafka)
- **Exactly once**—Messages are delivered once and only once (this is what everyone actually wants, but it's tricky)

Murphy's Law of Distributed Systems:

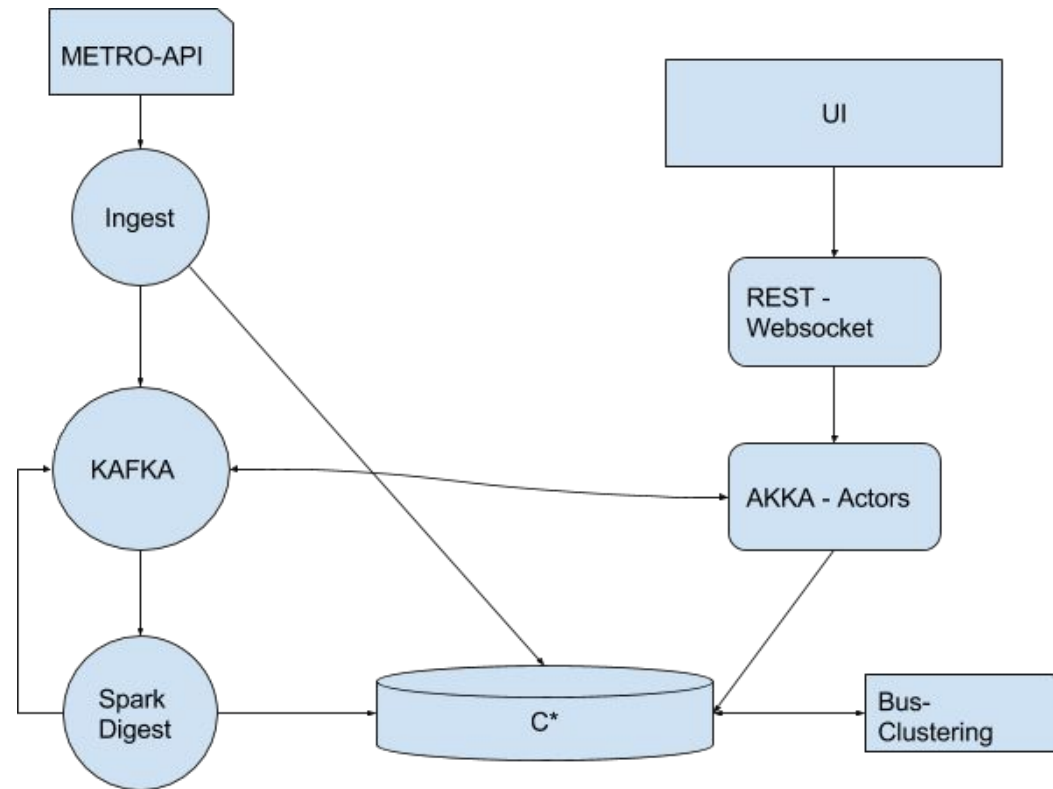
Anything that can go wrong, will go wrong ... partially!

Key Features of Kafka on DC/OS

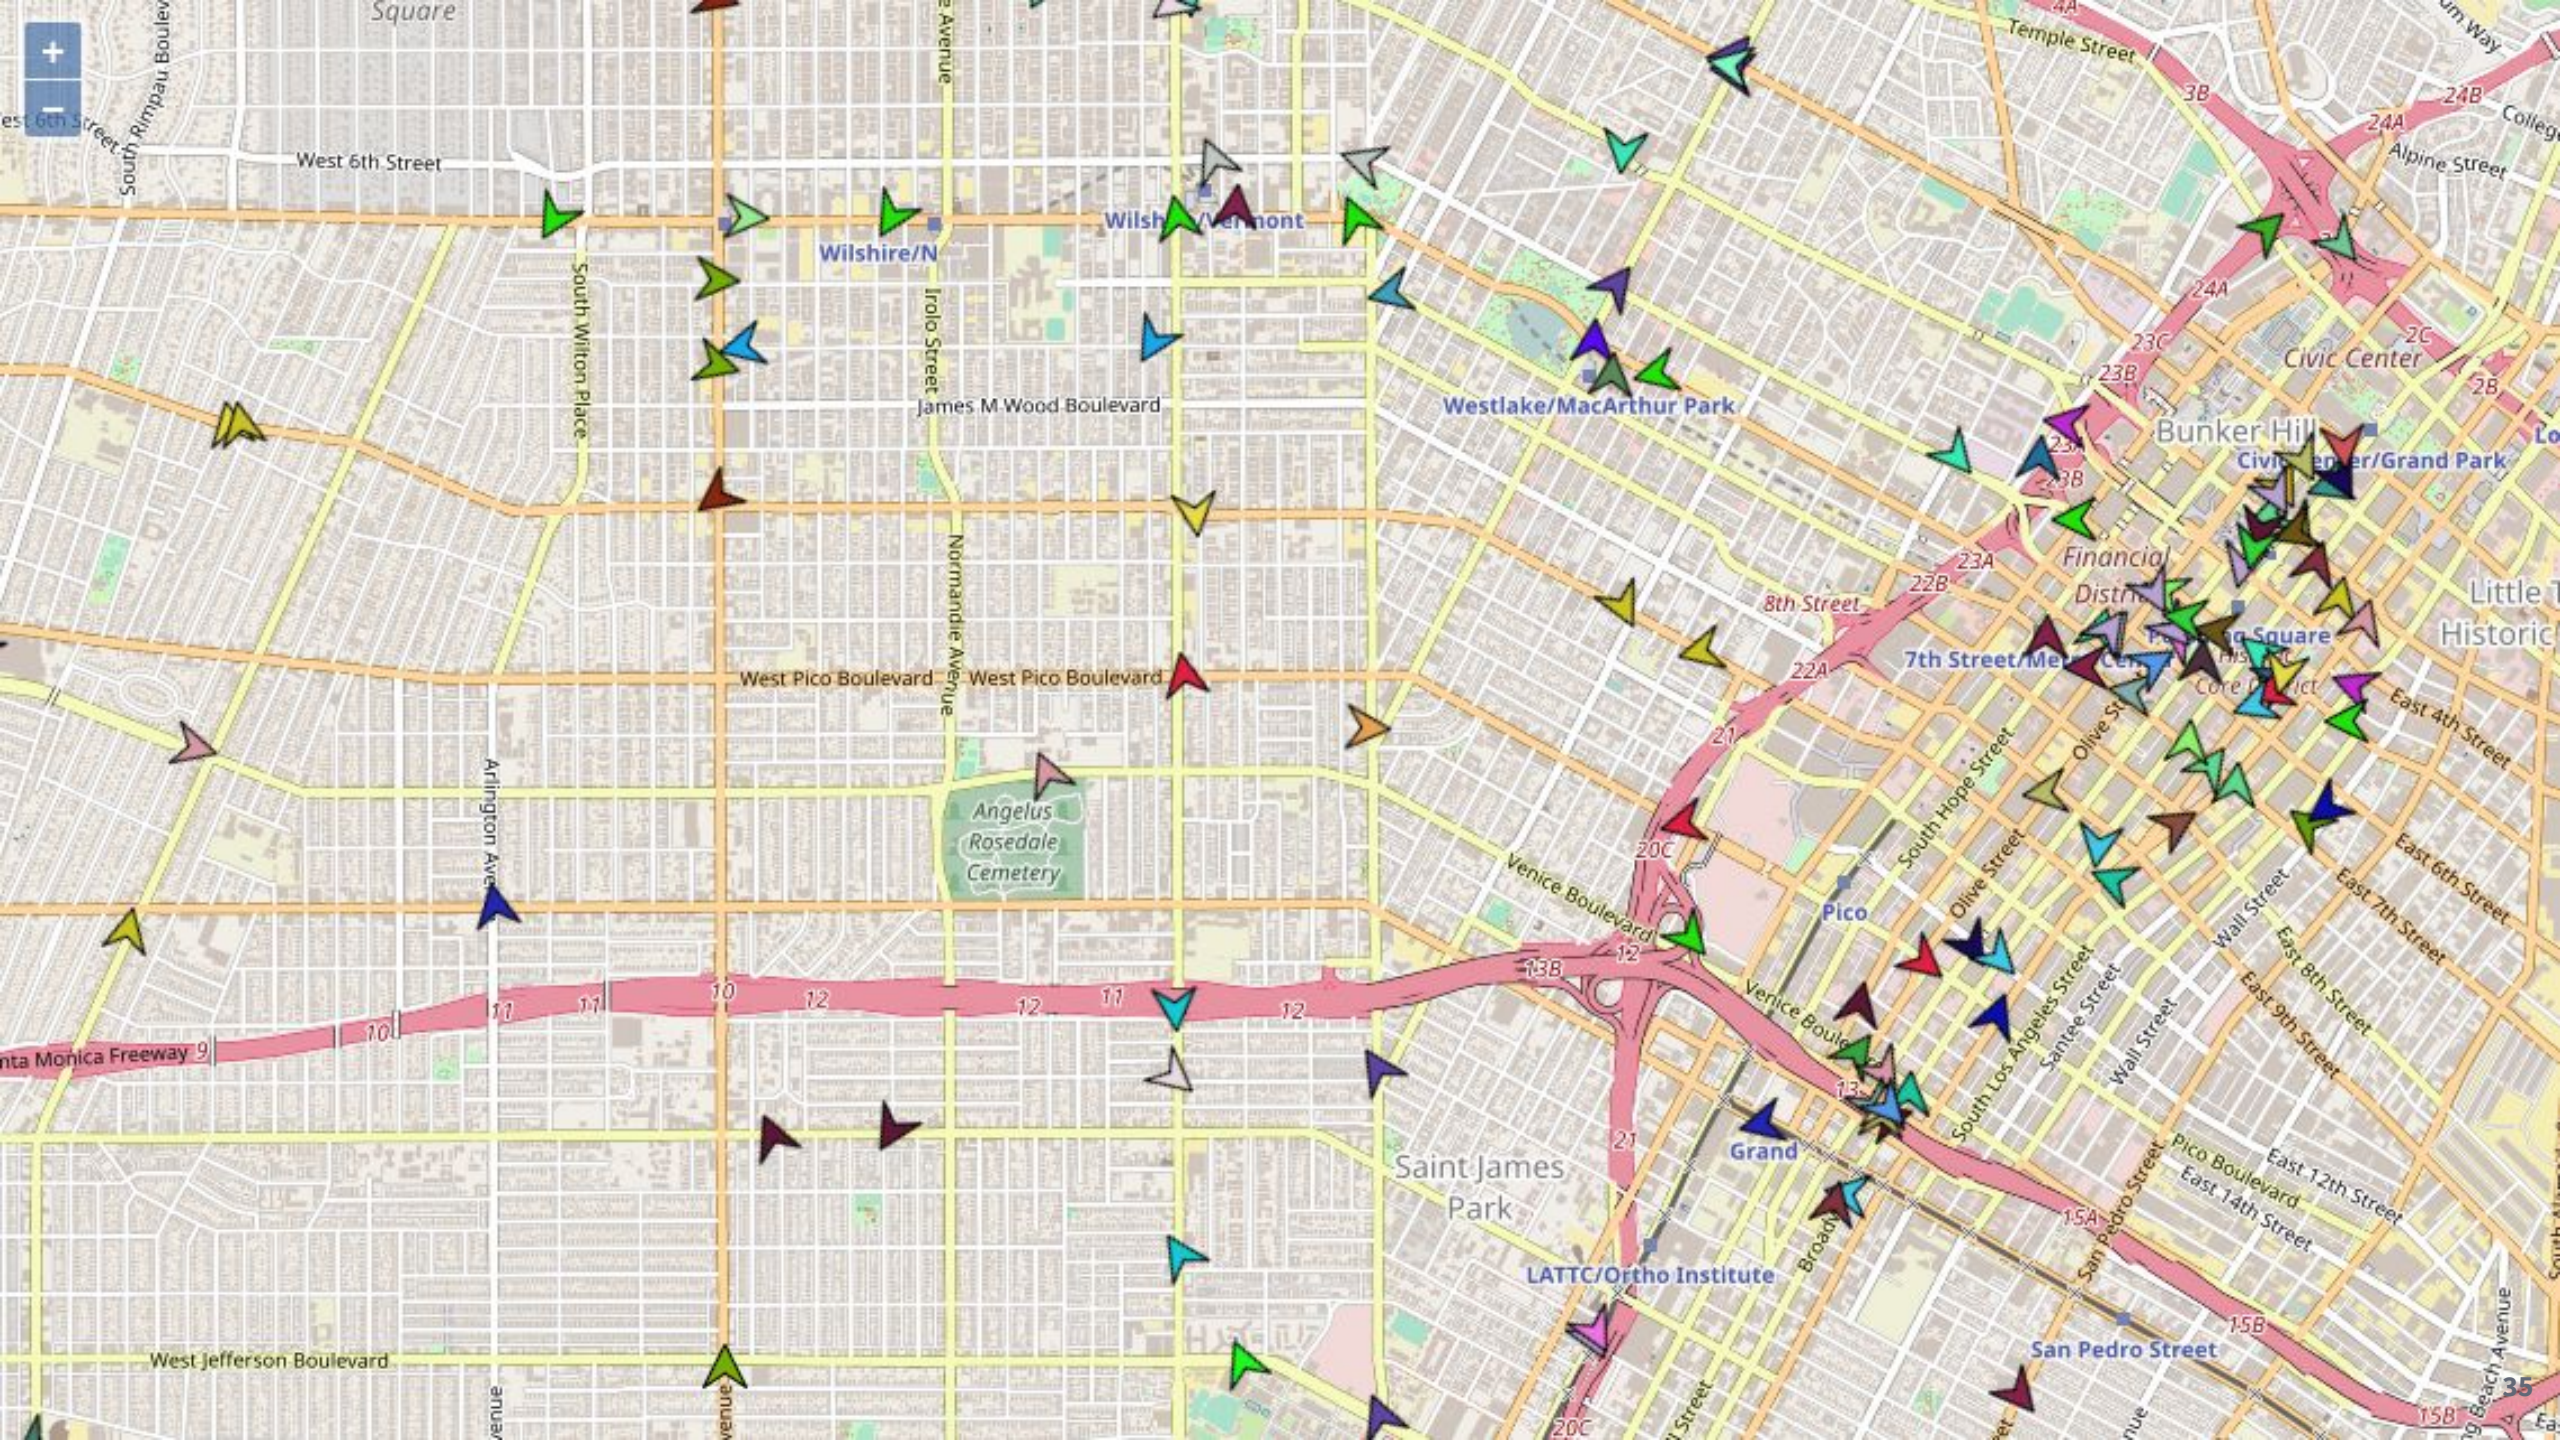
- Easy Installation
- Service discovery within the cluster
- Ability to run multiple Kafka clusters
- Elastic scaling of brokers
- Cluster and broker monitoring
- Support in the DC/OS CLI command

Learn more at <https://docs.mesosphere.com/service-docs/kafka/>

SMACK Stack Demo: Los Angeles Metro



Available for you to try at: <https://github.com/dcos/demos/tree/master/fastdata-iot/>



Questions?



@dcos



chat.dcos.io



users@dcos.io



/dcos

/dcos/examples

/dcos/demos

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