Resource Optimization with Charmander

Containers, Orchestration, Metrics

MesosCon 2015

Marcel Neuhausler, AT&T Foundry, Palo Alto CA

Motivation

Resource Optimization: Most Cost-Effective Match between a Client Request and Available Resources.

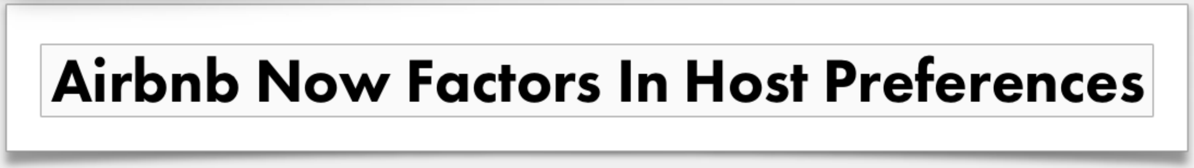

Airbnb Now Factors In Host Preferences

Resource Optimization: The Key Intellectual Property of a “Logistics” company.

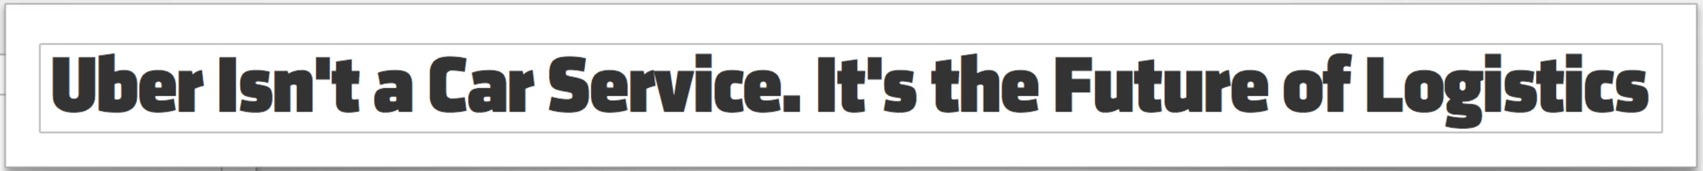

Uber Isn't a Car Service. It's the Future of Logistics

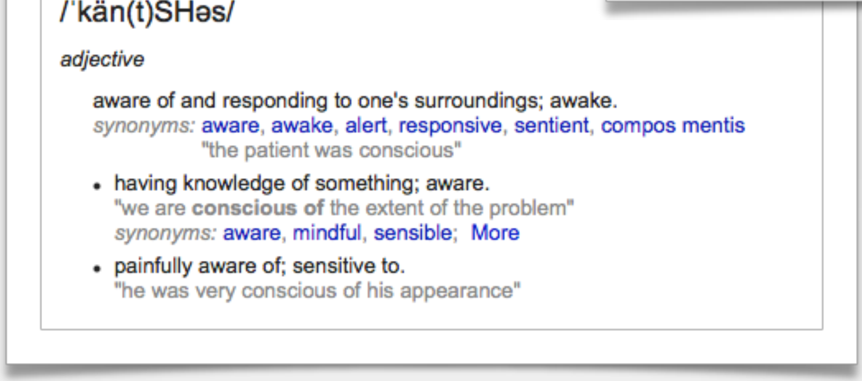

con·scious

/ˈkän(t)SHəs/

adjective

aware of and responding to one's surroundings; awake.

synonyms: aware, awake, alert, responsive, sentient, compos mentis

"the patient was conscious"

- having knowledge of something; aware.
  "we are **conscious of** the extent of the problem"
  synonyms: aware, mindful, sensible; More
- painfully aware of; sensitive to.
  "he was very conscious of his appearance"

https://encrypted.google.com/search?hl=en&q=conscious
http://recode.net/2015/04/15/airbnb-now-factors-in-host-preferences/
http://www.inc.com/adam-vaccaro/uber-isnt-a-car-service.html

Motivation

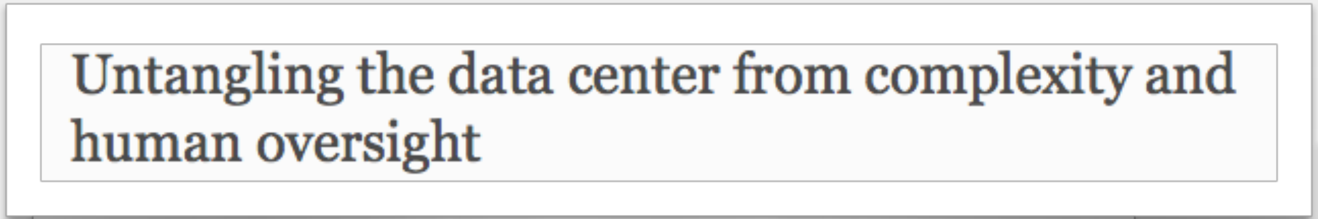

Untangling the data center from complexity and human oversight

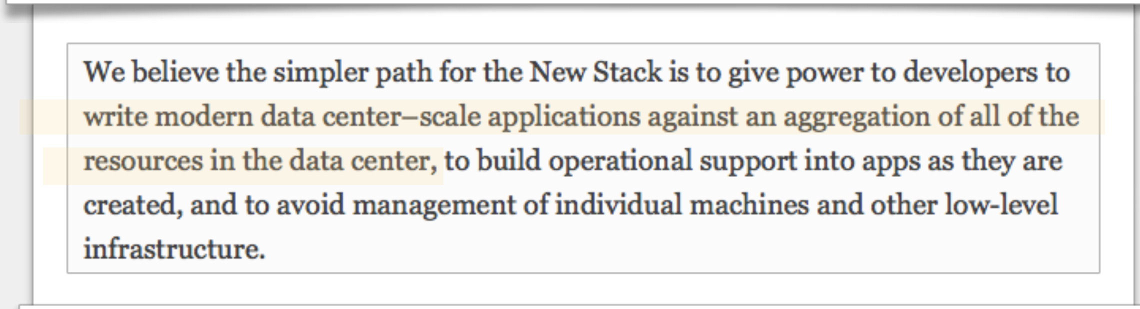

We believe the simpler path for the New Stack is to give power to developers to write modern data center–scale applications against an aggregation of all of the resources in the data center, to build operational support into apps as they are created, and to avoid management of individual machines and other low-level infrastructure.

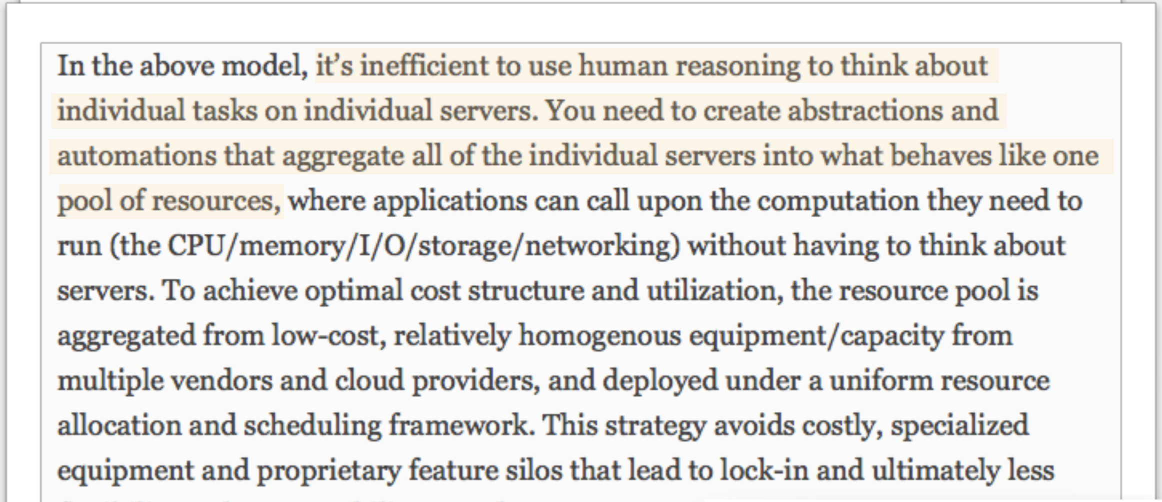

In the above model, it's inefficient to use human reasoning to think about individual tasks on individual servers. You need to create abstractions and automations that aggregate all of the individual servers into what behaves like one pool of resources, where applications can call upon the computation they need to run (the CPU/memory/I/O/storage/networking) without having to think about servers. To achieve optimal cost structure and utilization, the resource pool is aggregated from low-cost, relatively homogenous equipment/capacity from multiple vendors and cloud providers, and deployed under a uniform resource allocation and scheduling framework. This strategy avoids costly, specialized equipment and proprietary feature silos that lead to lock-in and ultimately less flexibility and manageability at scale.

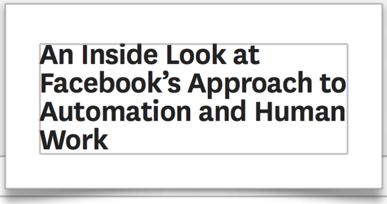

An Inside Look at Facebook's Approach to Automation and Human Work

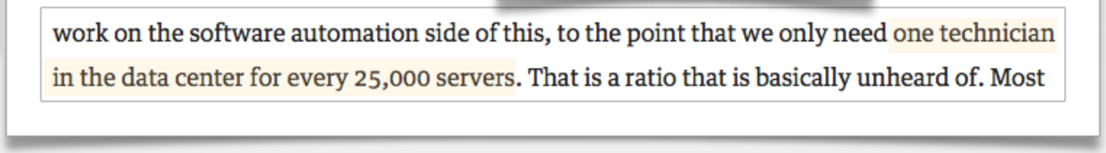

work on the software automation side of this, to the point that we only need one technician in the data center for every 25,000 servers. That is a ratio that is basically unheard of. Most

Scale, complexity and knowledge

Cyborg to Automation

Vinod Khosla: https://gigaom.com/2014/12/14/untangling-the-data-center-from-complexity-and-human-oversight/
Mark Burgess: "Thinking in Promises for the Cyborg Age"
Facebook: https://hbr.org/2015/06/an-inside-look-at-facebooks-approach-to-automation-and-human-work

Motivation

Anomaly Detection

Noisy Neighbors

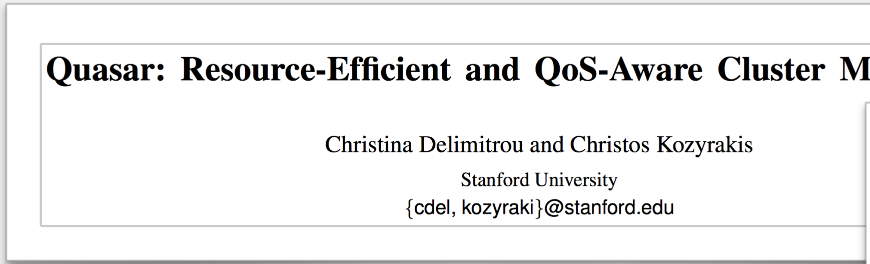

Quasar: Resource-Efficient and QoS-Aware Cluster Management

Christina Delimitrou and Christos Kozyrakis
Stanford University
{cdel, kozyraki}@stanford.edu

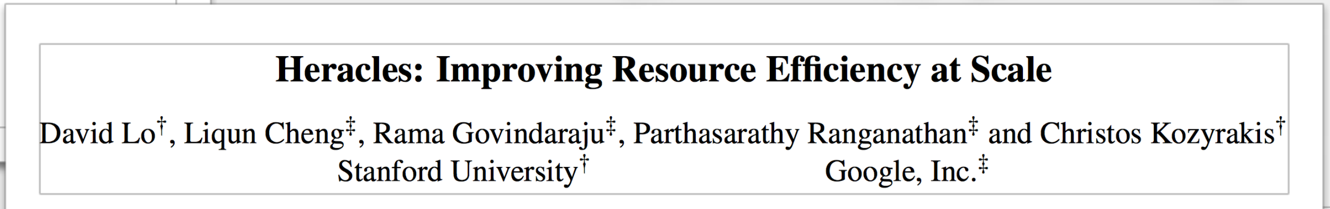

Heracles: Improving Resource Efficiency at Scale

David Lo[†], Liqun Cheng[‡], Rama Govindaraju[‡], Parthasarathy Ranganathan[‡] and Christos Kozyrakis[†]
Stanford University[†] Google, Inc.[‡]

Infrastructure Insights

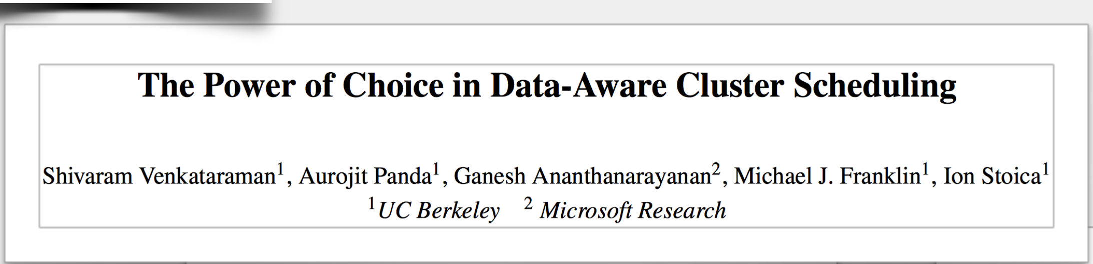

The Power of Choice in Data-Aware Cluster Scheduling

Shivaram Venkataraman[1], Aurojit Panda[1], Ganesh Ananthanarayanan[2], Michael J. Franklin[1], Ion Stoica[1]
[1]*UC Berkeley* [2] *Microsoft Research*

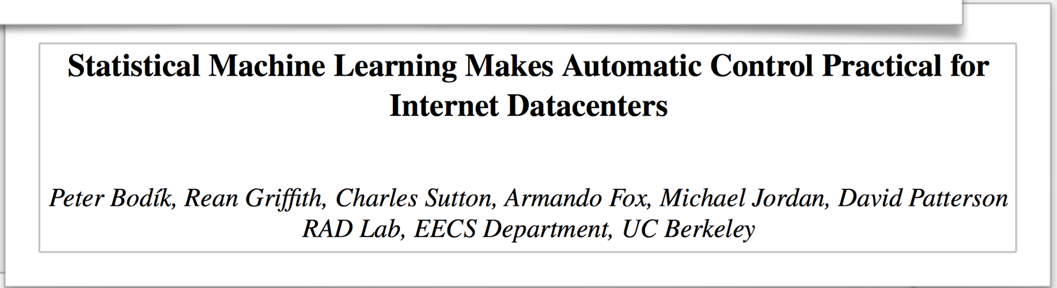

Statistical Machine Learning Makes Automatic Control Practical for Internet Datacenters

Peter Bodík, Rean Griffith, Charles Sutton, Armando Fox, Michael Jordan, David Patterson
RAD Lab, EECS Department, UC Berkeley

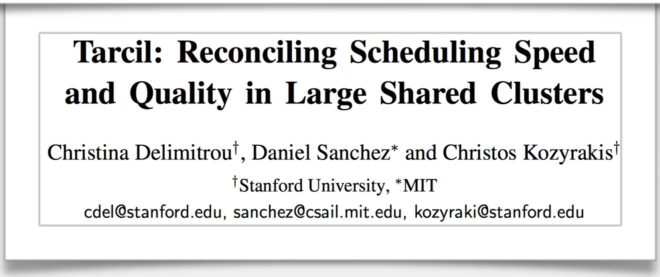

Tarcil: Reconciling Scheduling Speed and Quality in Large Shared Clusters

Christina Delimitrou[†], Daniel Sanchez[*] and Christos Kozyrakis[†]
[†]Stanford University, [*]MIT
cdel@stanford.edu, sanchez@csail.mit.edu, kozyraki@stanford.edu

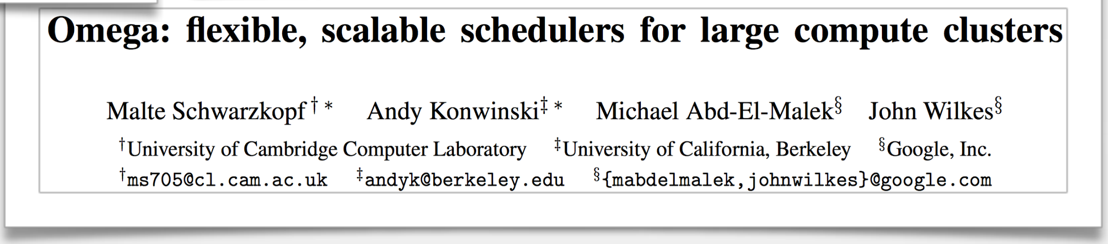

Omega: flexible, scalable schedulers for large compute clusters

Malte Schwarzkopf[† *] Andy Konwinski[‡ *] Michael Abd-El-Malek[§] John Wilkes[§]
[†]University of Cambridge Computer Laboratory [‡]University of California, Berkeley [§]Google, Inc.
[†]ms705@cl.cam.ac.uk [‡]andyk@berkeley.edu [§]{mabdelmalek,johnwilkes}@google.com

Auto-Profiling

Autoscaling, Machine-Learning

http://web.stanford.edu/group/mast/cgi-bin/drupal/system/files/2014.asplos.quasar.pdf
http://web.stanford.edu/group/mast/cgi-bin/drupal/system/files/2015.heracles.isca_.pdf
http://web.stanford.edu/group/mast/cgi-bin/drupal/system/files/Tarcil_socc15.pdf
http://amplab.cs.berkeley.edu/wp-content/uploads/2014/09/kmn-osdi-final.pdf
http://www.cs.berkeley.edu/~jordan/papers/bodik-etal-hotcloud09.pdf
http://eurosys2013.tudos.org/wp-content/uploads/2013/paper/Schwarzkopf.pdf

Charmander 1.0

- Use Charmander to:
  - Analyze more sophisticated scheduling algorithms.
  - Analyze interference between containers, “noisy neighbors”.
  - Investigate anomaly-detection algorithms.
  - Evaluate machine-learning for task-categorization, task-profiling, and performance prediction.

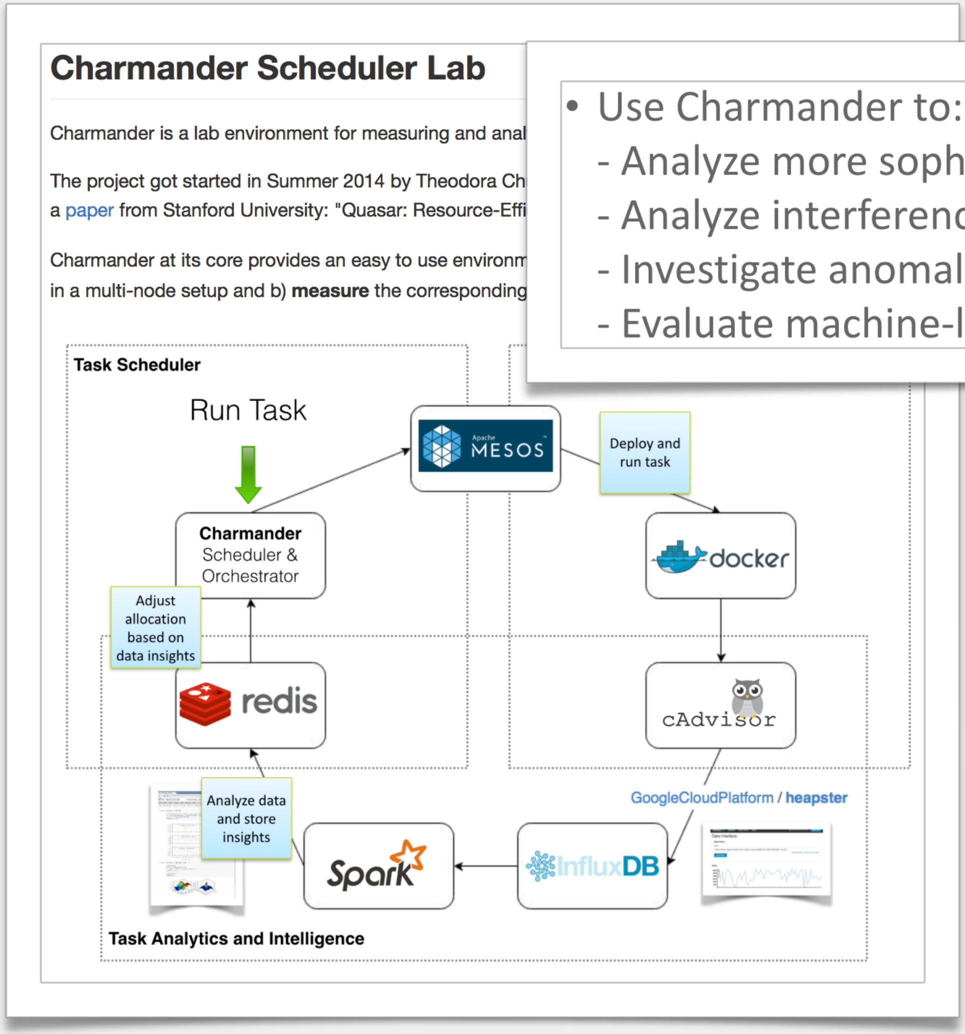

Charmander Scheduler Lab

Charmander is a lab environment for measuring and anal

The project got started in Summer 2014 by Theodora Ch
a paper from Stanford University: "Quasar: Resource-Effi

Charmander at its core provides an easy to use environm
in a multi-node setup and b) **measure** the corresponding

Task Scheduler
Run Task
MESOS
Deploy and run task
docker
Charmander Scheduler & Orchestrator
Adjust allocation based on data insights
redis
cAdvisor
GoogleCloudPlatform / heapster
Analyze data and store insights
Spark
InfluxDB
Task Analytics and Intelligence

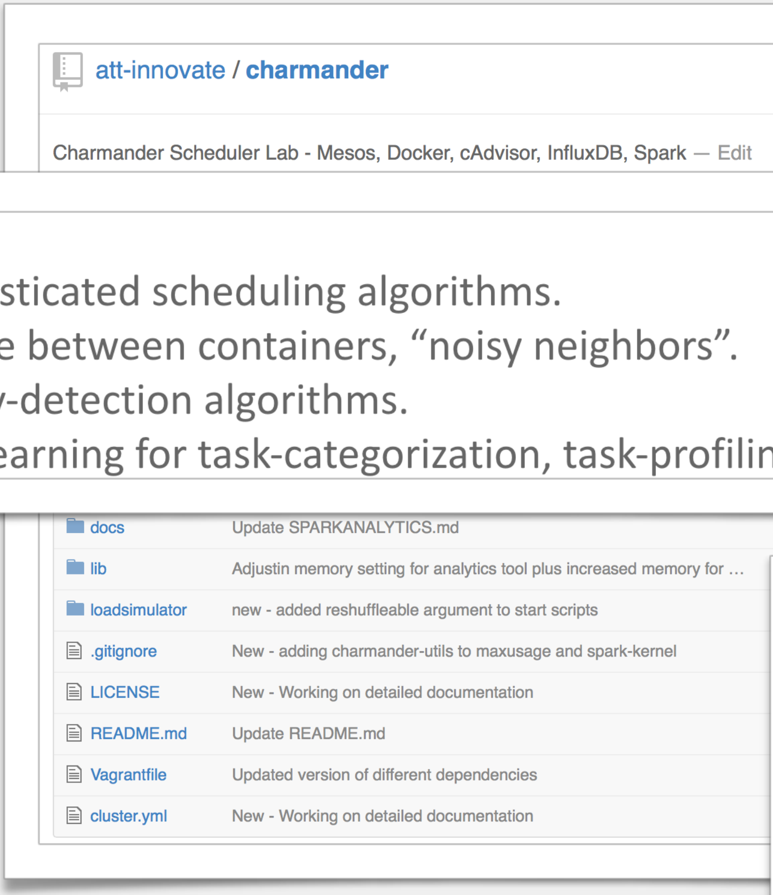

att-innovate / **charmander**

Charmander Scheduler Lab - Mesos, Docker, cAdvisor, InfluxDB, Spark — Edit

docs	Update SPARKANALYTICS.md	7 days ago
lib	Adjustin memory setting for analytics tool plus increased memory for …	
loadsimulator	new - added reshuffleable argument to start scripts	
.gitignore	New - adding charmander-utils to maxusage and spark-kernel	
LICENSE	New - Working on detailed documentation	
README.md	Update README.md	
Vagrantfile	Updated version of different dependencies	
cluster.yml	New - Working on detailed documentation	

Charmander
https://github.com/att-innovate/charmander

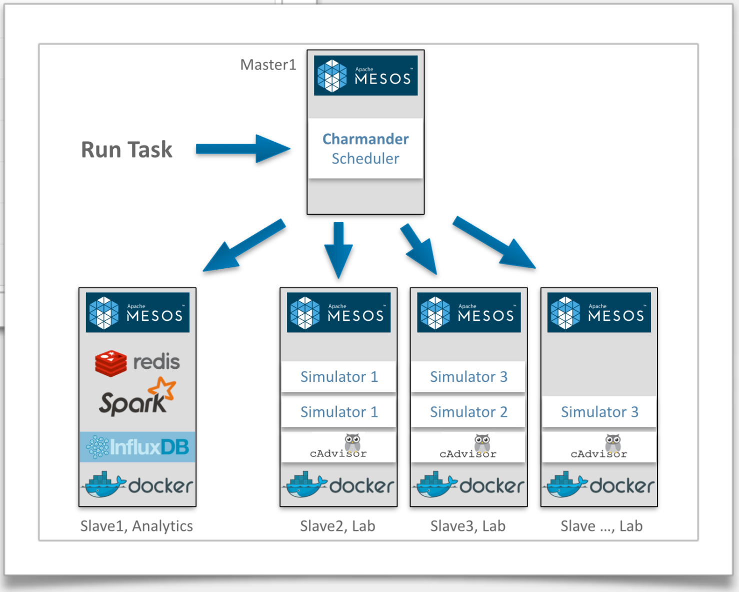

Demo

Technology Stack - Charmander 2.0

https://github.com/att-innovate/

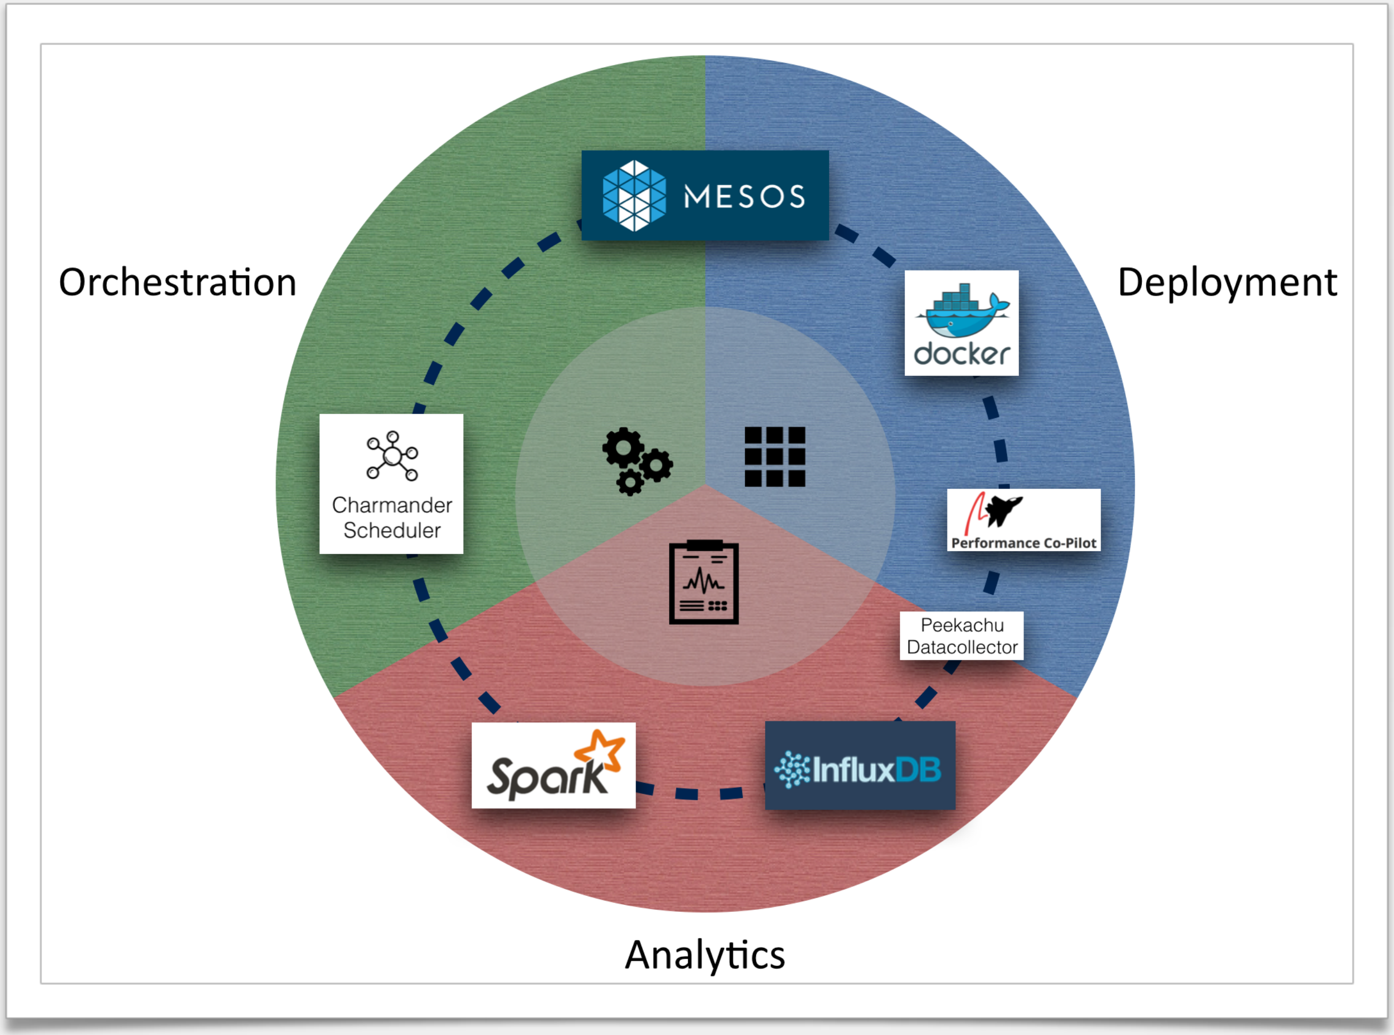

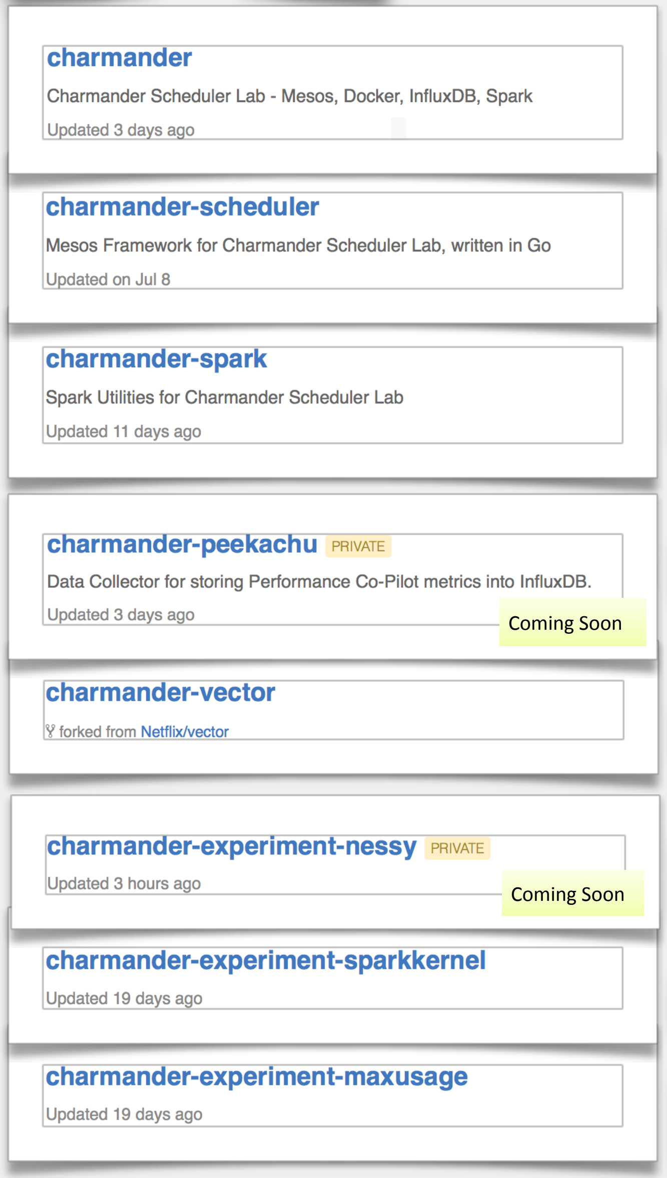

charmander
Charmander Scheduler Lab - Mesos, Docker, InfluxDB, Spark
Updated 3 days ago

Vagrant, Dockerfiles, Scripts ..

charmander-scheduler
Mesos Framework for Charmander Scheduler Lab, written in Go
Updated on Jul 8

Native Mesos Framework (Go)

charmander-spark
Spark Utilities for Charmander Scheduler Lab
Updated 11 days ago

RDD from InfluxDB

charmander-peekachu PRIVATE
Data Collector for storing Performance Co-Pilot metrics into InfluxDB.
Updated 3 days ago
Coming Soon

Configurable Collector (Go)

charmander-vector
forked from Netflix/vector

Optimized for Containers

charmander-experiment-nessy PRIVATE
Updated 3 hours ago
Coming Soon

DNS auto-scaling,
Spark streaming (Scala)

charmander-experiment-sparkkernel
Updated 19 days ago

SparkKernel notebook

charmander-experiment-maxusage
Updated 19 days ago

Spark SQL
(Scala)

Technology Stack - Charmander 2.0

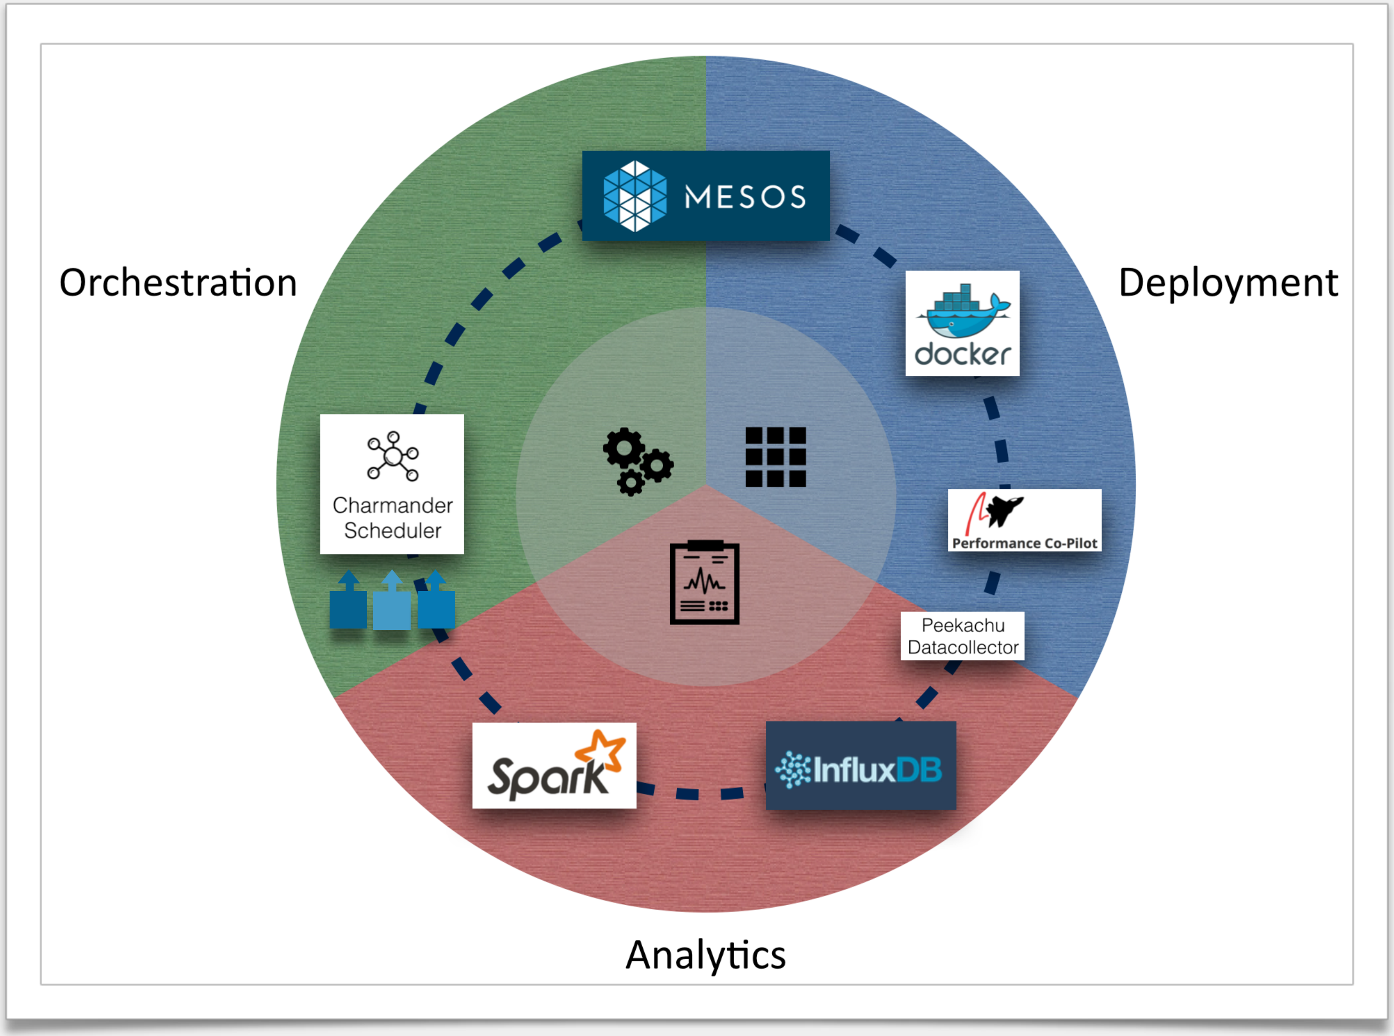

Resource-efficient

Resource-aware

Resource-conscious

Semi-Autonomous Infrastructure

What’s Next

- Many more hands-on Experiments
- Improved Performance Co-Pilot and Vector integration
- Auto-Scaling with autonomous Container-Activation (0 to n)
- Include Network Resources - Hardware (OCP) and Software (VNF)
- Federated metrics-collection platform
- Infrastructure Knowledge Management System
- Embrace Closed-Loop Feedback Systems

Q&A

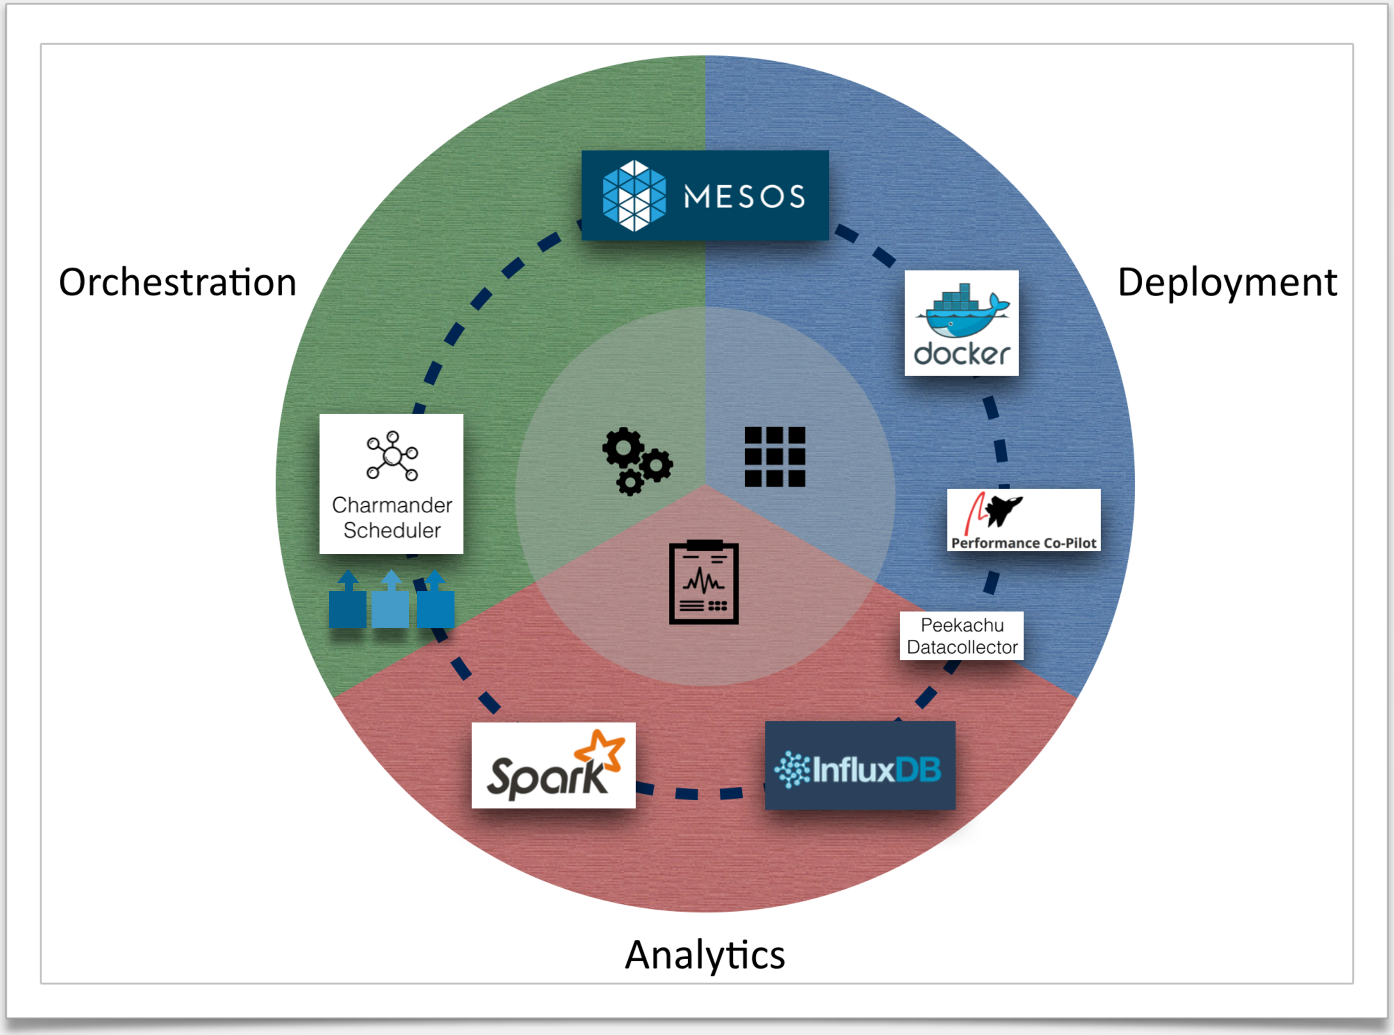

Resource-efficient

Resource-aware

Resource-conscious

Semi-Autonomous Infrastructure