
VC3

Virtual Clusters for Community Computation

US ATLAS Physics, Software and Computing Technical Planning Meeting,
Chicago, August 1, 2016

Rob Gardner
University of Chicago

DOE ASCR DDRM



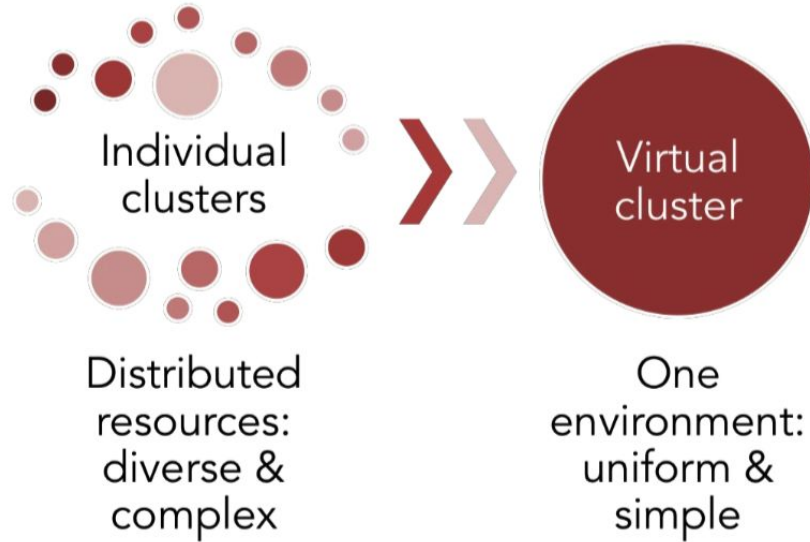
U.S. DEPARTMENT OF
ENERGY

Office of Science

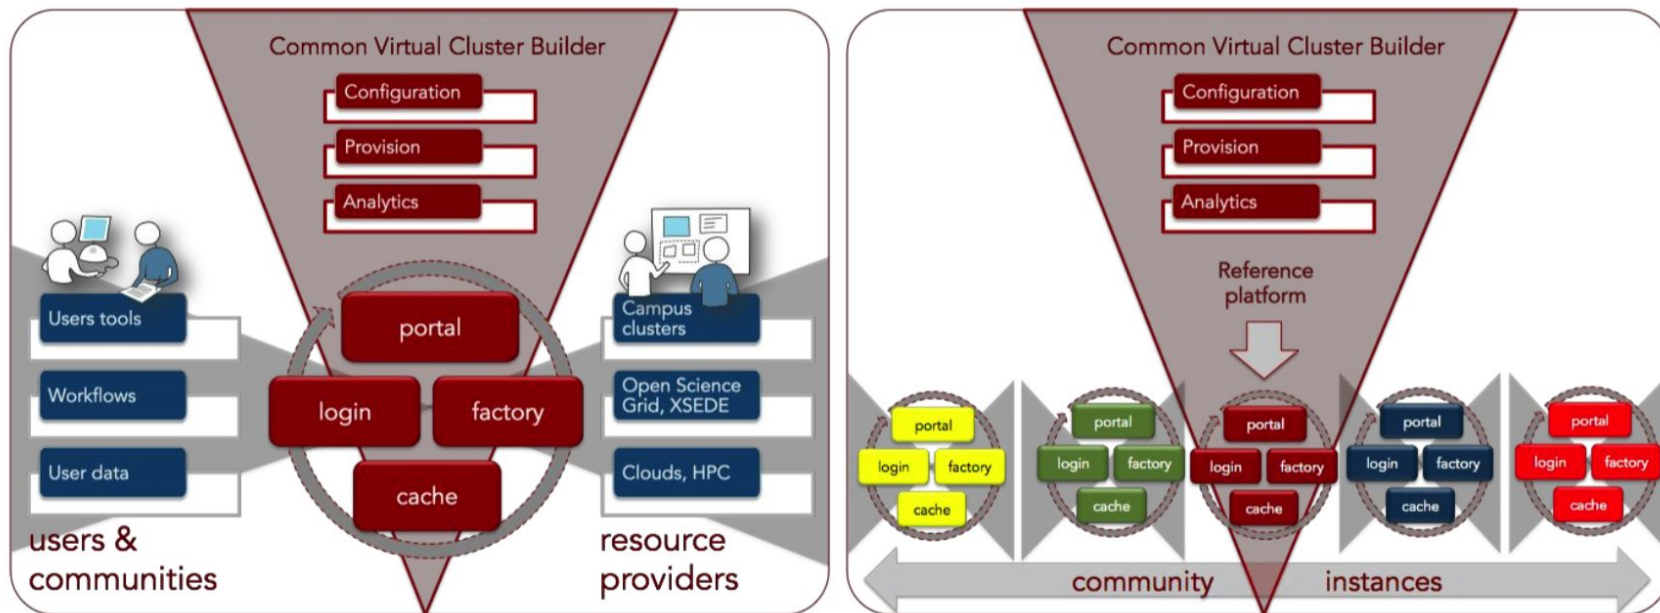
- Notre Dame
 - Thain (overall PI), Hildreth (co-PI), Lannon (co-PI), Brenner (co-PI)
- Chicago
 - Gardner (co-PI), Miller (co-PI)
- Brookhaven
 - Hover (co-PI)

Project kicked off June 2016

DOE ASCR DDRM



DOE ASCR DDRM



Approach

Distinguish the platform for the virtual cluster service from the physical resource targets


- We are not virtualizing data centers
- We are not building OpenStack virtual clusters
- We are building a platform that **delivers a virtual cluster abstraction, as a service**
- An “infrastructure as code” service platform to achieve scalability employing best of breed DevOps methods

VC3:


Virtual
Clusters
Community
Computation
from campuses






Campus elastic scale out
(enterprise hybrid)



campus data center





Transparently compute
over a superhybrid

VC3 Reference Platform & Terminology

- VC3 portal
- VC3 core services
- VC3 service platform
- VC3 resource targets

VC3 Summary

- New project to accelerate delivery of infrastructure to scientific applications
- Provide a virtual cluster abstraction as a service connecting user communities to resources
- Create “infrastructure as code” platform to rapidly build virtual clusters

Extra

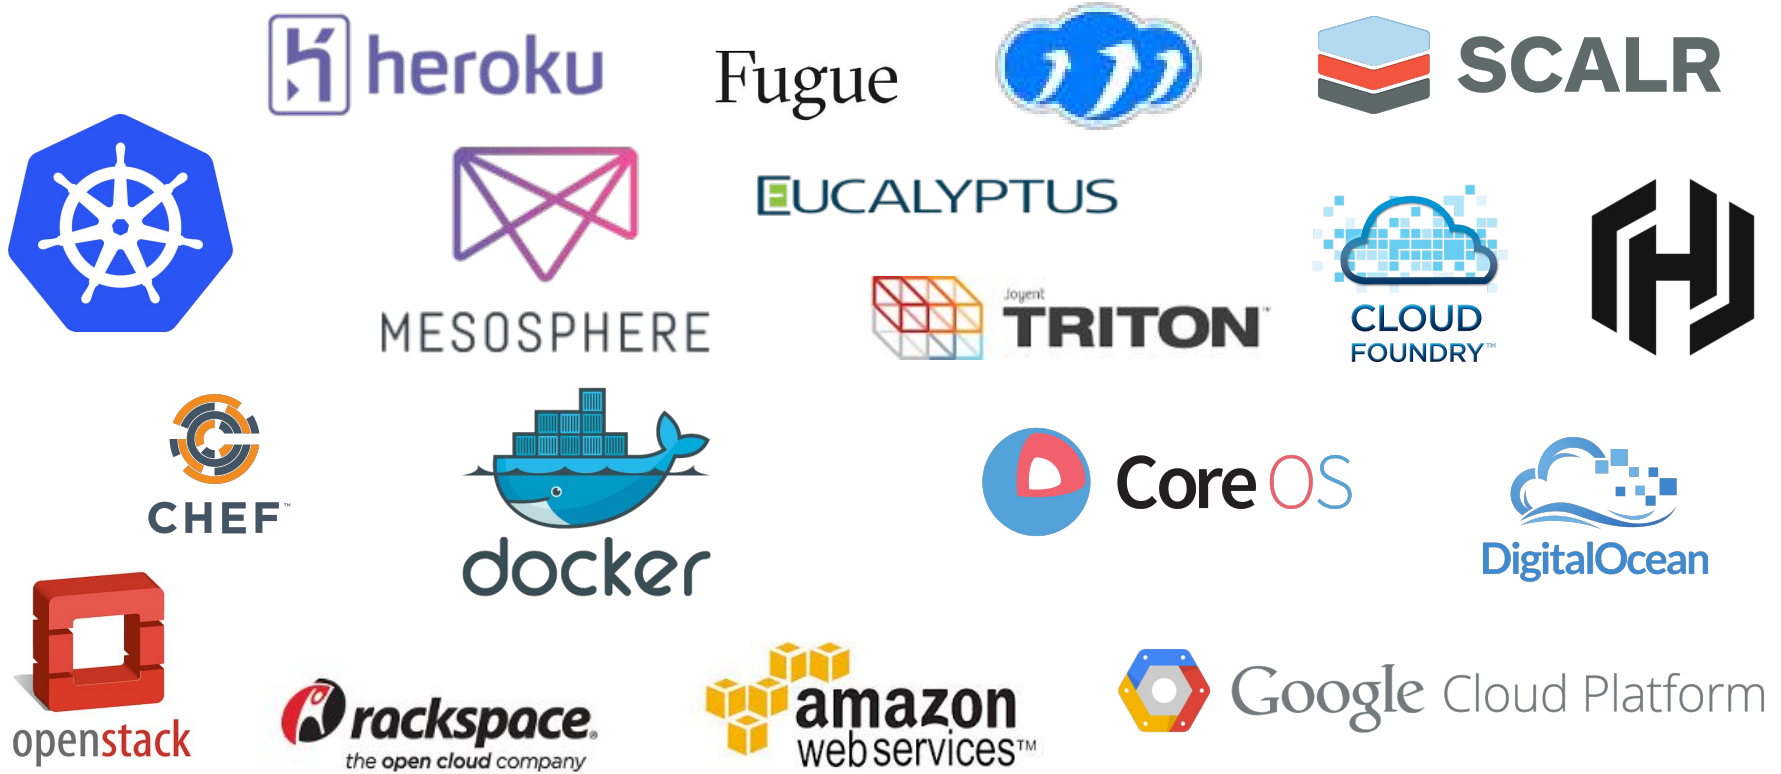
To that end, **we will deliver:**

- A reference virtual cluster platform service which describes the layered architectural components and the relationships and connections between them;
- Two or more production service instances, providing a core provisioning feature set, to be used to coordinate a heterogenous mix of institutional research computing clusters, HPC centers, public commercial cloud, and grid clusters;
- A selection of workflow tools and services that are capable of functioning effectively in the dynamically-configured environment of a virtual cluster. The effectiveness of these tools will be demonstrated at scale with CMS and ATLAS applications.

Work Packages

- WP1: Reference virtual cluster platform for community computation
- WP2: Adapting and enhancing provisioning factory
- WP3: Adapting existing middleware to virtual clusters
- WP4: Community experience portal
- WP5: Service instances
- WP6: Communities engagement

Leverage data center virtualization software & services



Atmosphere taxonomy

Karl Eisenburg (Mesosphere) on choosing the cloud tools:

1. No one tool solves all problems
2. Pick the right tool for the job
3. Pick tools that do one thing well
4. Prefer small APIs
5. Chain your tricks together

	Application Orchestration	CloudFoundry, Heroku, OpenShift, Deis
	Container Orchestration	Kubernetes, Marathon, Swarm, Fleet, Lattice, ECS
	Job Scheduling	Chronos, Kubernetes
DCOS	Containerization	Docker, Rkt, Garden, Mesos
	Resource Management	Mesos
	Provisioning	Vagrant, Otto, Puppet, Chef, Ansible, Salt
	Machine Management	AWS, Azure, GCE, OpenStack, vSphere, VirtualBox, Fusion

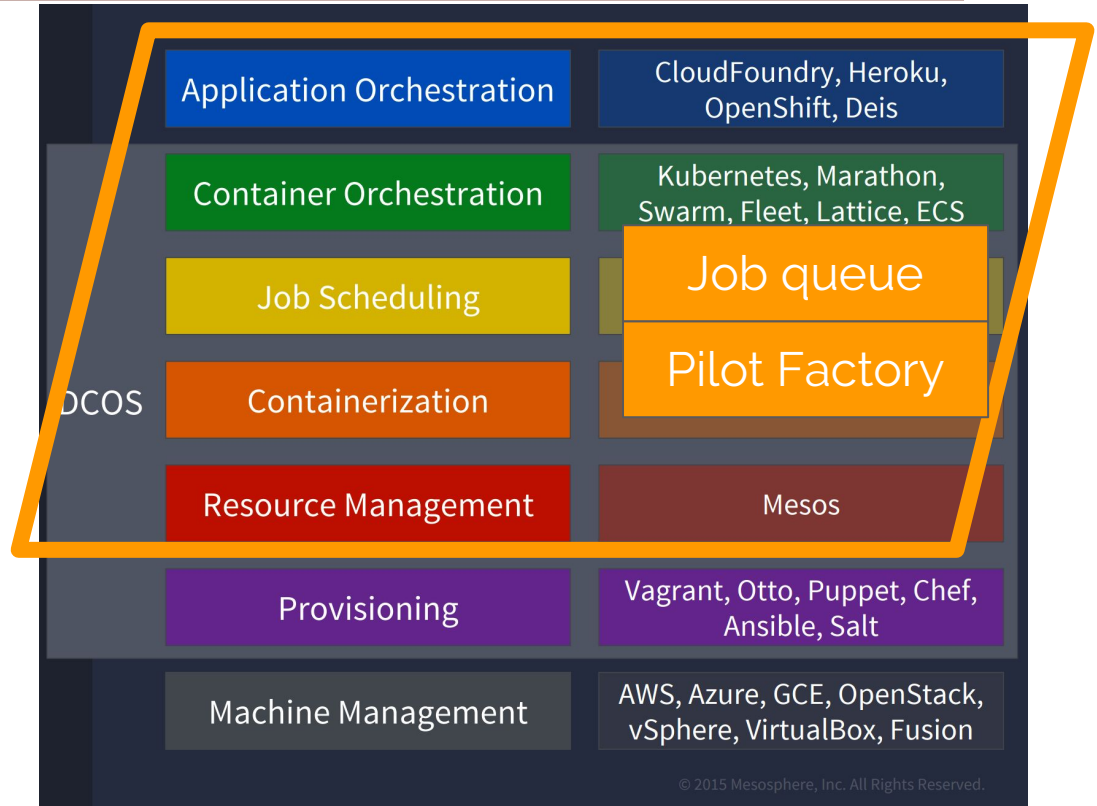
© 2015 Mesosphere, Inc. All Rights Reserved.

Placement by function

How does VC3 map?

Need more than two dimensions to express role by function

Perhaps “hovers above”, reaches for a service when needed

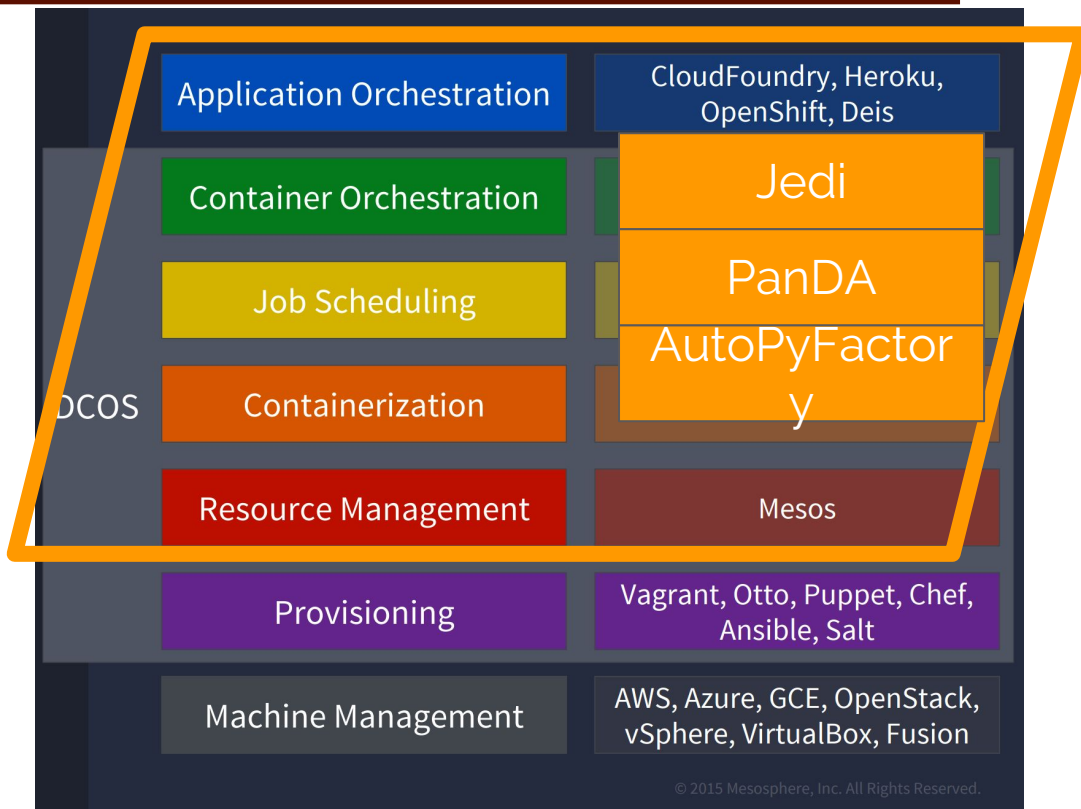


Placement by function

How does VC3 map?

Need more than two dimensions to express role by function

Perhaps “hovers above”, reaches for a service when needed



Placement by function

How does VC3 map?

Need more than two dimensions to express role by function

Perhaps “hovers above”, reaches for a service when needed

