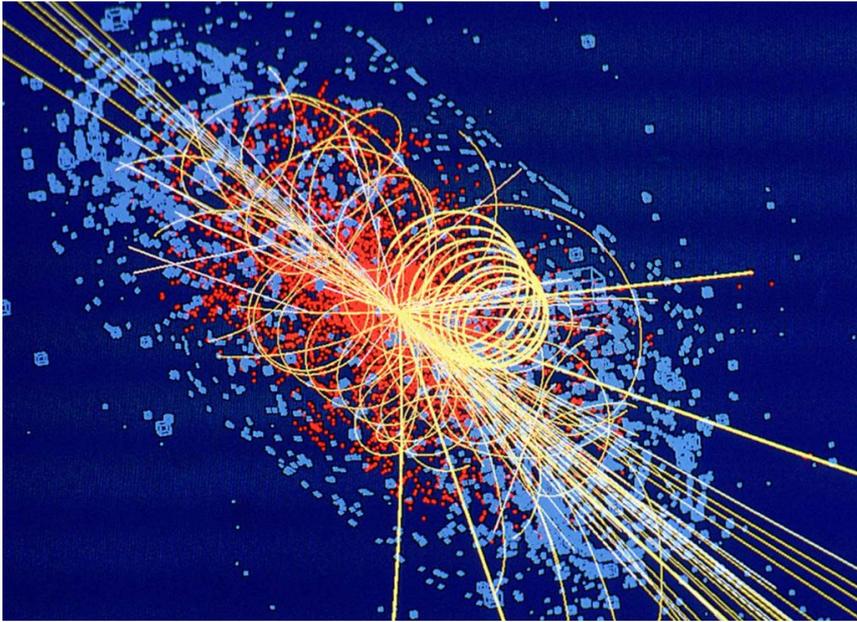


## After Discovering the Higgs Boson, CERN Integrates OpenStack with Kubernetes

12 May 2017 9:00am, by [Scott M. Fulton III](#)



1

The beauty of a scientific program, whether it be centered around a module orbiting the planet or a subatomic particle riding a narrow, ring-shaped corridor underground, is that it compels its participants to solve bigger problems in new and untried ways. The space program prompted NASA to create new ways to envision the way computers work. “Software” was a catch-phrase NASA engineers came up with to refer to digital programs that did the work of mechanisms. Shipping containers packed with servers outside Kennedy Space Center and the Jet Propulsion Laboratory became the seeds for the world’s first clouds.

In Switzerland, the [European Organization for Nuclear Research \(CERN\)](#) operates the [Large Hadron Collider](#). Already nearly five years ago, researchers at the LHC confirmed the existence of a theoretical particle, the [Higgs Boson](#), which was predicted by a theory of symmetry at the subatomic level.

Yet CERN finds time to innovate in the field of computing as well. The organization presently outsources its data center facilities to several firms, including one that utilizes a genuinely innovative cooling system that [involves immersing its servers completely in mineral oil](#).

Now, joining forces with Rackspace — the first commercial firm to emerge from the NASA cloud experiment, [from which OpenStack was born](#) — CERN has provided a testbed for [OpenStack](#) to integrate with the [Kubernetes](#) open source container orchestration software, by way of the former platform’s native [Magnum component](#). They’re solving the integration problem without splitting the atom.

At the end of 2015, CERN needed a way to effectively distribute workloads across its data centers (including those that it leases from other service providers) throughout Europe. As CERN staff member [Ricardo Rocha](#) told attendees at [OpenStack Summit 2017](#) in Boston on Wednesday, Magnum was one of the few options for integrating container orchestration with OpenStack at the time, so it began building a container service on top of Magnum in early 2016.

“Magnum also had the possibility for choosing the container engine. This was very important for us,” said Rocha. “We had groups of people who were pushing for Kubernetes; we had groups of people who were already using Mesos; and others who were just using plain Docker, and they wanted to rely on the Docker API, where Swarm has great potential. We wanted this to be easy to use, so people don’t have to understand complicated templates for how to configure their clusters.”

Helping CERN launch its project was Rackspace distinguished architect [Adrian Otto](#). In addressing attendees here in Boston Wednesday, Otto answered [one of the questions we raised](#) in our preview of this conference: What happened to Magnum?

“Magnum is an OpenStack service that allows you, as a cloud user, to produce a cluster that runs a container orchestration engine,” stated Otto. “It allows you to use your existing cloud credentials to produce those clusters. So if you’re already an OpenStack user,

That choice gave CERN the option of mixing container orchestrators among clusters. Rocha told attendees that, while the organization currently operates about 40 Kubernetes clusters, it also runs some 20 or so on Docker Swarm and five on [Mesosphere DC/OS](#).

Otto drove home another of the principal reasons why OpenStack may provide a more complete solution for containerization than just an orchestrator acting as an IaaS provider: Magnum enables multi-tenancy. ([The Hyper.sh project also seeks to provide multi-tenancy](#) for Kubernetes, but it's doing so by mixing in code from Magnum and other OpenStack components).

“So you can have clusters that are side-by-side, of the same time or of a different type,” Otto continued, “but that are guaranteed never to share the same kernel with each other. Which is important for security installation reasons. Because of the way that Magnum works, you get multi-tenancy not just at the control plane — which you might be accustomed to with your favorite orchestration system — but all the way down through the entire cloud.”

OpenStack veterans are accustomed to provisioning resources through a template using the [Heat component](#). That's a file artifact, as Otto explained, presented to the orchestrator to produce a stack. “The disadvantage of that model,” he continued, “is that it's not represented in a way that's reusable by all the users of your cloud. Every user needs to have his or her own file artifact, in order to produce the stack.”

That's because OpenStack did not originally envision the need for creating an environment within an environment, where the contained environment is a model for something common and ephemeral instead of unique, special, and indestructible. So for Magnum, OpenStack introduced an API resource which the administrator can expose as a public resource. That resource may then be borrowed by other users to produce containerized clusters.

Yes, this may not be new information, at least in terms of published literature. But for some who are investigating the question of why an orchestrator may need an infrastructure stack and vice versa, it's old information that points the way to a new solution. Though Kubernetes is frequently offered as an infrastructure solution, including from [some of OpenStack's best friends](#), many organizations need to stand up private clouds that work like public clouds and not like Google (specifically, like Google internally).

I asked Ricardo Rocha, what innovations has CERN made to its infrastructure, whose benefits may then be shared upstream with the rest of the community.

Through [CERN's OpenLab](#), produced in conjunction with Rackspace, the organization is now contributing upstream to the Magnum project, responded Rocha [pictured right with [Adrian Otto](#)]. “We actually have a person at CERN from this collaboration [[Tim Bell](#)] who became part of Magnum, and has been contributing all the features that we require internally.



Then there are these side projects, like integration with [Cinder](#), contributions to [libstorage](#), which is the library that Swarm is using to provide Cinder support. We've done a couple of patches for the OpenStack driver for Kubernetes. And we are reaching out to the communities to work with them. That worked really well for OpenStack, and we count on doing the same for other projects.”

It is a firm, solid, and evidently working example of an organization dedicated not just to the public interest but the betterment of humanity. Oh, yeah, and Rackspace. Answering users' inquiries with working code, as open source was designed to do. Maybe certain countries lack a space program, but it may not take — or need — a country anymore to drive the evolution of the data center.

Feature image: Higgs boson decay to four muons (CERN).



**A digest of the week's most important stories & analyses.****Subscribe**

We don't sell or share your email. Occasionally, we send updates and useful info.

## RELATED STORIES

PODCAST / TECHNOLOGY

**What Comes of Docker with a New CEO?**18 May 2017 3:00pm, by [Alex Williams](#)

ANALYSIS

**Microsoft Brings Container Orchestration to Azure Service Fabric, for Windows and (Soon) Linux**17 May 2017 1:00am, by [Mary Branscombe](#)[View / Add Comments](#)

## SPONSORED FEED

**vmware**

Hantera världens resekosystem med  
VMware NSX  
MAY 23, 2017

**Aporeto**

Cisco's first security GM, now a VC, talks  
about 3 big vulnerabilities  
MAY 22, 2017

**CoreOS**

Kubernetes sessions at CoreOS Fest  
MAY 22, 2017

**pagerduty**

How to Manage a Tier Zero Service

**THE NEW STACK**

How to warmly welcome new email subscribers  
MAY 22, 2017

---



Chef Community Engineering – Quarterly Update  
MAY 22, 2017

---



Diversity Scholarship Series: Inspired by Kubernetes and Its Community  
MAY 19, 2017

---



Securing Containers on OpenShift with Aqua  
MAY 18, 2017

---



Leveraging the dockernignore File to Create Smaller Images  
MAY 18, 2017

---



SD-WAN: Just the Ticket for Greater Flexibility and Lowered Costs  
MAY 18, 2017

---



Continuous Deployment with Wercker: A Tale of Python Bots  
MAY 18, 2017

---



Just released: Kubernetes CIS Benchmark for 1.6  
MAY 17, 2017

---



Deploying ASP.NET Core Applications on Apcera  
MAY 16, 2017

---



On the road to Channel 2.0  
MAY 15, 2017

---



South Korea Standardizes IT Service Delivery with Cloud Foundry  
MAY 11, 2017

---



DigitalOcean Monitoring: Insight Into Key Design Decisions  
MAY 10, 2017

---





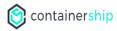
New Members

MAY 08, 2017



How to Thrive in a Post AWS World

MAY 07, 2017



A Pivotal Moment for the container industry

MAY 03, 2017



Gene Kim joins as Advisor to CloudMunch

FEBRUARY 21, 2017



Why we've moved from go.cd to gocd.io

JANUARY 04, 2017

TOPIC HUBS

Architecture

Data

Development

Do-It-Yourself

Operations

Security

Tech Culture

ECOSYSTEMS

AI

Cloud

Containers

Kubernetes

Microservices

Node.js

Serverless

CONTENT

Analysis

Ebooks

Events

Podcasts

Research

THE NEW STACK

About

Contact

Sponsors

Privacy Policy

© 2017 The New Stack. All rights reserved.