



NEWS

OpenStack Magnum on the CERN production cloud

Providing 2,000 users a self-service cloud that stands up to “extreme computing.”



**Nicole
Martinelli**
April 21, 2016

European particle physics laboratory CERN is on a mission to push the boundaries of knowledge in physics. Tim Bell, CERN infrastructure manager and OpenStack board member, recently brought to light another kind of innovation the research lab is

<http://superuser.openstack.org/articles/openstack-magnum-on-the-cern-production-cloud>



brought to light another kind of innovation the research lab is fronting.


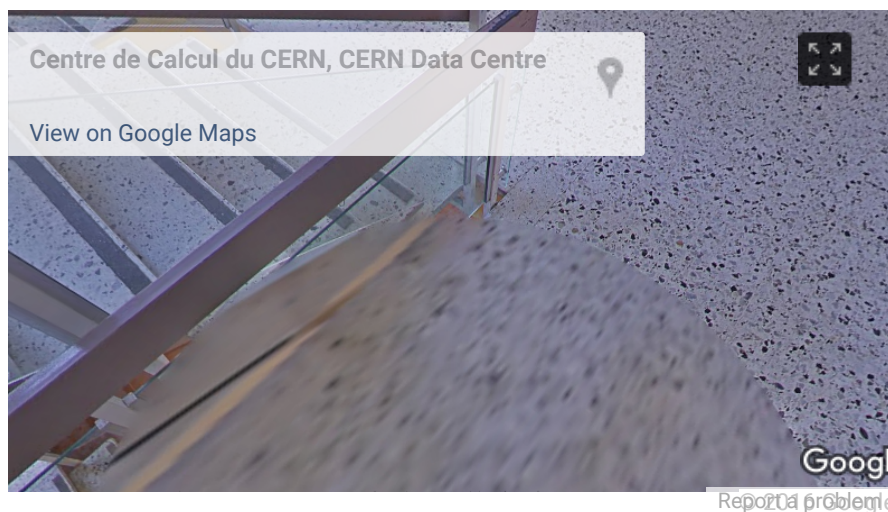
“My first bays [are] running on CERN production cloud with Magnum,” he reported on [OpenStack Successes](#).

The goal of using the OpenStack API service was to provide “agnostic container orchestration” for its 2,000 users as they crunch numbers in the quest to understand the mysteries of the universe. Researchers at four particle detectors (ATLAS, CMS, LHCb, ALICE) generate a mind-blowing 30 petabytes of data annually in the Large Hadron Collider (LHC.)

CERN and OpenStack go way back. CERN’s IT department started working on OpenStack in late 2011, moving most of its IT infrastructure to OpenStack into virtual machines in 2014. CERN has been using OpenStack in production since summer 2013. To stay on the cusp of new technologies, Bell and his team have been working to enable [Magnum](#) for container native applications since [late 2015](#).

So what are the takeaways from CERN about containers for more pedestrian OpenStack users? It turns out that even the experts at CERN find it challenging to parse the various options (Kubernetes, Docker Swarm, Mesos, plus Atomic, CoreOS and Rocket) to recommend the right ones for their users.

Superuser talked to [Bell](#) about how his team explores new technologies, the difficulty level for adopting Magnum and how operators can get involved in OpenStack.




DefCore isn't a punk band, but if you use OpenStack, listen up



Nicole Martinelli

February 20, 2015



Why OpenStack will change the way we live, work and play

Niki Acosta of Cisco on

April 6, 2015



Hong Nga Nguyen



A peek inside CERN's data center.

Can you sum up your experience deploying Magnum?

With a cloud the size of the CERN installation, we use Puppet to ensure the configurations of thousands of servers are consistent and can be changed regularly with new configuration options or software updates. The CERN approach is to use upstream Puppet modules where available from the Puppet-OpenStack project. In the case of Magnum, we worked with the Puppet community to create the puppet-magnum configuration now available at <https://github.com/openstack/puppet-magnum/>. With the Puppet configuration, we can easily deploy at scale, such as with the other 11 puppet components we use in the CERN cloud. We also worked with the [RDO distribution](#) and the [RPM Packaging project](#) to deploy the code itself.

Did you find documentation?

When working with a new OpenStack project, there are often parts of the documentation which need enhancing to cover installation and configuration. The existing developer documentation at <http://docs.openstack.org/developer/magnum/> and assistance on the #openstack-containers IRC channel was a great help but given that our focus is on production deployment, some work was needed to understand the options applicable to our cloud. Given our experience, we are now working on the [installation guide](#) to help others deploy more easily.

What's the difficulty level?

It is important to match the skills of the Organization to the demands of each OpenStack component. The [project navigator](#) can help to assess each OpenStack project and make an assessment if it is compatible with the skills in the team.

We had [Heat](#) already available as part of the CERN cloud service


Why OpenStack needs to keep it simple

Pundits say complexity and a "science project" reputation hinder

September 11,  Superuser 2015



OpenStack and OPNFV strengthen collaboration for telcos

 Kathy Cacciatore
November 23, 2015




Superuser Awards: your vote counts



catalog but for Magnum, we have added Neutron and Barbican to our cloud to handle the networking and secrets. We expect both of these components to be used in future for other cloud activities so sharing existing OpenStack projects is a benefit for the service. Many of these components are relatively newly packaged but the support from the community has been good when we encountered problems.

Technically, the largest challenges have been around understanding how containers would be used rather than deploying Magnum. With many options, such as Kubernetes, Docker Swarm and Mesos, along with Atomic, CoreOS and Rocket, it takes some time to understand the best approaches to recommend to the CERN users. Using [Keystone's endpoint filtering](#), we can expose a pilot service for specific projects only and use this to validate the approach and improve documentation before making it generally available.



CERN
@CERN

Follow

CERN congratulates the discoverers of [#gravitationalwaves](#) cern.ch/go/f6fs
(Image: W.Benger/ZIB/AEI/CCT-LSU)

10:01 AM - 12 Feb 2016

798 1,014



Five years in: Charting the OpenStack galaxy



Why OpenStack's drive for inclusivity matters



OpenStack User Survey Insights: November 2014



OpenStack + Kubernetes = More choice and flexibility for developers



Take these OpenStack Infrastructure tools and run!



OpenStack crowns "Liberty" its next release name



Getting my hands dirty on an OpenStack lab



Contributing to open cloud projects without losing your mind



How to craft a successful OpenStack

How long did it take?

With over 2,000 users of the CERN self-service cloud, our community is often exploring new technologies to solve the extreme computing challenges of the Large Hadron Collider and other CERN experiments. Several of our advanced users, such as <http://tiborsimko.org/docker-on-cern-openstack.html>, had been exploring Docker in 2015 using VMs so these made good pilot testers. The EU Indigo Dataclouds project (<https://www.indigo-datacloud.eu/>) was interested to develop an open source data and computing platform targeted at scientific communities, deployable on multiple hardware platforms and provisioned over hybrid, private or public, e-infrastructures. Rackspace, collaborating with [CERN openlab](#), have been exploring how to use containers for high throughput computing, building on our previous Rackspace collaboration on OpenStack cloud federation. For the production cloud, we started to look at Magnum in the second half of 2015 and have now started work with pilot users in the first half of 2016. We would expect to go into production for general users during 2016.

Any other things operators should be aware of?

As with CERN's previous Rackspace collaborations, many EU projects and the OpenStack community, the CERN team works using the upstream, open design processes. This allows easy sharing with other High Energy Physics laboratories using OpenStack such as IN2P3 in Lyon, SLAC in California, DESY in Germany and [INFN](#) in Italy. This code can be developed and available to other user communities within the open source umbrella. Operator meetups such as the recent one in Manchester, UK, allow the operator community to share experiences on deploying the new components and give input to decisions on when and what applications to consider a container deployment within their organization.

You can learn more about how CERN is deploying Magnum in this



Summit proposal



Working to bridge the developer/end-user divide in OpenStack



Superuser weekend reading



BBVA on OpenStack



Three cloud computing myths, debunked

recent talk from CERN's Bertrand Noel, Ricardo Brito Da Rocha and Mathieu Velten on [containers and orchestration in the CERN cloud](#). and hear from the team in sessions at the upcoming [Austin Summit](#).

Cover Photo of an an automated magnetic tape vault at CERN computer center. Credit: Claudia Marcelloni and Maximilien Brice, copyright CERN.



Comments **Community** **Login** ▾

Recommend **Share** **Sort by Best** ▾

Start the discussion...

Be the first to comment.

ALSO ON SUPERUSER

Unlocking the cloud: day one of the OpenStack

1 comment • a year ago



David H Deans — The keynotes and insightful sessions on the first day

OpenStack application developers share insights

4 comments • a year ago



David H Deans — The OpenStack User's Survey includes some

How to migrate from VMware and Hyper-V to

3 comments • 7 months ago



KLP — Could you script this up in Chef or Puppet?

October 2015 user survey highlights increasing

1 comment • 6 months ago



Raghav — rate of openstack deployment made industries better

Subscribe

SUPERUSER

[About](#)[Awards](#)[Contact us](#)[Subscribe](#)[User Survey](#)

COMMUNITY

[User Groups](#)[Events](#)[Jobs](#)[Contribute](#)

USER RESOURCES

[OpenStack Marketplace](#)[OpenStack Training Courses](#)[OpenStack Documentation](#)[OpenStack Operations Guide](#)[Application Developer Resources](#)[OpenStack Security Guide](#)[OpenStack Questions](#)

OpenStack Superuser is a publication built to chronicle the work of superusers, and their many accomplishments personally, professionally, and organizationally. The emphasis is on a blend of original journalism and user-generated content, ranging from technical to business-level issues with feature stories, case studies, tips and videos for OpenStack cloud architects and administrators.

