

How CERN uses OpenStack to drive its scientific mission

[D4MER](#) [BLOG, OPENSOURCE](#) [OCTOBER 23RD, 2015](#) [0 COMMENTS](#) [21 VIEWS](#)



One of the world's largest scientific organization is using OpenStack to understand what makes up everything in our universe. CERN runs one of the most collaborative scientific projects on Earth, responsible for producing enormous amounts of data on a routine basis to make Nobel prize winning discoveries such as the Higgs boson has some pretty unique computing requirements.

Belmiro Moreira takes us through how OpenStack is being used to enable this groundbreaking science at CERN, and how they have contributed back to this open source project. Satisfying the demanding needs of computer savvy scientists can be challenging, and CERN certainly runs the gauntlet of legacy applications that have been employed for decades through to new deployments written for newly conceived approaches. Moreira will be speaking on CERN's cloud architecture at the upcoming OpenStack Summit in Tokyo, Japan.



CERN and other facilities have traditionally used supercomputers. Why has CERN created such a large private cloud, and has this enabled new or different research to be done?

CERN provides compute resources to a large community of scientists and experiments. The Large Hadron Collider (LHC) experiments are the biggest scientific collaborations that use CERN compute resources, and their data analysis fits better with a distributed computing model. It entails the execution of a large number of loosely coupled jobs; this requires high throughput computing rather than high performance for a single program.

In the late 90s, the approach of using a large number of commodity x86 servers running Linux was becoming popular. To process the data from the LHC experiments, CERN developed a computing infrastructure based in that model to be the Tier-0 of the worldwide LHC computing grid. At that time there were very few tools available to manage thousands of nodes, and this led to the development of our own tools. Over the years maintaining these tools required substantial effort and it was becoming difficult to find new trained staff.

In 2012, CERN IT made the decision to deploy a private cloud based on commonly used open source tools like OpenStack, Puppet, and Elasticsearch. This allows us to manage in a consistent way a large pool of physical resources, share them between different scientific communities and services, improve resource utilization, and enable cloud interfaces for a self-service resource allocation model.



VIDEO



POPULAR POSTS




[Color Grading in Blender](#)
March 18th, 2013



[Bring out the GIMP!!!! natively for OSX](#)
December 18th, 2012



[What to do if a plugin breaks your wordpress install.](#)
August 2nd, 2012



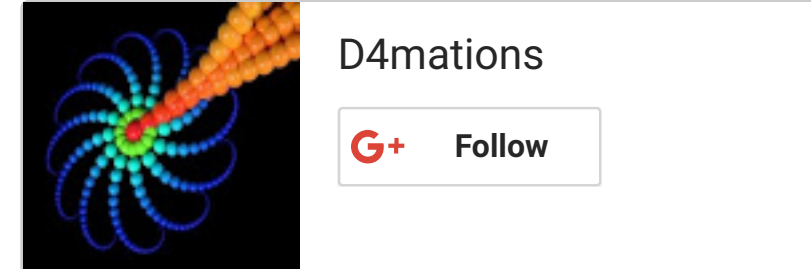
[StructureSynth to Luxrender Exporter v.025...Bug Fix](#)
May 27th, 2013



[Blenders 2.66 New features](#)
March 18th, 2013

FOLLOW US ON GOOGLE+

This represented a big step for all members of CERN scientific community, enabling them to access compute resources in minutes rather than waiting weeks/months for a physical server to be installed and allocated.



How large is CERN's deployment?

CERN's private cloud runs in two data centers (Geneva and Budapest) with a total of about 5,000 servers (about 130,000 cores). By summer 2016, we expect to grow to about 200,000 cores. For block storage, CERN runs Ceph with a capacity of 3.5PB.

What kind of applications are running on it?

The cloud infrastructure runs a very heterogeneous set of applications that ranges from IT services to scientific data processing applications to users' personal remote desktops. It's very challenging since these projects have different requirements and users communities. As an example, we run projects such as "ATLAS Analysis," "IT Videoconference," and "Personal Belmiro" in the same servers now, while in the past they would have had different dedicate resources running side by side.

Are there criteria for determining whether cloud or supercomputer are the best match for an application?

Applications that are highly parallelized and require low latency network run in special dedicated hardware, not in the Cloud infrastructure. However, those don't represent the majority of the scientific workload that we have at CERN.

The main scientific applications and frameworks running in CERN cloud infrastructure are for LHC data analysis. Those are very computational intensive, but the jobs can run independently. This is the perfect workload for a cloud environment where each job can have its own running environment in an ephemeral virtual machine (VM). Because the main analyses consist of crunching data that is stored in CERN storage systems, the job can be resubmitted in the event of a VM failure.

How has OpenStack changed the way big data analysis happens at CERN?

Scientific experiments use not only the compute resources available at CERN but also other smaller data centres around the world and even public clouds. OpenStack, being an open source project, allows CERN and other research facilities to expose a common set of cloud APIs to its users which gives significant productivity gains.

In the past resources were static and allocated to different scientific projects. It was extremely difficult to move that capacity if it was not being used. Now, since the resources are only created when needed, it results in more resources being available for all scientific communities which enables new opportunities.

How does CERN participate in the upstream development community?

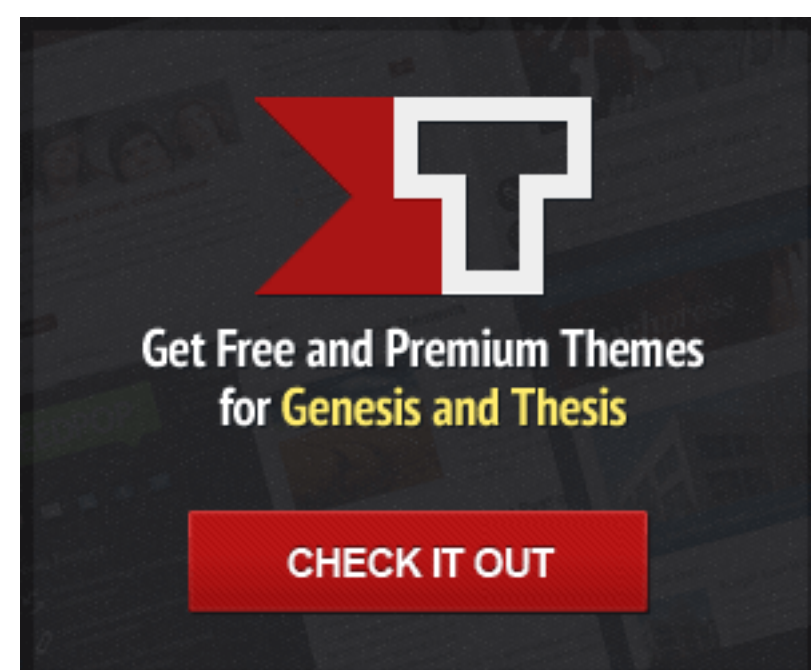
CERN collaborates actively with the upstream development community in several projects. For example, the initial integration with LDAP on Keystone, improvements to cloud federation as part of the CERN openlab/Rackspace collaboration, and lately the scalability enhancements in Nova cells v2.

CERN has been participating in OpenStack design summits since the Boston one (Essex Design Summit). We have very active discussions in the operators community, especially in the large deployments group, where we share experiences and best practice with other deployments.

Also, we participate actively in the RDO project, where we contribute to the packaging efforts. Finally, we share our experiences on our "OpenStack in Production" blog.

What are the largest challenges in getting scientific applications onto your deployment?

We have a few applications that were written many years ago when the cloud or virtualization model



RECENT TWEETS

Follow @d4mations

🐦 New post: 3 open source projects for modern COBOL development <http://t.co/b1be5JfPNn>
4 weeks ago

🐦 Get Yours! I will give you 110 #selfimprovement #boosts #ebook for \$5 on #Fiverr <https://t.co/XvRmvTfrMO>
1 month ago

🐦 Finally Here! I will give you 110 #selfimprovement #boosts #ebook for \$5 on #Fiverr <https://t.co/XvRmvTfrMO>
1 month ago

🐦 Out Now! I will give you 110 #selfimprovement #boosts #ebook for \$5 on #Fiverr <https://t.co/XvRmvTfrMO>
1 month ago

ADS

didn't exist yet. Usually they are monolithic—they can run only in one large server and don't tolerate failure. It's challenging to have these applications running in any cloud environment.

However, for most of these cases we have chosen to move the applications to the cloud infrastructure. We even moved the ones that require very large VMs. Having them running in OpenStack gives us much more flexibility because we can have a common management approach to all physical servers in the data centres.


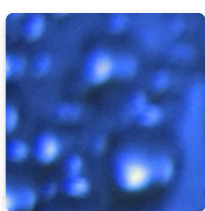

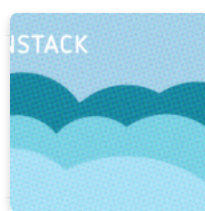
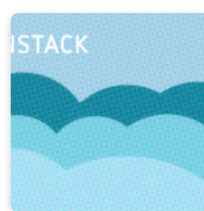
Satisfying the needs of a large, creative, and very compute-savvy scientific community is a challenge, and we closely track emerging technologies like containers. It's challenging, but great fun.

OpenStack Summit
Speaker Interview

This article is part of the Speaker Interview Series for OpenStack Summit Tokyo, a four-day conference for developers, users, and administrators of OpenStack cloud software.

Source link

You might like these too...

				
How do we keep track of ephemeral containers?	Accelerating OpenStack cloud authentication with Keystone tokens	How the Big Tent conversation changed OpenStack	Getting ready for the OpenStack Tokyo Summit	Dispatches from the Tokyo OpenStack Summit
(0)	(0)	(0)	(0)	(0)
Cloud-native	This July, the	Because "cloud" [...]	Interested in [...]	Interested in [...]
[...]	[...]	[...]	[...]	[...]

 ShareThis  ShareThis

Link to this post!

[How CERN uses OpenStack to drive its scient](http://d4mations.com/how-cern-uses-openstack-to-drive-its-scientific-mission/)


TAGS



[← Previous](#)

[Next →](#)

ABOUT THE : D4MER



RELATED POSTS



LEAVE A COMMENT

You must be logged in to post a comment.

