




Cloud

How open source supports CERN's Large Hadron Collider

by Tamlin Magee | Jun 5, 2018 7:10am



The 27-kilometer-long Large Hadron Collider (LHC) buried beneath the France-Switzerland border near Geneva is best known for helping to prove the existence of the Higgs' Boson particle - otherwise known as the God particle - crucial to the Standard Model of particle physics.

The LHC, which uses superconducting magnets to steer beams through its long pipes at just below the speed of light, is supported by open source IT systems at CERN to crunch through about 60 petabytes of data a year. These are built with Openstack, a free and open source software platform for building clouds.

The Openstack cloud first went into production at CERN in July 2013, marking the 13,000-physicist-strong laboratory as an early adopter. Today it has scaled to roughly 300,000 cores – and it's this kind of high-powered, scalable, open source cloud computing that got the attention of many private enterprises, now contributing to the code.

Tim Bell, compute and monitoring group leader at CERN, outlined to Computerworld UK just how the technology world is catching up to the open, collaborative model scientists have depended on for decades.

"It's certainly interesting to compare worldwide collaborations like the LHC, and the associated experiments with open source communities," he said during the Vancouver Openstack Summit last week. "At CERN, we have got 13,000 physicists from 100 different countries working to solve big problems, and it's very much a similar structure you get in open source. You share a common vision and you work together by consensus.



There have been situations recently where CERN has deployed a technology at scale and within months a large corporate has picked it up, and vice versa. **Huawei** for example took notice of CERN's work in putting spot market-type functionality (like the big three public cloud players all have) inside Openstack, and joined a partnership group called CERN Openlab to work on this.

"[They joined] in order to allow us to be working with them, getting this functionality through prototyping, into mainstream, upstream functionality – so they can deploy it as part of their product set," Bell said. "So they're benefiting from the research and we're benefiting from their assistance in doing that work."

On the other side, Bell sat in on a session from AT&T at the summit where the American networking company was talking about how it was running a cloud deployment with Openstack Helm and Openstack Helm Charts on top of containers.

"Because we deployed Openstack about five years ago in production, we were using the more traditional tools at the time," Bell said. "The Puppet deployment models. Since then the container world has come along but we hadn't seen a good reason to be re-architecting the way we had been deploying. But we're now seeing situations where people are genuinely coming up with large scale deployments, at the sort of scale of CERN or even bigger, and using even more agile, Agile 2.0-style tools.

CERN is also partnering with other large-scale scientific endeavors, such as the Square Kilometer Array – a project to model the shape of the universe from the Big Bang to the present. The SKA will have an even higher data rate than CERN when it's operational.

When the two organizations first talked they discovered there were many, many areas with common requirements, the most immediately obvious being operating at massive scale.

"What's good then, is that we are a fairly sophisticated user community, so we can not only come up with requirements, we can also contribute code to implement it," said Bell. "And it's a lot easier way of integration: it's not just a question of making a list of what you want, in some cases we can come along with the solutions and work them through with the design teams.

"SKA are coming online a little later than CERN is, but that means there is collaboration there at both the science levels, but also the compute level, with the team implementing their Openstack clouds," said Bell. "That sort of experience is shared."

But it does also share with other large commercial deployments, such as Oath – the media mega-merger between AOL and Yahoo – which scales "significantly larger" than the CERN installment.

"That sort of experience is also very useful for us," said Bell. "We needn't be the largest cloud and therefore have to find scalability issues, others are there breaking the boundaries as well."

HPC as a service

In the opening keynote sessions, Openstack Foundation COO Mark Collier said that if you want to see where technology is going to develop in the near future, take a look at what CERN is doing today. At the moment the Science Working Group, of which Bell is a member, is doing a lot of work around high performance computing 'as a service'.

That translates to "the ability to bring online bare metal or virtual environments" and then being able to quickly arrange those in "configurations that are appropriate for the application set" for the user.

"We have some applications that are very efficient with eight nodes, but make no sense with 64," Bell said. "What we want to be able to do is take a 64-node cluster and break it up into smaller chunks, configuring right for the application."

This could ultimately look like a user picking just another 'flavor' of cloud – HPC in an online marketplace with a clear user interface.

Bell thinks developments in this area will lead to a spike in HPC as-a-service. And because Openstack is software driven, it should be easier to push automation in this field. For example, by making use of automation, CERN has managed to scale from 30,000 cores to 300,000 with the same size of team.

"Working through with some of the bare metal teams around the [bare metal provisioning program] Ironic project, there's a lot of work there to get it so we can do not only deployment, but also things like when you give your machine back, we make sure it is completely cleaned and reset to its default configuration, so it can be allocated out to someone else.

"All of that process is aiming to be automated, so then you really get to a situation where staff can concentrate on handling things that are abnormal, rather than executing the standard steps for tickets."

Overcoming limitations

The big challenge for CERN in the future is making sure the experiments at the LHC aren't limited by technological constraints.

"The plan for the accelerator in 2023 will be a significant increase in the collision rate, the luminosity," Bell said. "And we're likely to need to the order of 60 times more compute and storage than we have now. Moore's law will get us maybe a quarter of the way there – and we have to work out how to get from one quarter to 100 percent, because it's very important the physics program is not affected by the ability of the computing to solve those problems."

He believes that this is a reflection on the success of the LHC - there was no guarantee at the start that it would be successful, but its performance has "exceeded all expectations".



find ways of expanding the compute capacity to meet that need.

"But equally to be ready in the situations where there are problems... and then we need to make sure the computing resources are used during that time for things like simulation of collisions, to be able to get more data, to validate the real experimental results when they come in."

Is all of that doable?

"Yes," he said.

Computerworld UK

Read More On

cloud computing Huawei

Cloud computing is the top transformative technology for HK IT pros | ComputerWorld Hong Kong

HKBN buys cloud MSP ICG for up to HK\$200m | ComputerWorld Hong Kong

Huawei launches CloudCampus solution | ComputerWorld Hong Kong

Huawei said to be developing mobile OS amid sanction investigation | ComputerWorld Hong Kong

From The Web

by Taboola

Sponsored Links by Taboola

A mechanical watch for a fraction of the price? A swiss start-up is taking on the challenge!

CODE41 Watches

Check if you are eligible for a U.S Green Card

U.S Green Card - Free check

Gima Group - 3D Factory

Gima Group 3D Factory

Father And Son Take Same Photo For 25 Years! Don't Cry When You See The Last One!

12Facts.com



FREE NEWSLETTER

Join 40,000+ of Hong Kong's IT professionals who read ComputerWorld Hong Kong's daily news and updates



Sign up now!

Email Address

SIGN UP



About the Author

Tamlin Magee



GENERAL

- [Home](#)
- [Privacy](#)
- [Terms Of Use](#)
- [Advertise](#)
- [About Us](#)

PARTNER SITES

- [Telecom Asia](#)
- [Networks Asia](#)
- [Enterprise Innovation](#)
- [CFO Innovation](#)
- [CMO Innovation](#)

NEWSLETTERS

- [Subscribe](#)
- [Manage Subscriptions](#)

CONNECT

© 2018 Questex LLC. All rights reserved. 13/F, 88 Hing Fat Street, Causeway Bay, Hong Kong

Reproduction in whole or part is prohibited.