

All the latest UK technology news, reviews and analysis

Cern uses OpenStack to scale as LHC compute requirements grow

by [Daniel Robinson](#)

04 Nov 2014

• Follow V3

Follow @V3_co_uk

- [LinkedIn](#)
- [Newsletters](#)



PARIS: Cern, the European physics laboratory and home of the large Hadron Collider (LHC), is using OpenStack as part of its key IT infrastructure, demonstrating how open source software can evolve to meet even the most extreme customer requirements.

Cern infrastructure manager Tim Bell told the keynote audience at the [OpenStack Summit in Paris](#)

how the organisation is using OpenStack to drive its key IT infrastructure supporting the LHC, and contributing to the open source project by making code available to the wider community.

Bell explained how the infrastructure to handle data from the LHC experiments currently has a 100PB archive, and is growing by 27PB per year.

This has so far been met with infrastructure comprising 11,000 compute nodes, but from 2017 the LHC will have been upgraded to deliver higher energy collisions and will need to collect more data.

"As we look to the future, we're expecting to see 400PB per year by 2023, and compute requirements are expected to be around 50 times the current budget," he said.

Furthermore, upgrading Cern's existing data centre to support these requirements would have been very difficult, as the cooling would not be able to cope with a large number of extra servers.

The solution was to ask the countries that contribute to Cern to propose an external data centre to share the load, and one in Budapest was chosen, especially as there are high-speed telecoms links to it, Bell said.

With the data centre online, Cern faced new challenges in operating across multiple clouds and handling a vastly increased number of compute nodes.

"There are now four OpenStack clouds at Cern, the largest comprising 7,000 cores on approximately 3,000 servers, but this is expected to pass 150,000 cores in total by the first quarter of 2015," he said.

This problem is being approached using a technique called Cells, according to Bell, where compute nodes are

clustered into a larger Cell that is used as a building block to allow the infrastructure to scale.

Meanwhile, the problem of working across multiple clouds is being handled using the federated cloud support that Cern has been developing in partnership with hosting firm Rackspace.

"Last year, [Rackspace joined Cern openlab](#) to help us with the federated identity project, and I'm pleased to say that as a result of this collaboration anyone can now deploy federated ID on OpenStack," Bell said.

This was touted by the OpenStack Foundation as more evidence of the benefits of the collaborative open source approach in delivering solutions.

Earlier at the OpenStack Summit, Bell and his IT team at Cern were announced as winners of the [OpenStack Superuser Awards](#) in recognition of their efforts.

Share this:

- Comment [1 Comment](#)
- Tweet [_](#)
- Google plus [_](#)
- Facebook [_](#)
- LinkedIn [_](#)
- Stumble Upon [_](#)



About [Daniel Robinson](#)

Daniel Robinson is technology editor at V3, and has been working as a technology journalist for over two decades. Dan has served on a number of publications including *PC Direct* and enterprise news publication *IT Week*. Areas of coverage include desktops, laptops, smartphones, enterprise mobility, storage, networks, servers, microprocessors, virtualisation and cloud computing.

- [Follow V3 on Twitter](#)

Sponsored Links by Taboola

An Easy Trick To Naturally Get Your Vision Back

Restore My Vision

7 Wonderful Foods That Will Give You Great Skin

Womanitely

11 Child Celebrities You Wouldn't Recognize Today!

StarFluff



[NetSuite warns businesses against 'fake clouds'](#)



[SAP poaches Microsoft exec for top technology role](#)



[Windows Server 2003 migrations dogged by complexity issues](#)