




Tim Bell on the importance of OpenStack for CERN

Last month, local OpenStack users from France and Switzerland held a meeting at CERN. iSGTW discusses the event with Tim Bell, a member of the OpenStack management board and leader of the Operating Systems and Infrastructure Services group within CERN's IT department. Bell explains the vital role OpenStack plays in enabling CERN to flexibly tailor its computing resources to meet demand, as well as its importance for other organizations across the globe.

Last month, local OpenStack users from France and Switzerland held a meeting at CERN. iSGTW discusses the event with Tim Bell, a member of the OpenStack management board and leader of the Operating Systems and Infrastructure Services group within CERN's IT department. Bell explains the vital role OpenStack plays in enabling CERN to

Posted on JAN 8
2014 10:06AM

flexibly tailor its computing resources to meet demand, as well as its importance for other organizations across the globe...



Andrew Purcell
European editor



OpenStack is a global collaboration of developers and cloud computing technologists producing the ubiquitous open-source cloud computing platform for public and private clouds. Last year, OpenStack held summits in [Hong Kong](#) and [Portland, Oregon](#) (pictured). Image courtesy [Aaron Hockley, Flickr](#) (CC BY-SA 2.0).

What does OpenStack enable you to do?

OpenStack provides a way for people to get computing resources in a very short time: in a few minutes, rather than waiting for months to have a machine delivered. It allows people to come along to a web interface and say 'give me a computer'. Then, in the time that it takes to get a cup of coffee, you can have a computer ready and waiting for you, hosted in the CERN Data Centre.

Share this story



[↻ Republish](#)

Tags

CERN

cloud

computing

interview

open-source

OpenStack

How exactly is OpenStack important to CERN then?

We need to be able to provide large-scale computing resources to physicists around the world. OpenStack gives us a framework through which we can provide that to people with less administration overhead and enables it to be done in a way that is consistent with other OpenStack clouds in other labs around the world.

What people at CERN, and in high-energy physics in general, are looking for is a more flexible, self-service way of obtaining computing resources. People are used to using public clouds and getting resources at the push of a button - albeit with a credit card. As such, what they're looking for from internal IT departments is a similar degree of flexibility. Providing an OpenStack base allows them to have something which is familiar to them, but which can be handled in the scope of allocations or pledges of computing resources, such as those coming from [the Worldwide LHC Computing Grid](#).

Each of [the experiments on the Large Hadron Collider \(LHC\)](#) takes different approaches to handling its data. We're now seeing people exploring exactly how best to use cloud computing. With the service coming into production this year, we're in a phase in which people are trying to understand the best way of using it. The high-level triggers farms at the experiment pits are also running OpenStack and that allows them to opportunistically use those spare resources in the

same way as they use the resources in the CERN Data Centre. So, when the accelerator isn't running, such as at the moment, the experiments can take advantage of that compute resource to do analysis and simulation.

You



"In the time that it takes to get a cup of coffee, you can have a computer ready and waiting for you, hosted in the CERN Data Centre."

mention OpenStack being used in labs around the world... what do you think is the key to its popularity?

I think OpenStack has become very popular because it's an open community, with a lot of very enthusiastic developers working on the code. Also, there's interest from companies, both large and small, in enhancing that code and providing additional services on top of the open stack clouds.

How was OpenStack first developed?

OpenStack initially started in 2010, when [NASA](#) and [Rackspace](#) found that they were commonly

interested in building up an open-source cloud solution. Soon after that, a large number of other companies joined in. There are now around 200 companies contributing and the open-stack code itself is owned by a foundation, rather than being owned by a particular company. The code is licensed and can be used by anyone, subject to open-source rules. Also, anyone can join the OpenStack foundation and become a member, which allows them to vote in the various elections. Equally, members can participate in mailing lists and can contribute code, documentation, and blogs. It's a very vibrant community, which people can freely join and to which people can equally contribute code that can be used around the world in OpenStack clouds.

So can anyone get involved then?

In order to join the OpenStack foundation, [you just go to a web page and fill out your name](#), which immediately allows you to vote in upcoming elections. After that there are different ways of contributing. People with knowledge in [Python](#) can contribute code, so we at CERN are significant contributors to the code, and so are many other organisations and universities. However, there are lots of other modes of participation, as well. For example, there's participation in user groups, sharing experience, *etc.* In terms of just trying things out, building an OpenStack cloud is a 15-to-20-minute activity. You can follow one of the standard guides and have a cloud up and running in that time.

Construction of the Wigner data center

Earlier this year, CERN inaugurated a remote extension to its primary data centre. Located near Budapest, Hungary, the extension hosts an additional 500 servers, 20,000 computing cores, and approximately 5.5 petabytes of storage. [OpenStack allows CERN to manage this additional capacity in a seamless manner.](#) Video © 2013 Prosecutor systems.

Last month, CERN played host to a combined meeting of [the Swiss and Rhône-Alpes OpenStack user groups](#). Could you explain how this model of user groups works?

The OpenStack structure is very global. It's probably split about 40% US, 40% Asia, and roughly 20% in Europe. There are around 100 user groups around the world. These are generally started up by people who are enthusiastic about the technology and they reach out to their communities forming teams ranging between, say, 20 to about 400 people, who get together on a regular basis to discuss their experiences with OpenStack. The

Swiss and Rhône-Alpes user groups have members from academic research, as well as a mixture of industries, including banking, telecommunications, and pharmaceuticals.

Perhaps you could tell the *iSGTW* readers a little more about such collaboration with industry?

Collaboration with industry is a key part of OpenStack. [Rackspace recently joined CERN openlab](#) and is now working with CERN to address some of the issues around getting clouds to interoperate, so that you can just log into one cloud and use resources from another cloud. This is a very good example of where research has a difficult problem and works with industry in a framework such as [CERN openlab](#). The code is then contributed back to the community, so it becomes available for any OpenStack cloud around the world to use.

Finally, at last month's user-group meeting, there was much discussion about a recent OpenStack summit in Hong Kong. How do these events work?

Every six months, the OpenStack foundation organizes summits. The aim of these events is to get together the people who are working on the project to define the roadmap for the next six months. It also allows users and companies who are offering services a chance to get together and better understand the latest solutions available. The last meeting in Hong Kong had about 3,800 people and

we're generally seeing an increase in attendance of about 50% every six months. The next meeting will be in Atlanta, Georgia, USA, and after that they'll be coming to Paris, France, in autumn 2014. It's a very global event.

Read more about the joint Swiss and Rhône-Alpes user group meeting at CERN [here](#).

Join the conversation

Contribute



Do you have story ideas or something to contribute? **Let us know!**

OUR UNDERWRITERS

Thank to you our underwriters, who have supported us since the transition from International Science Grid This Week (iSGTW) into Science Node in 2015. We are incredibly grateful.

CATEGORIES

Advanced computing
Research networks
Big data
Tech trends

CONTACT

Science Node
Email: editors@sciencenode.org
Website: sciencenode.org

[View all underwriters](#)

[Community building](#)



[GIVE NOW](#)

Copyright © 2022 Science Node™ | [Privacy Notice](#) | [Sitemap](#)

Disclaimer: While Science Node™ does its best to provide complete and up-to-date information, it does not warrant that the information is error-free and disclaims all liability with respect to results from the use of the information.