



# CERN Uses OpenStack

---

“With our OpenStack cloud, we’ve been able to offer physicists the resources they need to analyze and model data from the LHC [Large Hadron Collider].”

**Tim Bell, manager, infrastructure services, CERN**

---

## Unleashing Secrets of the Universe with Help from OpenStack

When one of the world's most prestigious research laboratories decided to embrace cloud computing, it chose OpenStack.

Using the world's largest, most complex scientific instruments, CERN, the European Organization for Nuclear Research, continues its study of the most basic constituents of matter – fundamental particles – to help scientists gain a better understanding of the very structure of the universe. Recently, after years of searching, calculations, speculation, and smashing atoms together, CERN together with the ATLAS and CMS experiments announced that the elusive Higgs Boson that provides mass to elementary particles had been found.

An exciting time in science, indeed. And as one might suspect, all of this research generates tremendous amounts of data, which require thousands of computers to process. Using highly-specialized algorithms, the data from the most interesting – and scientifically promising – events are transferred to CERN's data center for processing. The data are saved to tape storage and then distributed to the more than 150 sites worldwide that comprise the Worldwide LHC (Large Hadron Collider) Computing Grid (WLCG) for analysis.

The instruments CERN uses to smash these particles and collect data are as massive as they are complex. The accelerators stimulate beams of particles to high energies before the beams are made to collide. The LHC detectors observe and record the results of these collisions, which generate up to an astonishing one petabyte (that's a quadrillion bytes) every second before filtering. One of the detectors at the LHC is more than 7,000 tons measuring over 5 stories high.

For years, commodity computing devices kept pace with these massive data processing needs. For instance, the CERN data center in Geneva boasts 10,000 servers, 80,000 disks, and 100 petabytes of data stored in mass-storage systems– but it has reached its power and cooling capacity. In Budapest, Hungary, an additional similar data center is online with 200 Gbp/s network connections connecting both sites. Still, as CERN's data processing needs continued increasing while staffing remained fixed, the organization decided that cloud computing could provide a complementary approach to deliver services at scale to its physicists.

## Ready for Cloud

The timing for that decision proved ripe. Cloud computing technologies had begun to mature to the point that they could help handle such demanding workloads. Also, CERN users were growing increasingly familiar with using public clouds to provide additional capacity during peak times. “Providing a modern infrastructure-as-a-service private cloud that had similar capabilities to those available in public clouds would support the data and compute demands of their research,” says Tim Bell, manager of infrastructure services at CERN.

CERN began to investigate various virtualization and cloud platforms. CERN's IT team decided it would build a private cloud that would need to integrate well with a very heterogeneous environment. CERN's service consolidation

environment is based on the Microsoft Service Center Virtual Machine Manager and the Hyper-V hypervisor. The IT team also built a cloud test bed based on OpenNebula and KVM (Kernel-based Virtual Machine) hypervisor. “Multiple hypervisors are attractive with respect to support models, performance analysis, and flexibility. The choice of hypervisor should be a tactical one not determined by the cloud infrastructure, and we plan to run a mixed hypervisor environment in the future,” says Bell.

As CERN investigated potential components for new infrastructure tools and processes in 2011, it reviewed a number of candidates. In the end, CERN selected OpenStack. As a cloud platform, OpenStack controls and automates pools of compute, storage, and networking resources to turn standard hardware into a powerful cloud computing environment. Today, OpenStack is the fastest growing open cloud community, working to build software that powers public and private clouds for a growing number of organizations, including Cisco WebEx, Comcast, eBay, HP, Intel, MercadoLibre, NeCTAR and Rackspace.

“OpenStack’s technical architecture clearly addressed our needs to run at scale,” says Bell. “Also, the technology and developer ecosystem around OpenStack are very vibrant and would enable us to build the services we needed within the cloud. With an open community, we can benefit from the work of the active contributors but also use our engineering skills to enhance the product for others.”

A critical aspect of CERN’s move to cloud computing was the ability to interact with all of its existing IT services. “Our network infrastructure is based on a home-grown framework to meet the needs of running a laboratory with millions of IP devices. We were able to extend OpenStack to support dynamic allocation of network addresses and register those with our network management system. Further extensions for creating DNS entries, Kerberos, and X.509 certificates were also implemented,” he says.

## Building the CERN Private Cloud with OpenStack

CERN’s IT department started working on OpenStack toward the end of 2011, building test clouds for physicists to explore cloud technologies and test integration with CERN specific customizations. Using Scientific Linux, developed by CERN and Fermilab based on the Red Hat distribution, a cloud was rapidly built with Compute, Image, Identity and Dashboard services.

This environment has gradually grown to over 400 hypervisors with high availability on all OpenStack controllers using recipes from the community. Planning and testing is now ongoing to expand the cloud to 15,000 hypervisors with over 150,000 VMs in the next 18 months. “Seeing the size of several deployments in production that are already larger than our target validates our approach to be able to share and to benefit from others”, said Bell.

During 2012, a team at the CMS detector at CERN, was planning their activities for the 2 year upgrade of the LHC. Currently, a compute farm of 1,300 servers and around 13,000 cores is installed to filter the data from the detector before sending it for recording at the CERN computer centre. However, during the upgrade, this farm would not be required.

In less than two months, the team created a proof-of-concept cloud based on OpenStack with a controller, a distributed authentication service, two Compute nodes, and a node for Image Services. They used MySQL for their database, RabbitMQ for messaging, and because KVM had been used before in CMS, it was chosen as the hypervisor. “The deployed infrastructure was very stable and the OpenStack API layer makes it possible for us to fully manage the entire VM lifecycle,” says Jose Antonio who led the work to implement the cloud at CMS.

At the beginning of October 2012, the team performed a large-scale test with OpenStack controlling about 1,200 hypervisors. The OpenStack environment deployed 720 virtual machines in roughly an hour – about one every 5 seconds. “The infrastructure and networking were stable. The cloud controller, a server with four cores, was at the limit of CPU usage,” says Wojciech Ozga, who completes the team that deployed OpenStack on the CMS cluster. Since the initial deployment, the infrastructure controller has been upgraded to a 16 core, 48GB RAM server, allowing greater throughput and is now serving physicists of the CMS experiment.


## The Vibrant OpenStack Community

OpenStack attracts experts in various technologies and disciplines around the world to collaborate on the platform. “A

big part of that success, and anticipated future successes, is the deep, active, and highly collaborative OpenStack development community. An example came from our collaborations with PuppetLabs, which produced a set of Puppet configuration recipes to allow easy, yet very specific, OpenStack deployments. Building on top of those recipes, we provided the configuration options that we needed, such as SSL and Red Hat support. We incorporated these changes back into the PuppetForge community for others to use,” Bell says.

Another example is the significant identity management challenges at CERN. Currently, CERN has 44,000 users registered in its identity management system, and more than 400 users are added or removed each month. “Depending on their roles in the organization, some of these users will have rights to be administrators or project members of the private cloud,” Bell explains. While an earlier version of OpenStack provided basic LDAP support, a number of enhancements were required for Active Directory support at the scale CERN required. “Working with the OpenStack team, we enhanced the LDAP support to cover our use case – and these changes will be included in an upcoming version of OpenStack. These contributions illustrate the benefits of open communities and peer reviews that ensure both use case demand and software quality,” he says.

And the cloud deployment has been a smashing victory. Today, multiple clouds at CERN successfully run collision reconstructions on OpenStack. “Cloud technology has allowed us to be much more responsive to our user community allowing them to explore the frontiers of science without waiting for hardware to be delivered and configured”, says Bell.



## CERN

<http://home.web.cern.ch>

**Industry:** Academic / Research

**Headquarters:** GENEVA

**Size:** 5,001 to 10,000 employees

---

## OpenStack technologies CERN uses:

Openstack Compute (Nova)
Openstack Network
Openstack Dashboard (Horizon)
Openstack Identity Service (Keystone)
Openstack Image Service (Glance)

---

## Links About CERN

[🔗 What do the NSA and CERN Have in Common? \(http://www.youtube.com/watch?v=DaGDaltXDSg\)](http://www.youtube.com/watch?v=DaGDaltXDSg)

[🔗 Overlay Opportunistic Clouds in CMS/ATLAS at CERN \(http://www.youtube.com/watch?v=jRkTVh27XBQ\)](http://www.youtube.com/watch?v=jRkTVh27XBQ)

[🔗 OpenStack Users Speak Up \(including CERN\) \(http://www.youtube.com/watch?v=k1Gpltp4RQQ\)](http://www.youtube.com/watch?v=k1Gpltp4RQQ)