

## This content was printed from **Data Center Knowledge**



The team overseen by Tim Bell, infrastructure manager at CERN (pictured), is preparing to upgrade its OpenStack cloud from 115,000 to 150,000 compute nodes to handle an increase in data that will be generated by the Large Hadron Collider when its energy is doubled in 2015.

Feedback

[Big Data](#), [Cloud Computing](#), [Europe](#), [Featured](#)

### CERN's OpenStack Cloud to Reach 150,000 Cores by 2015

BY YEVGENIY SVERDLIK ON NOVEMBER 7, 2014

**PARIS** – Building the Large Hadron Collider itself was doubtless a massive feat, but the machine – a nearly 17-mile ring more than 300 feet underground on the Franco-Swiss border – is useless without the huge data storage and computing capacity needed to analyze the ungodly amount of data it generates when microscopic particles get recorded smashing into each other at extreme speeds.

That computational power at the European Organization for Nuclear Research (CERN) is delivered through four cloud environments the organization's IT team created using OpenStack, the suite of open source cloud software that is quickly becoming the industry standard for building clouds. Tim Bell, CERN's infrastructure manager who oversees that team, spoke about the organization's cloud during a keynote at this week's OpenStack summit in Paris.

CERN currently has four OpenStack clouds living in two data centers – one in Meyrin, Switzerland, where Bell's office is located, and the other in Budapest, Hungary, a remote business continuity site for the primary Swiss facility. Largest of the four has about 70,000 compute cores on about 3,000 servers. The other three comprise a total of 45,000 compute cores.

Bell's team started building its cloud environment in 2011 using the Cactus release of the open source cloud software. They went into production with the Grizzly release in July 2013. Today, all four clouds run the Icehouse release. Bell said he has about 2,000 additional servers on order to increase the cloud's capacity since the upcoming increase in energy of the particle beams inside the collider mean it will generate more data than the 1 petabyte per day it already generates when running at its current capacity.

CERN's cloud's architecture is a single system that scales across two data centers. Each data center, in Switzerland and in Hungary, has clusters of compute nodes and controllers for those clusters. Both controller "cells" speak to a master controller cell in Switzerland, and upstream from the master controller cell is a load balancer.

An OpenStack cloud is never built using just components of the OpenStack suite, and CERN's cloud is not an exception. Other tools in the box (all open source) are:

- Git: a software revision control system
- Ceph: distributed object storage that runs on commodity servers
- Elasticsearch: a distributed real time search and analytics system
- Puppet: a configuration management utility
- Kibana: a visualization engine for Elasticsearch
- Foreman: a server provisioning, configuration, and monitoring tool
- Hadoop: a distributed computing architecture for doing big data analytics on commodity-server clusters
- Rundeck: a job scheduler
- RDO: a package of software for deploying OpenStack clouds on Red Hat's Linux distribution
- Jenkins: a continuous integration tool

Bell's team recently devised a federated cloud system that currently works across CERN's cloud and the public cloud resources by Rackspace, whose cloud is also built on OpenStack. Users can deploy federated identity across Rackspace cloud and CERN's private cloud. In the future, Bell foresees more public clouds and other researchers' clouds to be able to gel with CERN's.

CERN's OpenStack environment is already massive, and it is going to become even bigger when the collider is upgraded to double its energy in 2015 so scientists who have already used LHC to confirm existence of the Higgs boson (50 years after Peter Higgs and five other physicists theorized about the "God particle's" existence) can continue looking for answers to some of the most fundamental questions about the universe. Why is gravity so weak? Are there dimensions we're not aware of? Do gravitons, the currently hypothetical guardians of gravity, exist?

Those are things physicists worry about when they wake up in the morning, Bell said. What fills *his* mornings with worry is how to make sure those physicists have an IT environment capable of crunching through the amount of data required to answer questions of that magnitude.

## About the Author



**Yevgeniy Sverdlik** (332 Posts)

San Francisco-based business and technology journalist. Editor in chief at Data Center Knowledge, covering the global data center industry.

### RESOURCE LINKS:

- [Disaster Safe Utah Data Center - 10Kw up to 10 MW.](#)
- [Download expert white paper – "Best Practices and Critical Considerations for Choosing the Right Colocation Solution"](#)

### Inside DCK

About Us  
Advertise  
Staff  
Contact Us  
Submit News  
Submit Guest Column  
Site Map

### Hot Topics

Data Center Infrastructure Management  
Downtime  
HPC / Supercomputing  
Facebook  
Apple  
Microsoft  
Google  
Northern Virginia  
New York  
Silicon Valley  
North Carolina

### News Channels

White Papers  
Events Calendar  
Virtualization  
Consolidation  
Open Compute  
Internet of Things  
Storage  
Convergence  
Disaster Recovery  
Managed Hosting  
Content Delivery  
Site Selection  
Investing

### Stay Connected

Daily Email Newsletter  
RSS  
Twitter  
Facebook  
LinkedIn  
Data Center Videos  
Humor  
Webinars

Visit our other properties:

**HOSTING AND CLOUD** Web Hosting Talk, HostingCon, Web Host Industry Review, Hosting Catalog, Host Voice

**WEB DEVELOPMENT** Hot Scripts, DB Forums

**PERFORMANCE MARKETING** ABestWeb

**CONSUMER TECH** Windows Secrets, Overclockers, Mac Forums

**DATA CENTER** Data Center Knowledge, Data Center World, AFCOM

Copyright © 2014 iNET Interactive. All rights reserved.

