



Organized by





Under the patronage of

AGID Agenzia per l'Italia Digitale Sponsored by









Operating OpenStack at CERN

Jan van Eldik / CERN IT

Organized by





Under the patronage of

GAGID Agenzia per l'Italia Digitale Sponsored by









CERN's primary mission: **SCIENCE**

Fundamental research on particle physics, pushing the boundaries of knowledge and technology



CERN

World's largest particle physics laboratory

ATLA

ALICE

Image credit: CERN

CERN Reformer

27 km



The Large Hadron Collider: LHC

dipole magnets 15 metres 35t EACH

230

27km



ATLAS, CMS, ALICE and LHCb







40 million

pictures per second **1PB/S**

> Run: 297041 Event: 59057181 2016-04-24 05:41:50 CEST



Data Flow to Storage and Processing



RUN 2 ALICE: 4GB/s ATLAS: 1GB/s CMS: 600MB/s LHCB: 750MB/s



CERN Data Centre: Primary Copy of LHC Data





WLCG: LHC Computing Grid

About WLCG:

- A community of 10,000 physicists
- ~250,000 jobs running concurrently
- 600,000 processing cores
- 700 PB storage available worldwide
- 20-40 Gbit/s connect CERN to Tier Is

Tier-0 (CERN)

- Initial data reconstruction
- Data recording & archiving
- Data distribution to rest of world

Tier-1s (14 centres worldwide)

- Permanent storage
- Re-processing
- Monte Carlo Simulation
- End-user analysis
- Tier-2s (>150 centres worldwide)
- Monte Carlo Simulation
- End-user analysis



CERN Data Centre: Private OpenStack Cloud



```
Back in 2011...
Managing all this became...
      ....very...
            ....very...
                  ....very...
                       ...tricky...
```



2012: Agile Infrastructure project

- Provisioning + Configuration + Monitoring
- Aim: virtualize all the machines
 - Unless really, really, really not possible
- Offer Cloud endpoints to users
- Scale horizontally
- Consolidate server provisioning
 - Yes, we use the private cloud for server consolidation usecases as well



CERN Tool Chain 🚯 git elasticsearch. 6 Grafana openstack[™] CLOUD SOFTWARE **CentOS** ceph RDO hadoop puppet 🤪 Jenkins FOREMAN = RUNDECK kibana











Settings

Sign Out

Logged in as: jveldik



System Panel Overview Instances Services Flavors Images



Overview

October

• 2012

- Submit

Active Instances: 1218 Active Memory: 685GB This Month's VCPU-Hours: 232987.37 This Month's GB-Hours: 180227.04

Usage Summ	Download CSV Summary				
Project Name	VCPUs	Disk	RAM	VCPU Hours	Disk GB Hours
boinc	871	20	433GB	217987.71	5039.11
tests	7	90	9GB	1920.48	3622.21
openstack	71	940	92GB	8834.04	171565.73
batch	300	-	150GB	4245.13	0.00
Displaying 4 items					

CERN

Production in Summer 2013



OpenStack production service since August 2013





Cell	Nodes	Cores	VMs
Geneva	373	10912	1359
Wigner	291	9312	368
Total	664	20224	1727



Rich Usage Spectrum ...



- Personal VMs
 - Development

... rich requirement spectrum!



CERN OpenStack Service Timeline

















OpenStack Magnum

An OpenStack API Service that allows creation of container clusters

- Use your OpenStack keystone credentials
- You choose your cluster type
- Multi-Tenancy



an OpenStack Community Project

 Quickly create new clusters with advanced features such as multi-master

Today, the CERN Cloud hosts 500 Kubernetes clusters









Ironic Overview

- Bare Metal Project in OpenStack
 - Provision 'physical' instances
 - Allows Compute service to manage and provide physical servers as if they were virtual machines
 - User interfaces with Nova (which also provides quotas, scheduling,)
- Hardware management possible
 via common interfaces (& vendor-specific ones)
 - PXE, IPMI
 - Allows for unified interface to manage

heterogeneous machine park (~50 h/w types at CERN for hypervisors only!)

IRON User Requests physical Nova instance API + Scheduler picks Nova Ironic API + Conductor Compute enrolls Glance Admin Neutron **Physical Servers**

3300 Servers are managed using OpenStack Ironic



A new use case: Containers on Bare-Metal

- OpenStack managed containers and bare metal so put them together
- General service offer: managed clusters
 - Users get only K8s credentials
 - Cloud team manages the cluster and the underlying infra
- Batch farm runs in VMs as well
 - 3% performance overhead, 0% with containers
 - Evaluating federated kubernetes for hybrid cloud integration
- Large scale testing starting now



Integration: seamless!

(based on specific template)

Monitoring (metrics/logs)?

- \rightarrow Pod in the cluster
- \rightarrow Logs: fluentd + ES
- \rightarrow Metrics: cadvisor + influx



Community Experience

- Open source collaboration sets model for in-house teams
- External recognition by the community is highly rewarding for contributors
- Reviews and being reviewed is a constant learning experience
- Productive for job market for staff
- Working groups, like the Scientific and Large Deployment teams, discuss wide range of topics
- Effective knowledge transfer mechanisms consistent with the CERN mission
 - >100 outreach talks since 2011
- Dojos at CERN bring good attendance
 - Ceph, CentOS, Elastic, OpenStack Day, ...







Compute needs growing...

- With the needs of LHC computing in future years, efficient and flexible delivery of compute resources will be key
 - Computing needs in 2023 estimated at 60x the current capacity (HL-LHC)





Commercial Clouds







Development areas going forward

- □ Spot Market / Pre-emptible instances
- Software Defined Networking
 - Introducing LBaaS this month
- Magnum rolling upgrades
- □ Enrolling all 15K servers in Ironic
 - Containers on Bare Metal



_ _ _

The next 5 years?

- □ LHC to Run 4
 - Computing must not limit the physics
- Open Infrastructure
 - OpenStack is a key part but lots of others too
 - Ceph, Tungsten Fabric, Grafana, Puppet, CI/CD, K8S, ...
- Open Source collaboration is the way forward
 - Natural culture fit for sharing and giving back
 - Work with the communities which are open and vibrant



Summary

- OpenStack has provided a strong base for scaling resources
- Additional functionality on top of pure Infrastructure-as-a-Service is now coming to production
- Community and industry collaboration has been productive and inspirational for the CERN team
- □ Some big computing challenges up ahead...



Thank you!





Further Information

- CERN OpenStack blog at <u>http://openstack-in-production.blogspot.com</u>
- Recent CERN OpenStack talks at Vancouver summit at

https://www.openstack.org/videos/search?se

arch=cern

CERN Tools at <u>https://github.com/cernops</u>

