
CERN

The CERN OpenStack Cloud

Compute Resource Provisioning for the Large Hadron Collider

Jan van Eldik

for the CERN Cloud Team

Cloud Service 101

- **Cloud Computing**
- **Why a CERN Private Cloud?**
- **Cloud Components**
- **Usage**

Cloud Computing

- Cloud provide Infrastructure-as-a-Service
  - Compute, Storage, Network resources
- Cloud implies:
  - APIs
  - Elasticity
  - Accounting
  - …
- Clouds are ubiquitous:
  - Public clouds: Amazon, Microsoft, Google, …
  - Private clouds on many sites

Why a CERN private Cloud?

Advantages for Service Providers:

- Consolidate resource provisioning
  - Offer single pane-of-glass
- Optimize resource utilization
  - Pack Virtual Machines on physical nodes
  - Thinly provision / overcommit resources
- Reduce operational cost
  - Single team to look after resource life-cycle
  - Migrate workloads to facilitate upgrades/repairs

Why a CERN private Cloud?

Advantages for customers:

- Rapid provisioning, based on agreed quota
- Scale-out possibilities
- Easy to build redundant services
- Possibility to deploy specific images
- No hardware to worry about

In general: loose coupling between IaaS and Application layers

The sky is cloudy

- Many HEP sites offer Cloud endpoints
- WLCG experiments consume cloud resources
  - Public, HEP, CERN private, Online farms, …
- Hybrid clouds
  - Size private cloud for regular load, buy public resources in times of peak demand
- Federated clouds
  - Accept federation credentials to allow seamless resource sharing.

CERN private cloud details

- Based of OpenStack Cloud software
  - Open-source project, started in 2011
  - Enormous traction in industry
- Production service at CERN since 2013
  - 2600 users, 600 shared projects
- Growing list of components
  - Compute: 8000 servers, 250K cores, 28K VMs
  - Block Storage
  - Container Orchestration
  - Networking
  - …

Project types

- Personal
  - Limited resources, tied to individuals
  - Intended for test/dev work
  - Resources deleted once person leaves
- Shared
  - Agreed resource quota
  - Non-standard flavors
  - Tied to experiments, services
  - Production workloads
  - Possibility for optimization
    - CPU optimized, higher IOPS, Critical power, Wigner, …

Interfaces

- API access
  - OpenStack API
  - EC2 API
- CLIs
  - Installed on lxplus
  - And on lxplus-cloud for newer versions
- Web interface
  - https://openstack.cern.ch

Compute - Nova

Manage Virtual Machines

- Create VMs in a Project
- Based on images
  - Linux, Windows, CernVM, private,…
- “Flavours” to describe virtual hardware
- VNC console access
- Contextualize VMs at boot time
  - Install experiment software, run Puppet, …

Compute – Nova - II

Advanced operations

- Snapshot VM, boot new “clones”
- Deploy VMs across multiple Availability Zones
  - Increase resilience against infrastructure failure
- Upon request:
  - Special flavours
  - Placement on Critical Power or Wigner DC
  - CPU optimized configurations
  - …

Block Storage - Cinder

Manage Volumes

- Create volumes, attach to VMs
- Boot VMs from volumes

Upon request

- Higher-IOPS volumes
- Volumes in Critical Power or Wigner DC

Notes

- Volume service backed by Ceph
- FSaaS in beta testing

Other interesting features

- Container Orchestration
  - Kubernetes, Docker Swarm, Mesos, …
  - See next talk ☺
- VM Orchestration
  - Auto-scaling deployments, healing
- Networking
  - Very much WIP
- Bare Metal provisioning

CERN Cloud in Numbers (1)

- **7000** hypervisors in production + **~2000** being added (5.8K 1yr ago)
- **220K** cores + **86K** to be added (155k)
- 3.7K volumes with 1.2 PB allocated (Cinder) (2.8K)
- 3.8K images/snapshots (Glance) (2.7K)
- 27 fileshares with 18 TB allocated (Manila) (new)
- 71 container clusters (Magnum) (new)

CERN Cloud in Numbers (2)

- ~ **400** operations/min in Nova API
  - Mainly creation/deletion of VMs
- **~27K** VMs at the moment (**+5k** 6 months ago)

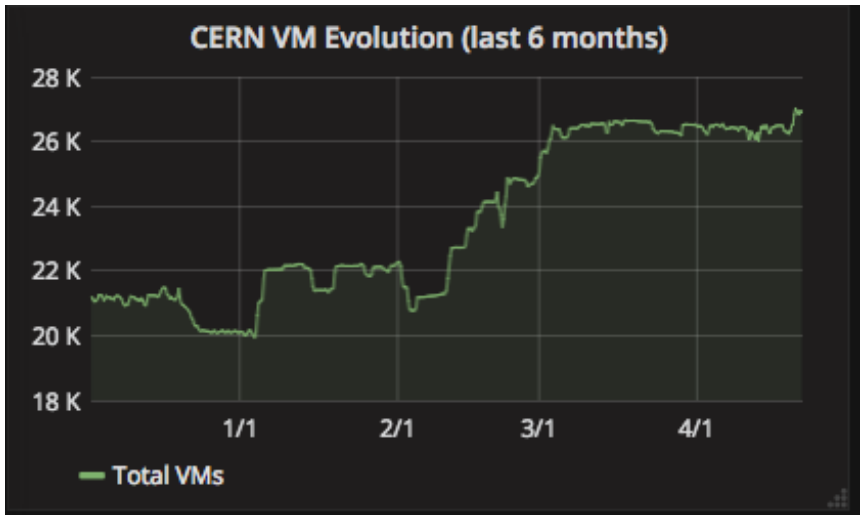

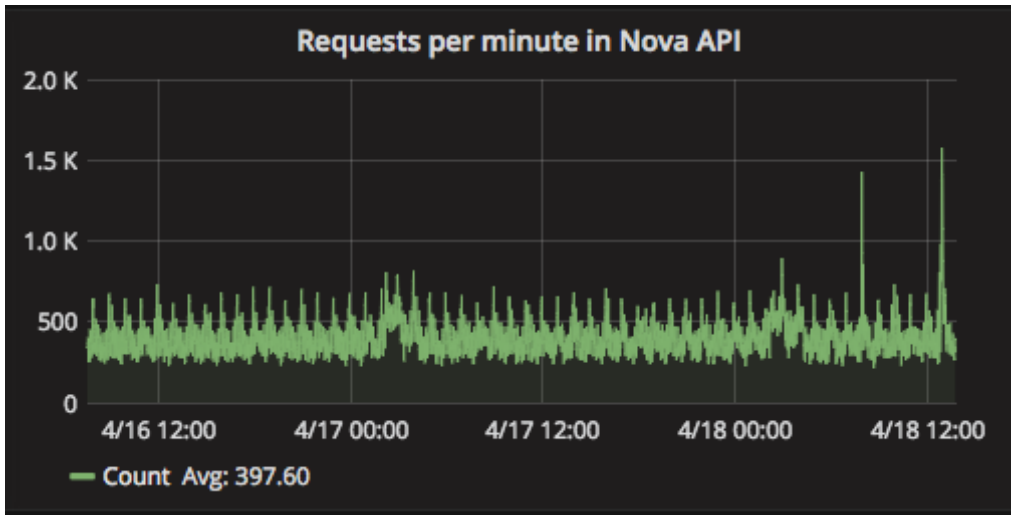

LHCb projects

Shared Tenant	Instances			CPUs		
	Used	Quota	Used%	Used	Quota	Used%
LHCb Build servers	39	47	82	245	248	98
LHCb CLOUD.CERN.ch	365	502	72	365	502	72
LHCb Cloud Workers	0	100	0	0	12	0
LHCb ICE-DIP ESR2	1	50	2	4	50	8
LHCb LHCb Analysis preservation	12	25	48	24	25	96
LHCb LHCbDirac services	14	14	100	196	208	94
LHCb Lc2pXX analysis	1	1	100	8	8	100
LHCb Reproducible Experiment Platform	4	8	50	12	12	100
LHCb VOboxes	30	90	33	105	216	48
LHCb Vcycle	5	20	25	5	20	25

In summary

- Cloud service is open to all CERN users
- Hosting Lxplus, Lxbatch, Terminal Servers, Web servers, Build servers, …
- As well critical services for experiments and organization
- Container services are taking off!

Homepage, user guide:

https://openstack.cern.ch


CERN

OpenStack at CERN

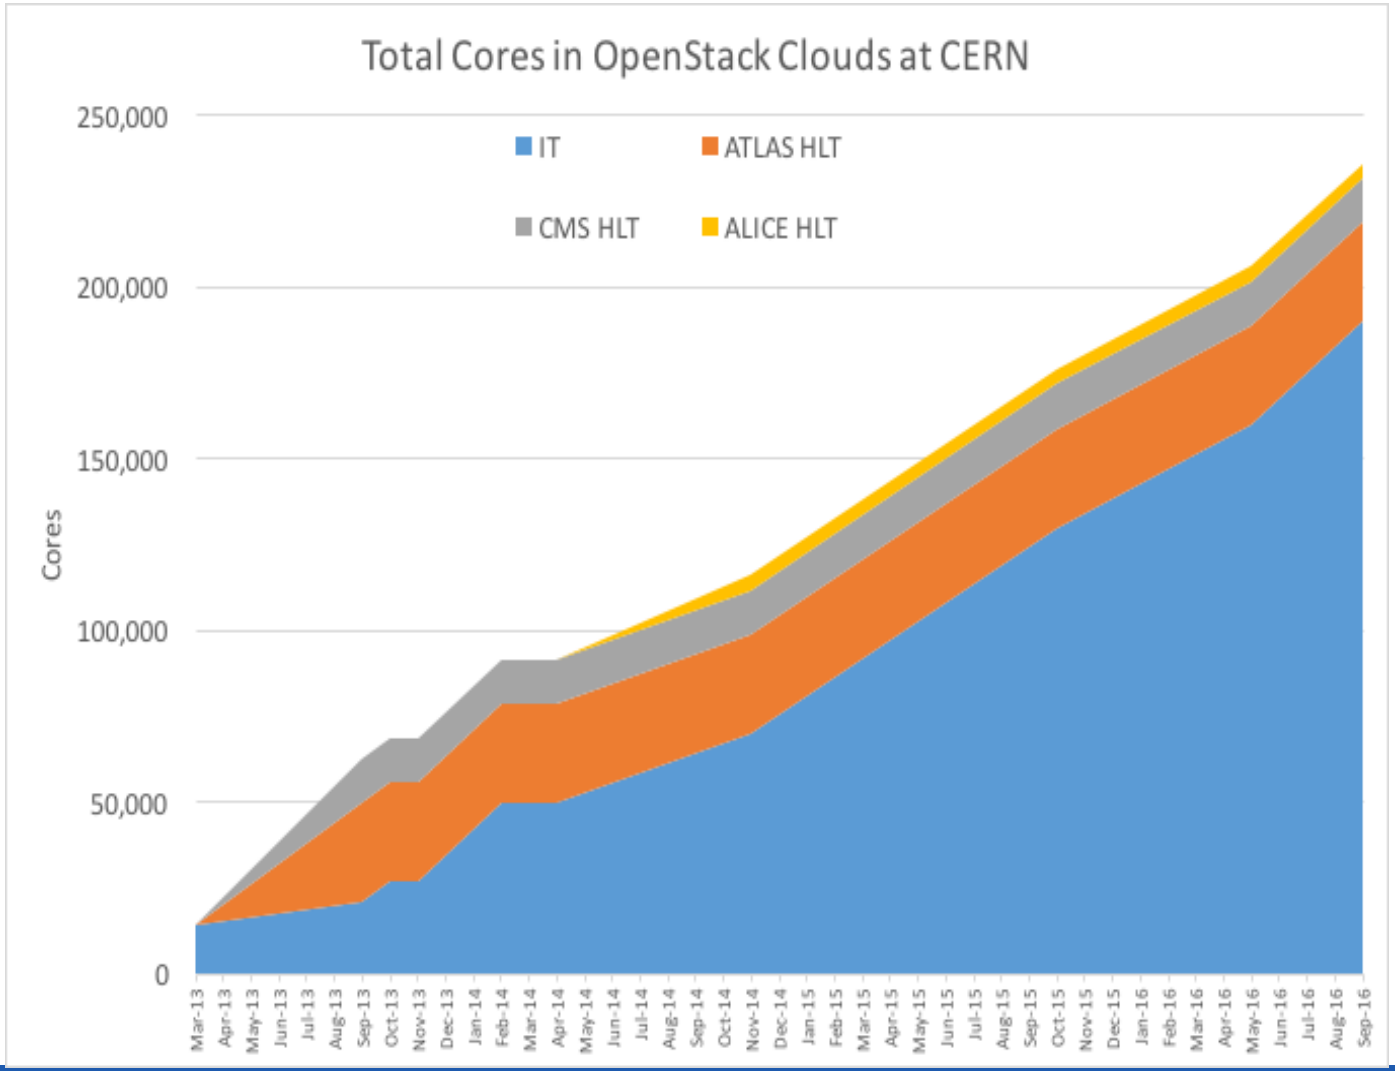

In production:

- 4 clouds
- >200K cores
- >8,000 hypervisors

90% of CERN's compute resources are now delivered on top of OpenStack

Rich Usage Spectrum …

- Batch service
  - Physics data analysis
- IT Services
  - Sometimes built on top of other virtualised services
- Experiment services
  - E.g. build machines
- Engineering services
  - E.g. micro-electronics/chip design
- Infrastructure services
  - E.g. hostel booking, car rental, …
- Personal VMs
  - Development

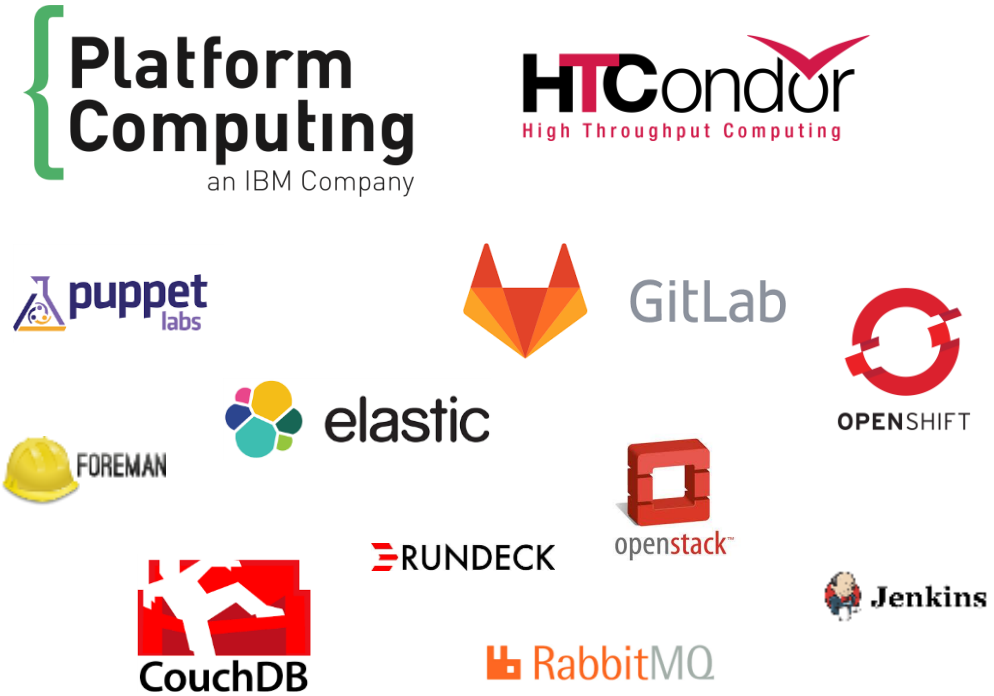

… rich requirement spectrum!