


Openlab Compute Provisioning Topics

Tim Bell
1st March 2017

Agenda

- Review of past Openlab activities
- Some ongoing cloud work
- Research areas
- Open discussion

Current Private Cloud Status




- >200K cores in production under OpenStack
- >90% of CERN compute resources are virtualised
- >60K cores to be added in the next 6 months

Past Activities with Rackspace

- Provide identity federation
 - Use resources on one cloud being authenticated in another
- Code is part of the standard OpenStack release
 - Other companies have been enhancing it
- Used in production at CERN and other labs
 - Horizon 2020 Indigo Datacloud project with 28 partners

Ongoing : Rackspace Containers

- OpenStack Magnum provides containers as a service
 - Kubernetes
 - Mesos DC/OS
 - Docker Swarm
- Single pane of glass for containers and virtual machines
 - Simplified user experience
 - Consistent accounting, permissions
- Openlab fellow now a core contributor




Ongoing : Private Cloud Service

- Investigating further functions
 - Filesystem as a service
 - Secrets as a service
 - Bare metal management
 - Ensure clean systems on return
 - Clear disks
 - Reset BIOS
 - Automate provisioning from spare pool
 - Single inventory for all compute resources with accounting and ownership

Ongoing : Public Cloud Tests

- Various tests of public clouds have been performed over the past years
 - CERN IT co-ordinated
 - Experiments directly
- Current activities run from CERN IT
 - Oracle Bare Metal cloud using Terraform
 - HNSciCloud
 - T-Systems, Huawei, Cyfronet, Divia
 - IBM
 - RHEA Group, T-Systems, exoscale, SixSq

Research : Tetris



- Over 20 different hardware configurations
- Need to give users a simple experience
 - “Give me 2 cores, 4GB memory and 20GB SSD”
- Need to ensure good utilisation (but not too much)
- Potential topics
 - Spot market instances – pre-emptible
 - I/O intensive, Memory intensive flavors
 - Rack Scale allocation under Bare Metal management

Research : Scalable Clouds

- Aim to expand to HL-LHC scale within the staffing envelope
- Investigate process and technical improvements
 - OpenStack Cells
 - Software defined networking (Neutron)
 - Quota
 - Capacity planning
 - Accounting and monitoring

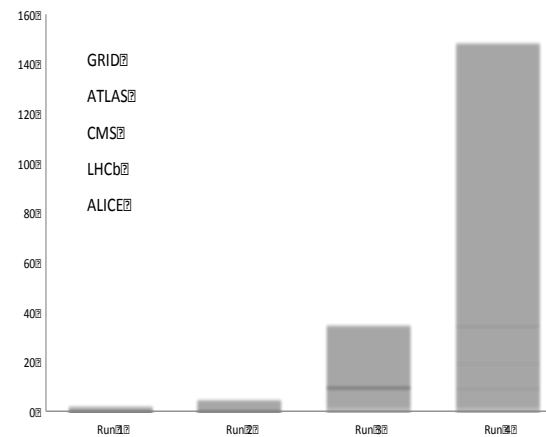


Research : Global Scientific Clouds

- OpenStack Scientific Working Group initiated a discussion at the Barcelona summit in Fall 2016
- “Boston Cloud Congress” workshop in May 2017 (after OpenStack summit) to identify the barriers
- Wide range of participants
 - Universities and research labs (US, Asia, Europe)
 - NRENs / EGI / NSF / Internet 2
 - Public cloud, distributions and on-premise providers
- Early days to understand potential impact on LHC computing models and industry

Thoughts?

Needs : LHC compute growth

The outline LHC schedule out to 2035 presented by Frederick Bordry to the SPC and FC June 2015 can be found [here](#)



Compute: Growth > x50

What we think is affordable unless we do something differently

Research : Global Scientific Clouds (II)

- Is this something that could be exploited?
- Challenges to be covered
 - Reference architecture
 - Security Considerations (digital rights management, IP protection, privacy sensitive data management, etc...)
 - Authorized Shared Use Facilitation
 - Data and Resource Federation (network/storage/compute)
 - Sharing and Business Facilitation
 - Interoperability and Shared-use Policy

Links

- [CERN OpenStack in Production blog](#)
- [HNSciCloud](#)
- [OpenStack Scientific Working Group](#)
- [Boston Cloud Congress](#)
 - [Documents](#)
- [Openlab technical workshop \(Dec 2016\)](#)