



CERN
openlab



Huawei and Rackspace cloud project highlights

IT R&D Meeting

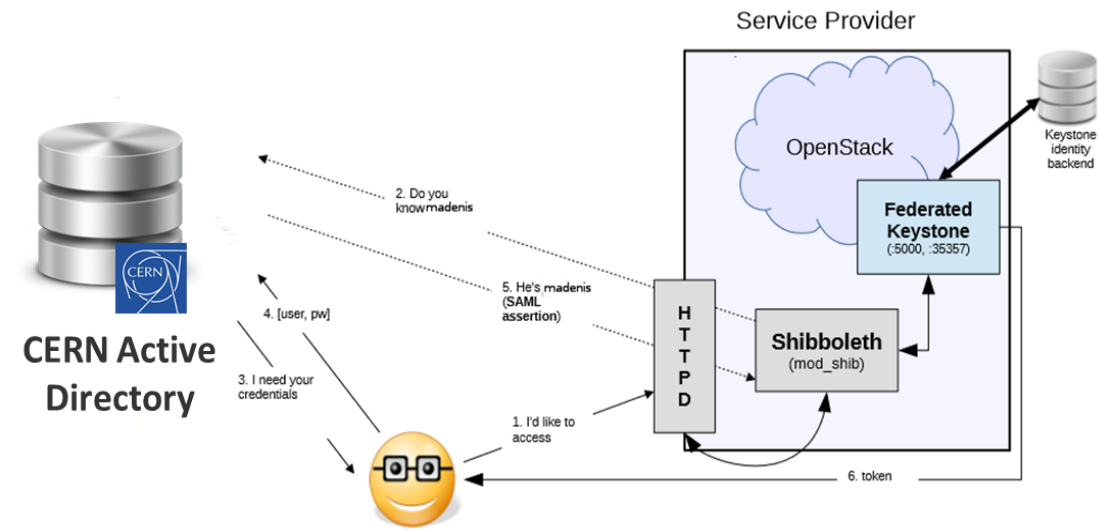
Tim Bell

02 July /2019


Cloud Federation

Rackspace (2013-2015) – 1 fellow in IT-CM

- Support Multi-cloud federated identity of CERN on-premise OpenStack cloud
- Maps between user in cloud A to project in cloud B
- First released in OpenStack Kilo release
- Used in production at CERN for Indigo Dataclouds and to add Kerberos support
- Results fed into the AARC and GEANT identify federation projects
- Contributions also from Red Hat and IBM
- Fellow now at Facebook



LOGIN: madenis
LANGUAGE: EN
DEPARTMENT: IT/OIS
FULLNAME: Marek Denis



ID: madenis
GROUP: openlab
Project: development



Containers as a Service

Rackspace (2015-2017) – 1 fellow in IT-CM

- Develop Magnum, a one-click Container clusters in OpenStack
- Quota, accounting and permissions consistent with VMs and storage
- In production at CERN from 2016
- Over 400 clusters running today at CERN
 - Kubernetes (324)
 - Docker (67)
 - Mesos (9)
- Widely deployed (including Blizzard!)
- Fellow is now a CERN LD

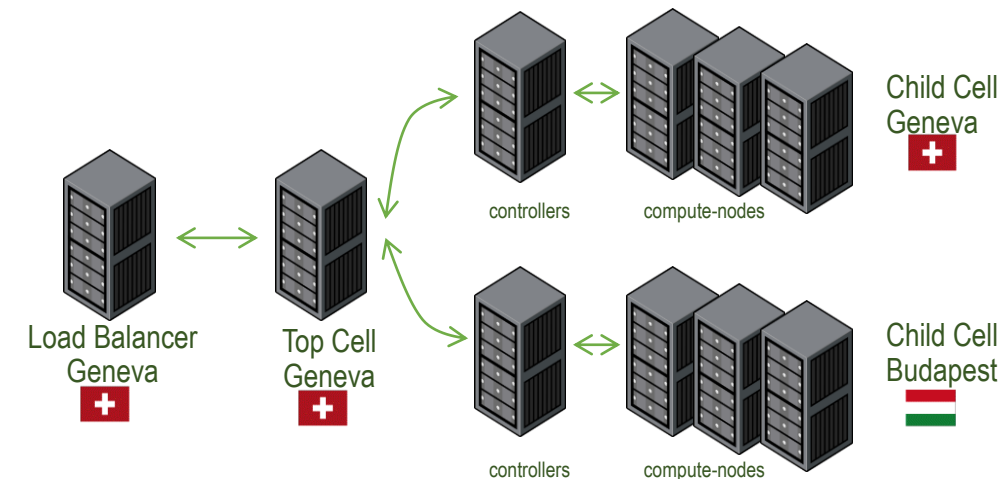
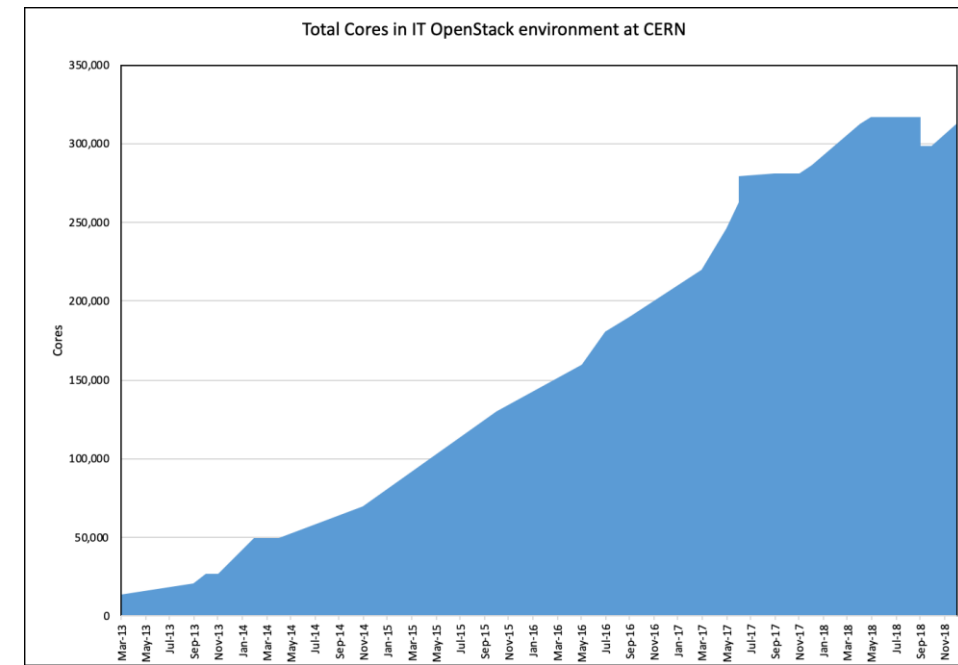
#	Company	Commits
1	NEC	577
2	IBM	415
3	Huawei	326
4	CERN	261
5	Intel	259
6	Rackspace	168
7	Red Hat	157
8	EasyStack	136
	*independent	112
9	Fujitsu	110



Cloud Scaling

Huawei (2017-2018) – 1 fellow in IT-CM

- OpenStack clouds scale by sharding smaller cells
- Cells was an experiment feature from 2013
- Project with Huawei to make this function fully supported and the default
 - Large scale clouds such as Deutsche Telekom
- Deployed in production in 2018
 - Now >80 cells
- Significant tuning post-production supported by openlab fellow
- Fellow extended by IT following project completion



Spot Instances


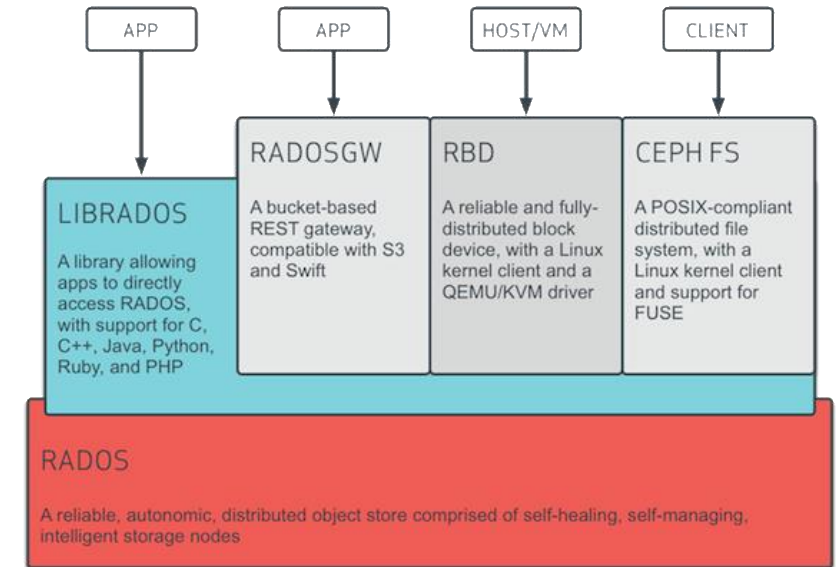
Huawei (2017-2018) – 1 fellow in IT-CM

- Some cloud resources are available opportunistically
 - Hardware before/after production use
 - Buffer from annual purchases for services
- Project to develop Amazon-like Spot instances
 - Can be terminated with little notice
- Project developed in collaboration with SKA, OpenStack Scientific SIG and Huawei
- Production usage at CERN planned during 2019 for experiment restartable workflows
- Fellow extended by IT following project completion

Ceph Performance

Rackspace (2017-2018) – 1 fellow in IT-ST

- Cloud storage at CERN is provided using Ceph
 - Block, Object, File shares
- Project with Rackspace was to optimize performance for CERN's use cases
 - Rackspace Private Cloud product uses Ceph
- Achieved 500K IOPS on 20 node cluster with tuning and code enhancements
- Improved Ceph tools for monitoring and tracing
- Fellow extended by IT following project completion



Conclusions

- Cloud collaboration projects in openlab have funded 9 person years of effort addressing common needs
- Projects have allowed enhancements for open source software to be developed in collaboration with industry
 - Often attracts further companies and labs
- Communities receptive to joint development given CERN's challenging use cases
 - Additional contributions from other companies
- Upstream collaboration with open source foundations simplifies agreements
 - NDA and IP handled by foundation bylaws
- Cost effective model for CERN
 - Community helps with mentoring and future maintenance
- Affordable and flexible for companies to acquire talent
 - Risk of non-renewal such as management re-organisation

References

- CERN Technical Blog
 - <https://techblog.web.cern.ch/techblog/>
- Openlab project pages
 - <https://openlab.cern/project/openstack-clouds>
 - <https://openlab.cern/project/cloud-storage>
- Final talks to conferences
 - **Spot Instances** - <https://www.openstack.org/summit/denver-2019/summit-schedule/events/23187/improving-resource-availability-in-cern-private-cloud>
 - **Cells V2** - <https://www.openstack.org/summit/denver-2019/summit-schedule/events/23287/whats-new-in-nova-cellsv2>
 - **Containers** - <https://www.openstack.org/summit/denver-2019/summit-schedule/events/23713/magnum-project-update>