





CERN Cloud Overview

From 0 to 300k cores

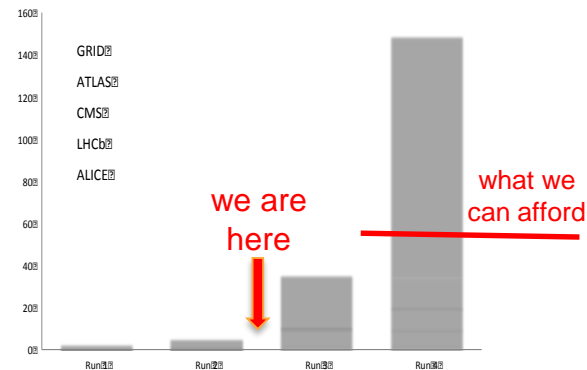
Arne Wiebalck

Compute Canada / MeerKAT Visit
September 16, 2019

2012: A Turning Point for CERN IT



- EU projects for IT tools finished in 2010: decreasing development and support
- LHC compute and data requirements increasing
 - Moore's law would help, but not enough
- Staff would not grow with managed resources
 - Standardization & automation, current tools not apt
- Other deployments have surpassed the CERN one
 - Mostly commercial companies like Google, Facebook, Rackspace, Amazon, Yahoo!, ...
 - We were no longer special! Can we profit?



LS1 (2013) ahead, next window for change would only open in 2019 ...

The Agile Infrastructure Project



- Resource provisioning (IaaS)

- Based on OpenStack



- Centralized Monitoring

- Based on Lemon (sensor) + 'ELK' stack



- Configuration Management


- Based on Puppet



Why IaaS at CERN?

“one server, one application”

- Low Infrastructure Utilization
 - Typically one application per server to avoid affecting the availability of another application
- Increasing Physical Infrastructure Costs
 - Power consumption, cooling and facilities costs
- Increasing IT Management Costs
 - Spend disproportionate time and resources on maintenance, and thus require more personnel
- Insufficient Failover and Disaster Protection
 - The threat of security attacks, natural disasters, and the importance of business continuity



Public Procurement Purchase Model

Step	Time (Days)	Elapsed (Days)
User expresses requirement		0
Market Survey prepared	15	15
Market Survey for possible vendors	30	45
Specifications prepared	15	60
Vendor responses	30	90
Test systems evaluated	30	120
Offers adjudicated	10	130
Finance committee	30	160
Hardware delivered	90	250
Burn in and acceptance	30 days typical 380 worst case	280
Total		280+ Days

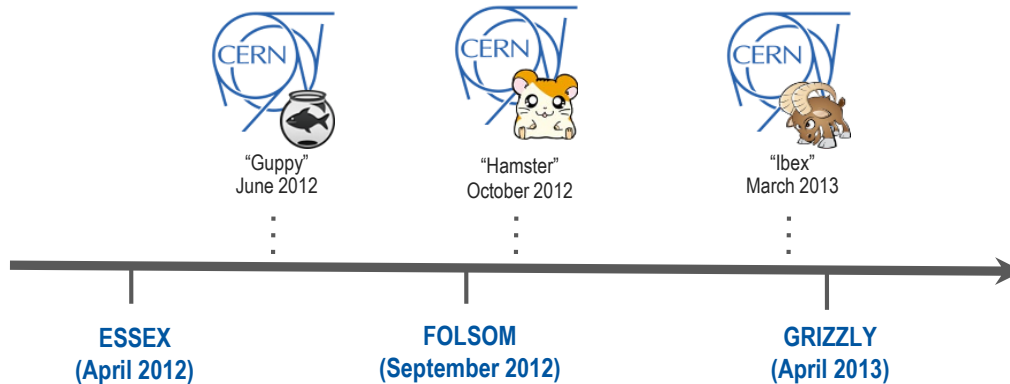
How can we address these challenges



Virtualization
Cloud Computing

Prototyping CERN OpenStack Cloud

- Iterate Fast...
 - Build test infrastructures and open them to early adopters
 - Few hundred nodes available
 - 2 different virtualization technologies (KVM, Hyper-V)
 - Integration with other Agile Infrastructure projects (puppet, monitoring, ...)



CERN OpenStack Cloud - 2013

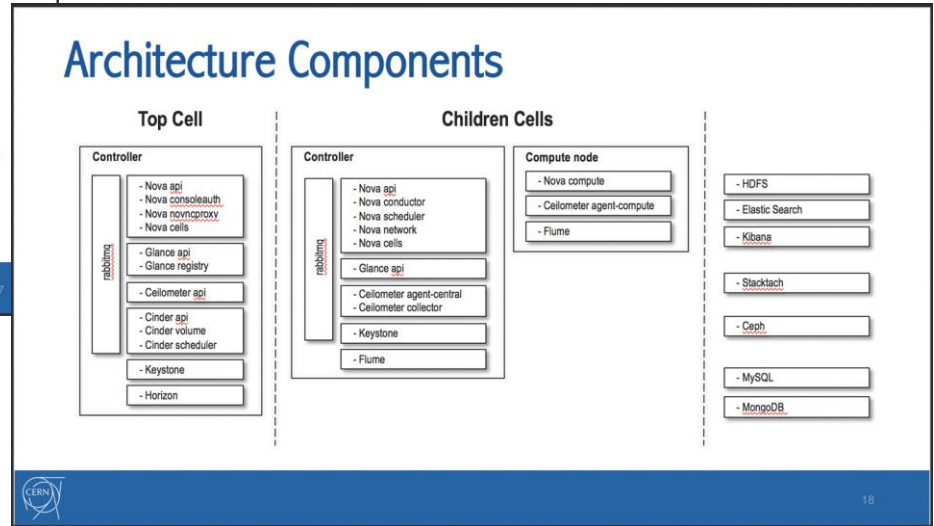
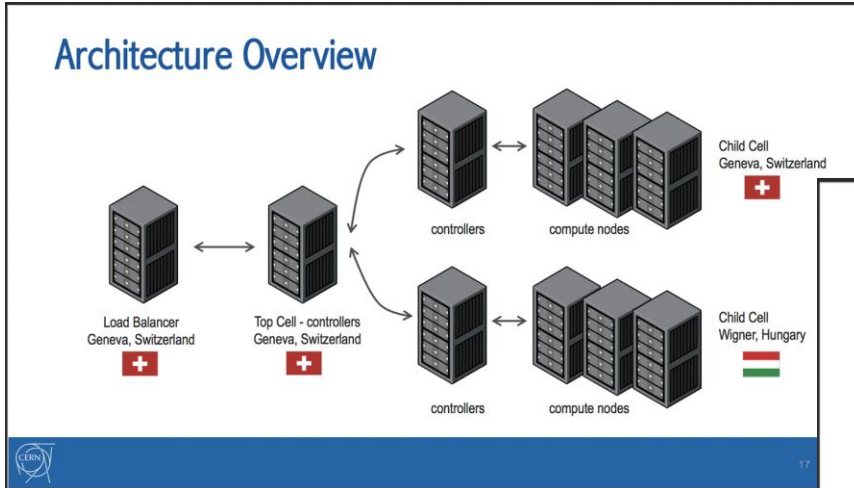
OpenStack at CERN - grizzly release

- +2 Cells – Geneva and Wigner Computer Centers
- HA+1 architecture
- Ceilometer deployed
- Integrated with CERN accounts and network
- Monitoring OpenStack components status
- Glance - Ceph backend
- Cinder tests - Ceph backend

Infrastructure Overview

- HAProxy as load balancer
- Master and Compute nodes
 - 3+ Master nodes per Cell
 - O(1000) Compute nodes per Cell (KVM and HyperV)
 - 3 availability zones per Cell
- Rabbitmq
 - At least 3 brokers per Cell
 - Rabbitmq cluster with mirrored queues


CERN OpenStack Cloud - 2013




CERN OpenStack Cloud - 2013

CERN Cloud Infrastructure adoption

Number of VMs



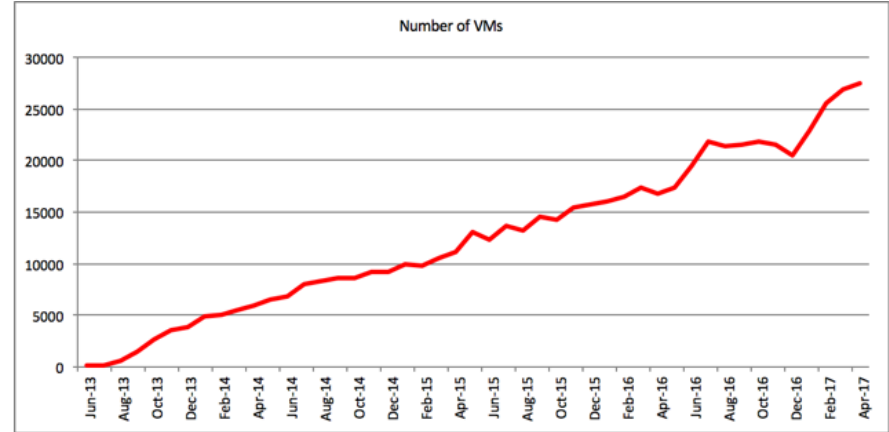
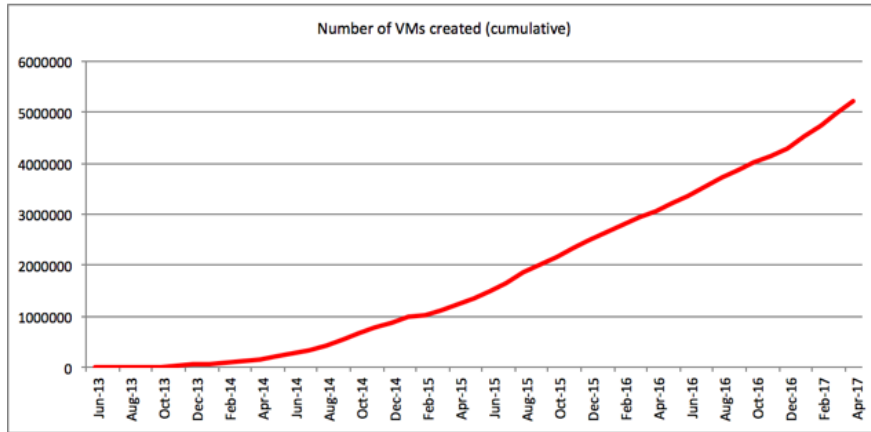
Number of Users



Cell	Nodes	Cores	RAM (GB)	Disk (TB)	VMs
Geneva	375	10976	21662	632	1438
Wigner	291	9312	18296	491	674
Total	666	20288	39958	1123	2112

(13/10/2013)

CERN OpenStack Cloud - Growth





CERN OpenStack Cloud - Growth

- OpenStack projects available in the CERN Cloud over releases

Grizzly	Havana	Icehouse	Juno	Kilo	Liberty	Mitaka	Newton	Ocata	Pike	Queens	Rocky
Nova	Nova	Nova	Nova	Nova	Nova	Nova	Nova	Nova	Nova	Nova	Nova
Glance	Glance	Glance	Glance	Glance	Glance	Glance	Glance	Glance	Glance	Glance	Glance
Horizon	Horizon	Horizon	Horizon	Horizon	Horizon	Horizon	Horizon	Horizon	Horizon	Horizon	Horizon
Keystone	Keystone	Keystone	Keystone	Keystone	Keystone	Keystone	Keystone	Keystone	Keystone	Keystone	Keystone
Ceilometer*	Ceilometer*	Ceilometer	Ceilometer	Ceilometer	Ceilometer	Ceilometer	Ceilometer	Ceilometer	Ceilometer	Ceilometer	Ceilometer
		Cinder	Cinder	Cinder	Cinder	Cinder	Cinder	Cinder	Cinder	Cinder	Cinder
			Heat*	Heat	Heat	Heat	Heat	Heat	Heat	Heat	Heat
			Rally*	Rally	Rally	Rally	Rally	Rally	Rally	Rally	Rally
					EC2API	EC2API	EC2API	EC2API	EC2API	EC2API	EC2API
					Magnum*	Magnum	Magnum	Magnum	Magnum	Magnum	Magnum
					Barbican*	Barbican	Barbican	Barbican	Barbican	Barbican	Barbican
					Neutron*	Neutron	Neutron	Neutron	Neutron	Neutron	Neutron
						Ironic?	Ironic?	Ironic?	Ironic*	Ironic	Ironic
						Mistral?	Mistral?	Mistral?	Mistral*	Mistral	Mistral
						Manila?	Manila?	Manila*	Manila*	Manila	Manila
								Trove?		Qinling?	Qinling?
								Murano?		Watcher?	Watcher?

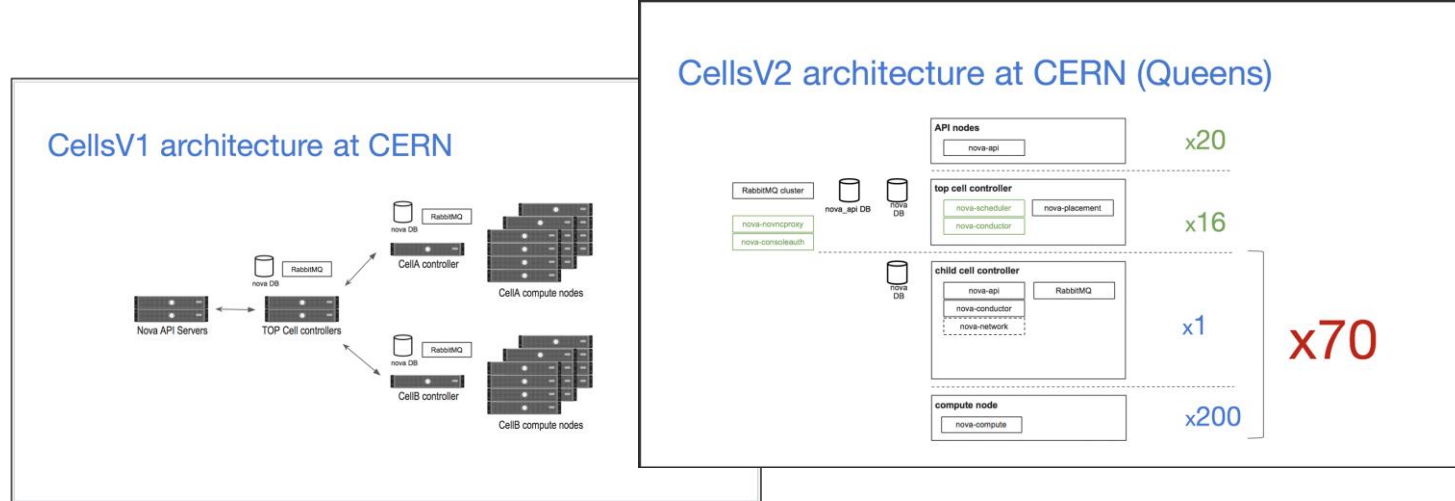
* - Pilot service

? - Trial service

Current Infrastructure

Nova - Cells

- Allows Nova to scale to thousands of compute nodes
- Biggest Nova Cells deployment
- Moved from 2 cells to +70 cells
- Upgrade from CellsV1 to CellsV2 in 2018




Storage - Cinder, Manila, S3

- OpenStack Cinder with Ceph backend (2014)
 - Several volume types available
- OpenStack Manila (Fileshare service). Backed by CephFS (2017)
- S3 available (end 2018)

Cinder

- Ceph and NetApp backends
- Extended list of available volume types (QoS, Backend, Location)
- Cinder nodes are VMs
- Active/Active?
 - When a volume is created a "cinder-volume" node is associated
 - Responsible for volume operations
 - Not easy to replace cinder controller nodes
 - DB entries need to be changed manually
- More about CERN storage infrastructure for OpenStack:
 - <https://www.openstack.org/summit/vancouver-2015/summit-videos/presentation/ceph-life-of-a-petabyte-scale-block-storage-service>

Cinder Deployment at CERN



Manila

- Fileshare service Manila
- Pilot since Q4 2016
- CephFS as the backend
- Off-the-shelf integration with Kubernetes, Swarm...
- Need for a highly available FS (to replace NFS filer service)
- Collaboration with FILER service
- Share configuration, certificates, etc




Container Orchestration - Magnum

- OpenStack Magnum service available since 2016
- Extremely popular service, +500 clusters

What's new? Magnum

- OpenStack project to treat orchestration engines as 1st class resources
 - Docker Swarm, Kubernetes, Mesos and DC/OS
- Current release is in Newton (with cherry-pick from master)
- Timeline



HEPIX Spring 2017 – CERN Cloud Service Update

Containers

Container service (magnum):

- Support many versions of kubernetes (1.9.x and 1.10.x in prod)
- Simplify user interface
- Support for traefik ingress
- RBAC for kubernetes, possible to federate with external clusters

Use cases:

- REANA/RECAST for reusable analysis
- Continuous Integration
- Spark on Kubernetes
- Interactive analysis

HEPIX Spring 2018 – CERN Cloud Service Update

Container improvements

Lifecycle operations on container clusters:

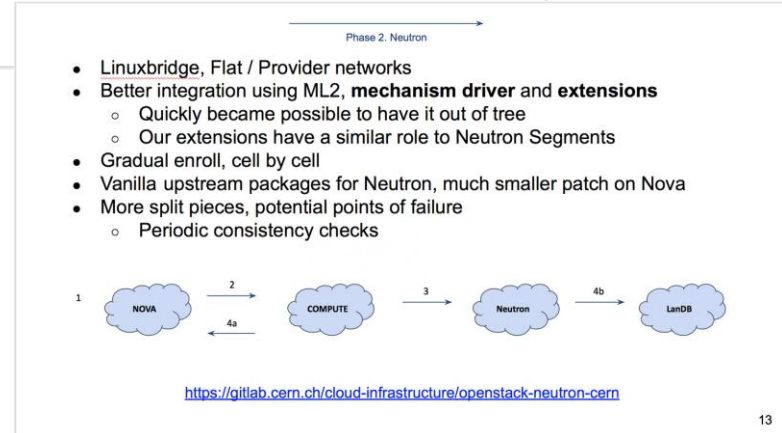
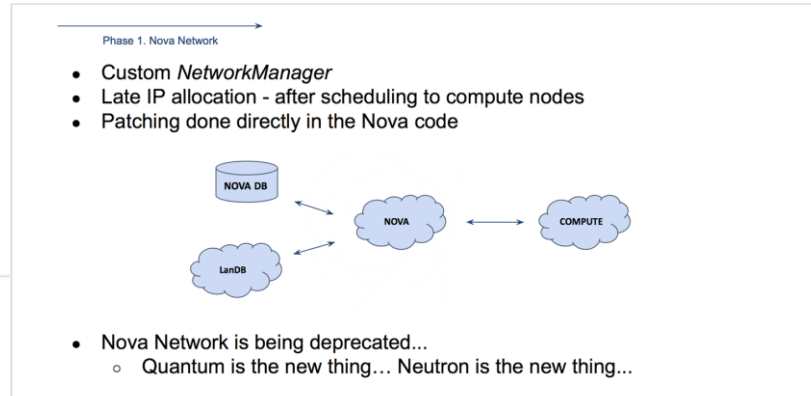
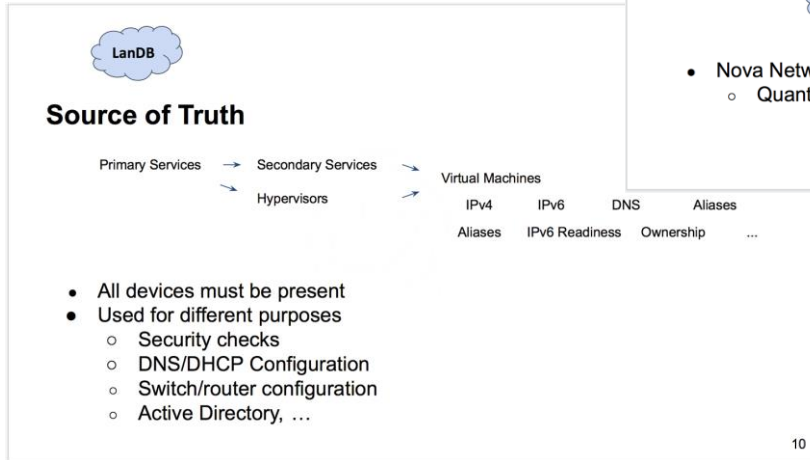
- Host upgrades, OS and container orchestrator
- auto-healing of faulty nodes

Storage and containers:

- csi-cephfs integration, users will be able to create and mount cephfs volumes to kubernetes pods (create only with admin creds)
- manila provisioner, end users will:
 - create shares with cephfs as backend
 - mount them to pods with csi-cephfs

HEPIX Spring 2018 – CERN Cloud Service Update

Networking - Nova-network to Neutron



Baremetal Provisioning - Ironic

- In production since 2018
- All new hardware is enrolled using Ironic. +1700 nodes managed by Ironic
- Existing hardware will be enrolled into Ironic during 2019

Why Bare-Metal Provisioning? (1)



- VMs not suitable for 100% of our use cases
 - Benchmarking, storage nodes, boot strapping, critical network equipment, specialised network setups, HPC clusters, ...
- Complete our service offerings
 - Physical nodes (in addition to VMs and containers)
 - OpenStack UI as the single pane of glass
- Simplify hardware provisioning workflows
 - For users: `openstack server create/delete`
 - For procurement: initial on-boarding, server re-assignments



Why Bare-Metal Provisioning? (2)



- Consolidate accounting & bookkeeping
 - Resource accounting input will come from less sources
 - Machine re-assignments will be easier to track
- Enable new use cases
 - Containers on bare metal



Doesn't change the overall policy ☺
The reasons why we introduced virtual machines have not gone away!



Meltdown/Spectre/L1TF

- Reboot campaigns and performance impact

Patching the Cloud

- Upgrade/reboot of hypervisors *and* virtual machines required
 - ◆ ~8'500 hypervisors, ~36'000 virtual machines
- Patching of non-batch hypervisors (a.k.a. shared/service cells)
 - ◆ ~1'400 hosts, ~18'000 guests (⇒ your personal and service VMs)
 - ◆ Hypervisor reboots will be staged by [availability zone](#)
 - ◆ Your VMs will be rebooted during the hypervisor patching campaign!
- Given the severity of the vulnerability and the size of the cloud deployment, there won't be much room for schedule discussions
- Schedule and updates will be available from the [Cloud SSB entry](#)

KEEP CALM REBOOT

Arne Wiebalck: Spectre/Meltdown - Impact on CERN Cloud Users (ASDF, 11 Jan 2018) 3

Patching the Cloud: Service Hypervisors

- All availability zones have been rebooted! (~1'100 hypervisors with ~11'500 virtual machines)
- ◆ 'cern-geneva-a'
- ◆ 'cern-geneva-b'
- ◆ 'cern-geneva-c'
- ◆ critical area
- ◆ 'cern-wigner-a'
- ◆ 'cern-wigner-b'

- Wed Sep 12: AVZ 'cern-geneva-a'
- Mon Sep 17: AVZ 'cern-geneva-b'
- Tue Sep 18: AVZ 'cern-geneva-c'
- Wed Sep 19: AVZs 'critical', 'cern-wigner-a,b'

- Batch not affected this time
 - ◆ Trusted VMs

KEEP CALM REBOOT

Review: Performance Impact

- 'l1tf=full' means unconditional flushing and SMT off
- Performance impact assessed for "HS06-like" workloads
 - ◆ "Impact on services unknown, but we have enough head room."
- Overcommit depends on hardware type
 - ◆ Hypervisors with a lot of RAM severely impacted → "SMT on" again

KEEP CALM REBOOT

Arne Wiebalck: Rebooting the Cloud (ASDF, 20 Sep 2018) 7


Operations - Rundeck and Mistral

- OpenStack Mistral
- RunDeck

RUNDECK

- Friendly and easy interface from where we can organize and launch jobs on our hosts
- Sharing of sensitive tasks to other groups without exposing credentials or procedures
- Use Cases
 - **SysAdmins:** Workflows related to hypervisor maintenance (h/w intervention, notify users...)
 - **Cloud-Operations:** Project creation, Health reports, Quota update

Rundeck Integration



Mistral

- Workflow service Mistral **MISTRAL**
an OpenStack Community Project
 - Will simplify operations
 - Already deployed with testing workflow prototypes
 - Will play along with Rundeck for workflows

Operations

- Experience growing/managing the Infrastructure during the last 6 years
- Several upgrades during this journey
 - OpenStack release cycle is every 6 months!
 - SLC6 to CC7 upgrade
 - CC7 upgrades
- Supported for few years KVM and HyperV in the same infrastructure
 - Migrated CVI VMs to OpenStack HyperV and then to OpenStack KVM
- Security updates required reboot of all cloud
- Most user management operations are automated
 - project creation; quotas; ...
 - VM expiration

What's next?

Splitting the Infrastructure into Regions

<https://techblog.web.cern.ch/techblog/post/region-split/>

CERN Cloud Infrastructure

- (2013) We decided to offer only one region!
 - Wigner datacentre was exposed to users as 2 AVZs
 - Direct project mapping for the compute use-case
- (2013) Why?
 - At that time was important to offer only one endpoint to users (Still is...)
 - **It's more simple to manage one small cloud than 2 small clouds**
 - Cells allows to scale Nova to thousand of nodes
 - No real advantage in having another region...

CERN Cloud Infrastructure

- What changed?
 - **It's more simple to manage two small clouds than 1 large cloud**
 - Deploy a new configuration change
 - Upgrades
 - High impact/visibility when something goes wrong
 - Nova-network -> Quantum -> Neutron
 - Neutron is not Nova cell aware
 - Neutron relies in a single RabbitMQ cluster
 - Challenge to scale!
 - Use cases are now very well defined
 - Compute VS services

Preemptible Instances

- Public Clouds
 - Based on different pricing/SLA considering resource availability
 - Reserved instances vs spot-market
- Private Clouds
 - Quotas are hard limits. Leads to a reduction in resource utilization
 - Preemptible instances
 - Projects that exhausted their quota can continue to create instances
 - Opportunistic workloads
 - Low SLA
- Preemptible Instances Workflow in OpenStack Nova
 - The creation of a non preemptible VM fails because there aren't available resources
 - Instances that fail with "Nova Valid Host", go to "PENDING" state instead of "ERROR"
 - The Reaper service is notified and it tries to free the requested resources
 - Rebuild the instance
 - Or change instance state to "ERROR"


Other Challenges

- Leveraging Container Orchestration to deploy OpenStack control plane
- Re-enroll existing physical resources into OpenStack Ironic
- Introduce SDN
- Dynamic resource provisioning based on Compute Nodes' load

Cloud resources



Openstack services stats



Resource overview by time

