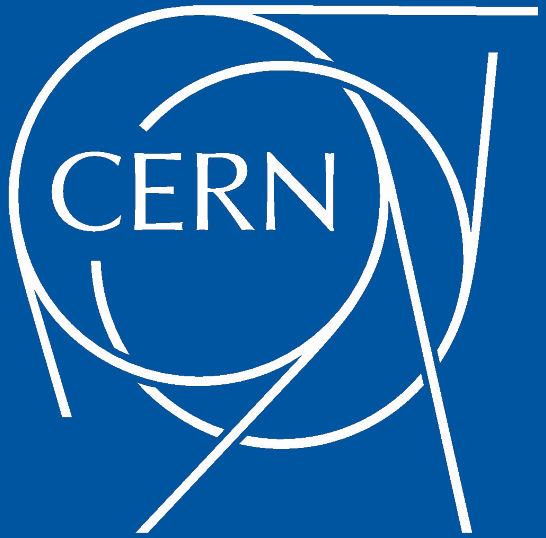
CERN

6 years of CERN Cloud

From 0 to 300k cores

HEPIX - San Diego - 2019

Belmiro Moreira

on behalf of the CERN Cloud Team

belmiro.moreira@cern.ch @belmiromoreira

Outline

- Early Virtualization attempts
- Finding the right Cloud Orchestrator
- CERN OpenStack Cloud
- Some Infrastructure Highlights
- What’s next?

2009 - 2011

Virtualization and Server Consolidation

Why IaaS at CERN?

“one server, one application”

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

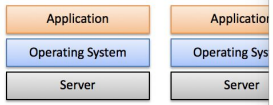

Public Procurement Purchase Model

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

LxCloud - Virtualize Batch Infrastructure

PES

LxCloud – advantages...

CERN IT Department

- The usual cloud advantages:
  - Maintenance
    - Decoupling from the underlying hardware
    - Easy migration between OS / configurations
  - Dynamic Environment
    - The composition of compute nodes can be rapidly adapted to optimize resources utilization
  - Encaps
    - Virt
      envi

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

PES

What is LxCloud?

CERN IT De

- LxCloud architecture:

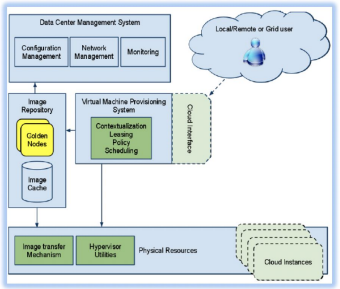

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

CERN - IT- PES - PS

6

Platform ISF

CERN IT Department

- Platform Computing;
  - Commercial license;
- Supported hypervisors: XEN, KVM, ESXi;
- Strong coupling with LSF (with Adaptive Cluster);
- Needs VMO agent in each computer node;
- Professional support;

- Scalability tests:
  - About 10000 VMs in abo
- In production since Dec
  - 6 Computer Servers run
  - 48 VMs;

OpenNebula

CERN IT Department

- Open source
  - Apache license;
- Supported hypervisors: KVM, XEN, ESXi;
- Small installation footprint;
- Only requires installation of the cloud controller;
- Professional and community support;

- Scalability tests:
  - About 16000 VMs in about 500 computer servers;
- In production since December 2010 for lxbatch:
  - 6 Computer Servers running KVM;
  - 48 VMs;

9

CVI - CERN Virtual Infrastructure

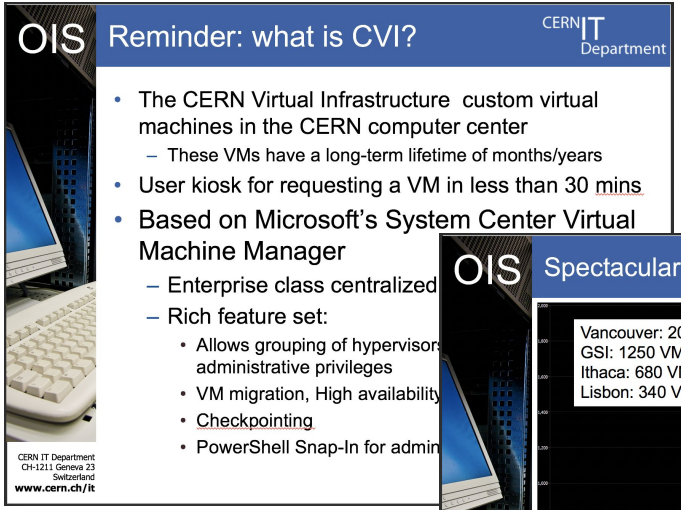

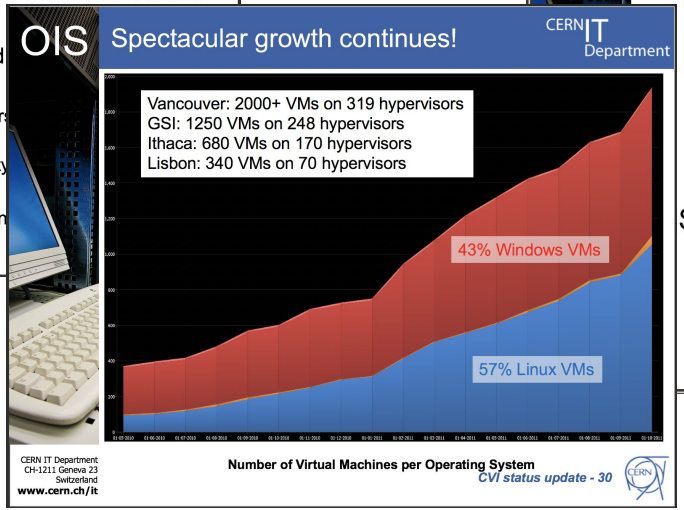

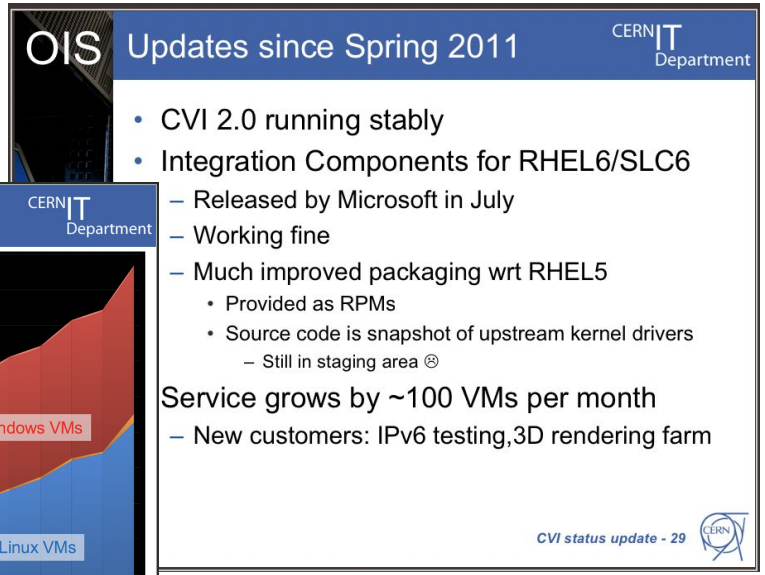

2011 - 2013

Cloud Prototype

Agile Infrastructure Project

Year	What	Actions
2011		Agree overall principles
2012		Prepare formal project plan Establish IaaS in CERN CC Production Agile Infrastructure Monitoring Implementation as per WG Migrate lxcloud Early adopters to Agile Infrastructure
2013	LSD 1 New Data Centre	Extend IaaS to remote CC Business Continuity Support Experiment App re-work Migrate CVI General migration to Agile with SLC6 and Windows 8
2014	LSD 1 (to November)	Phase out Quattor/CDB/…

OpenStack at CERN - Early days

- Released on October 2010
- Created by Nasa and Rackspace
- Austin release (Nova, Swift)
- OpenSource (Apache 2.0)

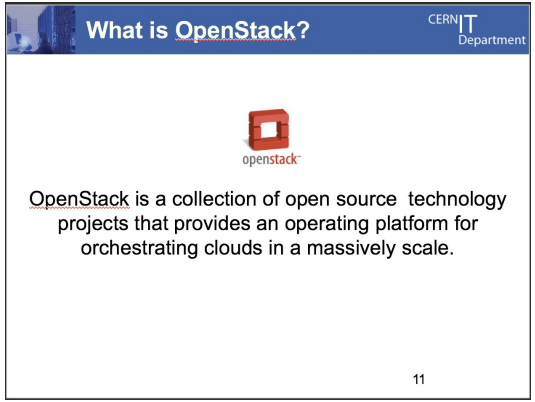

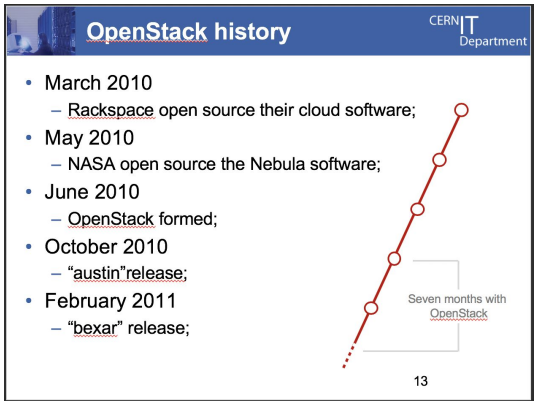

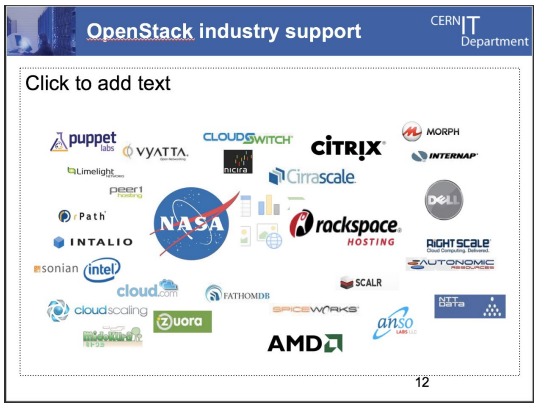

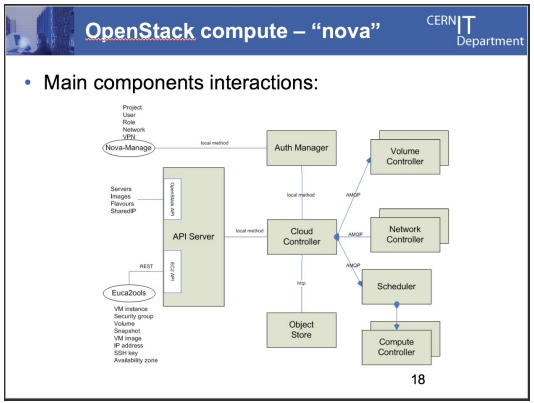

OpenStack at CERN - Early days

- First version of OpenStack Horizon

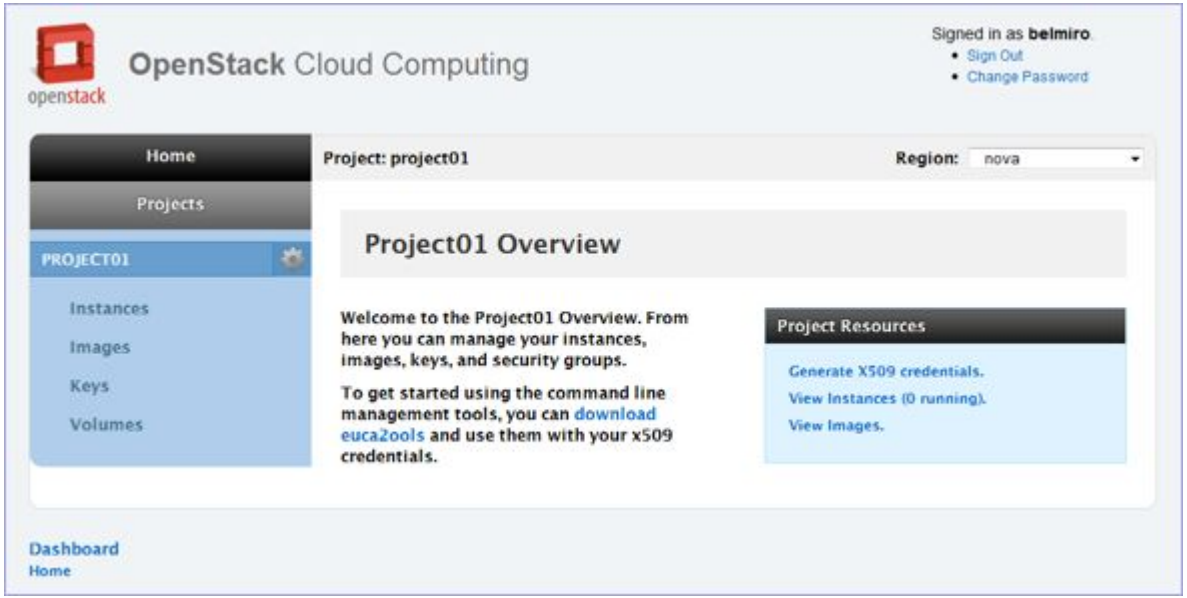

Prototyping CERN OpenStack Cloud

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

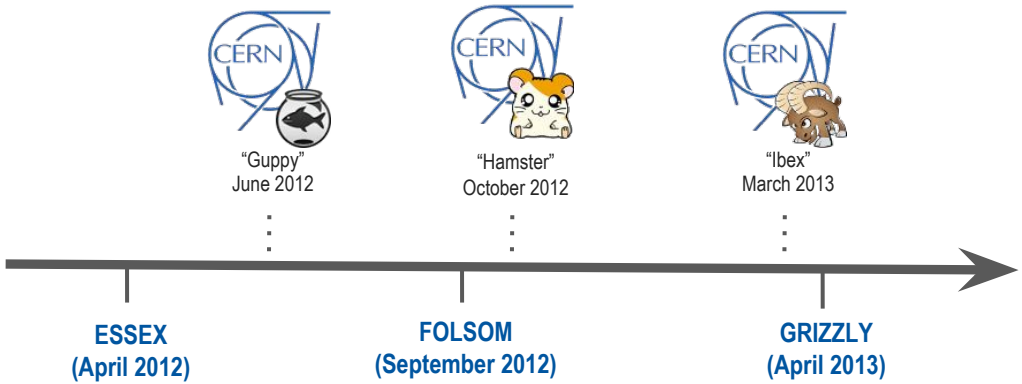

2013 - 2019

From 0 to +300k cores

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

19

CERN OpenStack Cloud - 2013

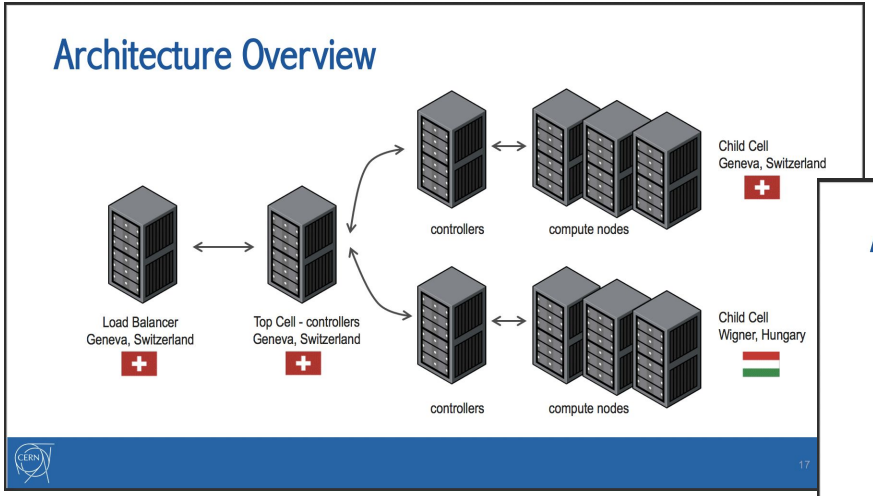

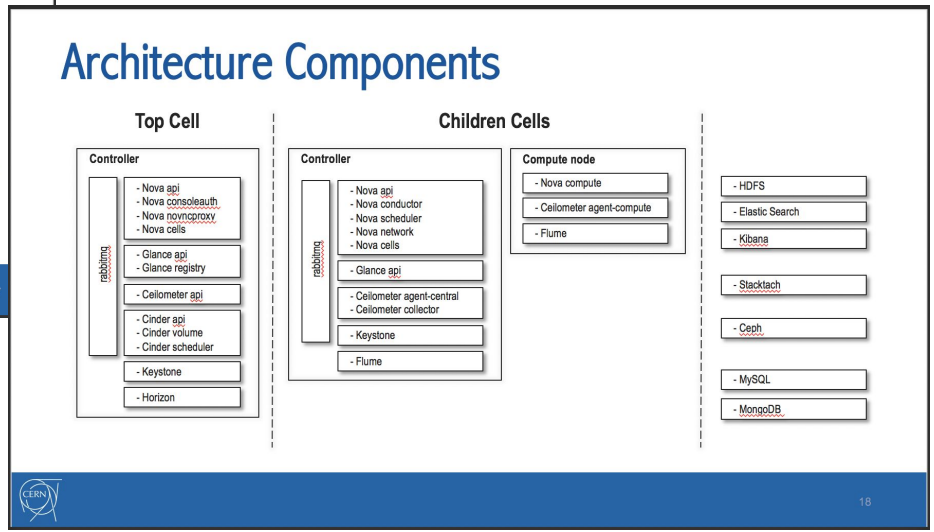

CERN OpenStack Cloud - 2013

CERN Cloud Infrastructure adoption

Number of VMs

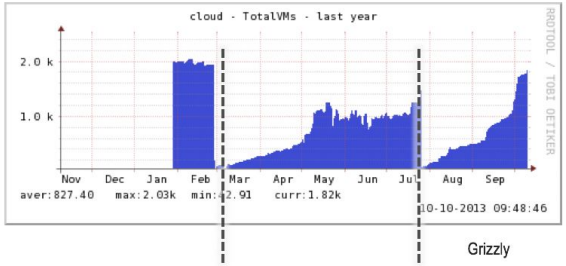

Grizzly

Number of Users

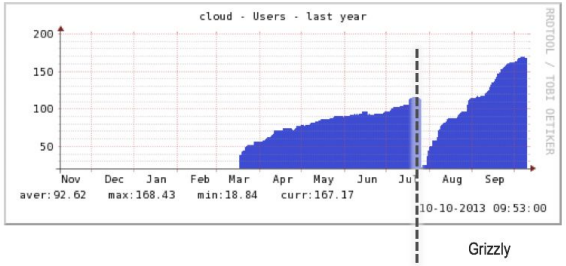

Grizzly

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

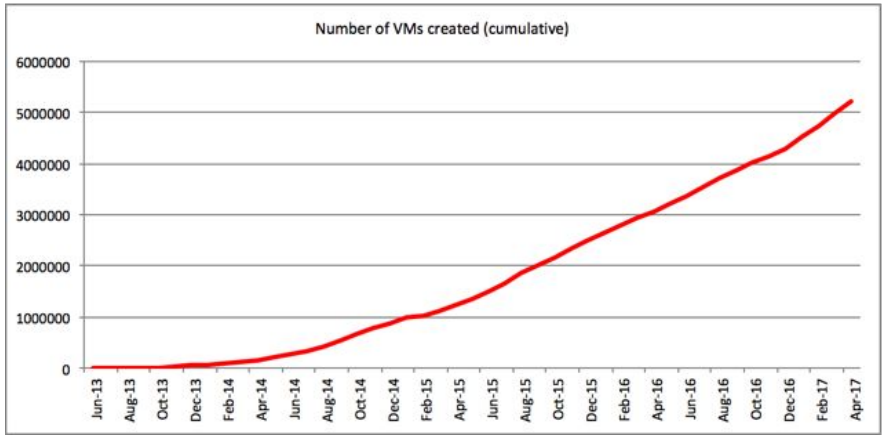

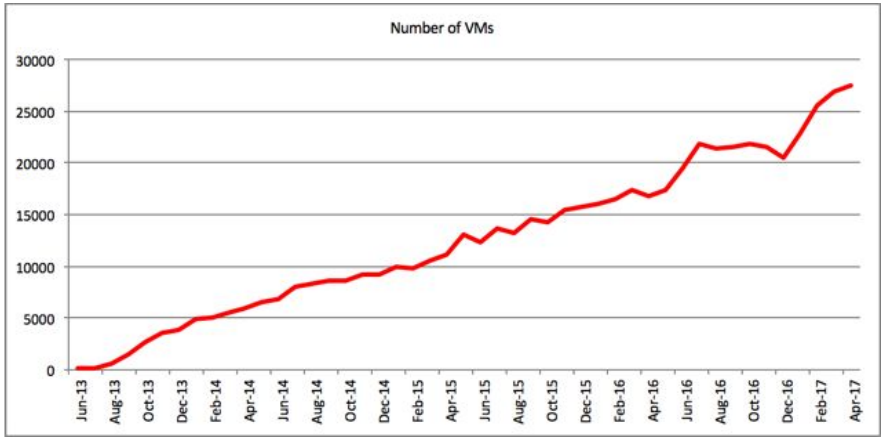

CERN OpenStack Cloud - Growth

- OpenStack projects available in the CERN Cloud over releases

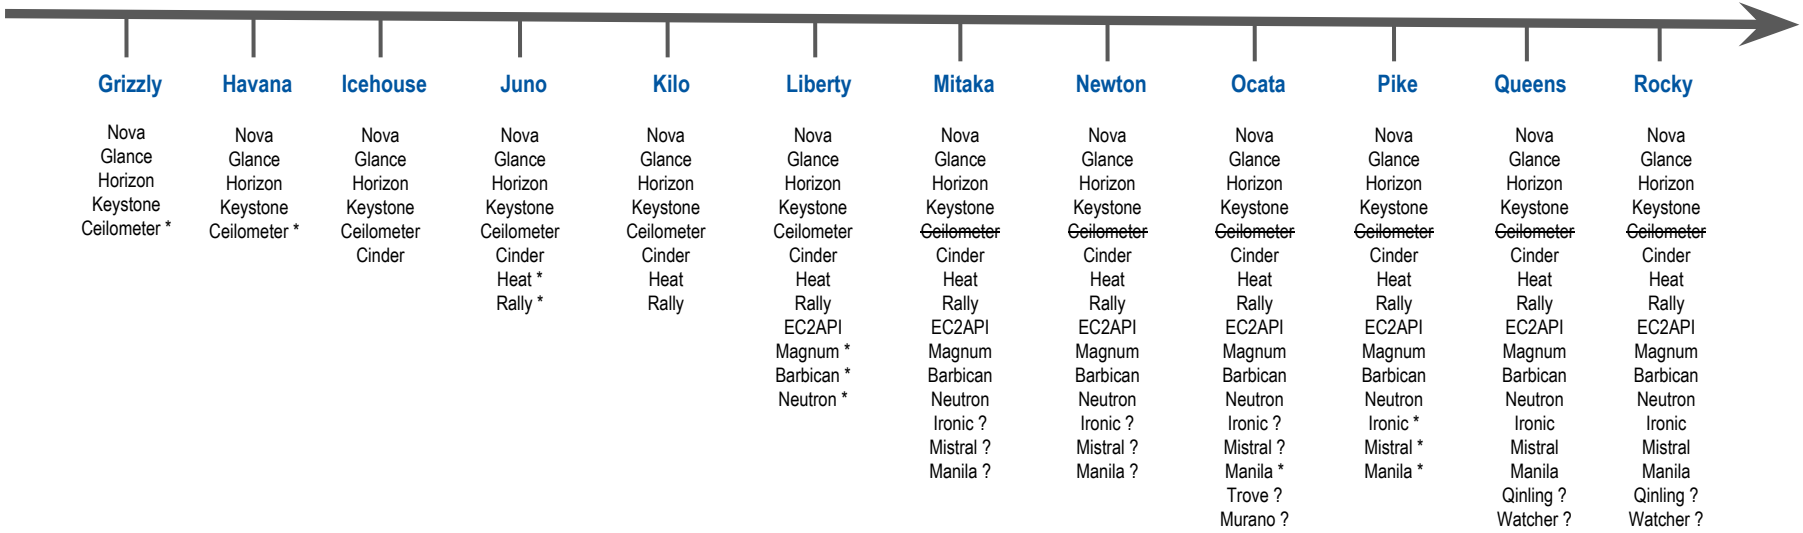

* - Pilot service
? - Trial service

Milestone Highlights

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

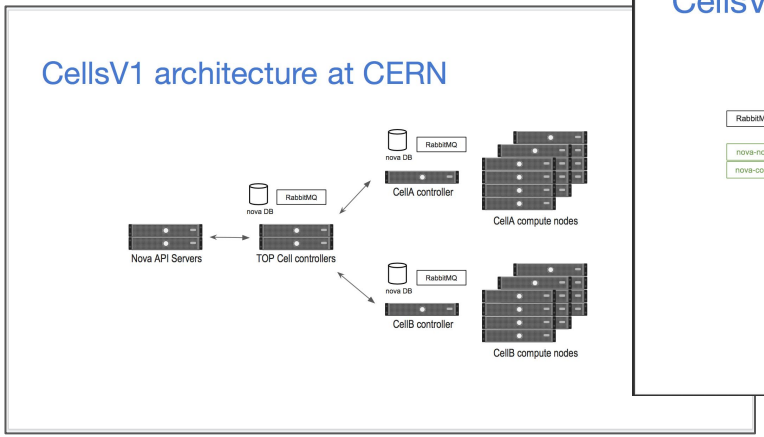

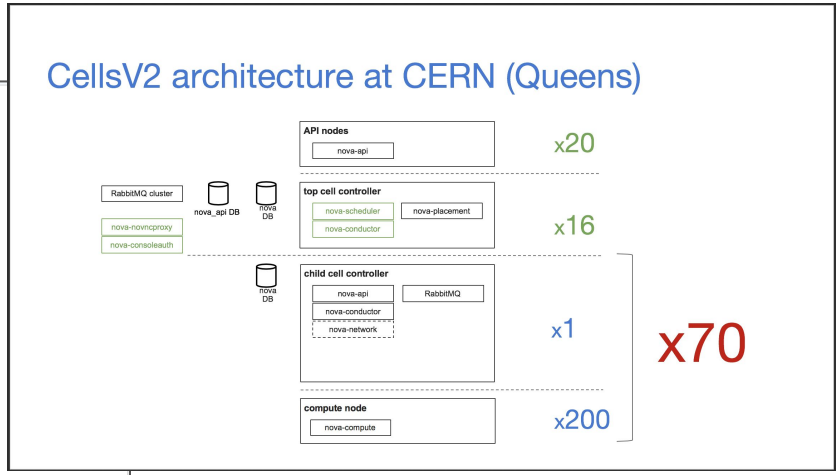

Ceilometer - The Rise & Fall

- OpenStack Ceilometer deployed
- Removed after run it for 3 years. Not scalable and difficult to retrieve data

Ceilometer

- Users are not directly billed
  - Metering needed to adjust Project quotas
- mongoDB backend – sharded and replicated
- Collector, Central-Agent
  - Running on “children” Cells controllers
- Compute-Agent
  - Uses nova-api running on “children” Cells controllers

Ceilometer Deployment at CERN

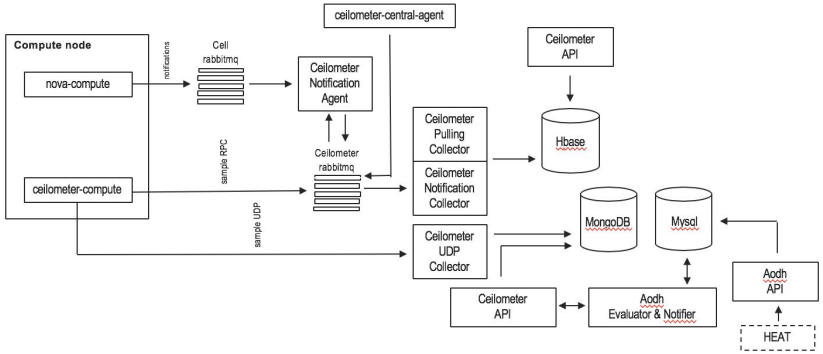

Storage - Cinder, Manila, S3

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

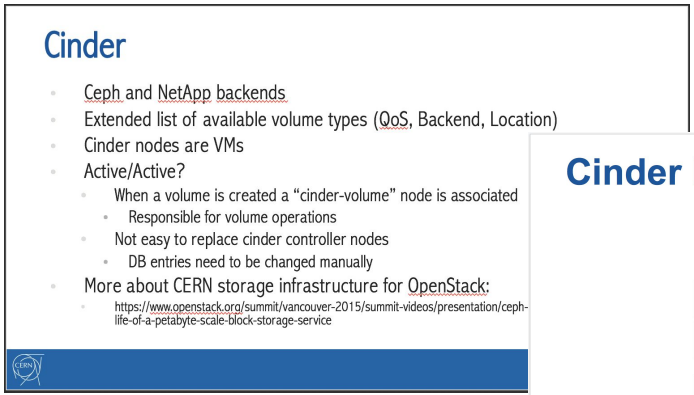

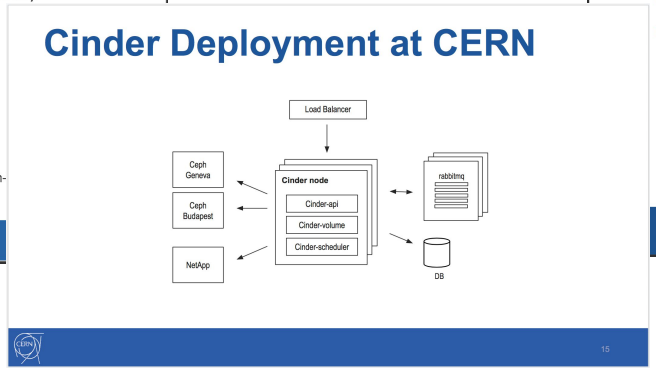

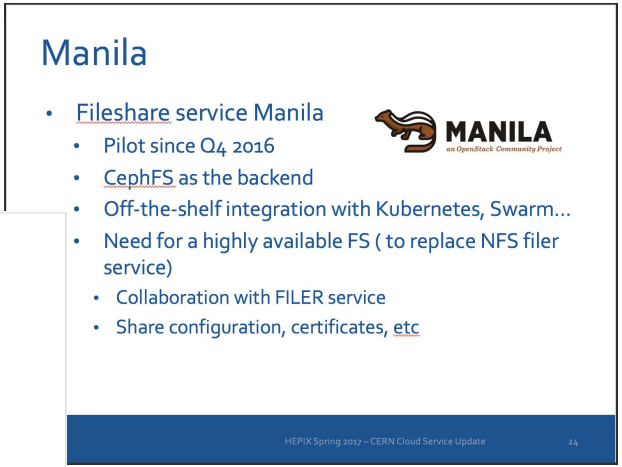

Container Orchestration - Magnum

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

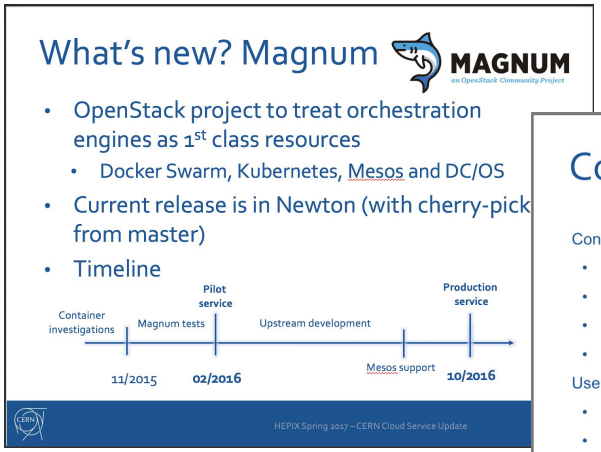

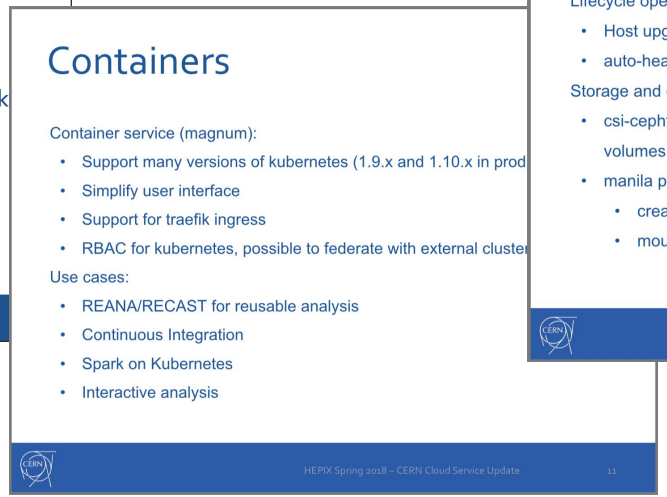

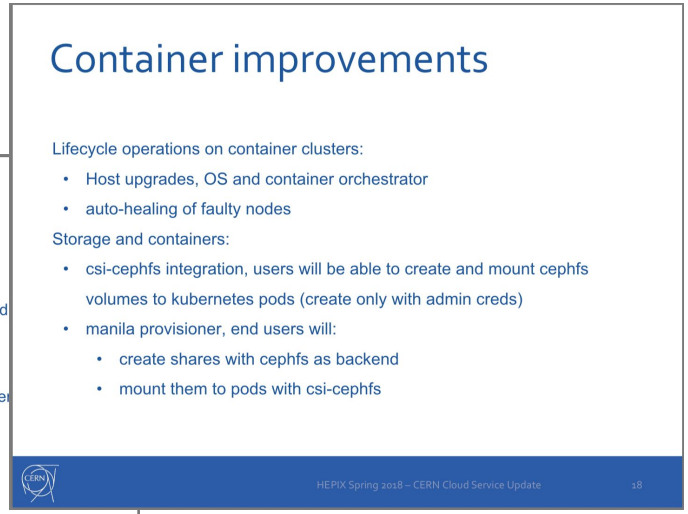

Networking - Nova-network to Neutron

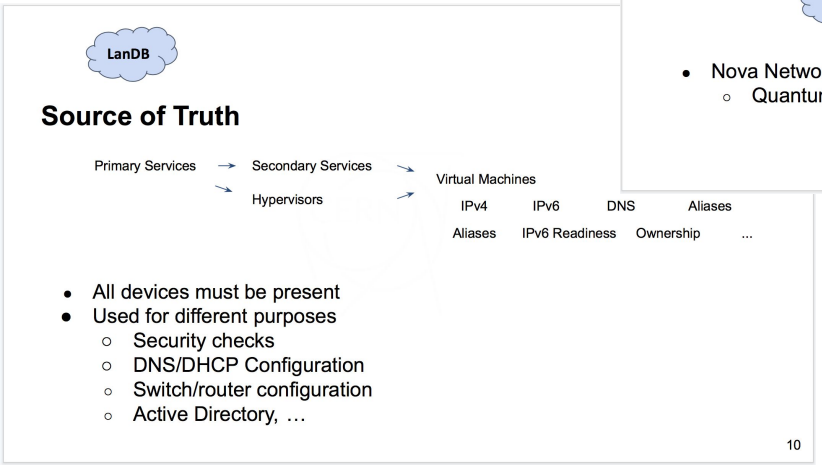

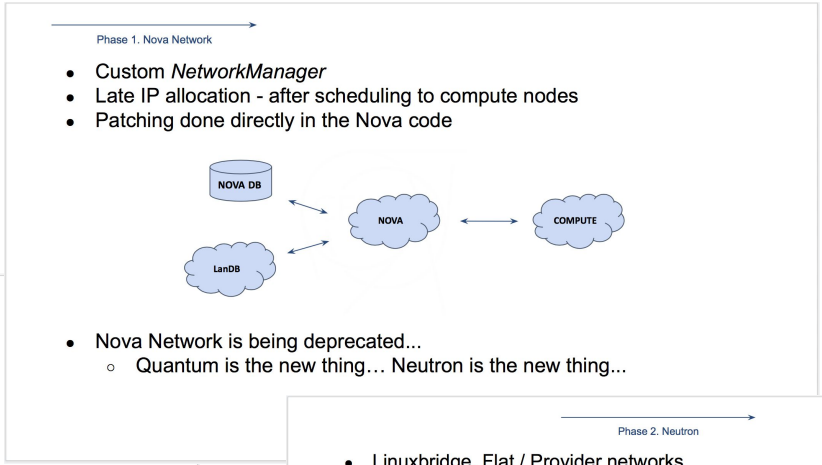

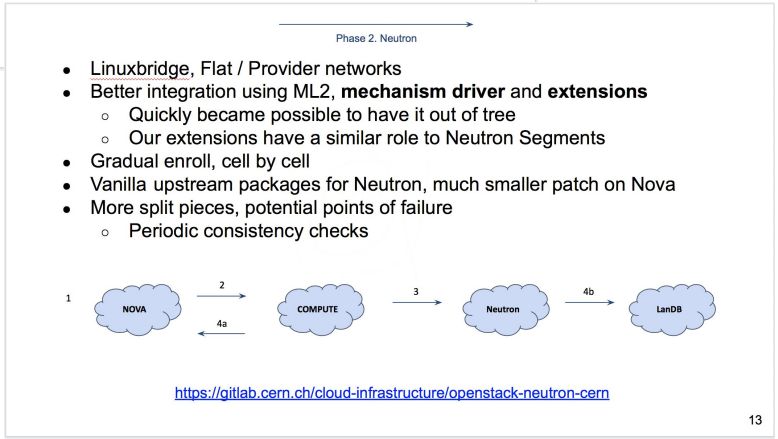

Baremetal Provisioning - Ironic

- 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

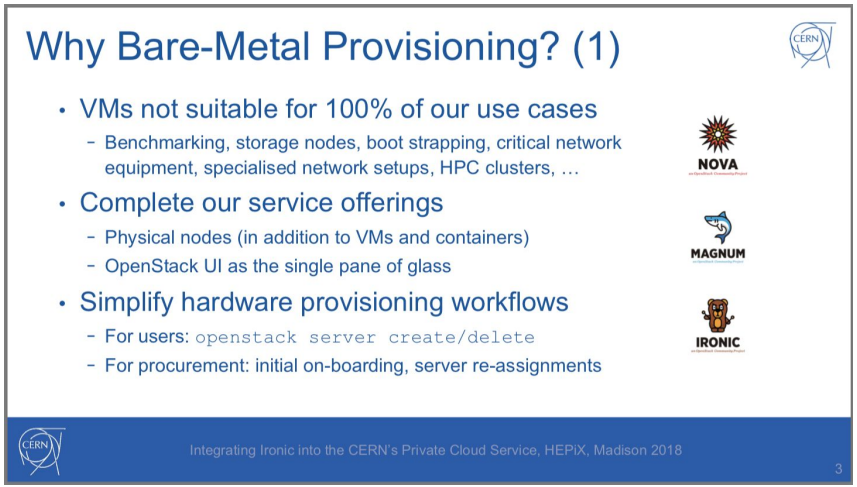

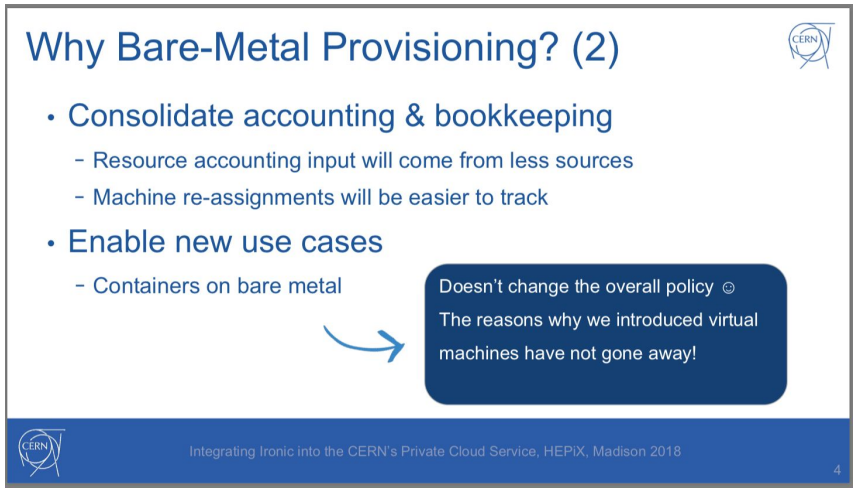

Meltdown/Spectre/L1TF

- Reboot campaigns and performance impact

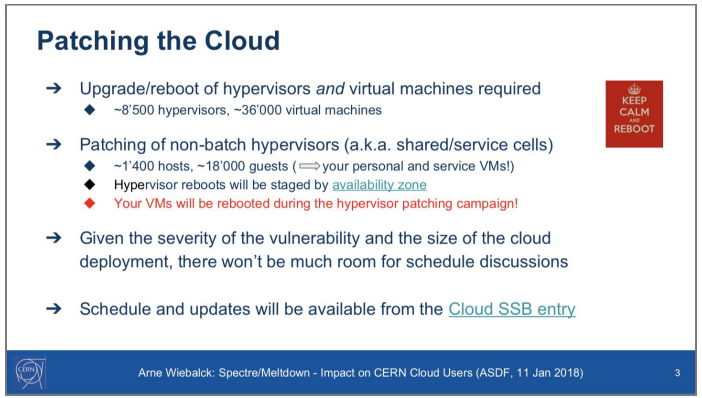

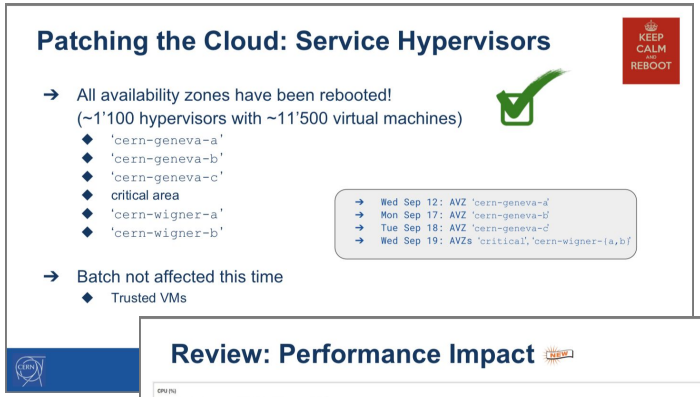

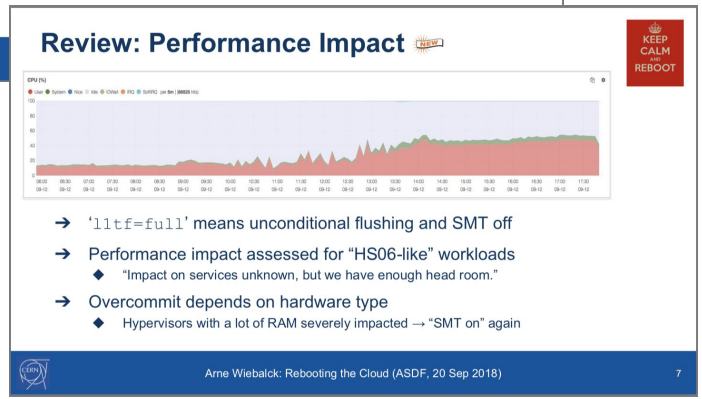

Operations - Rundeck and Mistral

- OpenStack Mistral
- RunDeck

Mistral

- Workflow service Mistral
  - Will simplify operations
  - Already deployed with testing workflow prototypes
  - Will play along with Rundeck for workflows

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

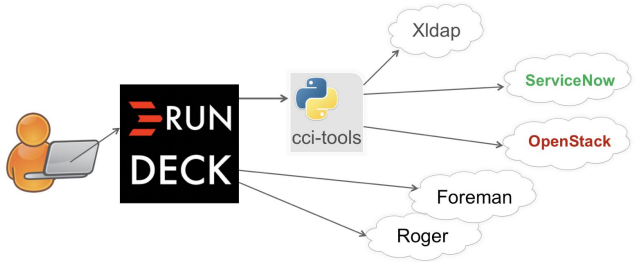

Operations

- 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

2019 - …

What’s next?

Splitting the Infrastructure into 2 Regions

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

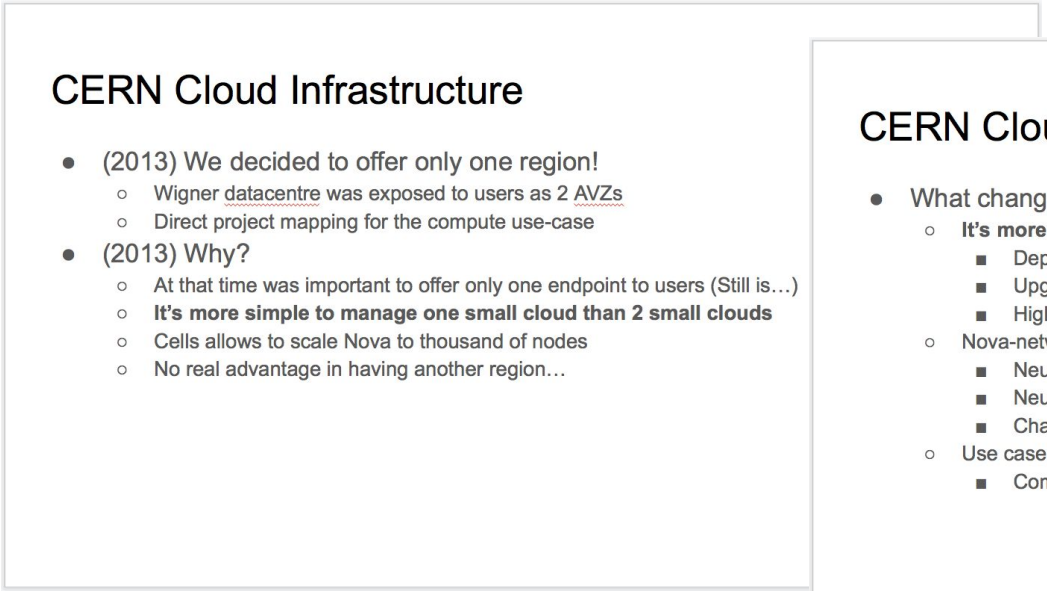

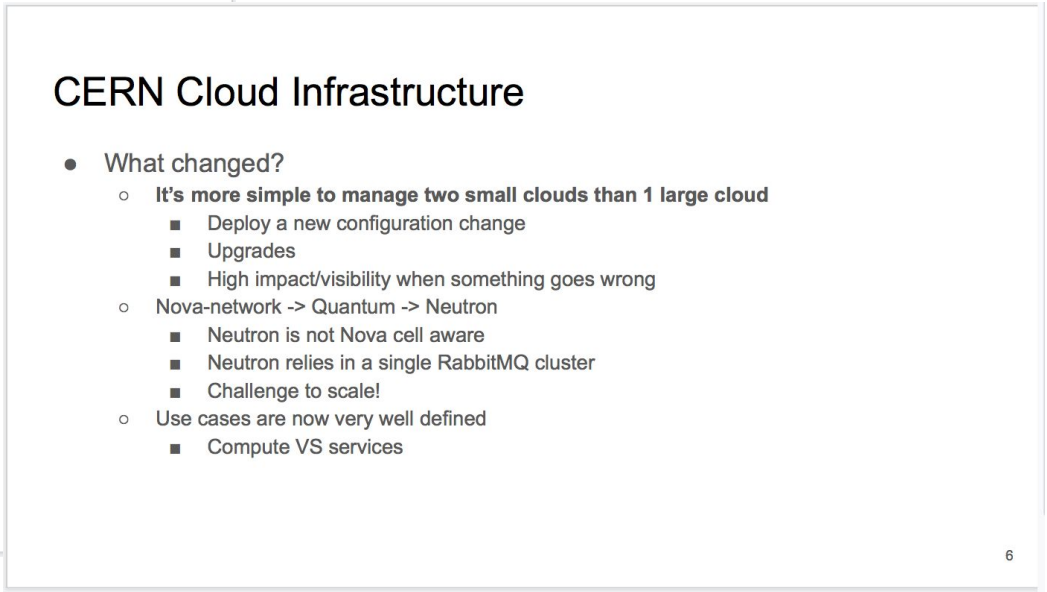

Preemptible Instances

- 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 in Compute Nodes load

CERN OpenStack overview

Last 30 days Refresh every 15m

Cloud resources

Used	Available	Used	Available	Used	Available
289.0 K cores	274.3 K cores	822.1 TiB RAM	903.1 TiB RAM	9.4 PiB disk	14.6 PiB disk

Openstack services stats

Users	Projects	VMs	Magnum clusters	Hypervisors	Images	Baremetal nodes
3591	4711	34713	525	8380	3360	1616

Volumes	Volume size	Fileshares	Fileshares size
6389	1.97 PiB	441	232 TiB

Resource overview by time

VMs created/deleted

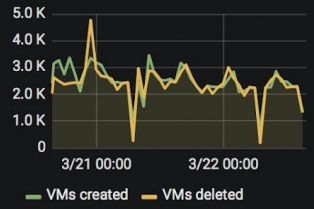

Shared cells availability

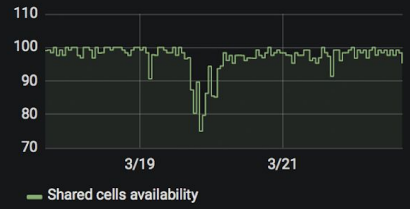

Total VMs

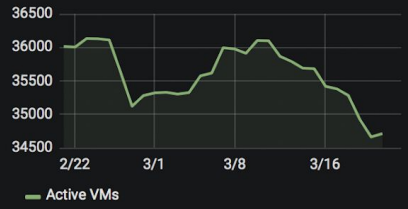

Average VM boot time

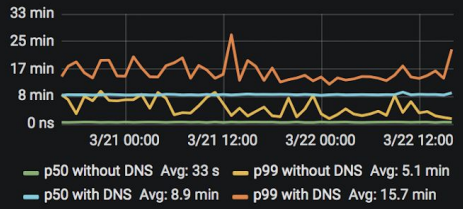

VM changes

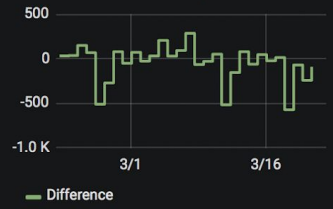

Hypervisors

Magnum clusters

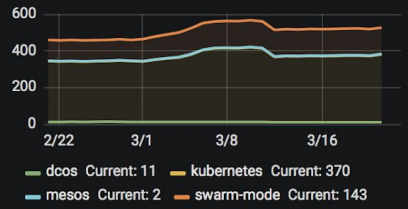

Projects and users

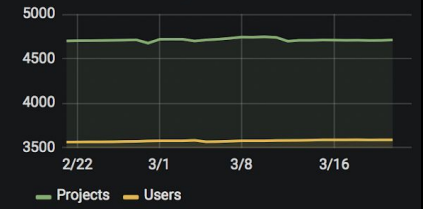

Summary

- During the last 10 years, resource management and deployment model changed completely
  - From Virtualization and Server consolidation to a Cloud Infrastructure
  - From Baremetal to VMs, to managed Baremetal to Containers
- Continue to adapt the Infrastructure to the new technologies and requirements
  - Iterative approach to introduce new services, new functionality
  - Continue to explore new approaches to deploy/manage a large infrastructure
    - Control Plane managed by kubernetes
    - New regions
    - SDN

Is serverless the new model?

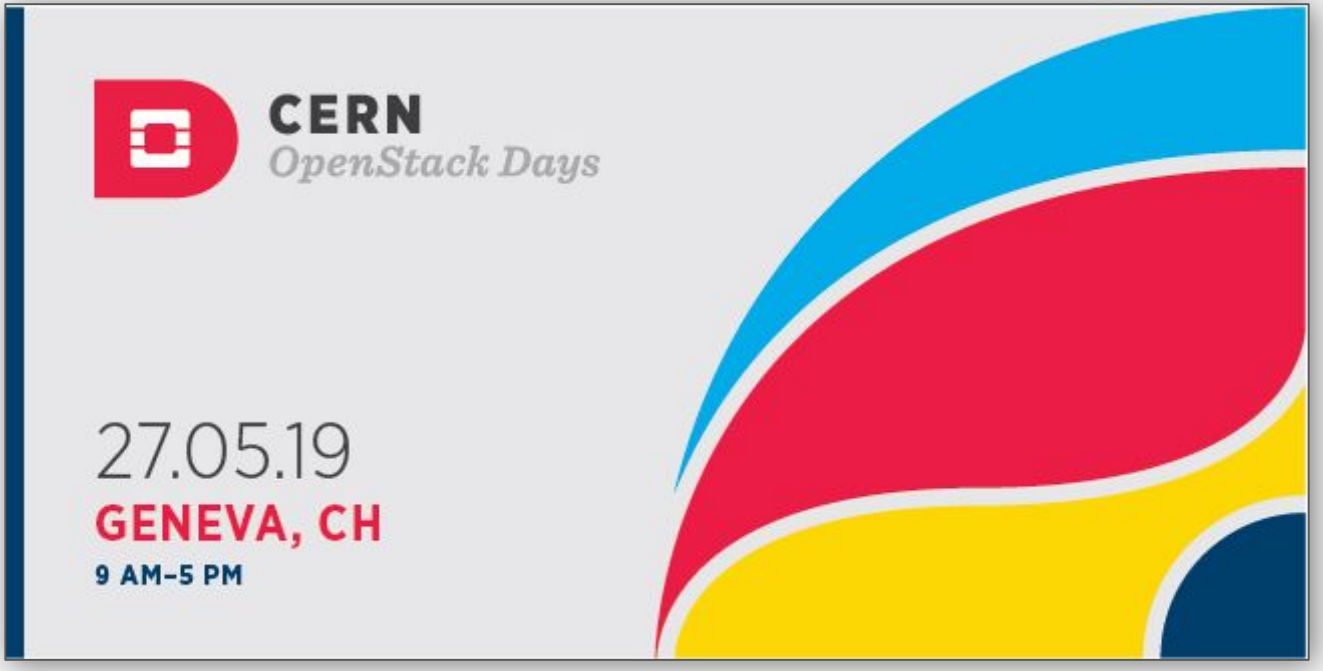

https://openstackdayscern.web.cern.ch

@belmiromoreira

www.cern.ch

Credits

- Used slides from several authors
  - Arne Wiebalck
  - Belmiro Moreira
  - Domenico Giordano
  - Jan Van Eldik
  - Luis Pigueiras
  - Ricardo Rocha
  - Spyridon Trigazis