



Cloud Infrastructure Update: Operations, Campaigns and Evolution

Maryna Savchenko

HEPiX Spring
28 March 2023



Outline

- Cloud Service Overview
- Operation updates
- Bare metal service updates
- Future plans
 - *Preveessin Data Center*
 - *Software Defined Network*
 - *CPU steal*






CERN Cloud Service

- Infrastructure as a Service
- Production since July 2013
- Highly scalable architecture
- From Train to Xena releases
 - *Train, Ussuri, Victoria, Wallaby, Xena, Yoga, Zed*
- Geneva Computer Center
 - *LHCb container*
 - *Preveessin Computer Center*



Cloud Components







User visible



Openstack services statistics

Users 3360	Projects 4621	Loadbalancers 297	Images 6000	Volumes 7354	Volumes si... 3.96 PB	File Shares 5235	File Shares... 1.79 PB	Object Stor... 490	Object Stor... 88.0 TB			
Servers		Cores			RAM			Batch				
Physical 8846	Physical in use 8387	Hypervisors 2020	Virtual 13523	Physical 485 K	Hypervisors 60.2 K	Virtual 88.6 K	Physical 2.07 PB	Hypervisors 393 TB	Virtual 202 TB	Servers 4535	Cores 255270	RAM 998 TB


Time series



Openstack services statistics

Users 3360	Projects 4621	Loadbalancers 297	Images 6000	Volumes 7354	Volumes si... 3.96 PB	File Shares 5235	File Shares... 1.79 PB	Object Stor... 490	Object Stor... 88.0 TB			
Servers				Cores			RAM			Batch		
Physical 8846	Physical in use 8387	Hypervisors 2020	Virtual 13523	Physical 485 K	Hypervisors 60.2 K	Virtual 88.6 K	Physical 2.07 PB	Hypervisors 393 TB	Virtual 202 TB	Servers 4535	Cores 255270	RAM 998 TB







Time series



Openstack services statistics

Users 3360	Projects 4621	Loadbalancers 297	Images 6000	Volumes 7354	Volumes si... 3.96 PB	File Shares 5235	File Shares... 1.79 PB	Object Stor... 490	Object Stor... 88.0 TB			
Servers		Cores			RAM			Batch				
Physical 8846	Physical in use 8387	Hypervisors 2020	Virtual 13523	Physical 485 K	Hypervisors 60.2 K	Virtual 88.6 K	Physical 2.07 PB	Hypervisors 393 TB	Virtual 202 TB	Servers 4535	Cores 255270	RAM 998 TB

Time series



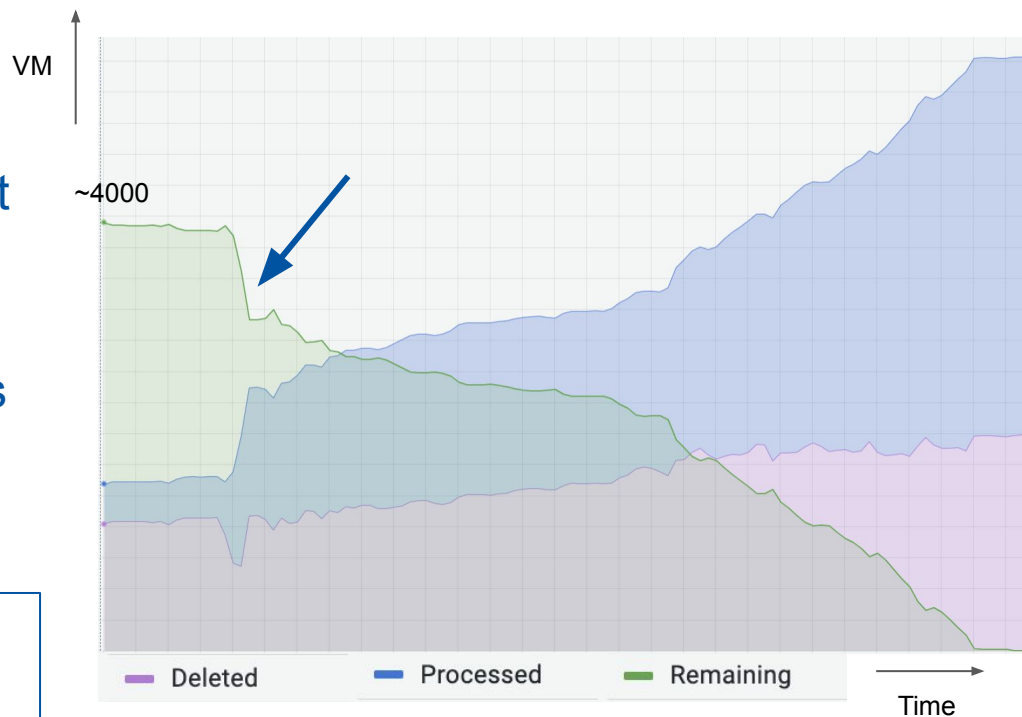
Cloud operation updates (1/3)

VM migration campaign

- Replace legacy network component
- Cold migration of ~4000 VMs
- Unblocking Nova upgrade
- Allowed advanced network features

More information:

[Cloud Cold Migration campaign retrospective](#)



Cloud updates (2/3)


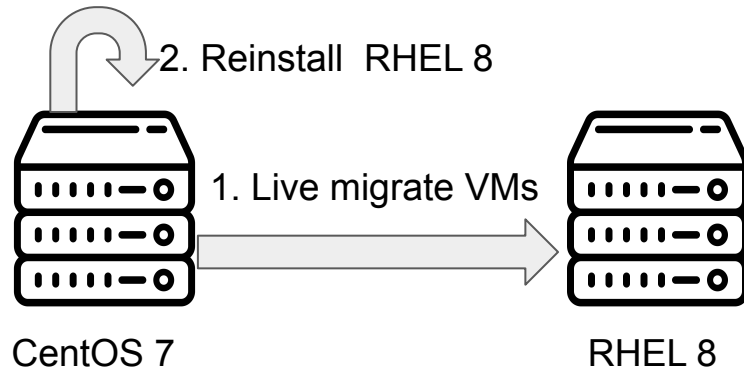
Upgrade Nova to Train release

- Without API downtime
- 48 cells one by one
- Previous Nova upgrade 3 years ago
- Code cleanup

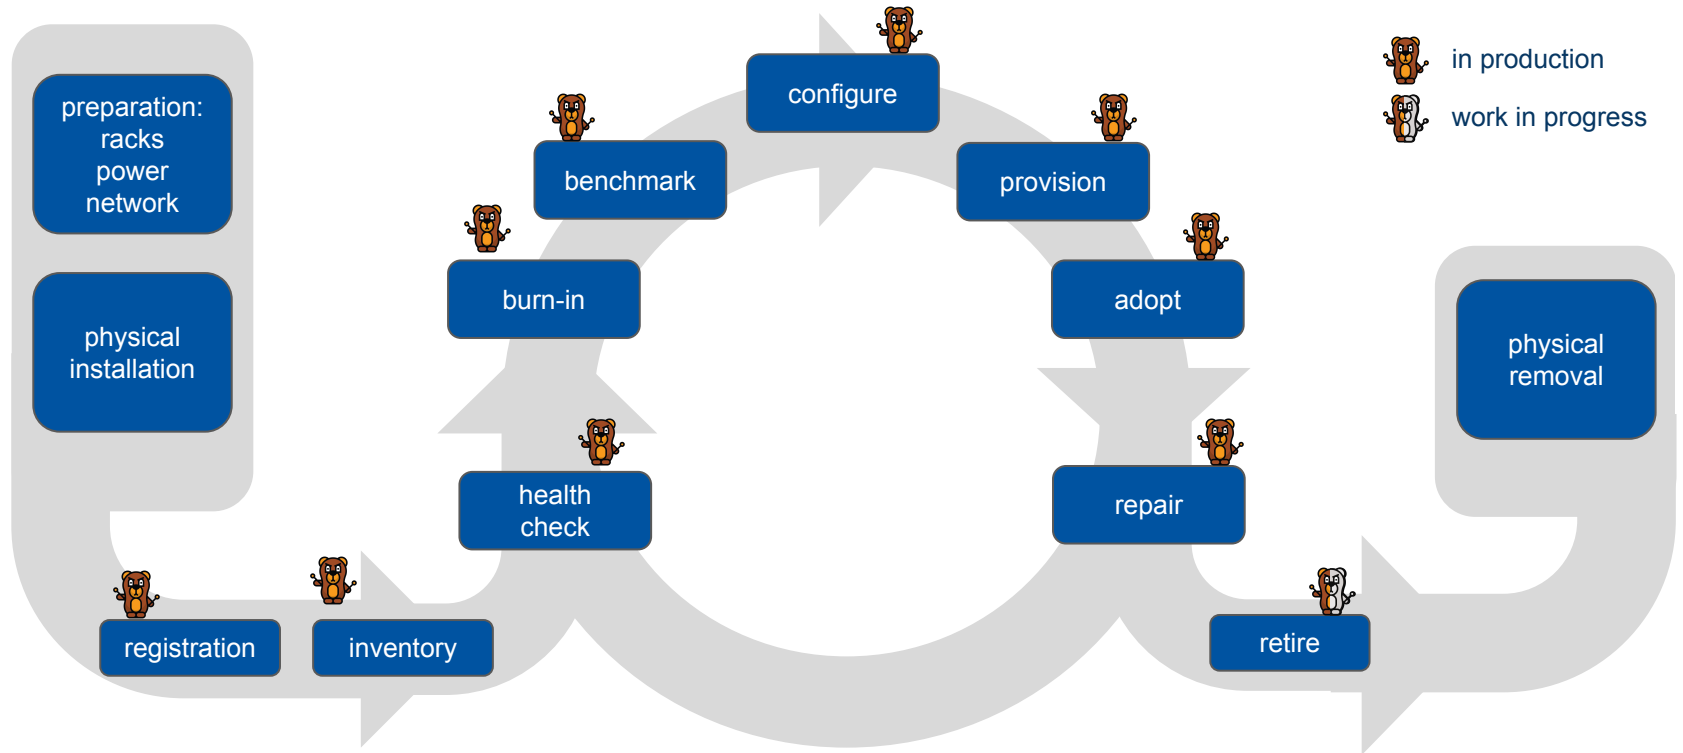
Cloud updates (3/3)

Hypervisor OS update

- EOL CentOS 7 - June 2024



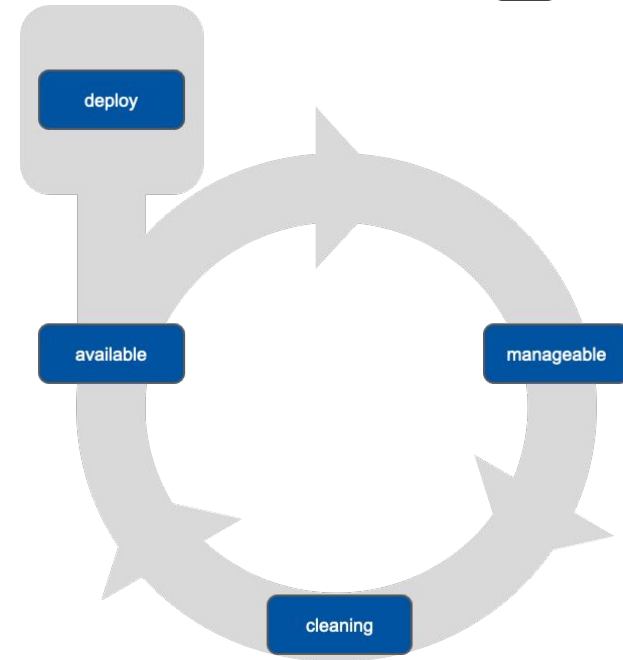
Irony: Server life-cycle management





Ironic updates

- Completed autoregistration
- Enrolled ARM and GPU servers, offering as VMs
- Introducing safeguards
 - *Skip list of devices*
 - *Node cleaning limiter*
- Upgrade to Xena release



28.03 at 10.25: *Providing ARM and GPU resources in the CERN Private Cloud Infrastructure*



Ironic future plans

- Retirement process fully with Ironic
 - *Full cleaning*
- [Deploying with Anaconda](#)
- Enrolment pipeline
- GPU benchmark
- GPU burn-in


New Data Centre in Preveessin

- Currently under construction, delivered by the end of 2023
- Introduce Software Defined Networking
- Fully independent region vs stretched API
- Availability zones by design
- VM boot from volume
 - ⊕ *Data stored on CEPH cluster*
 - ⊕ *Simplified live migration*
 - ⊖ *Rely on CEPH*
 - ⊖ *Latency*



Software Defined Network

- Static network configuration with linuxbridge
 - *Recabling to change network or IP*
 - *VMs cannot be moved across broadcast domains*
- Full SDN deployment in the Preveessin Data Center
 - *Virtual Networks, Floating IPs, Security Groups, QoS, Layer-3 LBaaS*
- Current setup based on TungstenFabric




LoadBalancing-as-a-Service



- Migrate from TungstenFabric to **Octavia**:
 - *Easier to maintain with components already used*
 - *Active community project with significant adoption*
 - *Superset of TungstenFabric's features*
 - *Hooks to the cloud network backend - can support ACTIVE-ACTIVE, Layer-3 LBaaS*
- General Network and Technical Network instances available

CPU steal

- Overcommit → fight for resources
- Live migration postpones but does not solve
- 3 aggregates in every cell:
 - *Overcommit*
 - *No overcommit*
 - *Spare*
- Increase flexibility



Thank you!

All our **open source** code is available on <https://gitlab.cern.ch/cloud-infrastructure>

My email: maryna.savchenko@cern.ch





home.cern