
CERN

Cloud Infrastructure Update: Operations, Campaigns and Evolution

Domingo Rivera Barros

HEPiX Autumn, 2 November 2022

Outline

- Cloud Service Overview
- Virtual Machines migration campaign
- Bare metal Service updates
- New features
  - GPU
  - SDN
- Future

CERN Cloud Service

- Infrastructure as a Service
- Production since **July 2013**
- **CentOS 7** based (updating to 8 soon)
- CERN Data Centre (adding a DC with a new region)
- **Highly scalable** architecture
  - 48 cells on 5 regions
- Most of the components in **Xena** release

  … Stein, Train, Ussuri, Victoria, Wallaby, Xena, Yoga, Zed

CERN Cloud Service

- Infrastructure as a Service
- Production since **July 2013**
- **CentOS 7** based (updating to 8 soon)
- CERN Data Centre (adding a DC with a new region)
- **Highly scalable** architecture
  - 48 cells on 5 regions
- Most of the components in **Xena** release

  … Stein, Train, Ussuri, Victoria, Wallaby, Xena, Yoga, Zed

CERN Cloud Service

- Infrastructure as a Service
- Production since **July 2013**
- **CentOS 7** based (updating to 8 soon)
- CERN Data Centre (adding a DC with a new region)
- **Highly scalable** architecture
  - 48 cells on 5 regions
- Most of the components in **Xena** release

  …Stein, Train, Ussuri, Victoria, Wallaby, Xena, Yoga, Zed

Cloud Components

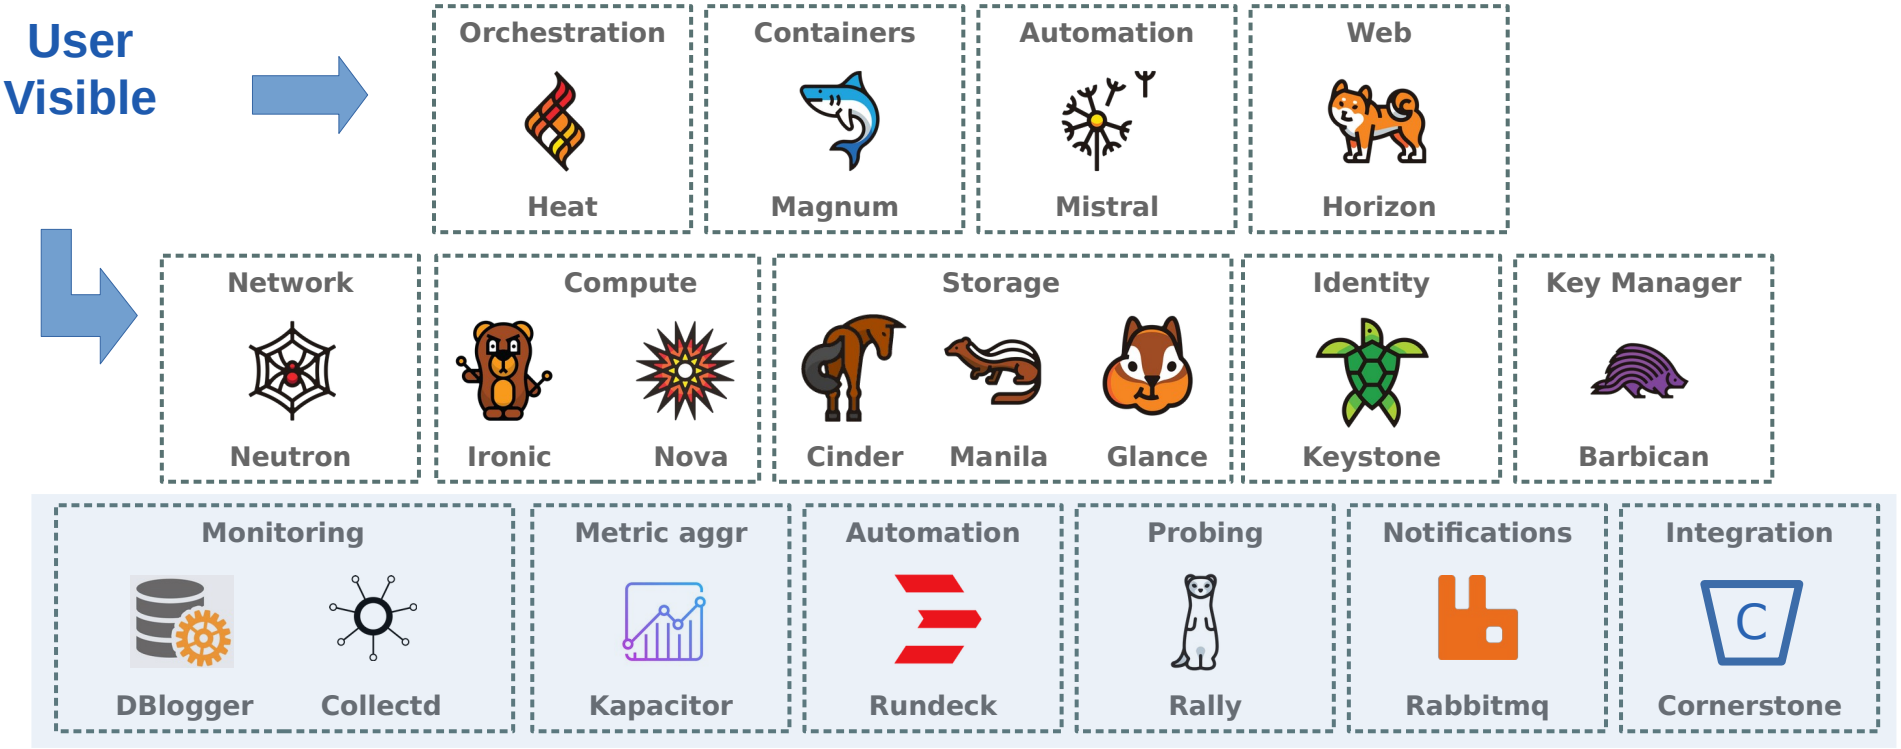

Last 60 days

Openstack services statistics

Users	Projects	Loadbalancers	Images	Volumes	Volumes...	File Shar...	File Shar...	Object St...	Object St...
3397	4629	339	5162	7325	3.80 PB	6196	1.13 PB	486	75.5 TB

Servers				Cores			RAM			Batch		
Physical	Physical in use	Hypervisors	Virtual	Physical	Hypervisors	Virtual	Physical	Hypervisors	Virtual	Servers	Cores	RAM
9060	8810	2030	13730	487 K	60.1 K	88.4 K	2.02 PB	393 TB	211 TB	5193	280011	1.06 PB

Time series

Physical servers

Last 10 months

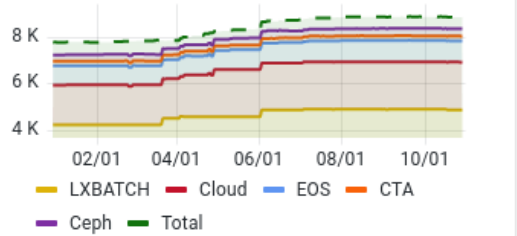

Virtual Machines

Virtual Machines created/deleted

Last 7 days

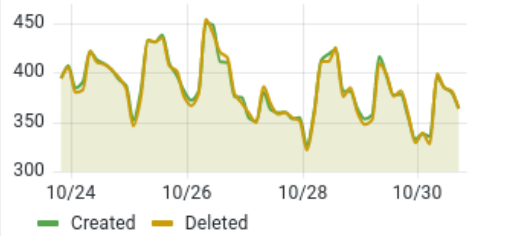

Shared cells availability

Last 30 days

Loadbalancers

Volumes

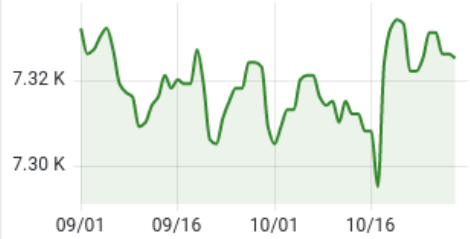

File Shares

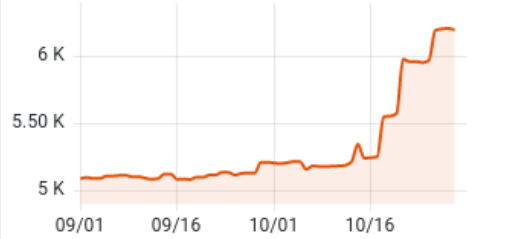

Object Store size

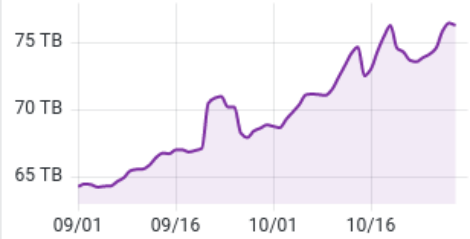

Last 60 days

Openstack services statistics

Users	Projects	Loadbalancers	Images	Volumes	Volumes...	File Shar...	File Shar...	Object St...	Object St...
3397	4629	339	5162	7325	3.80 PB	6196	1.13 PB	486	75.5 TB

Servers

Physical	Physical in use	Hypervisors	Virtual
9060	8810	2030	13730

Cores

Physical	Hypervisors	Virtual
487 K	60.1 K	88.4 K

RAM

Physical	Hypervisors	Virtual
2.02 PB	393 TB	211 TB

Batch

Servers	Cores	RAM
5193	280011	1.06 PB

Time series

Physical servers

Last 10 months

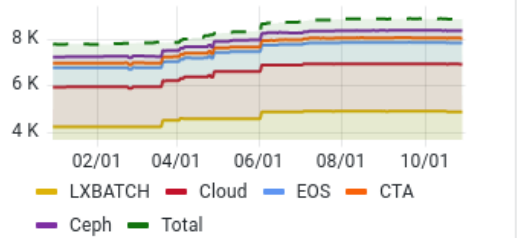

Virtual Machines

Virtual Machines created/deleted

Last 7 days

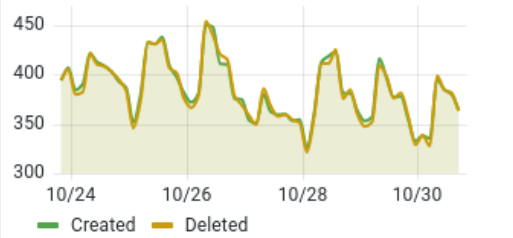

Shared cells availability

Last 30 days

Loadbalancers

Volumes

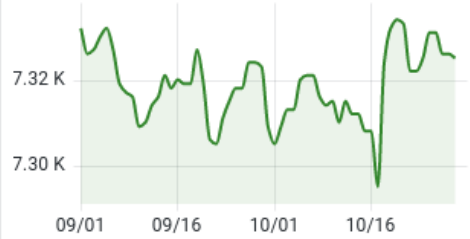

File Shares

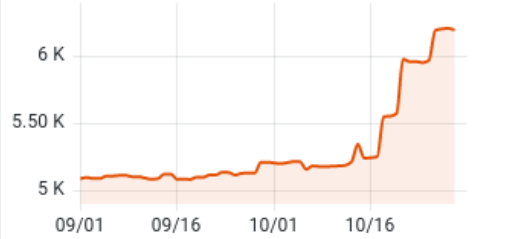

Object Store size

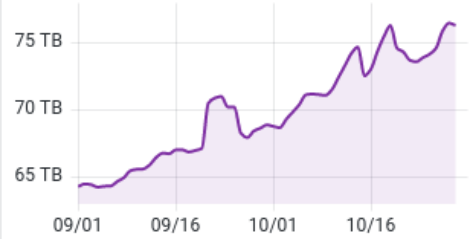

Last 60 days

Openstack services statistics

Users	Projects	Loadbalancers	Images	Volumes	Volumes...	File Shar...	File Shar...	Object St...	Object St...
3397	4629	339	5162	7325	3.80 PB	6196	1.13 PB	486	75.5 TB

Servers				Cores			RAM			Batch		
Physical	Physical in use	Hypervisors	Virtual	Physical	Hypervisors	Virtual	Physical	Hypervisors	Virtual	Servers	Cores	RAM
9060	8810	2030	13730	487 K	60.1 K	88.4 K	2.02 PB	393 TB	211 TB	5193	280011	1.06 PB

Time series

Physical servers

Last 10 months

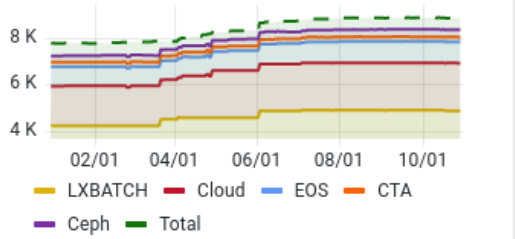

Virtual Machines

Virtual Machines created/deleted

Last 7 days

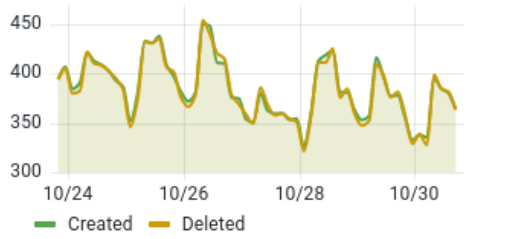

Shared cells availability

Last 30 days

Loadbalancers

Volumes

File Shares

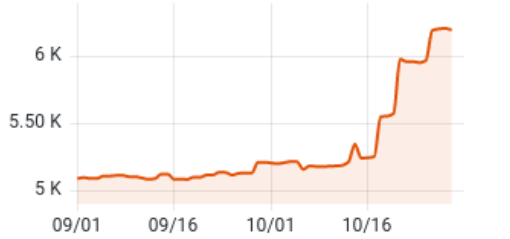

Object Store size

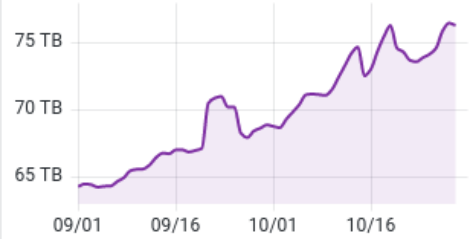

Last 60 days

Openstack services statistics

Users	Projects	Loadbalancers	Images	Volumes	Volumes...	File Shar...	File Shar...	Object St...	Object St...
3397	4629	339	5162	7325	3.80 PB	6196	1.13 PB	486	75.5 TB

Servers				Cores			RAM			Batch		
Physical	Physical in use	Hypervisors	Virtual	Physical	Hypervisors	Virtual	Physical	Hypervisors	Virtual	Servers	Cores	RAM
9060	8810	2030	13730	487 K	60.1 K	88.4 K	2.02 PB	393 TB	211 TB	5193	280011	1.06 PB

Time series

Physical servers

Last 10 months

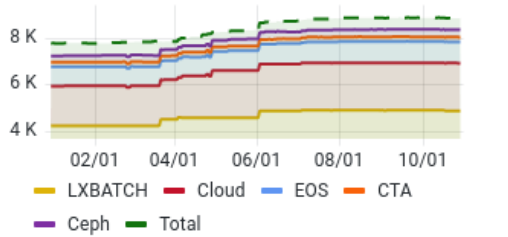

Virtual Machines

Virtual Machines created/deleted

Last 7 days

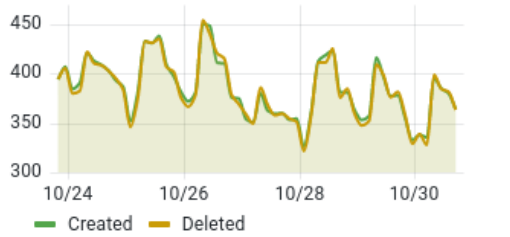

Shared cells availability

Last 30 days

Loadbalancers

Volumes

File Shares

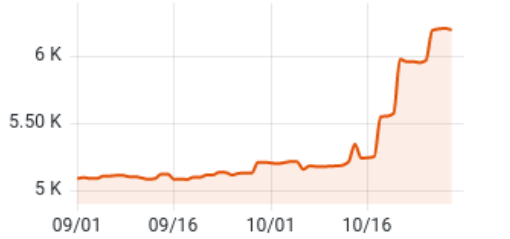

Object Store size

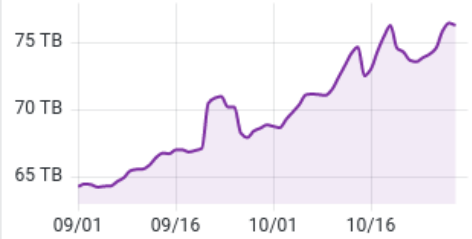

Last 60 days

Openstack services statistics

Users	Projects	Loadbalancers	Images	Volumes	Volumes...	File Shar...	File Shar...	Object St...	Object St...
3397	4629	339	5162	7325	3.80 PB	6196	1.13 PB	486	75.5 TB

Servers				Cores			RAM			Batch		
Physical	Physical in use	Hypervisors	Virtual	Physical	Hypervisors	Virtual	Physical	Hypervisors	Virtual	Servers	Cores	RAM
9060	8810	2030	13730	487 K	60.1 K	88.4 K	2.02 PB	393 TB	211 TB	5193	280011	1.06 PB

Time series

Physical servers

Last 10 months

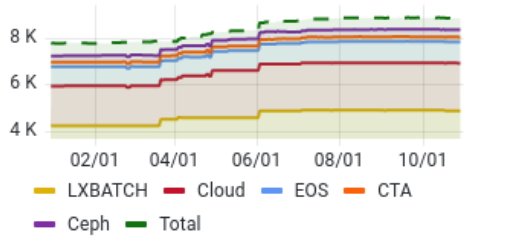

Virtual Machines

Virtual Machines created/deleted

Last 7 days

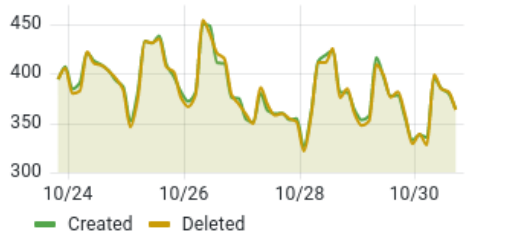

Shared cells availability

Last 30 days

Loadbalancers

Volumes

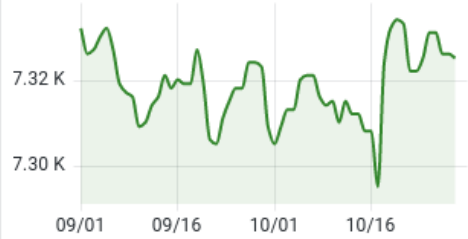

File Shares

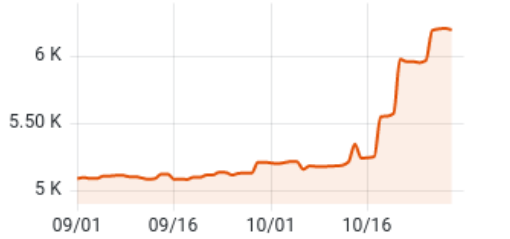

Object Store size

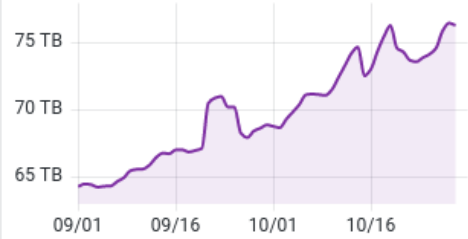

Last 60 days

Openstack services statistics

Users	Projects	Loadbalancers	Images	Volumes	Volumes...	File Shar...	File Shar...	Object St...	Object St...
3397	4629	339	5162	7325	3.80 PB	6196	1.13 PB	486	75.5 TB

Servers				Cores			RAM			Batch		
Physical	Physical in use	Hypervisors	Virtual	Physical	Hypervisors	Virtual	Physical	Hypervisors	Virtual	Servers	Cores	RAM
9060	8810	2030	13730	487 K	60.1 K	88.4 K	2.02 PB	393 TB	211 TB	5193	280011	1.06 PB

Time series

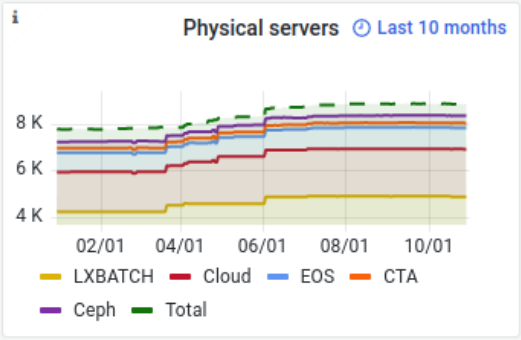

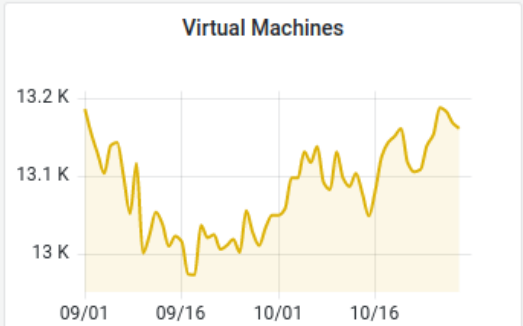

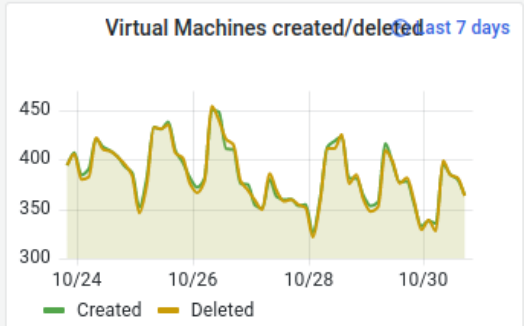

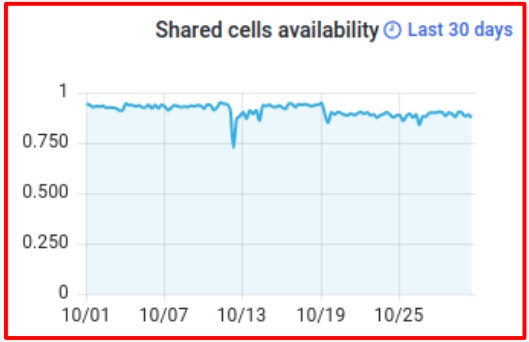

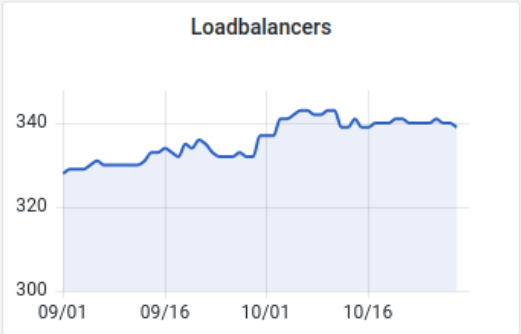

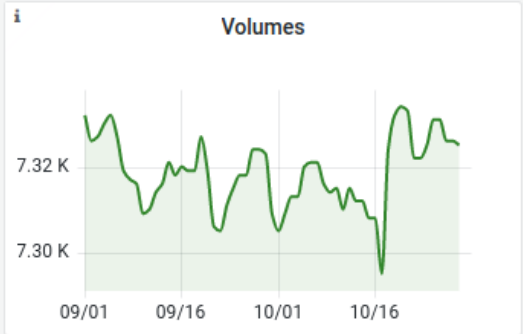

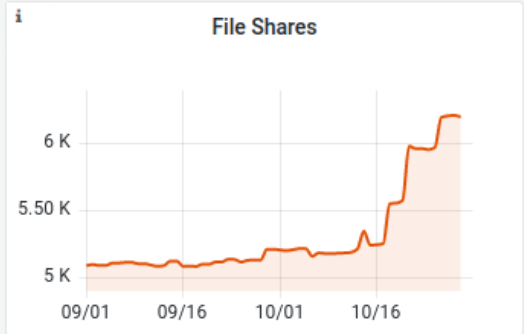

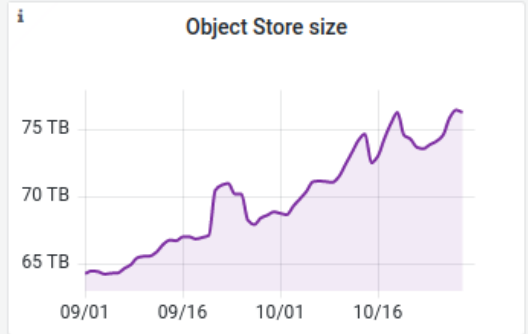

VM Migration Campaign

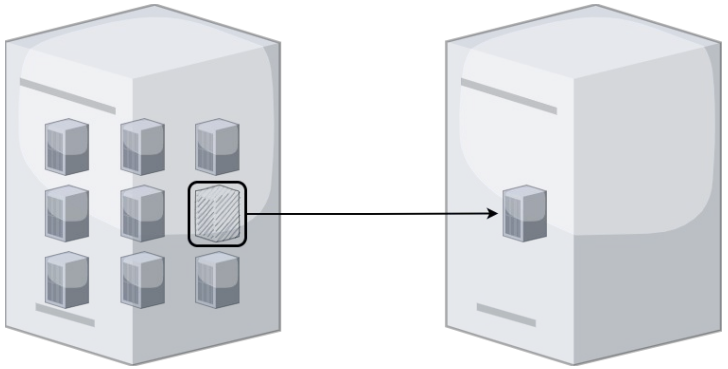

- **Replace legacy network component (Nova network to Neutron)**
- **Unblock OpenStack and hypervisor's OS updates**
  - Nova and Neutron in Stein release
  - Will allow to use advanced network features (e.g. Security groups)
- **Complex operation**
  - Around 4000 VMs
  - Stop, snapshot and migrate (cold migration)
  - 1-2 hours downtime per VM (depending on instance size and flavor)
  - Preserving VM properties

VM Migration Campaign

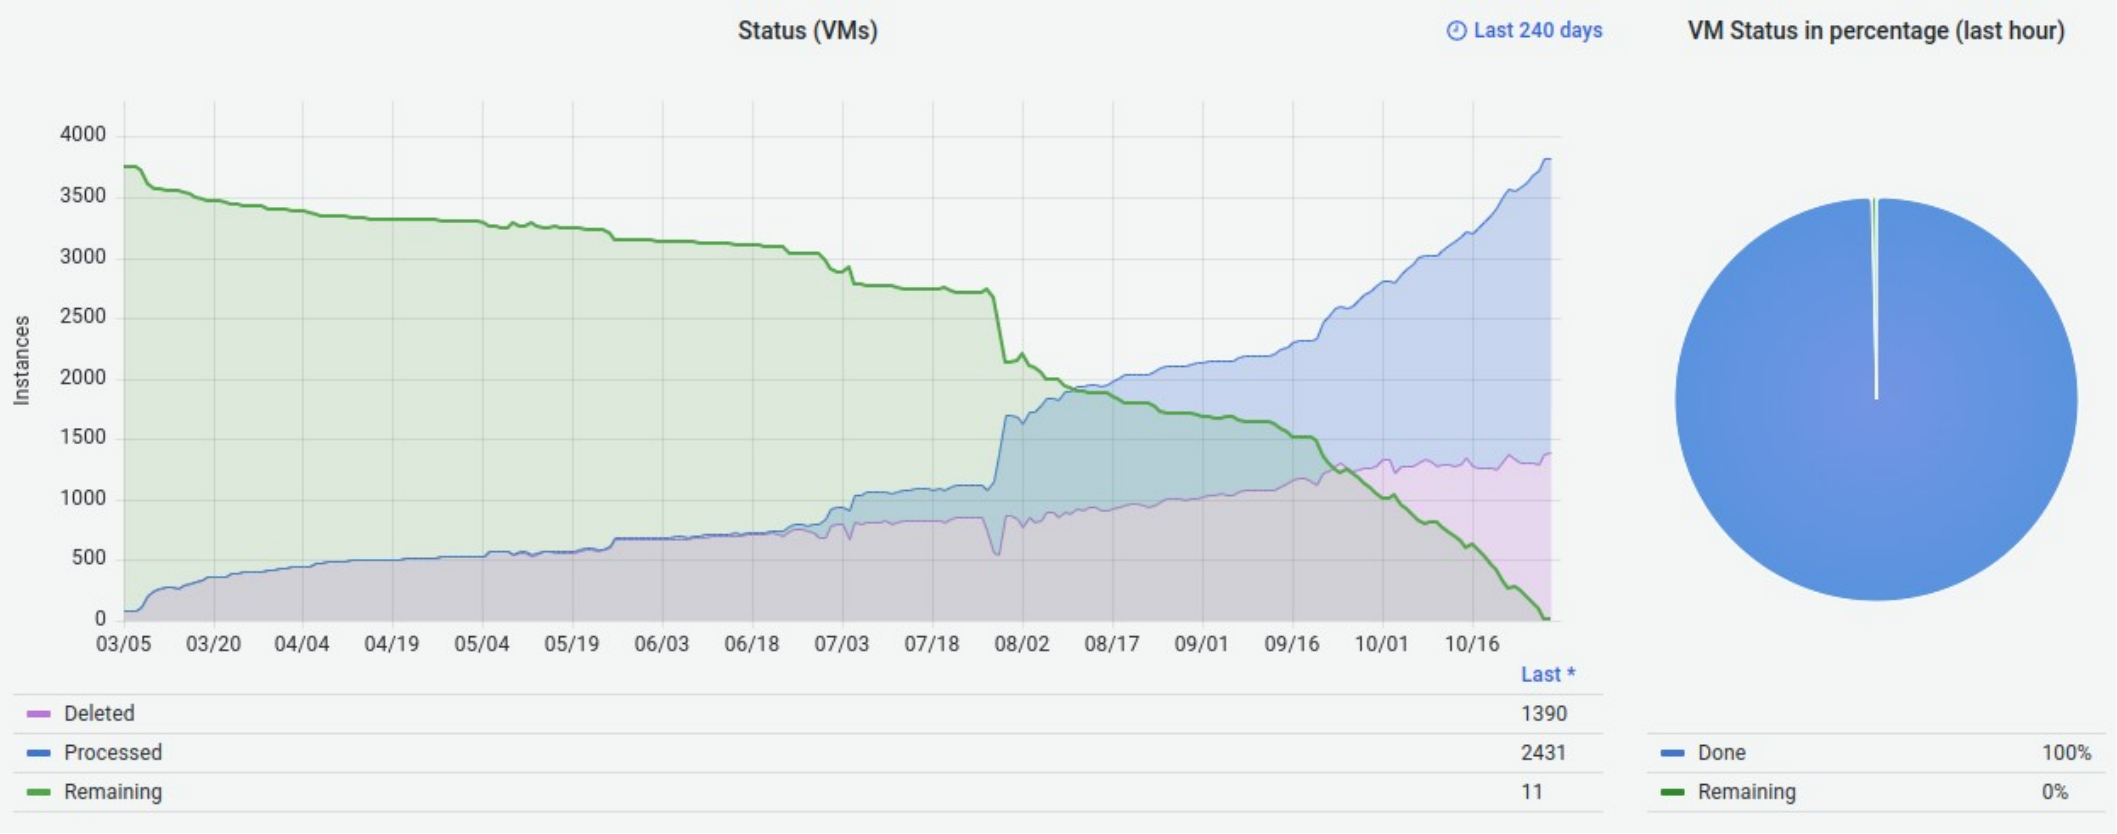

VM Migration Campaign

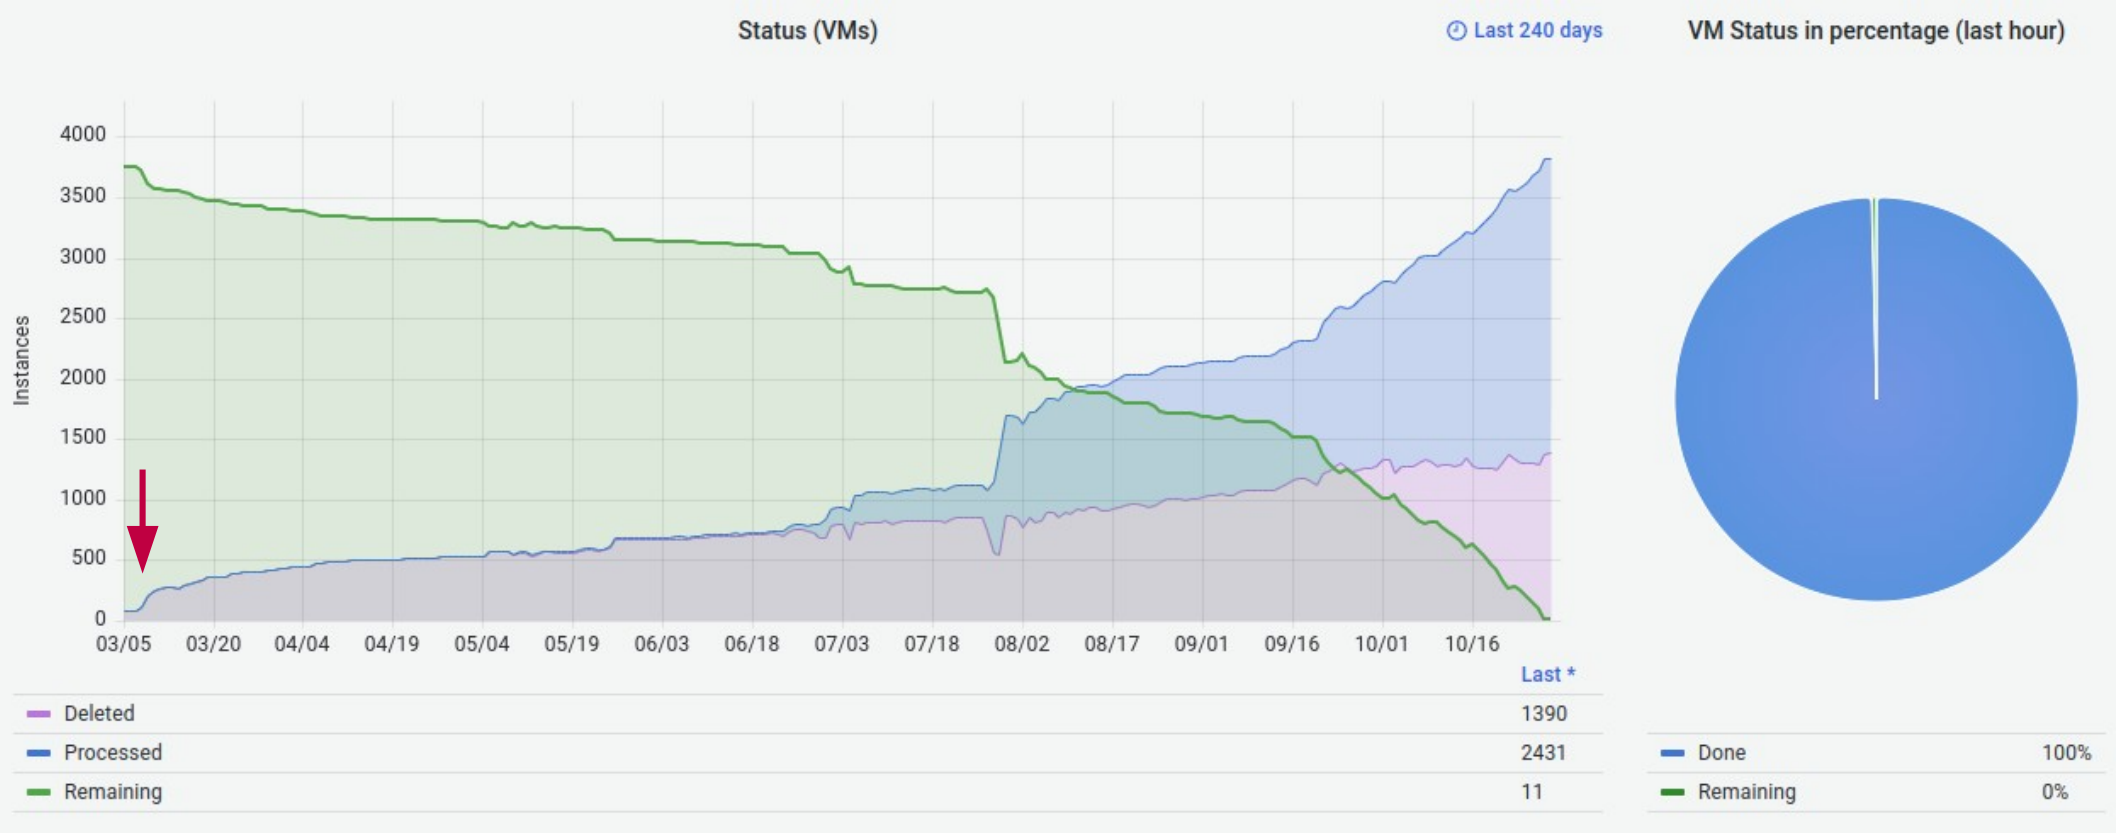

VM Migration Campaign

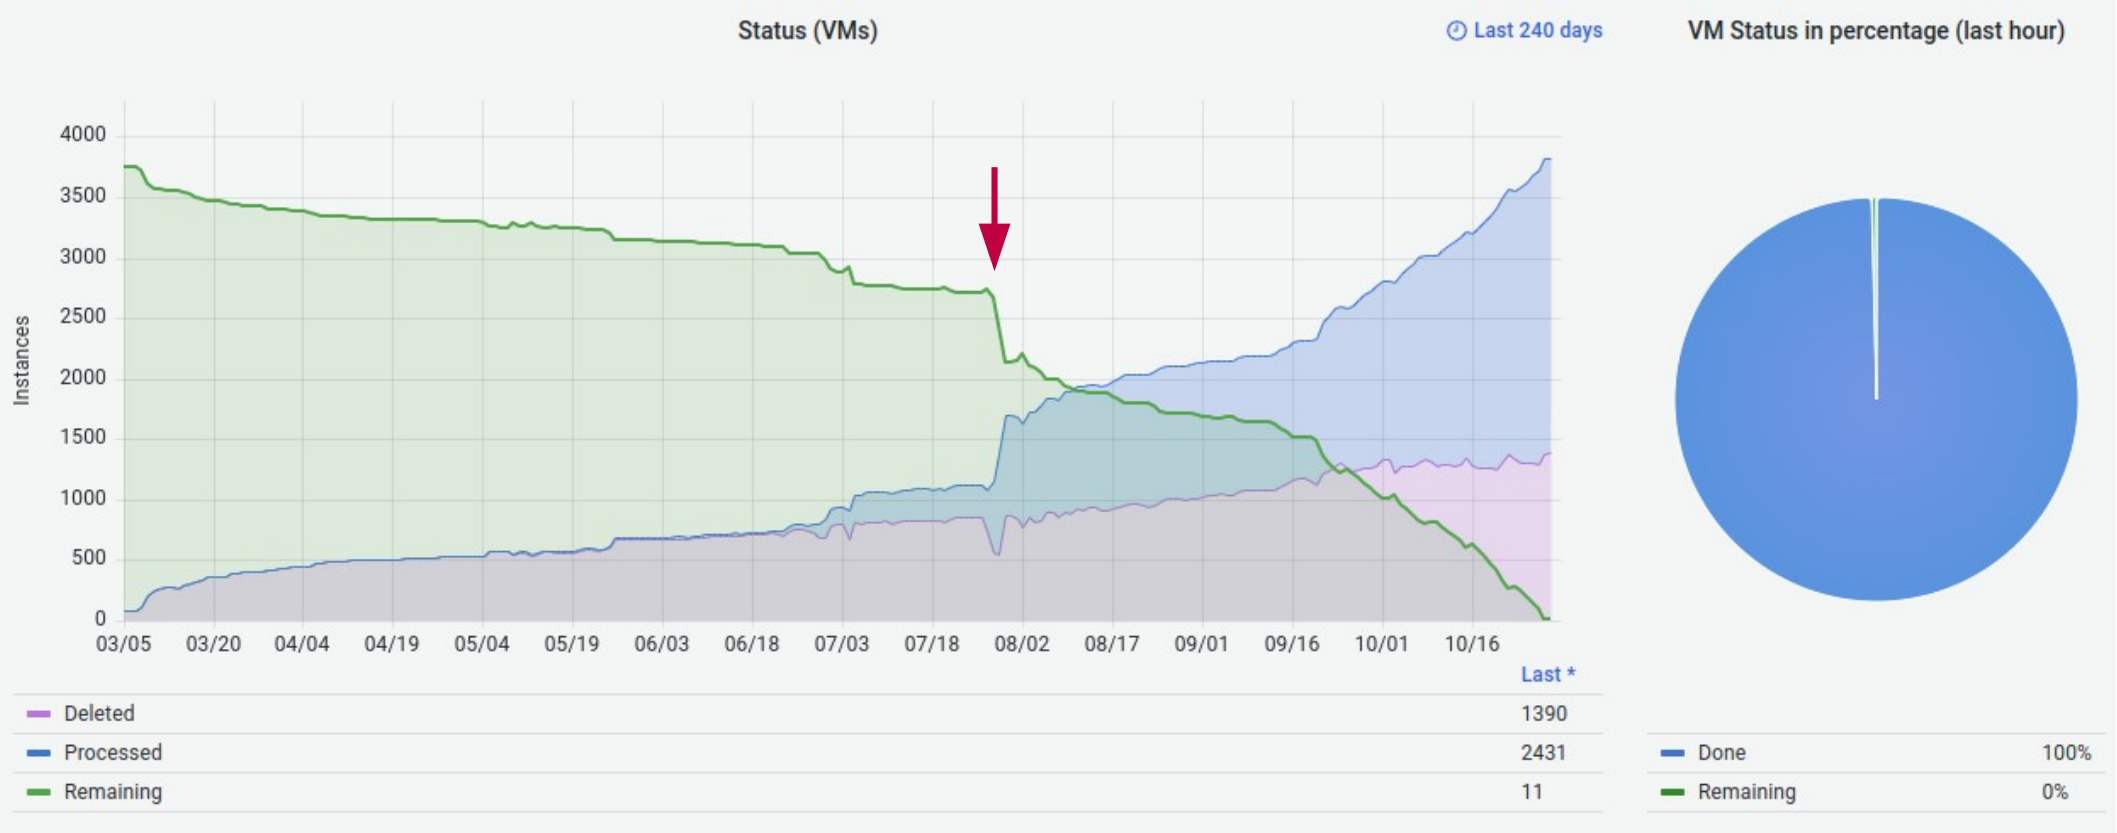

VM Migration Campaign

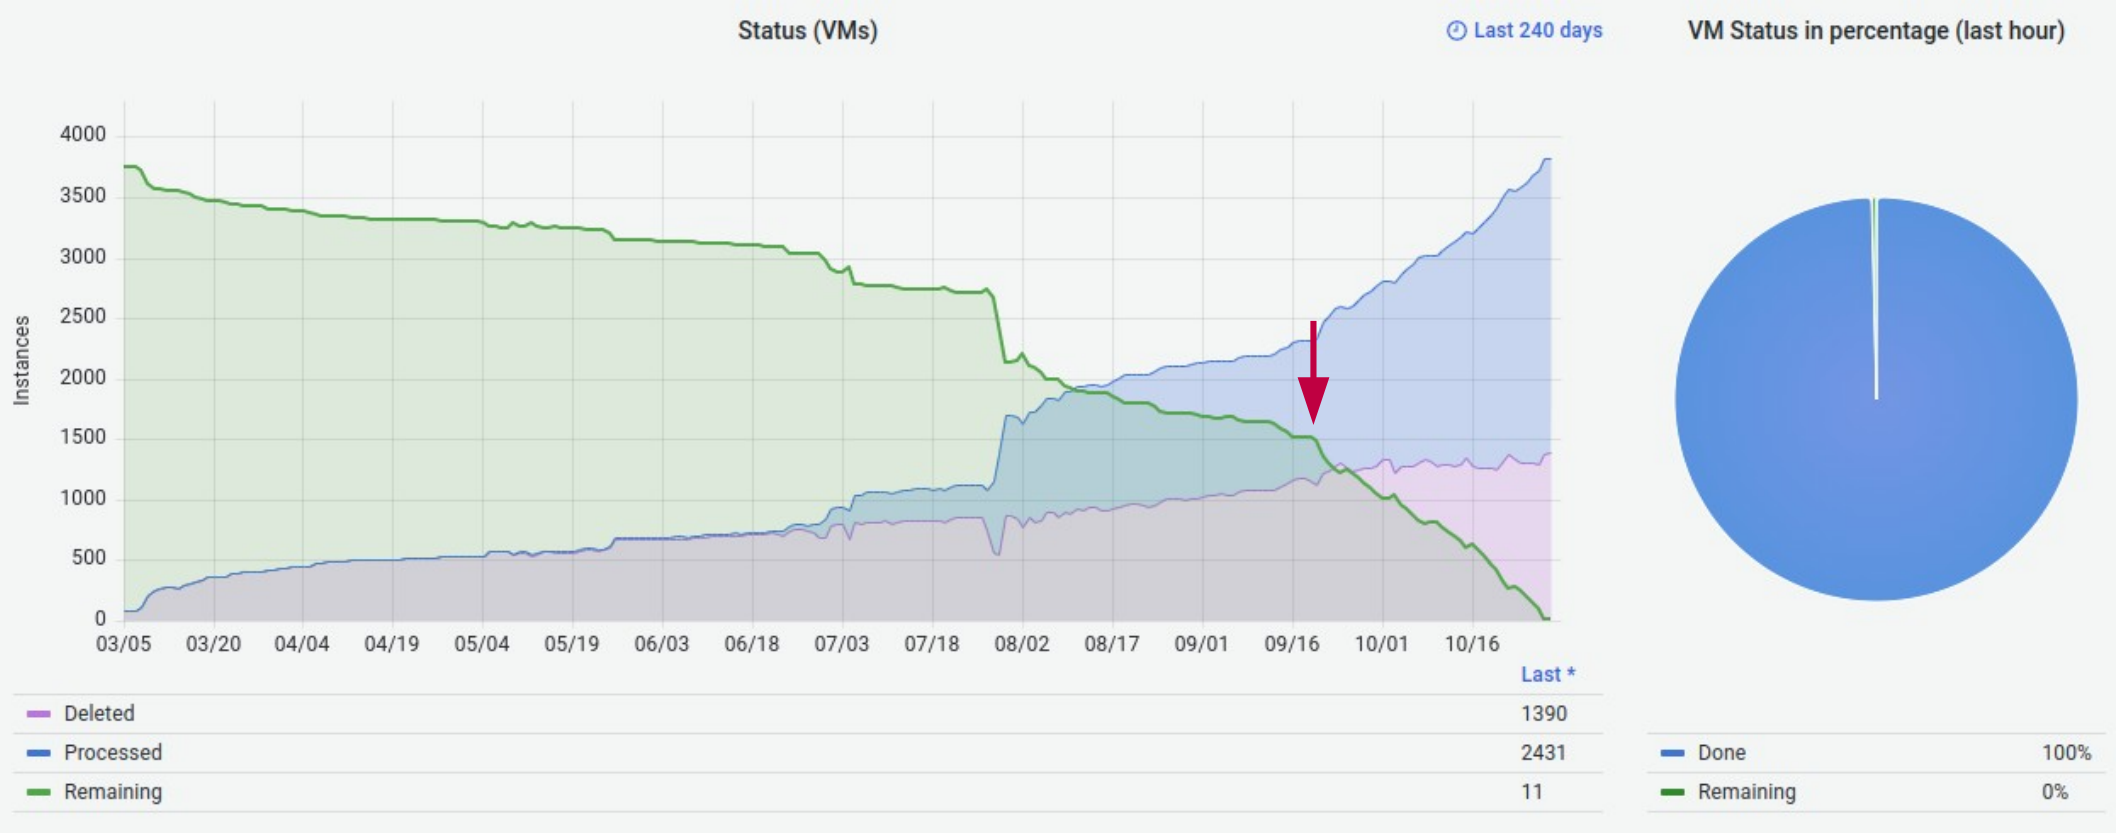

(OpenStack) Ironic

- **Bare metal provisioning Service**
- **Offers physical servers using the same interface as for VMs**
- **Stand-alone or integrated with OpenStack**
- **Facilitates and consolidates provisioning, auditing and inventory procedures**

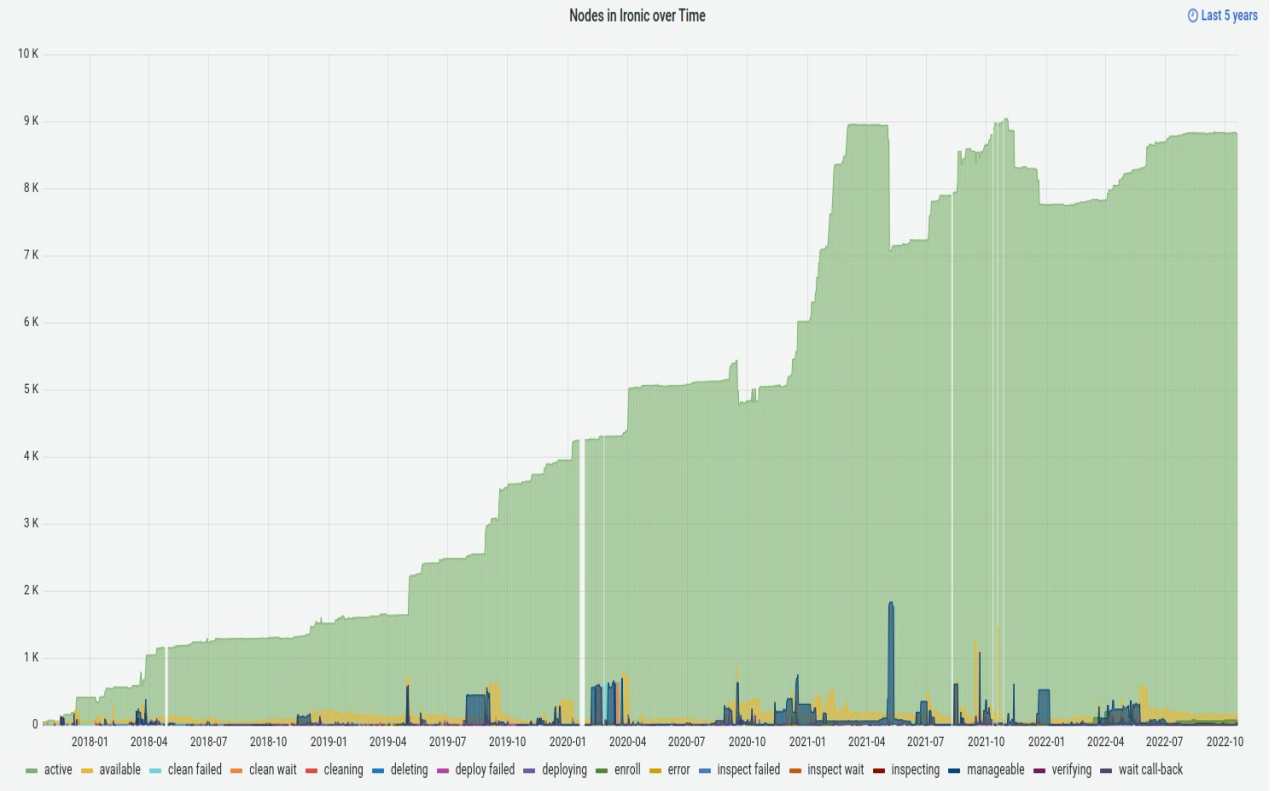

Ironic updates I: Server life-cycle management

Full provisioning life-cycle moved to production. Latest additions:

- Auto-registration for new deliveries
- Ironic's burn-in is the default

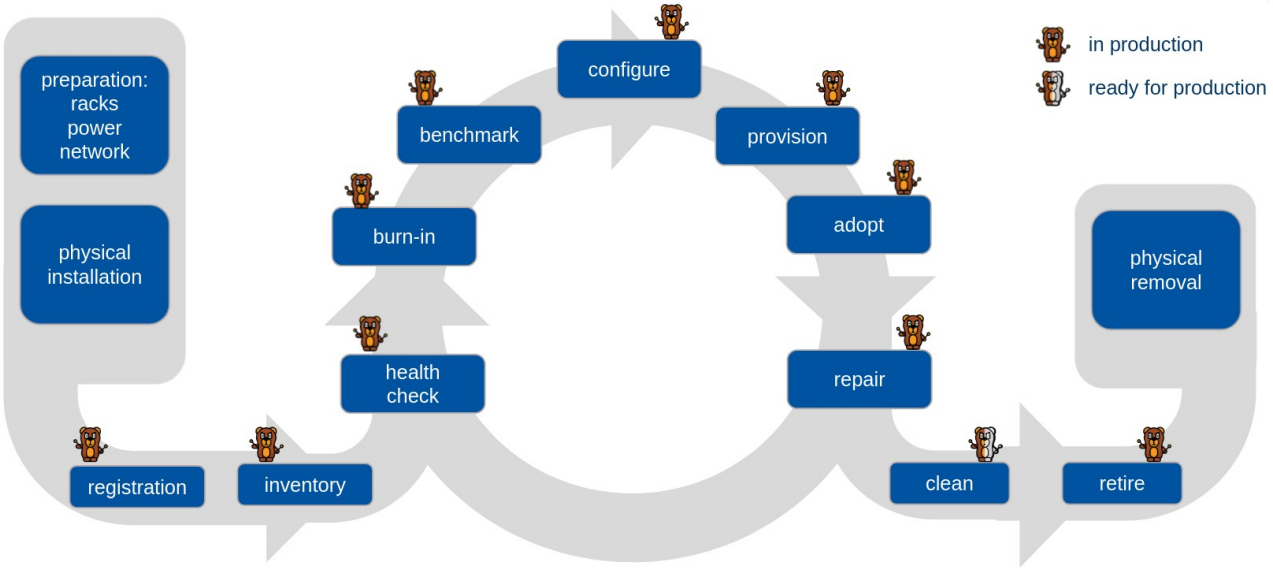

Ironic updates II: ARM nodes support

- **Successfully enrolled ARM servers with Ironic for the first time. This required adapting:**
  - Create an AArch64 CentOS image
  - Ironic (Python Agent) image
  - Auto-registration process
  - Benchmarking

- **Next step is to convert the servers into hypervisors to offer ARM based VMs**

GPU provisioning

- **Many different use cases require access to GPUs with different utilization**
  - deep learning, inference, analysis, simulations…
- **4 different Nvidia models available (T4, V100, V100s and A100)**
- **Available as vGPU or pci-passthrough (currently looking at Multi-instance GPU)**
- **Really low resources, preparing a lease model**
- **Quota handling to be improved**

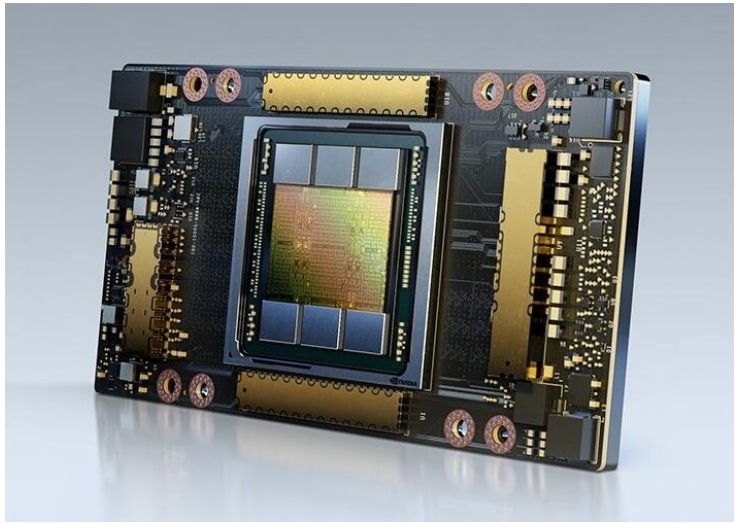

New Data Centre in Prevessin

- Currently under construction, delivery by end of 2023
- Provide **extra capacity** for the upcoming LHC runs
  - 3 floors with up to 4 MW per floor (**12 MW**).Current DC 3.5 MW
- Greenfield deployment
  - **AvZs isolation by design**
  - Dedicated OpenStack control plane and Ceph Clusters
  - May change **hypervisor disk layout**
  - Introduce **Software Defined Networking**

Software Defined Networking

- **Current networking model tightened to the infrastructure**
  - VMs cannot be moved across broadcast domains, rooms or switches
  - Physical machines need to be re-cabled if they need to change IP or network
- **Several technologies evaluated or under evaluation**
  - OpenDaylight, OpenContrail/Tungsten Fabric, OVN
  - Currently offering LBaaS
- **Full SDN deployment on new Data Centre**
  - Virtual Networks, Floating IPs, LBaaS ...
  - Provide maximum flexibility to end users

Future

- OpenStack components upgrades
- Hypervisors OS upgrades
- Provide ARM VMs
- GPUs offering
- Address CPU steal
- Test AMD processors
- Preparation for new Data Centre

Thank you

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

home.cern