

IN2P3-CC cloud computing (IAAS) status

FJPPL

Feb 9-11th 2016

- Use cases
 - R&D
 - Internal core services
 - Computing
 - Hosting
- Some implementation details
 - Deployment
 - Existing clusters and usage
- What's next

Destinated to IN2P3 labs, experiments and personnel for non production services

Need for self service →

- host lifecycle directly handled by end users
- accurate sizing and environment specification

Examples :

- Preproduction
- Testings
- Developments
- Trainings
- ...

Use cases : private cloud for infrastructure core services (2)

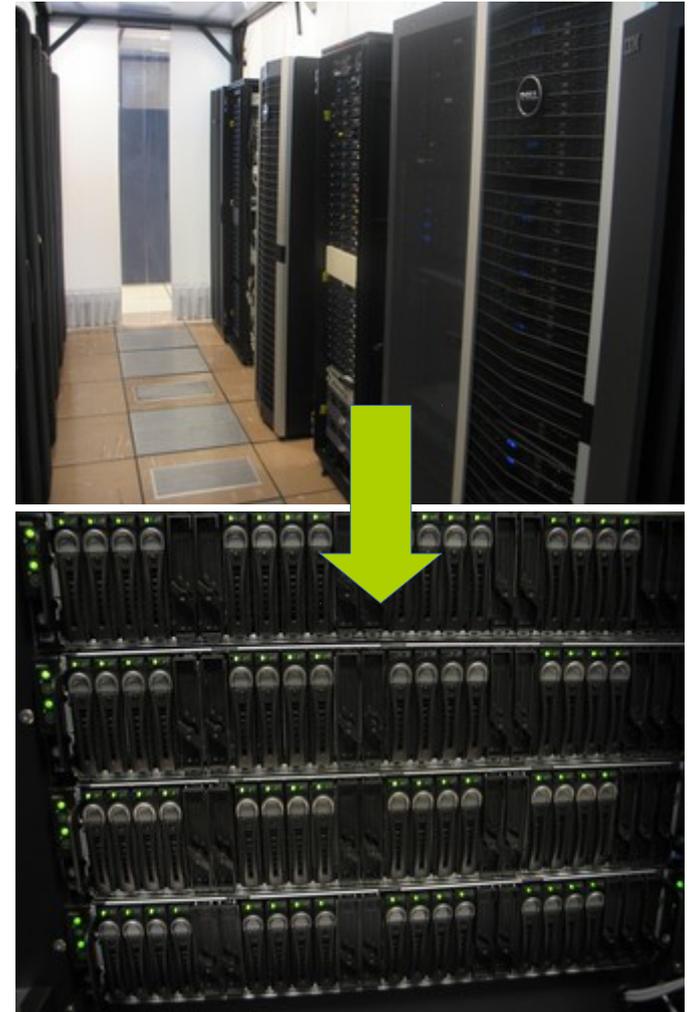
For whom : IN2P3-CC SA team provisioning infrastructure services

Why ?

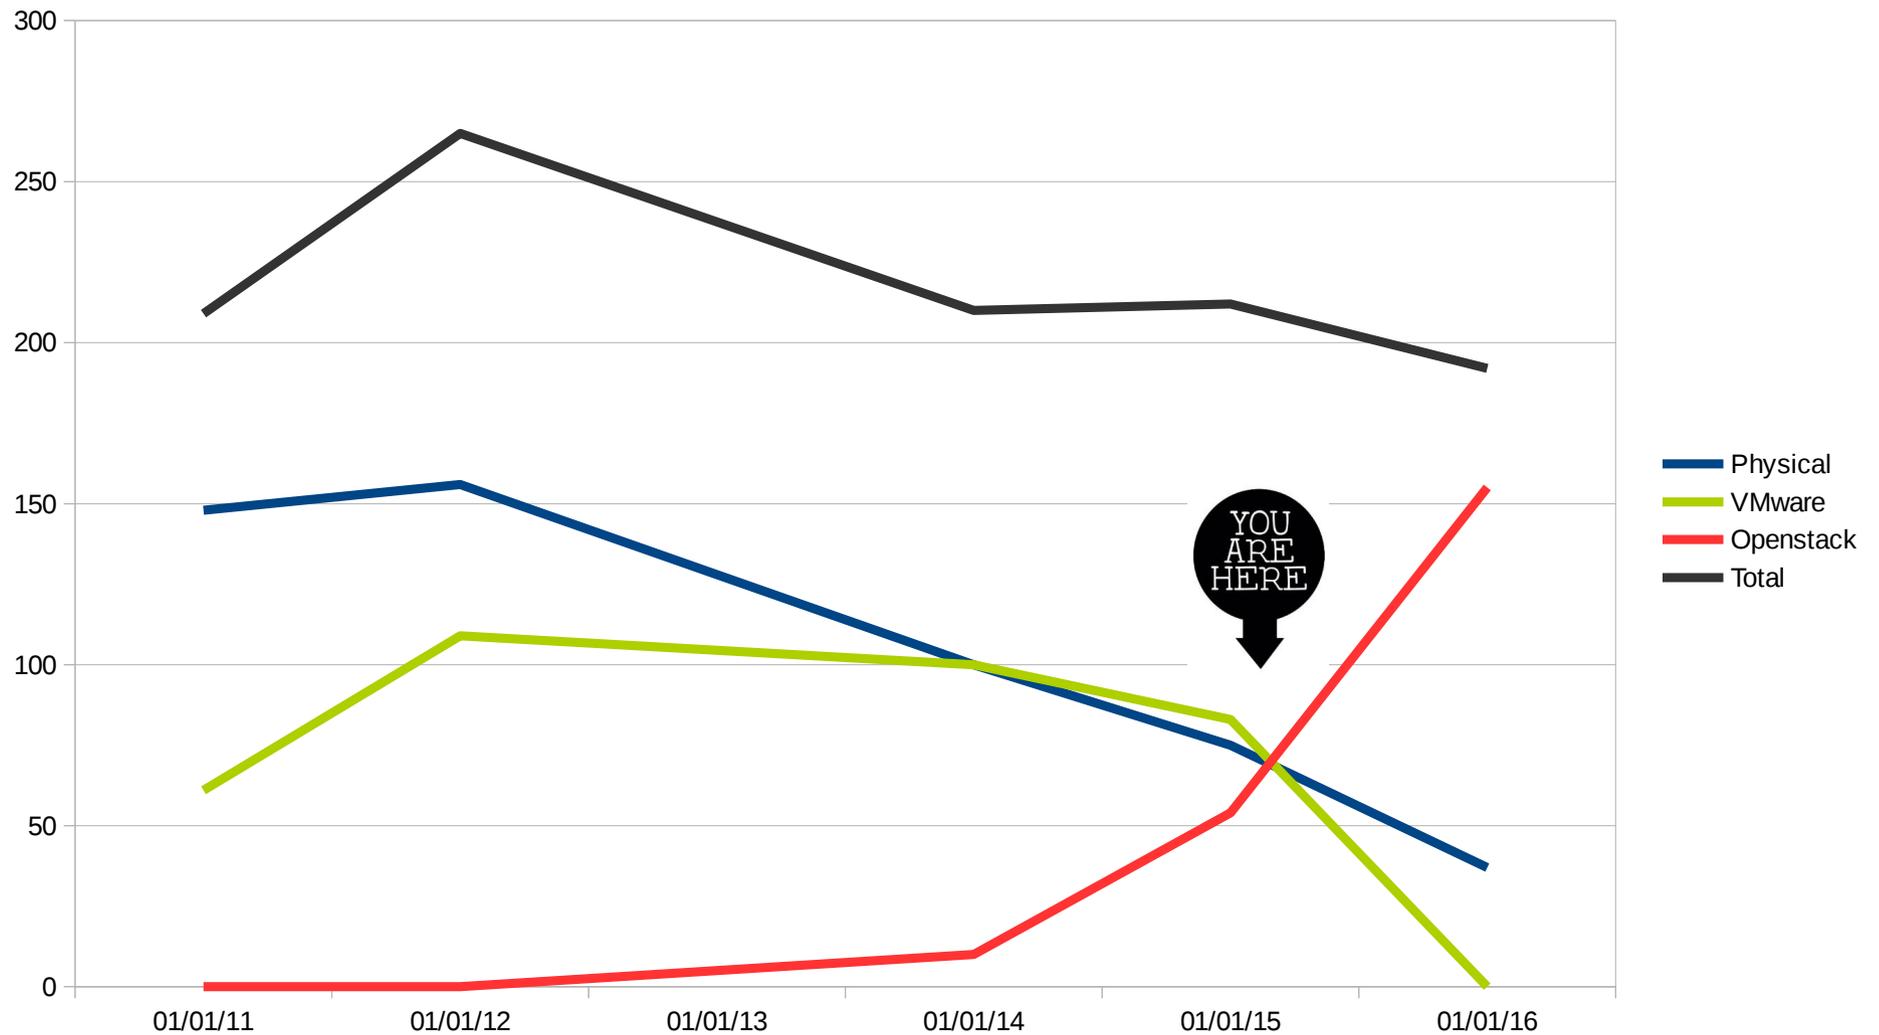
- All usual virtualization motivations
- Save VMware licences fee
- Satisfying service availability and performance :
 - No SPOF allows more than a few minutes DT
 - GPFS backend : horizontal capacity and IOPS scaling
 - LACP rr 2x10Gbps NICs
 - Live migration
 - Instances image automatic snapshots

Examples

- elastic search
- kerberos
- grid services
- puppet
- ... (80 different kinds)



Core services deployment



Motivations :

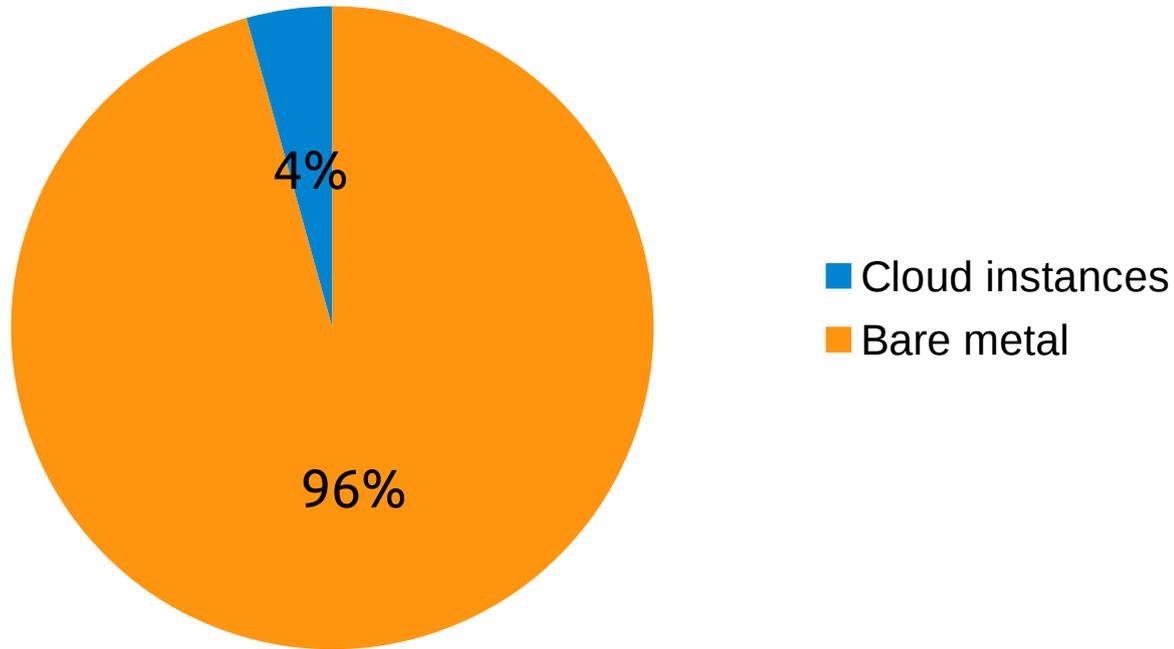
- Gain access to opportunistic resources
- Use specific environment (OS, softwares...)
- Implement their own scheduling
- Ease software deployment
- Enable and POC cloud resources for their computing models

Current alternatives to compute on cloud resources :

- Evaluation of HTCondor
- Direct use of cloud interfaces (EC2/Nova...), FSS enabling to share resources
- Tools to smooth WN instantiation for HTC administrators

Experiments using the cloud for computation :

- Large Synoptic Survey Telescop (<http://www.lsst.org>)
- Euclid (<http://www.euclid-ec.org/>)
- Atlas (MC simulation)
- Bioaster
- ELISA (www.elisascience.org)



Computing resources :

- ~1k HT cores in regard with 25k on HTC farm

Community cloud for whom ?

- IN2P3-CC institutional users (labs/experiments)

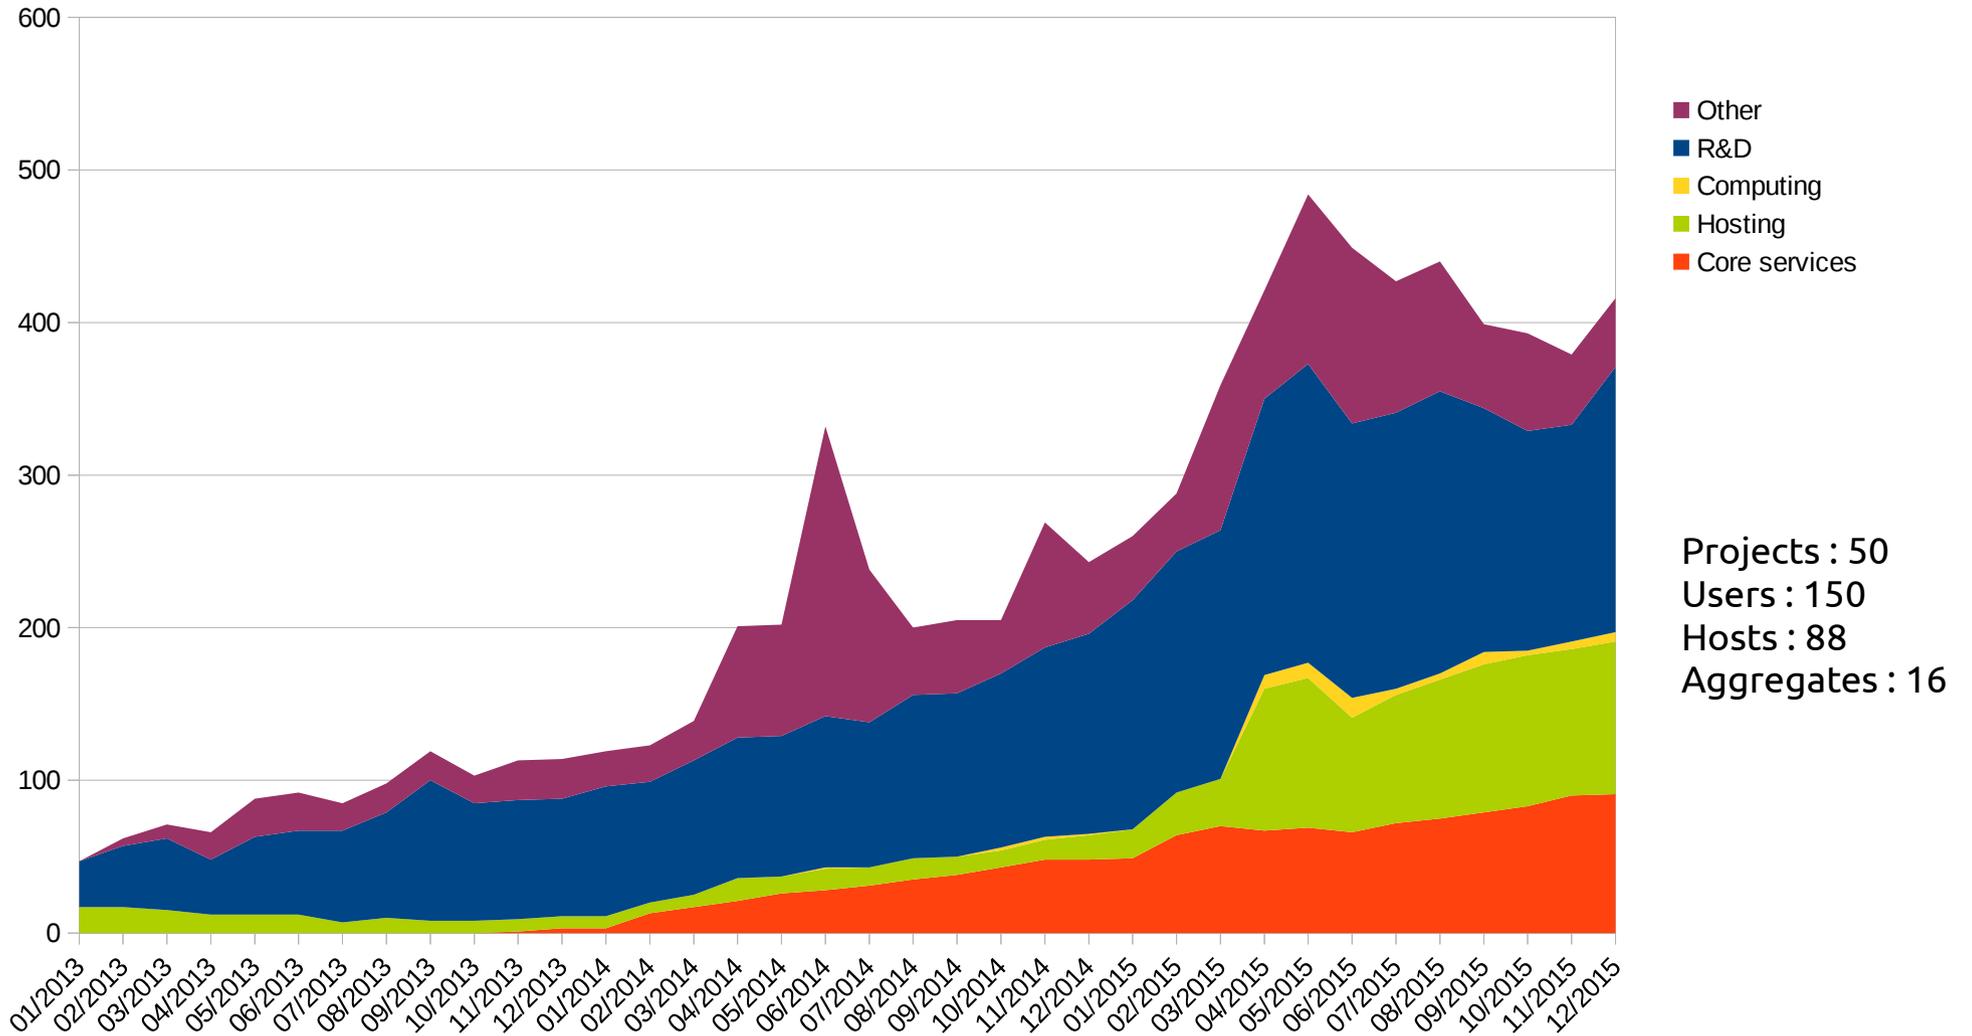
Why ?

- Users benefit from an highly available, optimized and cost effective environment

Examples :

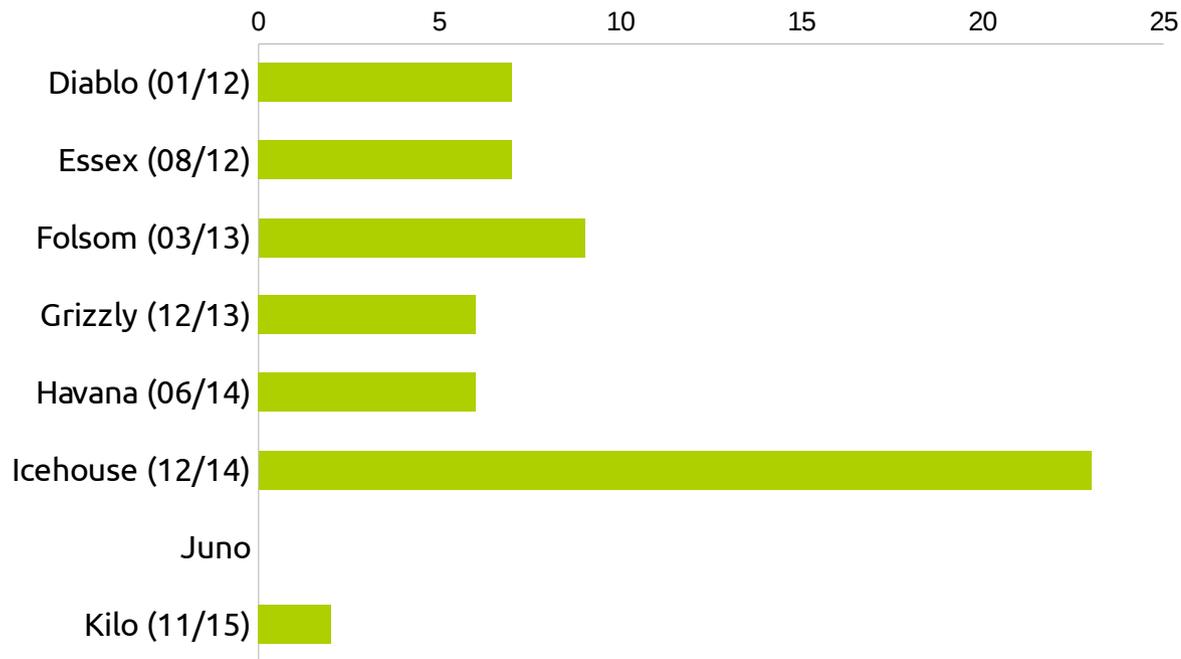
- AMI
- eTRIKS
- Bioaster

Overall usage



Some implementation details

Openstack deployments



Operational components :

Keystone
Glance
Nova
Horizon
Cinder
Ceilometer
Swift

Currently deployed :

Neutron

Deployment :

- CentOS 7
- RDO packages
- Puppet configuration

Computing clusters :

Hosting : 72 cores (DELL C6100)

R&D : 312 cores (DELL C6100)

Core services : 160 cores (DELL R630)

Computing : 1128 cores (DELL C6100)

→ **1672 cores in total**

Storage backend for HA usages:

- 24TB cluster (GPFS)

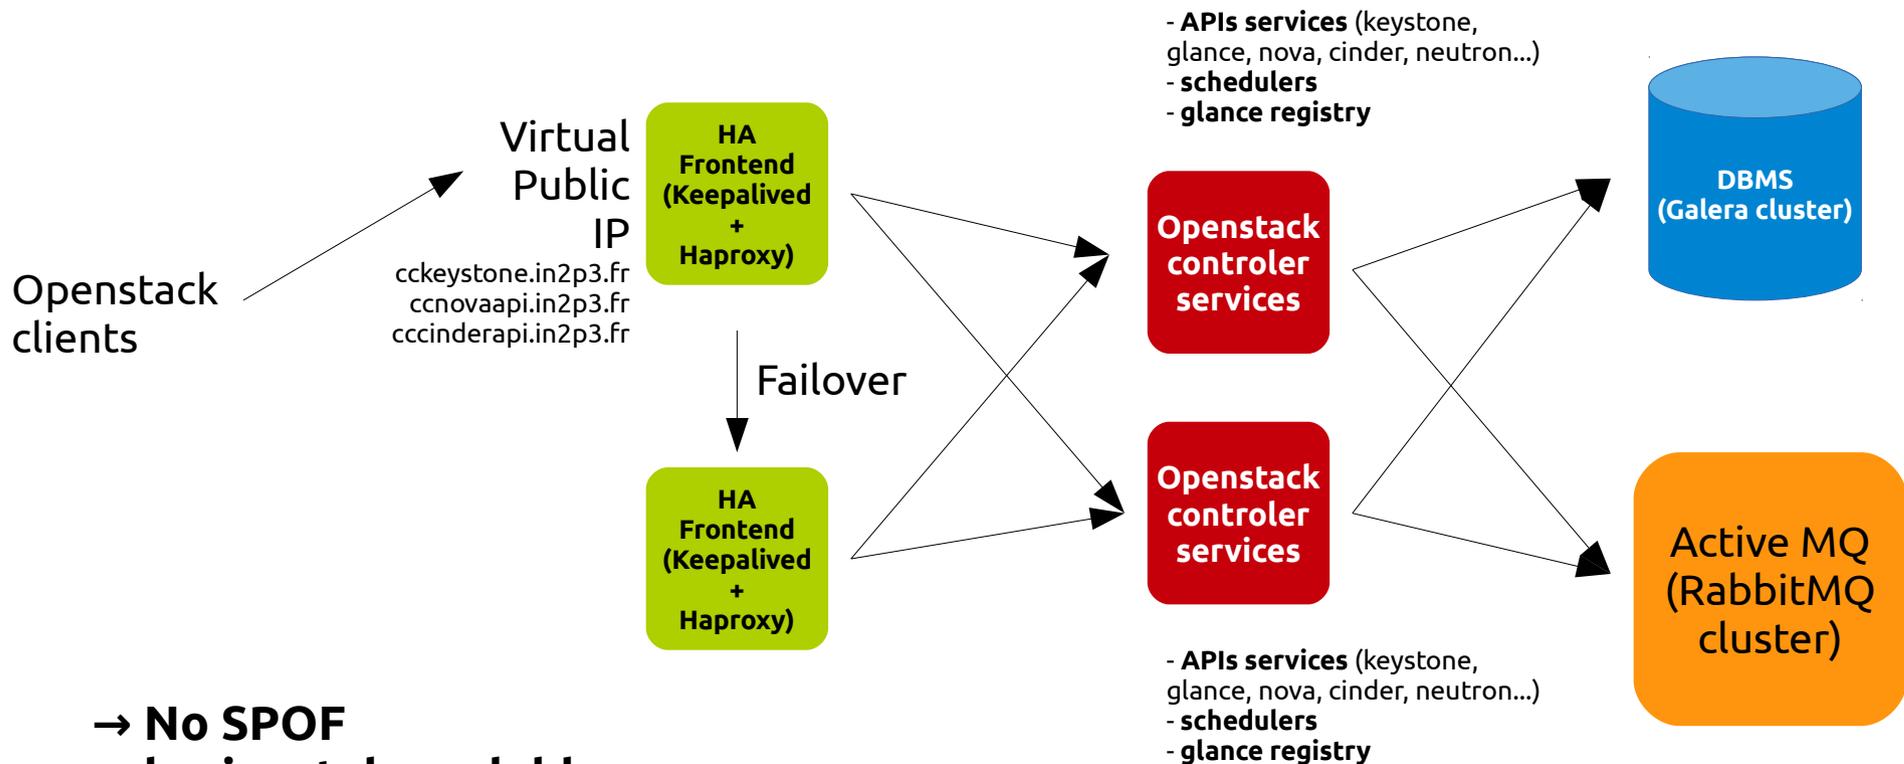
S3 Storage (Swift):

- 24TB DAS



Service availability

Keepalived → frontend high availability failover (using IPVS/VRRP)
Haproxy → load balancing amongst redundant services

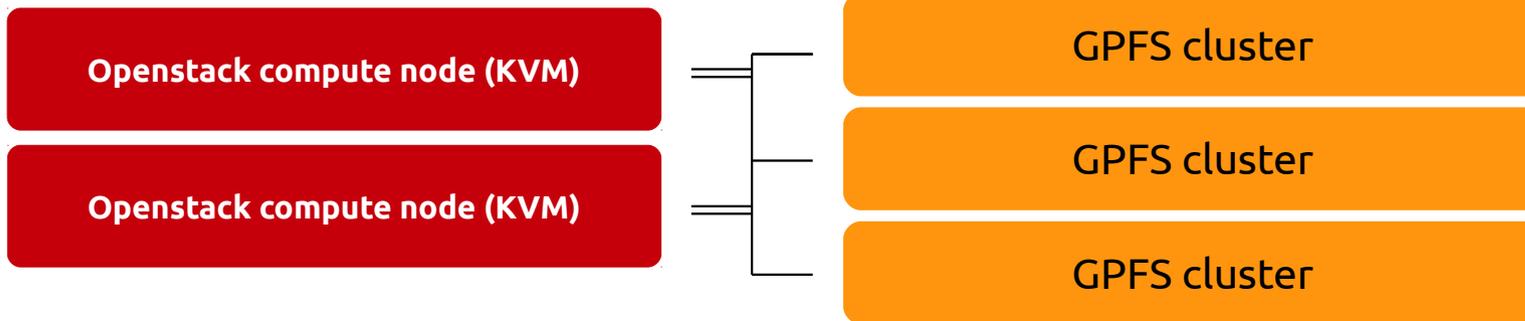


- No SPOF
- horizontally scalable
- Load balanced

RAM+CPU
horizontal
scaling

10 Gbps dedicated
network

Storage
capacity
& throughput
horizontal
scaling



3 way data replication
10 Gbps connection per host (12 drives) → maximum throughput

- only SPOF : compute node shutdown, service recovery within a VM boot
- horizontaly scalable

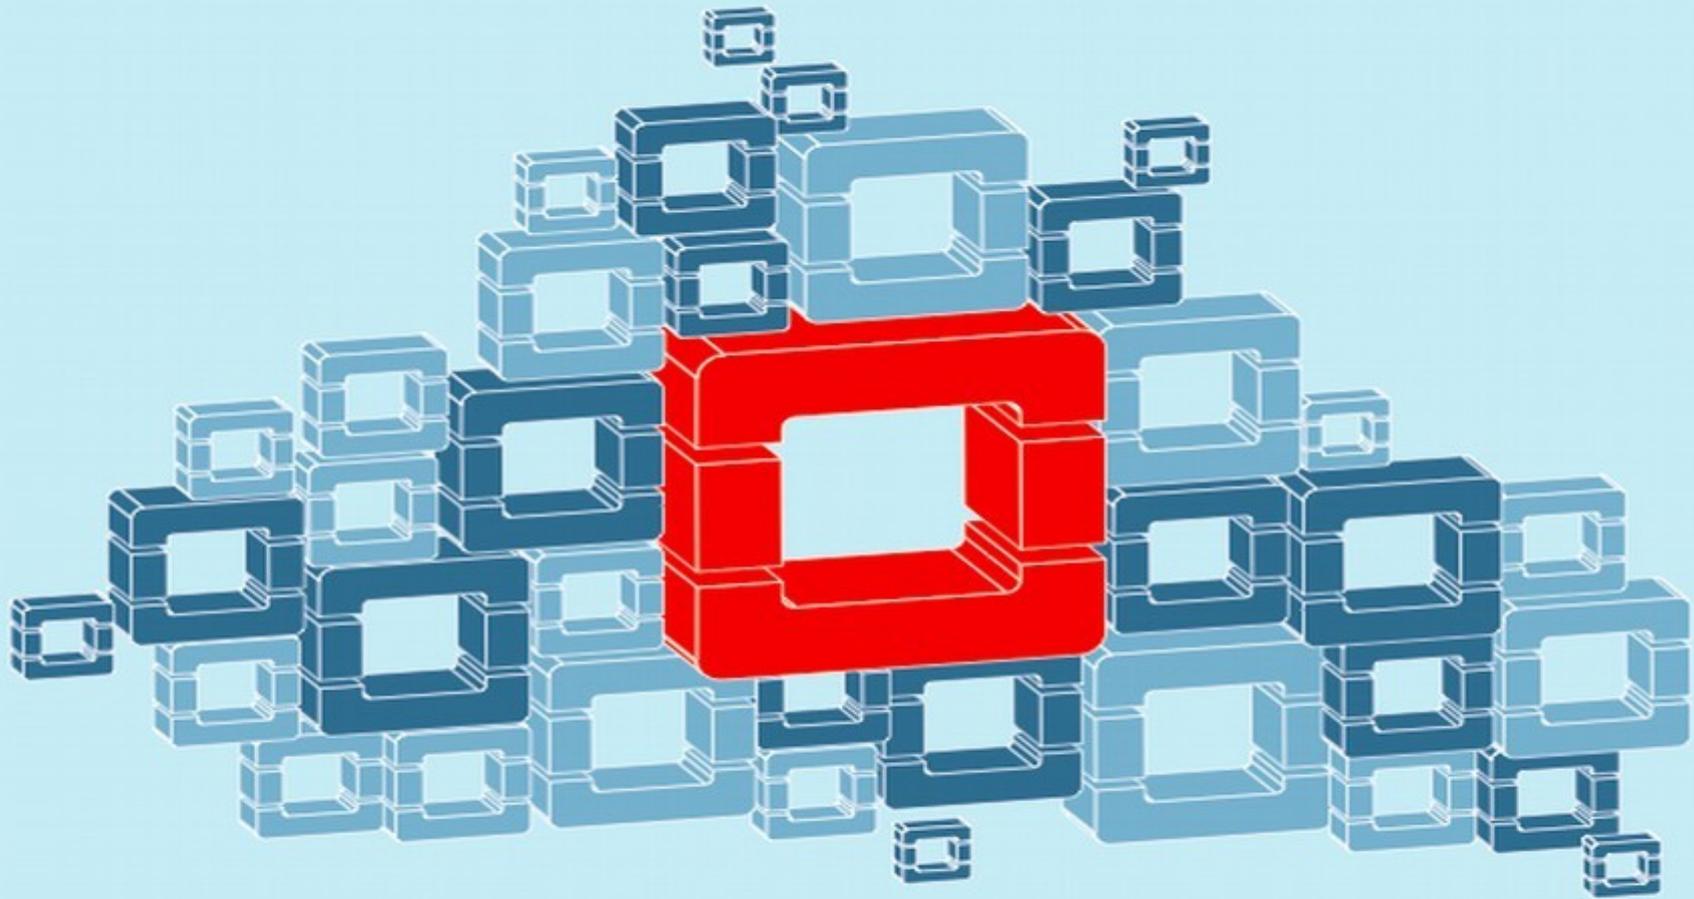
What's next ?

Short term milestones :

- FSS for computing
- Neutron migration

Longer term objectives :

- Migrate most computing farm



(thank you)

Questions ?