

IN2P3-CC cloud computing (IAAS) status

Feb '17

R&D : academic labs and experiments desiring to develop and evaluate new solutions.

- average availability, mainly private networking, average performance

Public cloud : offload IT services on a public IAAS cloud

Main characteristics :

- GPFS backend : horizontal capacity and IOPS scaling
- LACP rr 2x10Gbps NICs
- Live migration

- highest availability, public networking, requested performance

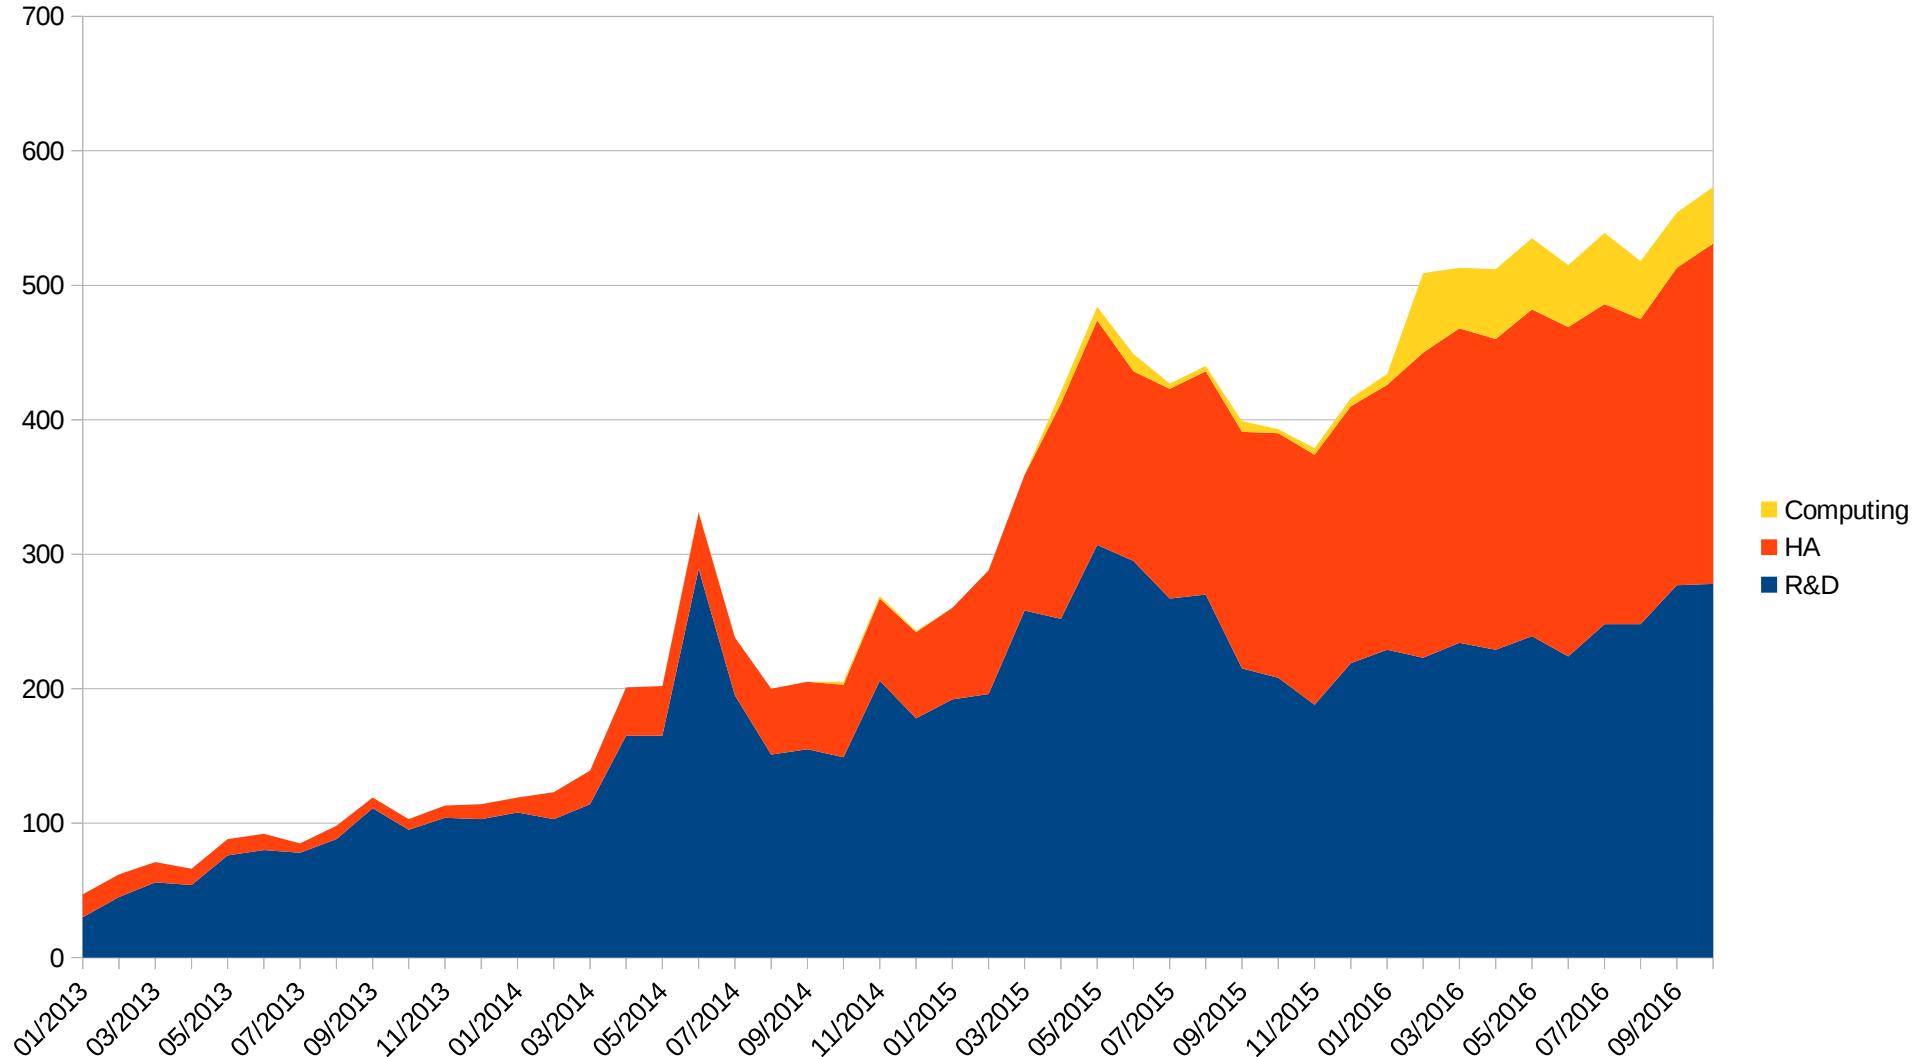
Computing : designed to handle HTC workloads

Main characteristics :

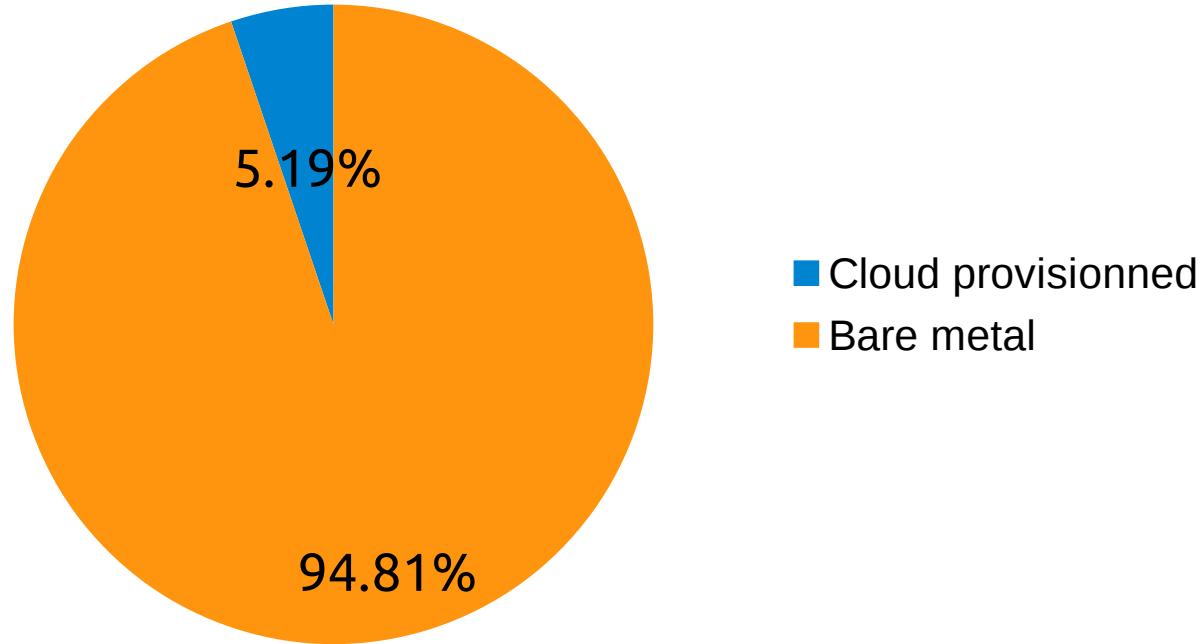
- ease software deployment and allows specific environment
- specific job management & computing model implementation

- low availability, private networking with outgoing internet connectivity, huge CPU cycle accesses (high latency & throughput)

Overall usage



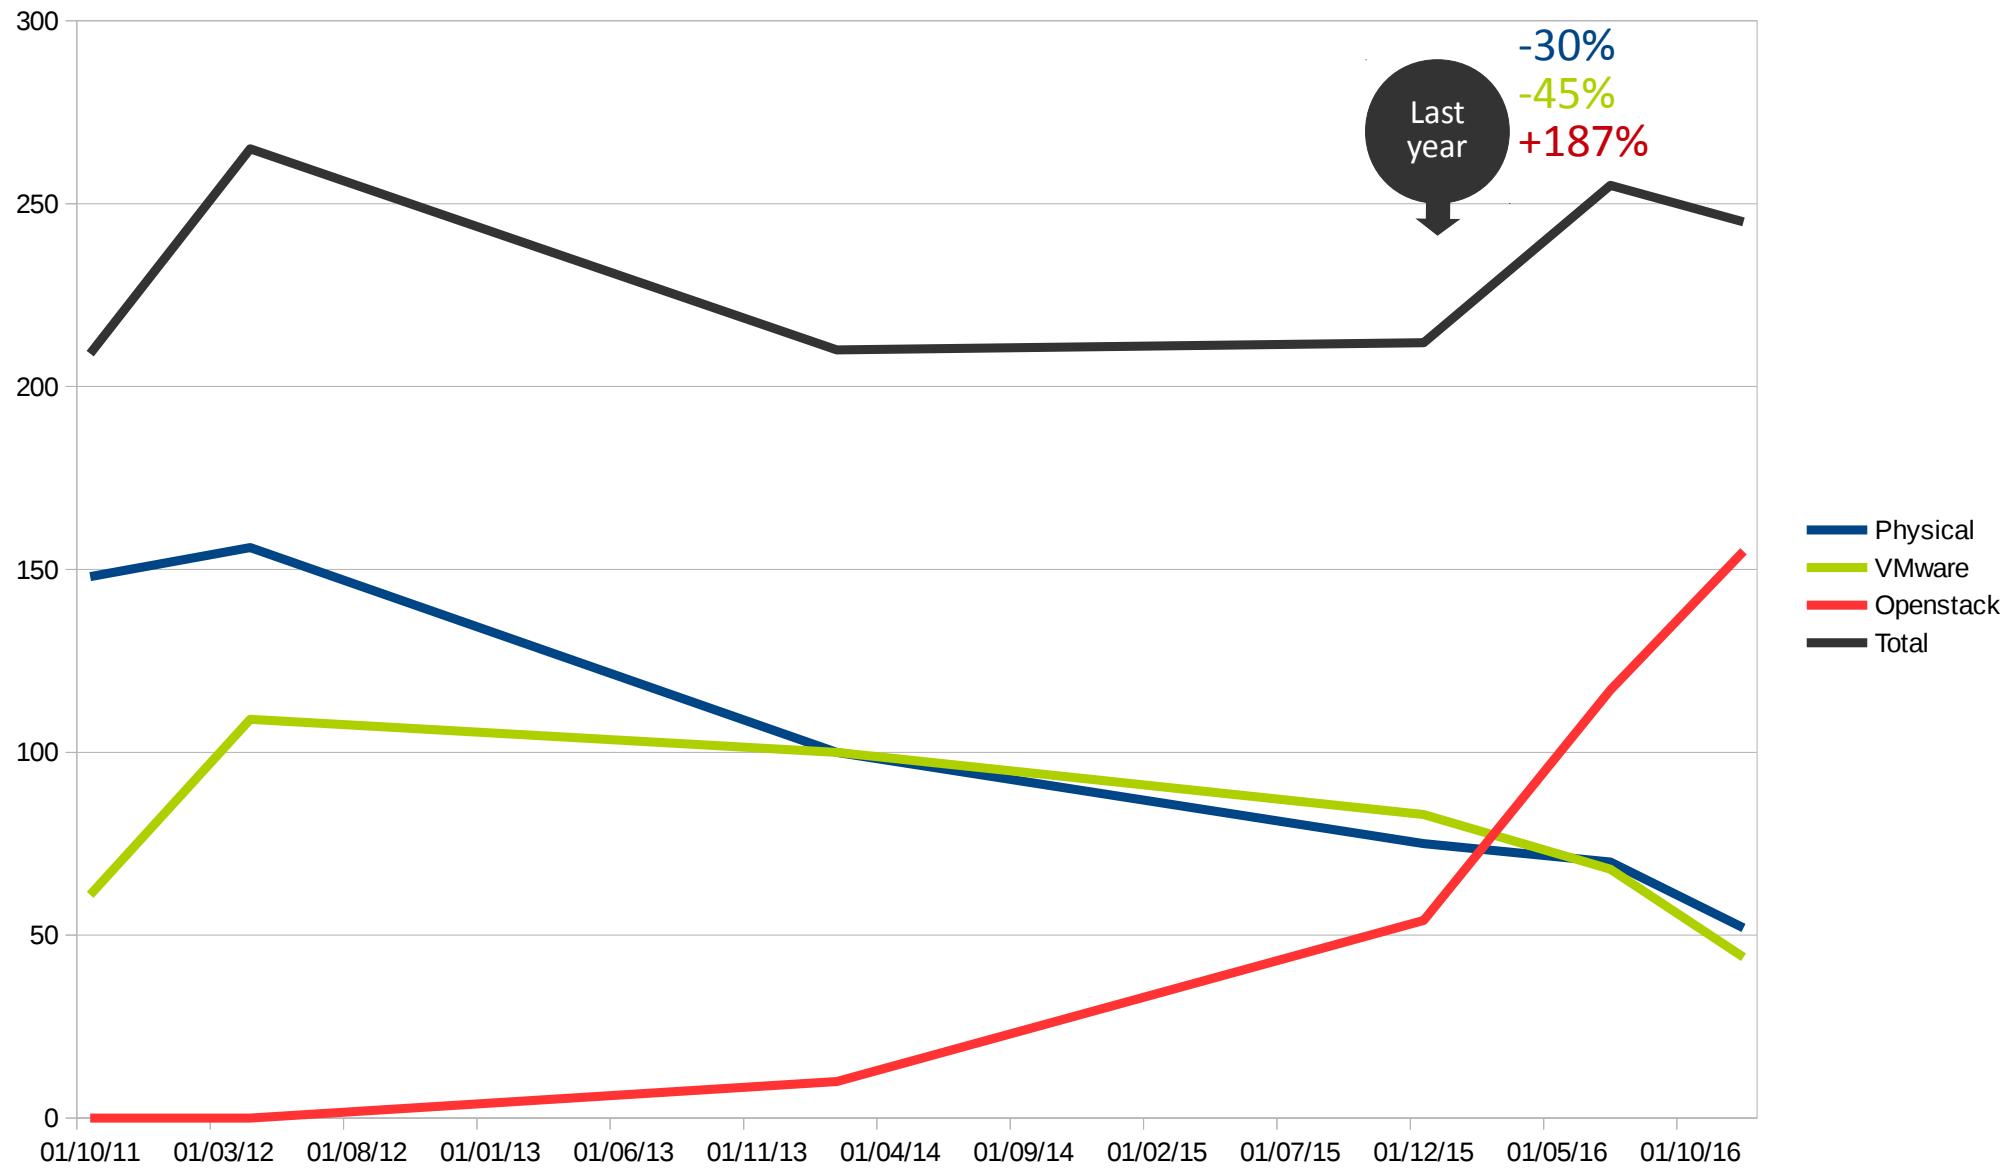
Projects : 50
Users : 150
Hosts : 80
Aggregates : 16



Computing resources :

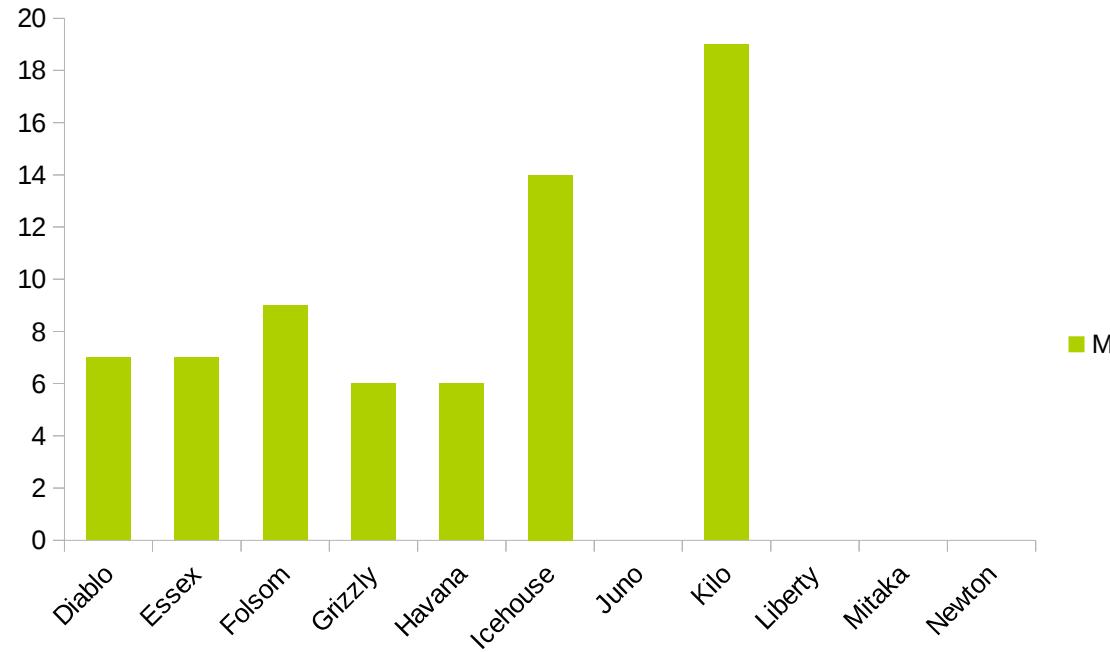
- ~1k HT cores in regard with 25k on HTC farm

Core services deployment (HA private cloud)

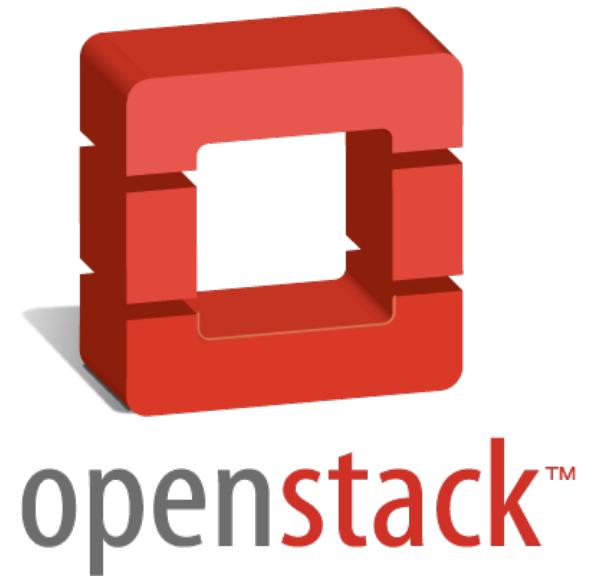


Some implementation details

Openstack deployments



■ Months



Operational components :

Keystone
Glance
Nova
Neutron
Horizon
Cinder
Ceilometer
Swift

In deployment/evaluation :

Heat
Magnum
Rally
Manila

Hardware resources (as of feb '17)

Deployment :

- CentOS 7
- RDO packages
- Puppet configuration

Compute clusters (may '16) :

	Cores	RAM	Storage
HA	544	3.6 TB	36 TB
R&D	288	1.2 TB	24 TB
Computing	1216	5.2 TB	40 TB

→ 2048 cores in total, 10TB RAM, 100 TB storage

S3 Storage (Swift):

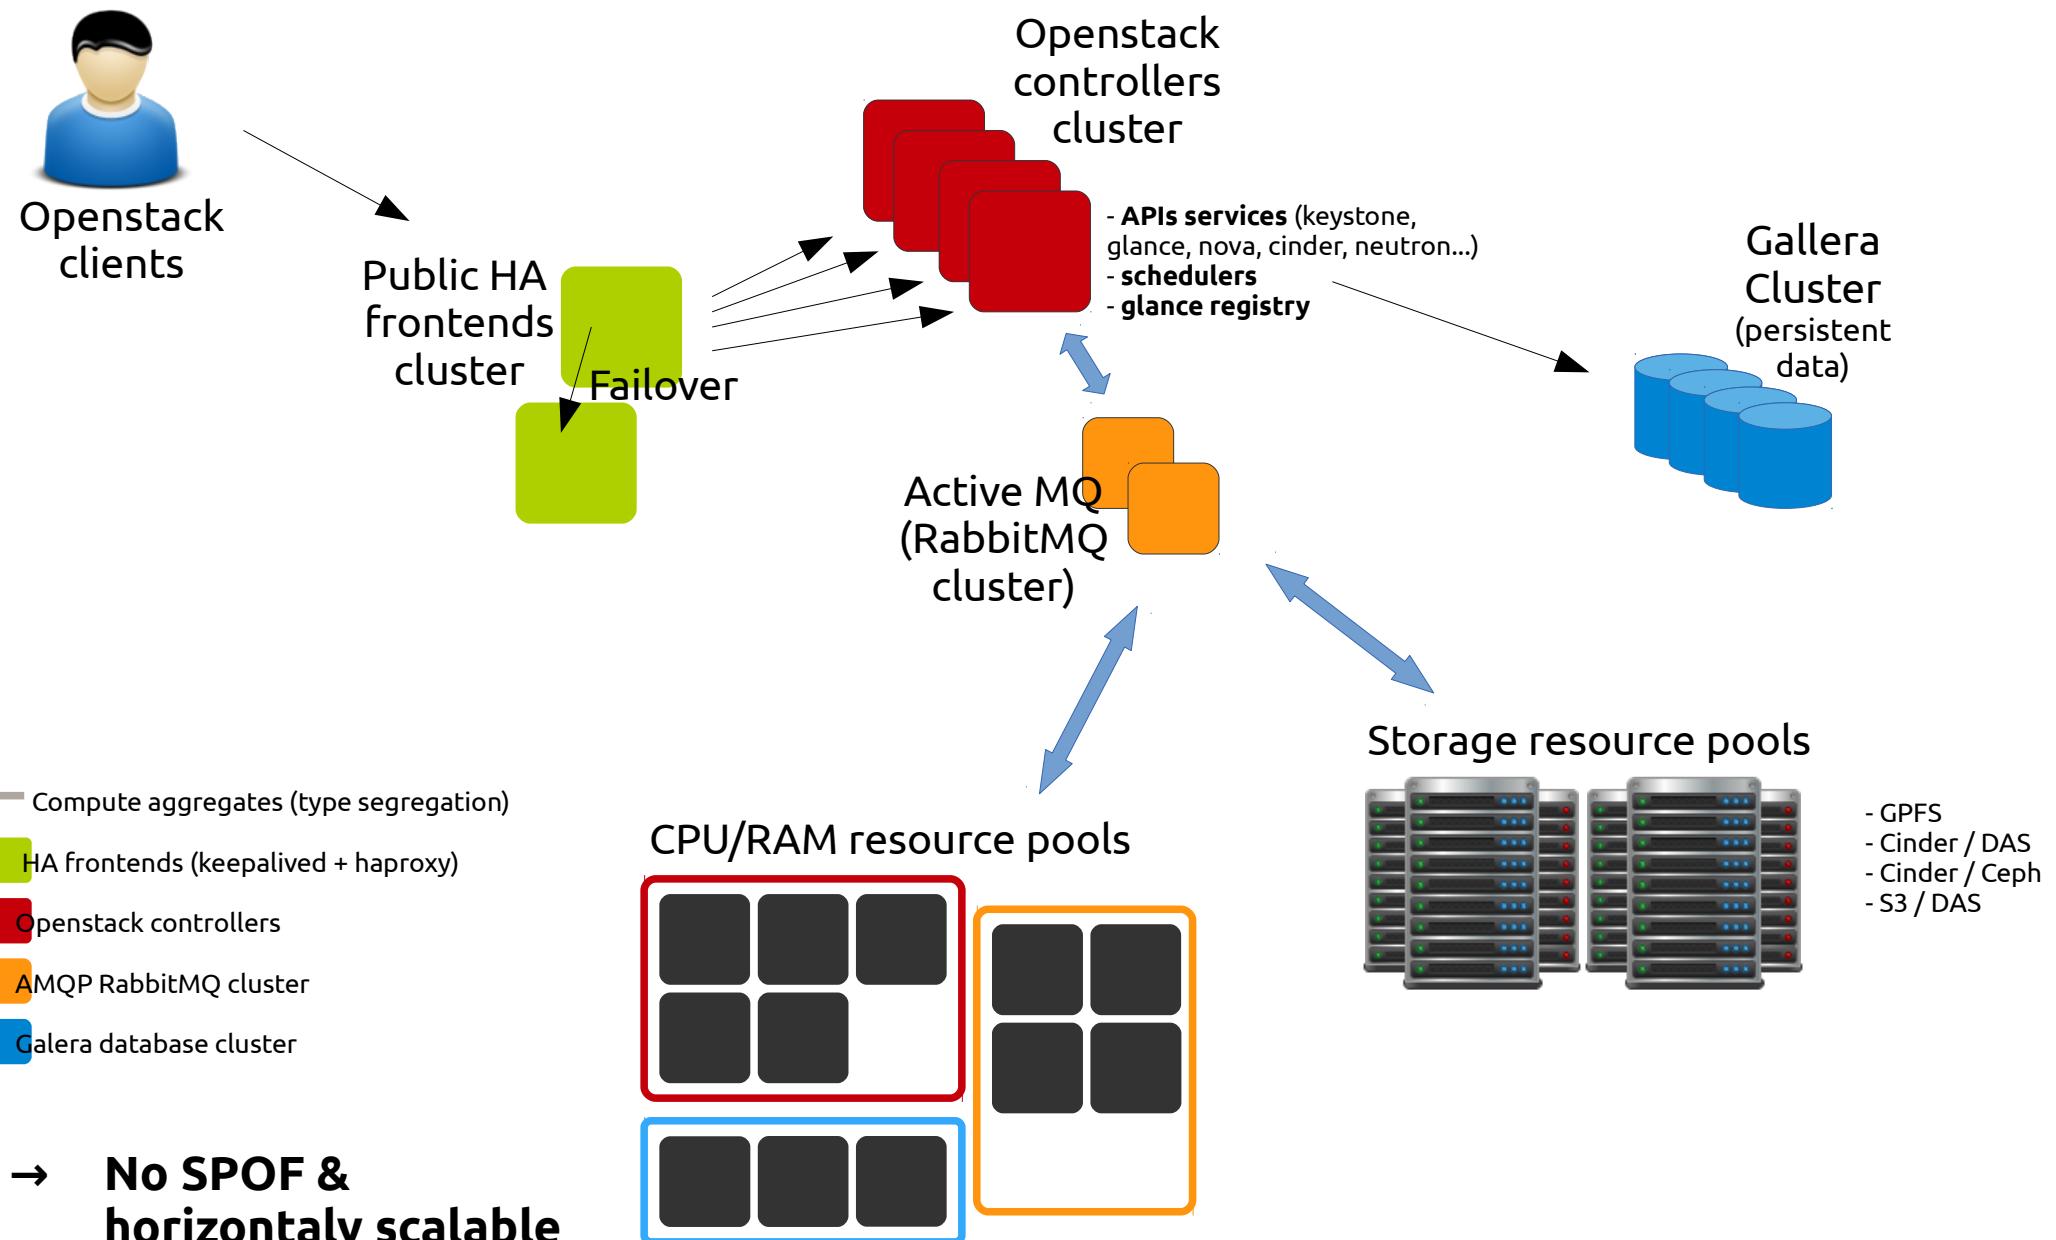
- 24TB DAS

Ceph Storage (Cinder):

- 480TB DAS



Service architecture



Neutron migration

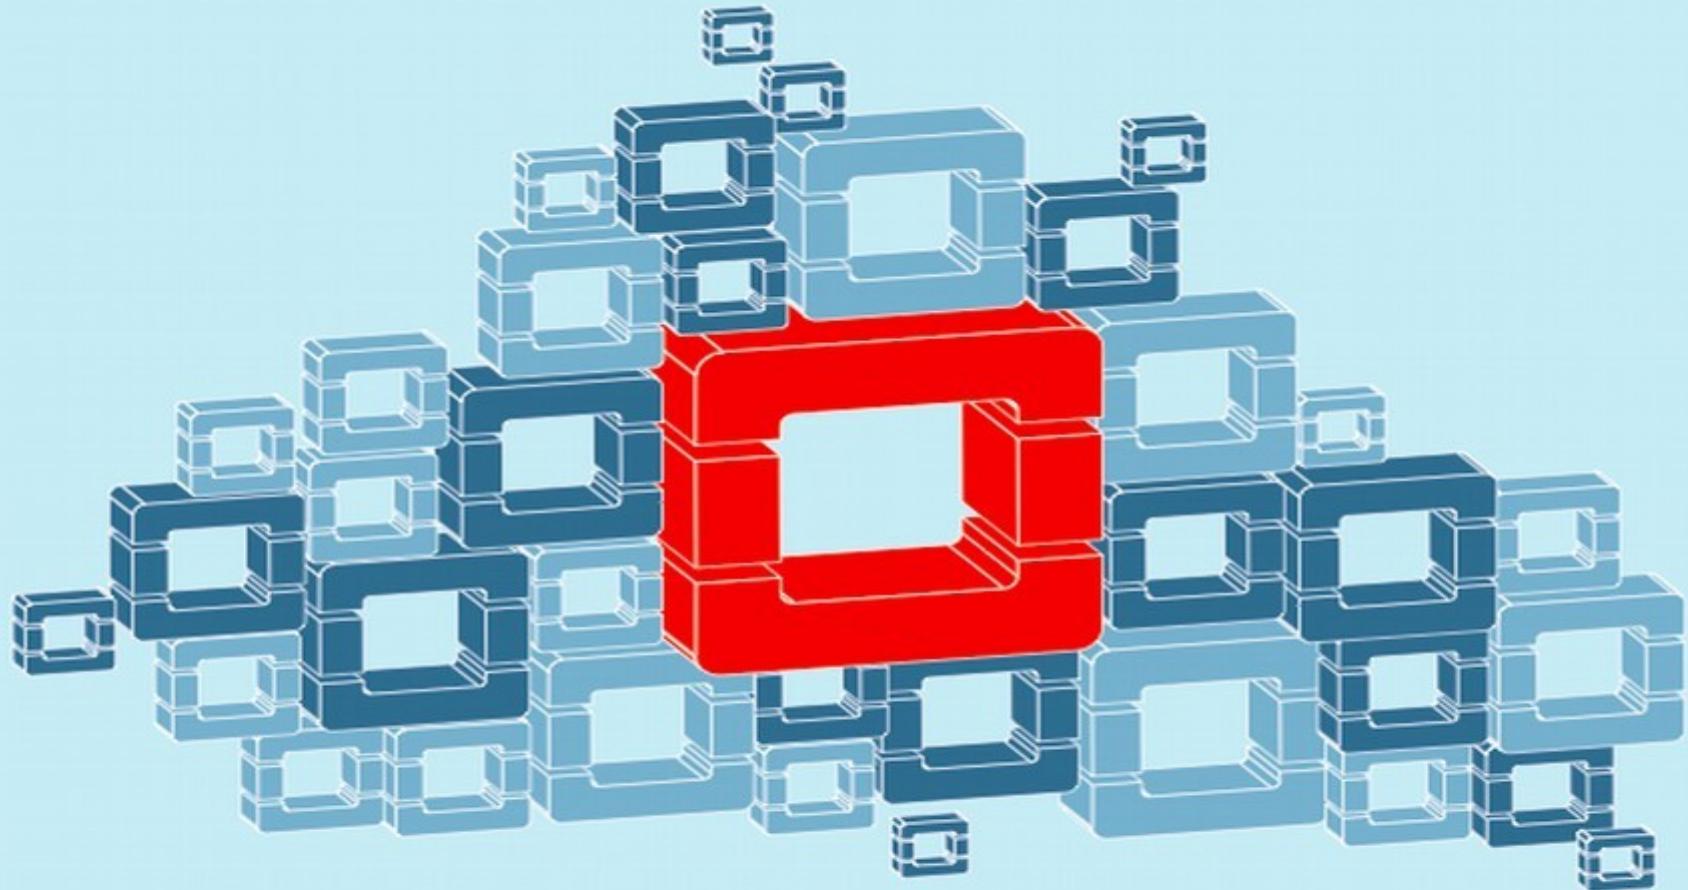
From legacy networking to Neutron, what's the challenge ?

- 50 different VLANs served (1 ↔ 1 projects)
- 6 different modes
 - Private isolated
 - Private internally routed
 - Private SNATED to the outside world
 - Floating IPs (SNAT/DNAT)
 - Public
 - Public distributed into existing networks
- Virtualized (non cloudish) production services implemented
- No supported migration process : you just rely on what's been done by others
- Linux bridges to OVS
- Static network provisionning to SDN
- HA/DVR

What's next ?

What's next in the mid-term

- Public cloud platform enhancements (storage & CPU) : targeted usage within 3 years : 300 instances, 1.4TB RAM, 800 vCPUs, 50 TB storage
- FSS with Synergy for computing
- Finalize resources orchestration (Heat)
- Container orchestration (Magnum/Kubernetes)
- Cinder refactoring with Ceph backend
- AAI integration
- Interest in shared FSaaS (Manila)



(thank you)
Questions ?