IN2P3-CC cloud computing (IAAS) status

FJPPL Computing Workshop

Mar 11th 2015 / Mattieu Puel

Outline

- Use cases
  - Testing, development and preproduction systems
  - Core services
  - Computing
  - Community cloud / hosting
- Implementation details
  - Deployment
  - Platform usage
- Impacts on IT management
- What's next

Centre de Calcul de l'Institut National de Physique Nucléaire et de Physique des Particules

Good for ?

Use cases : testing, development and preproduction systems

For whom : IN2P3-CC people to provision their testing and development systems

Goals

- Self service =
  - host lifecycle handled by end users
  - accurate sizing and environment specification
- Test systems are most of the time resources thrifty
  → efficient mutualization
- On top of it : puppet enables reproducability for the whole, moving from preproduction to production is handy

Use cases : core services

For whom : IN2P3-CC SA team provisionning core services

Goals

- Usual virtualization motivations
  - lowers hardware management
  - optimizes resources (mutualization)
  - soft servers sizing (cpu/ram/disk/net)
  - unbind servers from the hardware (higher availability, maintenance eased)

We formerly used VMware

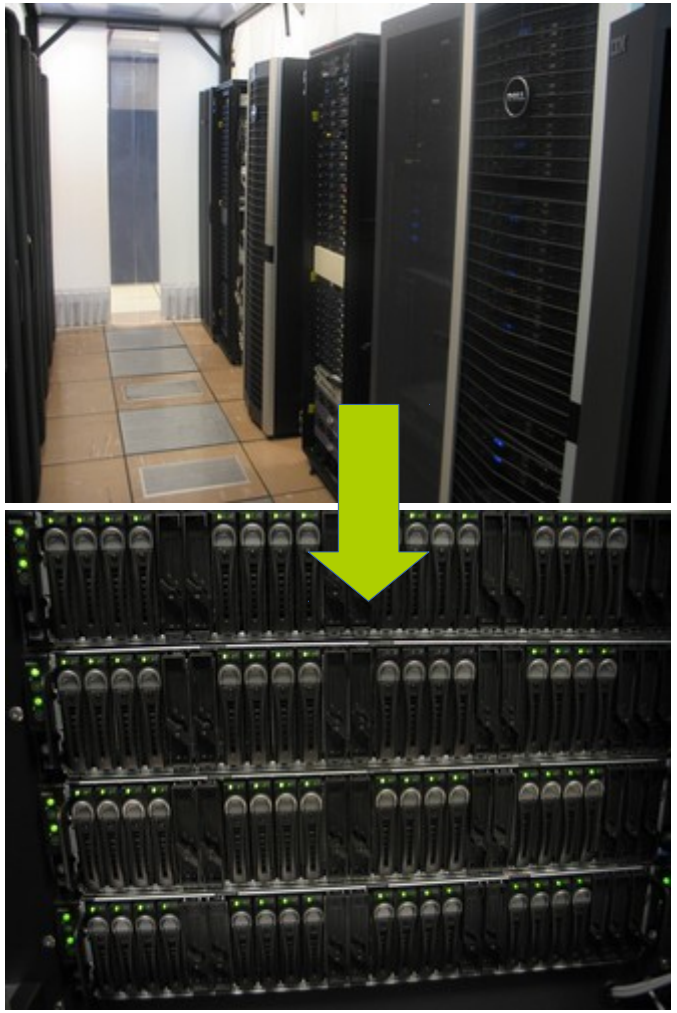

Use cases : computing

Users motivations :

- Gain access to opportunistic resources (fast access during activity bursts)
- Use specific environment (OS, softwares...)
- Implement their own scheduling
- Ease software deployment

Different models :

- The local batch system is aware of the cloud (statically, dynamically)
- Grid middleware leveraging cloud interfaces (Dirac...)
- Direct use of cloud interfaces (EC2/Nova...)

First productions since late '14:

- Atlas for MC jobs (opportunistic, simulation)

Currently evaluated cloud computing models :

- Large Synoptic Survey Telescop (http://www.lsst.org)
- Euclid (http://www.euclid-ec.org/)
- Bioaster

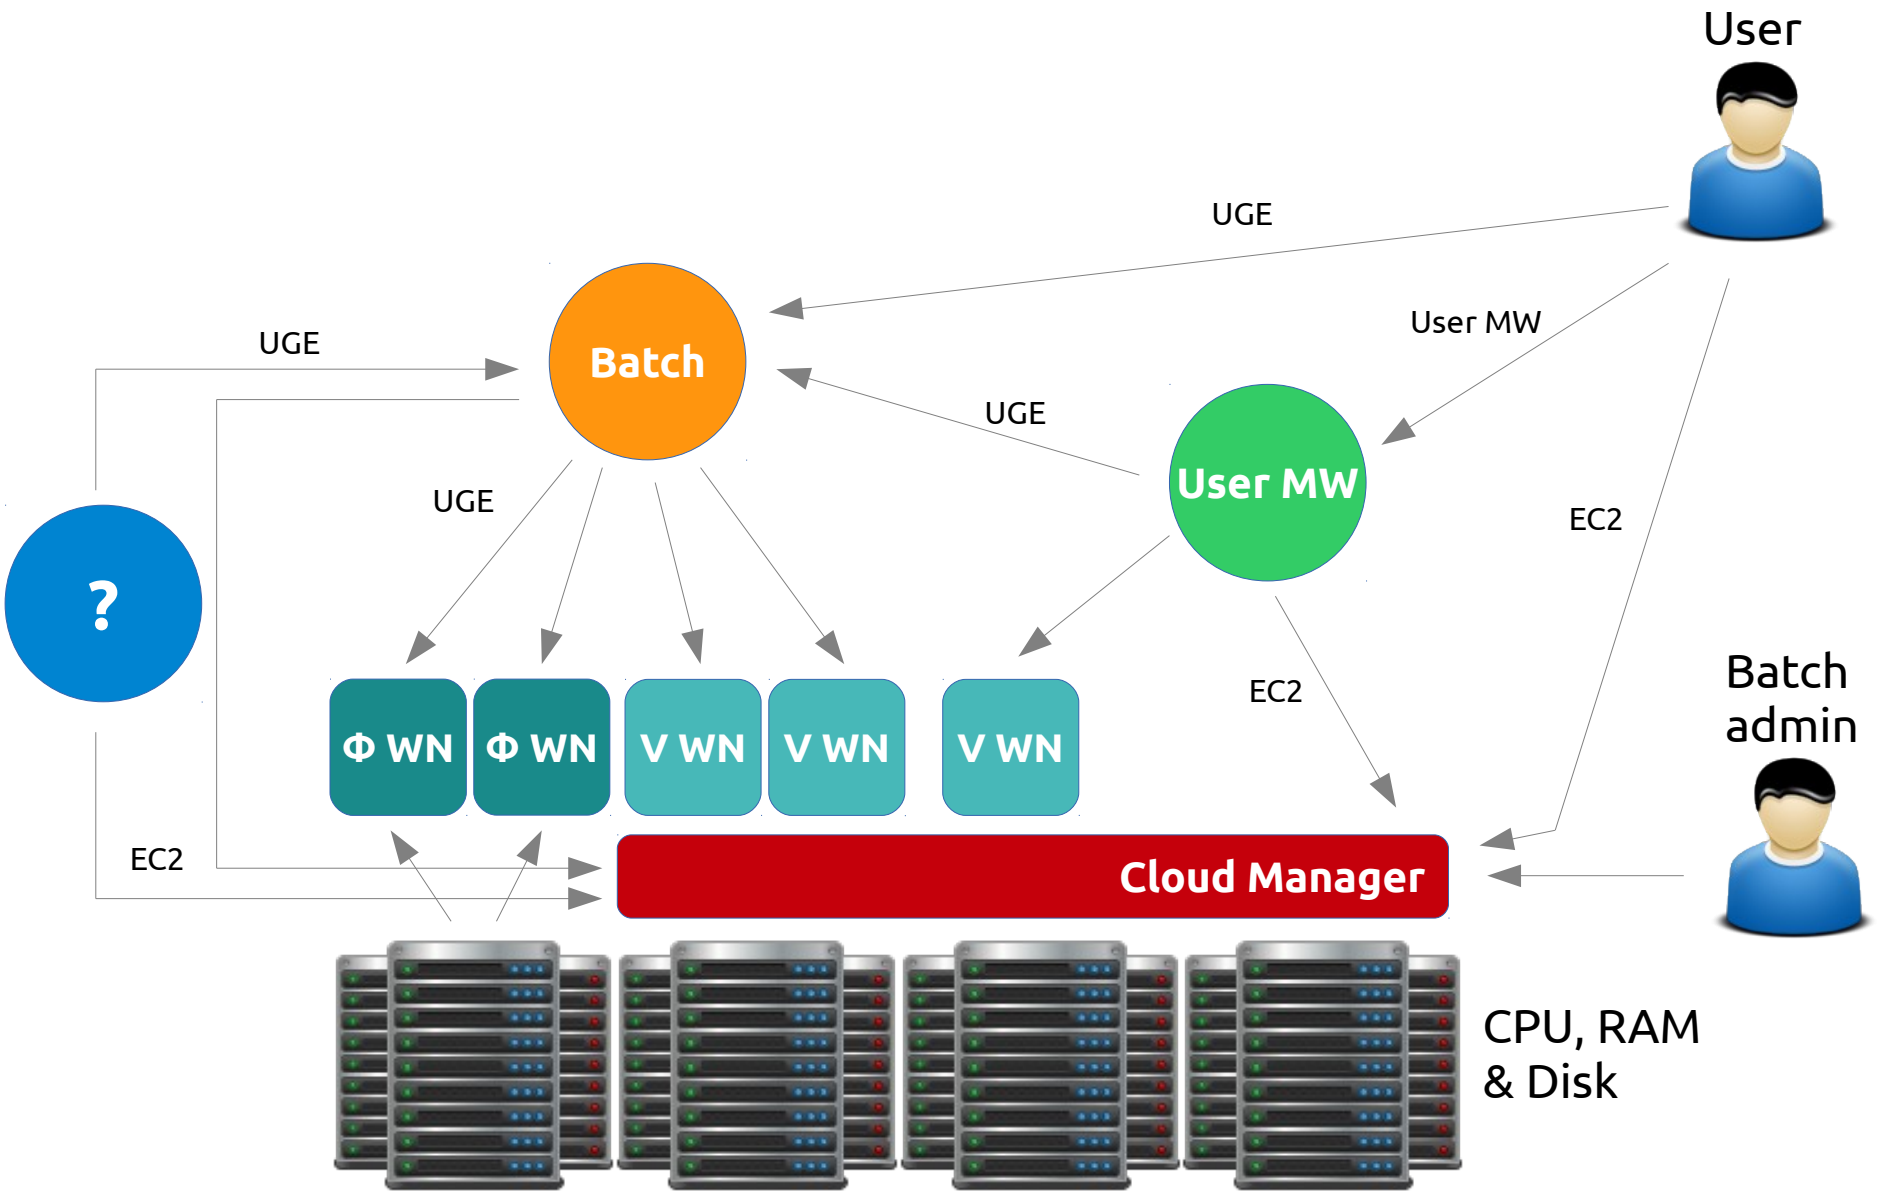
User
UGE
Batch
User MW
UGE
UGE
User MW
?
UGE
EC2
Φ WN
Φ WN
V WN
V WN
V WN
EC2
Batch admin
EC2
Cloud Manager
CPU, RAM & Disk

Use cases : community cloud, hosting

Community cloud for whom ?

- IN2P3-CC institutional users (labs and experiments)

Goals

- Help users to implement their services by themselves on top of an highly available infrastructure
- Centralize institutional resources → maximize mutualization, get bigger extra capacity
- Examples : Bioaster, AMI, eTRIKS

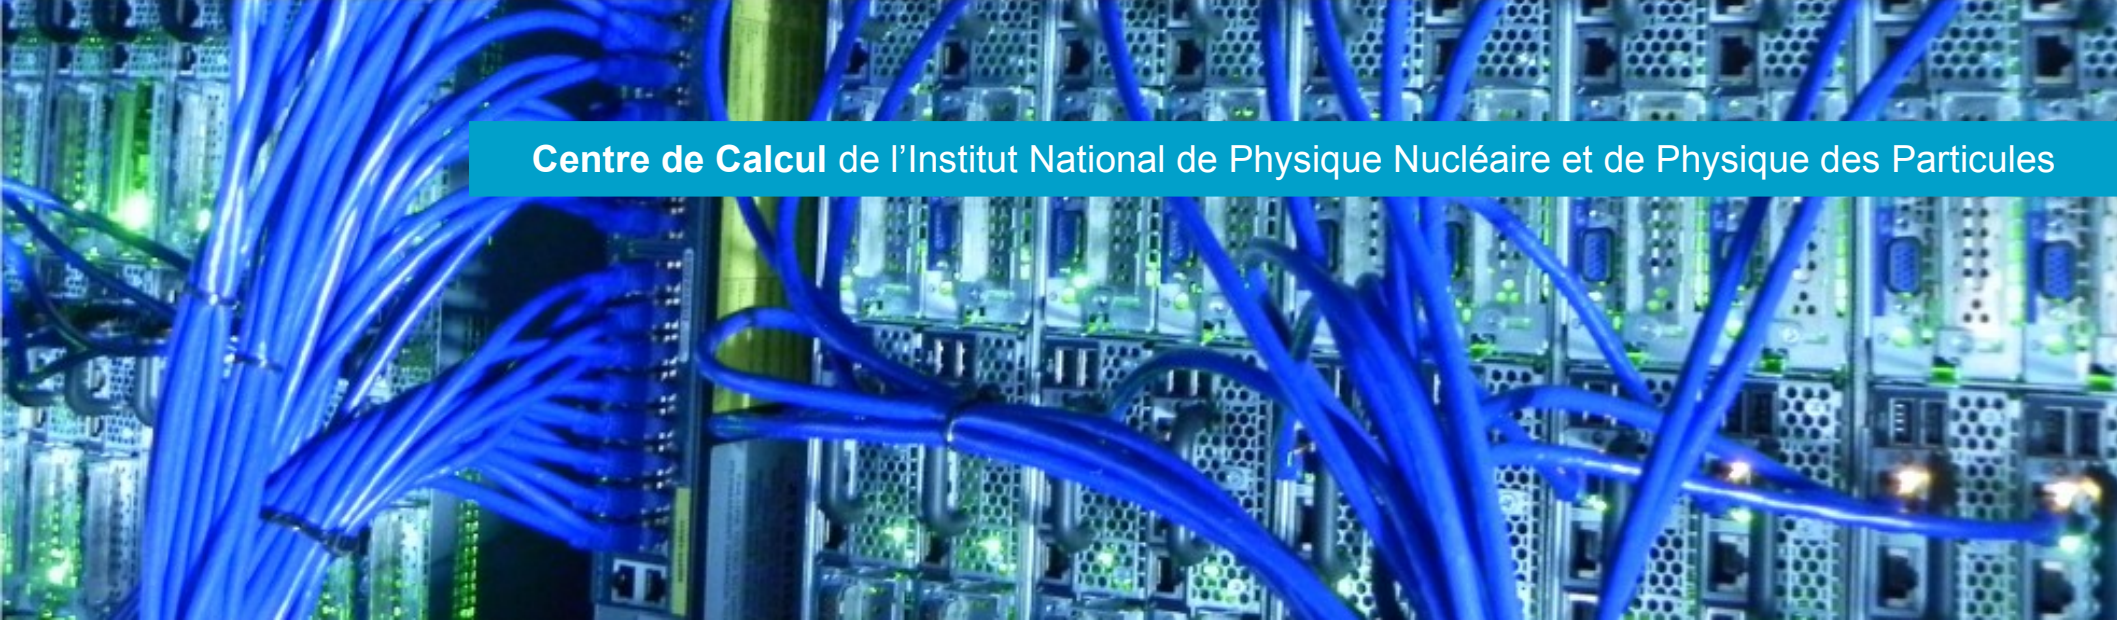

Some implementation details

Openstack deployments

Openstack components :

Operational :

Nova
Horizon
Cinder
Ceilometer
Swift

Still in evaluation :

Neutron

Openstack cluster

Deployment :

- Scientific Linux 6 (requirement for 7 in Kilo release)
- Griddynamics, then EPEL and now RDO
- Configured with Puppet

Resources :

Bunch of C6xx, R6xx, M6xx DELL Poweredge servers

Core services

- 150 CPUs
- 300GB RAM
- 5TB storage

Hosting

- 50 CPUs
- 200GB RAM
- 4TB storage

Computing

- 500 CPUs
- 3TB RAM
- 9TB storage

Preproduction/testing cluster :

- 300 CPUs
- 1TB RAM
- 28TB storage

+30 TB Cinder volumes

+24 TB Swift S3 storage

Grand total:

- 1k CPUs
- 4.5TB Memory
- 100TB storage

Some facts

Platform usage : 306 systems

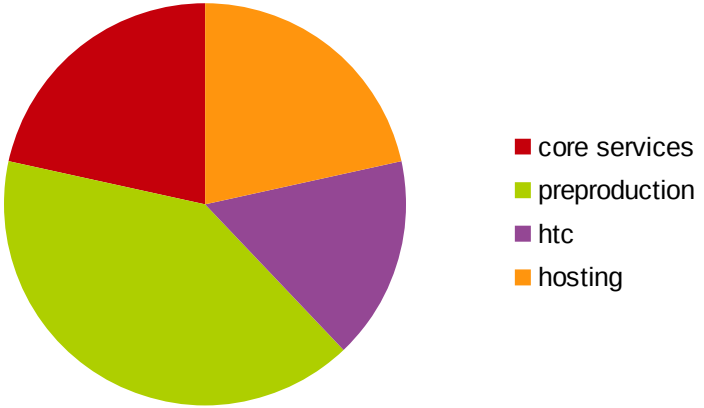

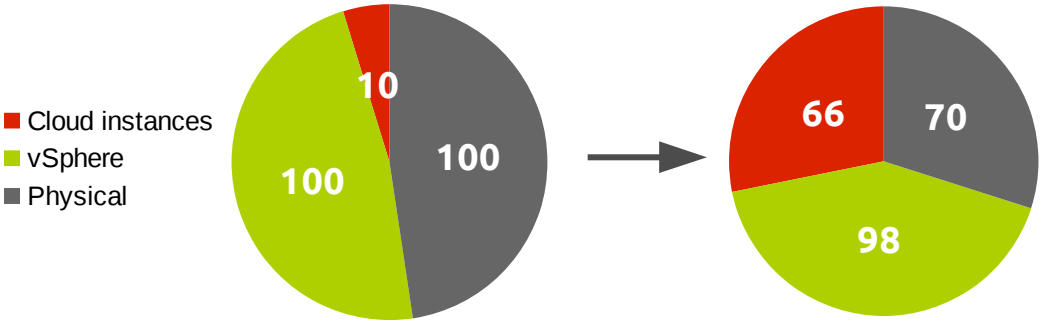

Apr '14
210 hosts

Today
244 hosts

Core services migration

Computing resources :

- ~500 HT cores in regard with 17k for HTC platform

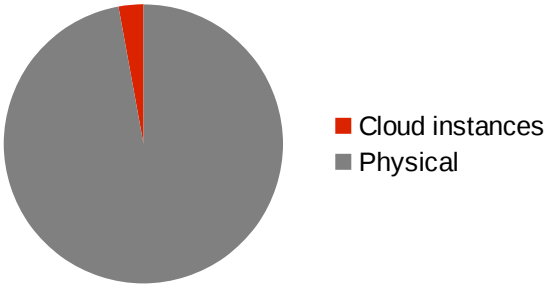

Impacts on IT organization

Impacts on IT organization

Hardware provisionning :

- Use commodity hardware : less vendors, OSes, configs...
- Capacity planning becomes a global consideration
- The more mutualization, the more optimized resources, the more money saved.

Roles shift :

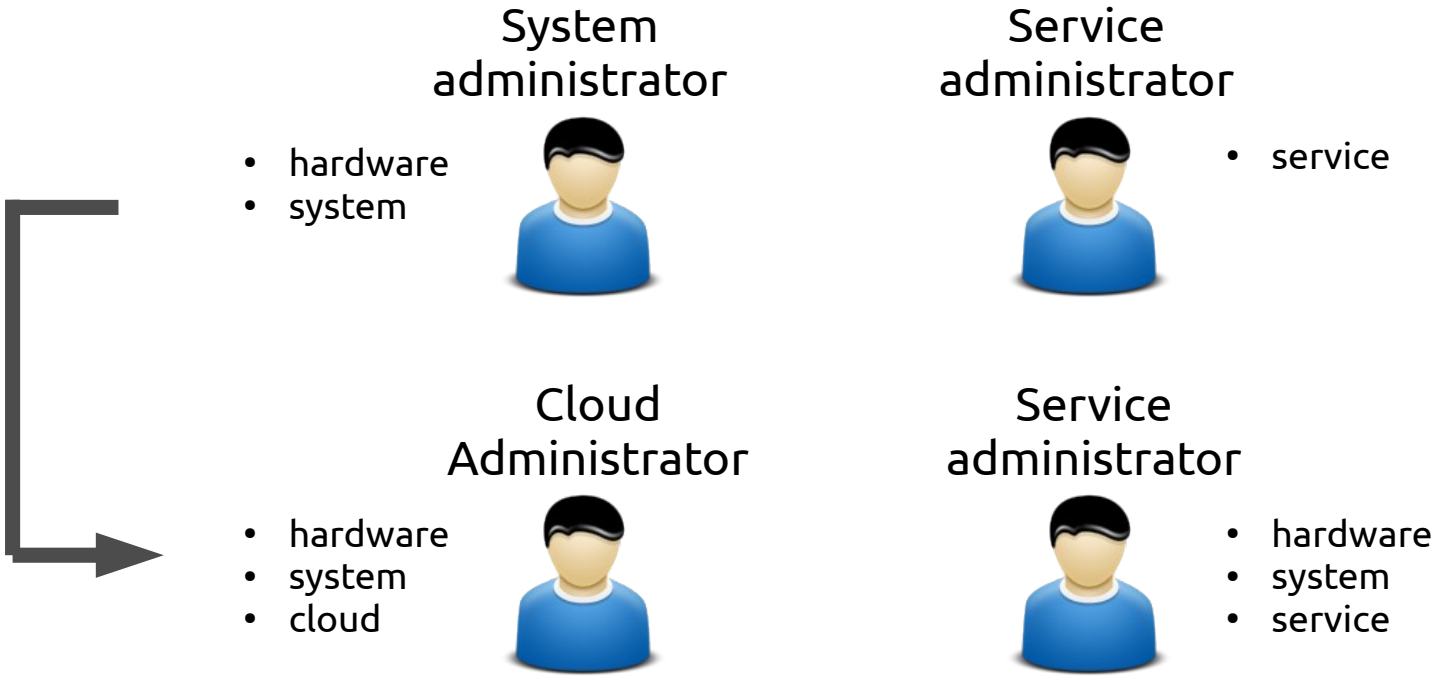

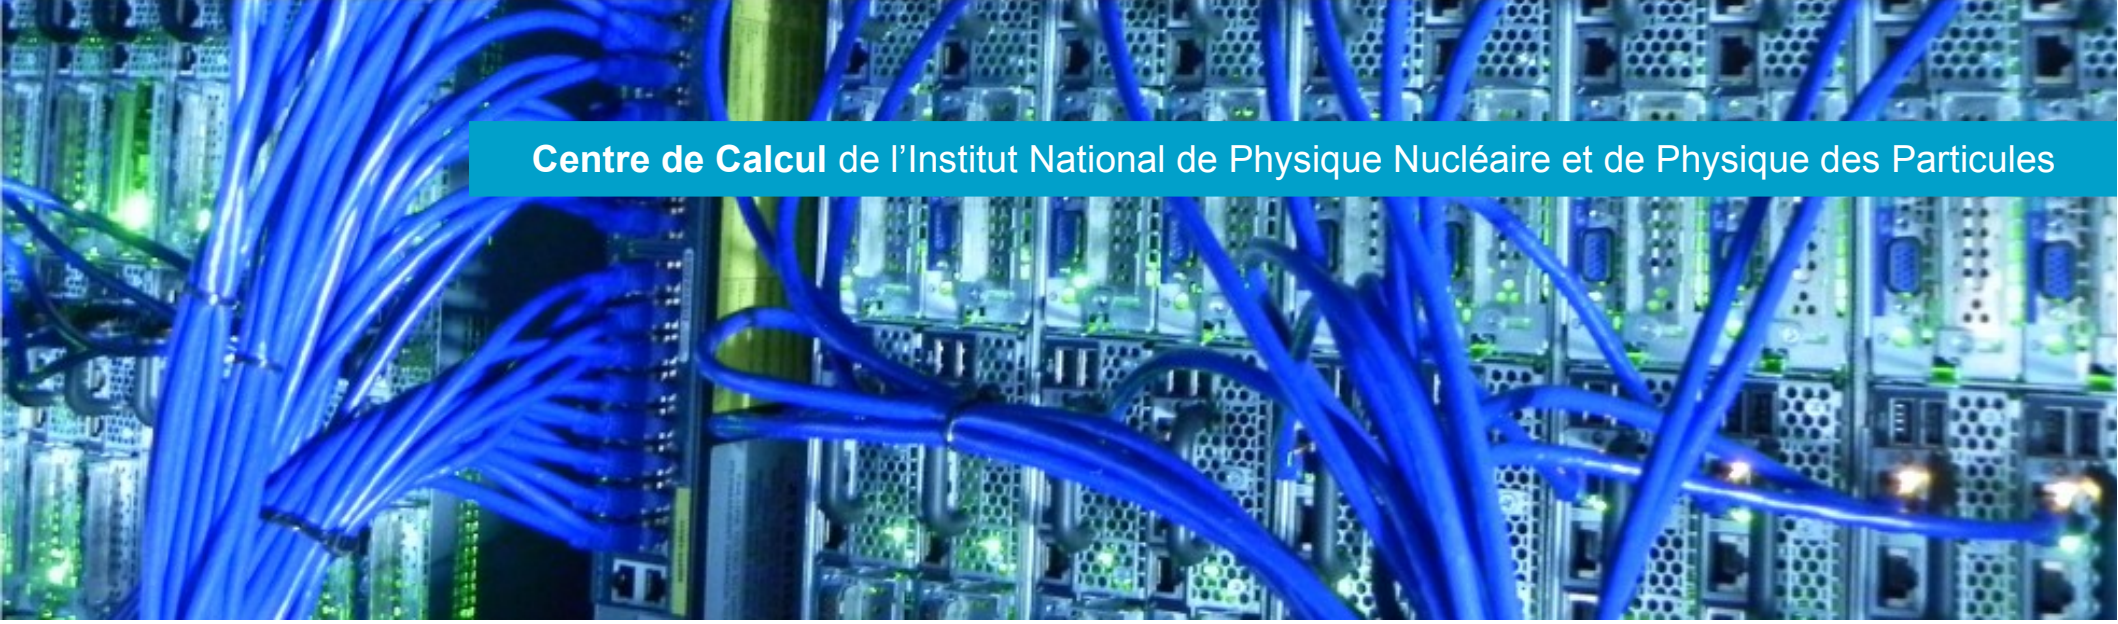
Centre de Calcul de l'Institut National de Physique Nucléaire et de Physique des Particules

What's next ?

What's next ?

Short term :

- End VMware migration
- Openstack Neutron (NaaS)

Longer term objectives :

- More PaaS offerings
- Need for lots more IPs, IPv6
- Production grade computing for diversified job classes

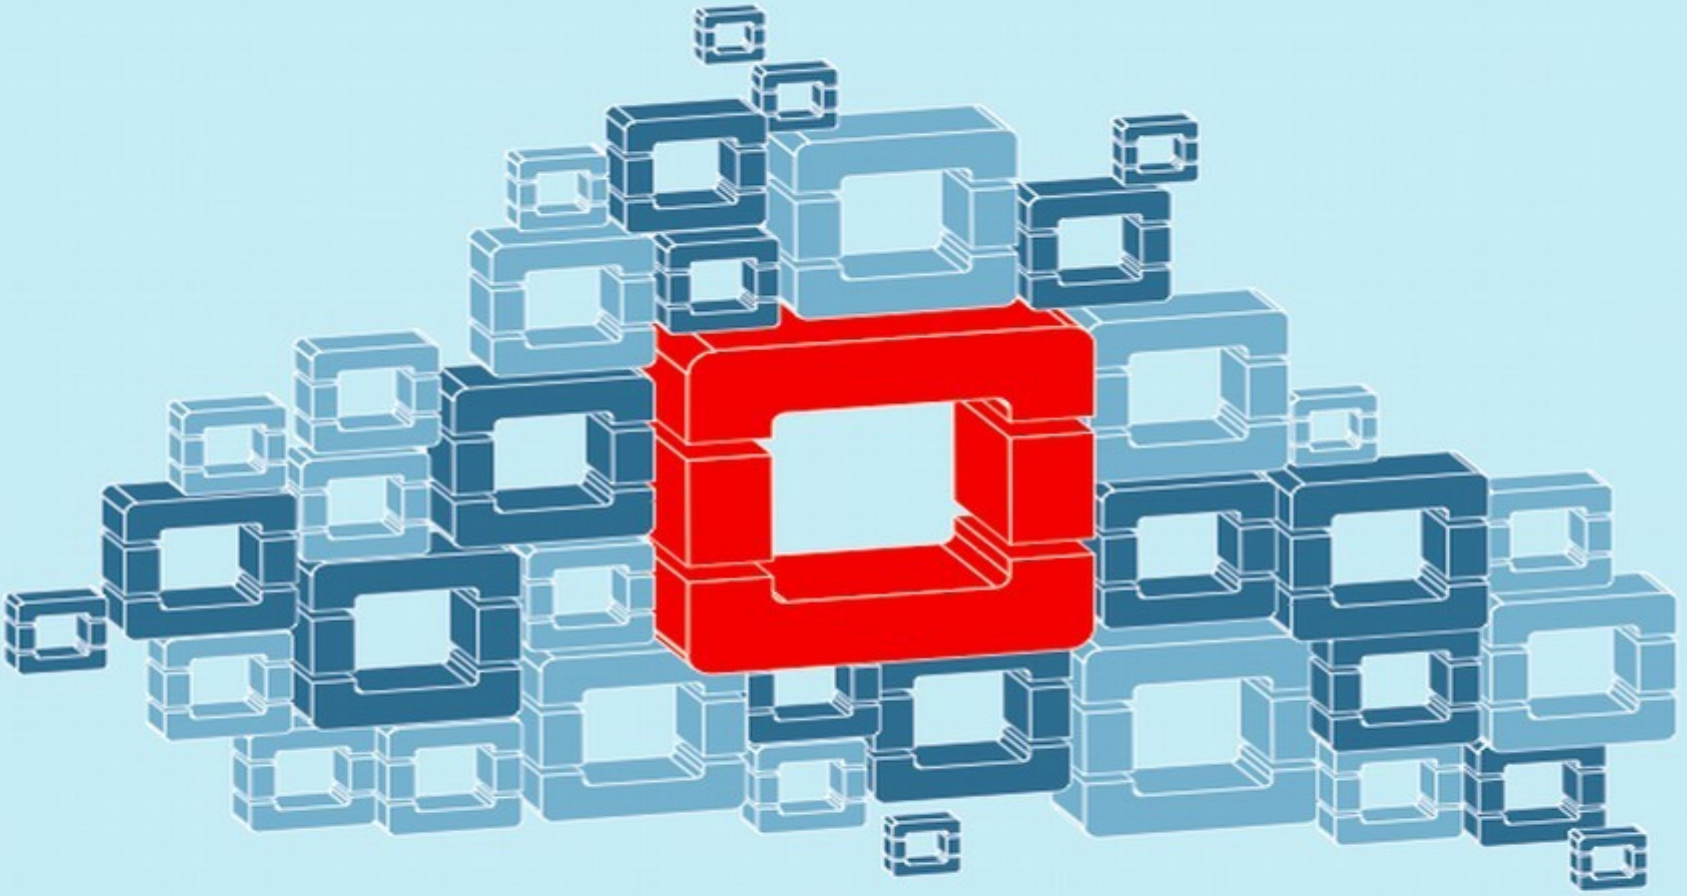

(thank you)

Questions ?