

IN2P3-CC cloud computing (IAAS) status

IHEP meeting

Jun 11th 2014

- Use cases
 - Testing systems
 - Infrastructure services
 - Computing
 - Community cloud
- Some implementation details
 - Deployment
 - Existing clusters and usage
 - Openstack pros and cons
- Impacts on infrastructure management
- What's next

IN2P3-CC IAAS cloud

Good for ?

For whom : IN2P3-CC people (~80) to provision testing and development systems

Why ?

- Self service =
 - host lifecycle directly handled by end users
 - accurate sizing and environment specification
- Test systems are most of the time resources thrifty
→ efficient mutualization


For whom : IN2P3-CC SA team provisioning infrastructure services

Why ?

- All usual virtualization motivations
 - less hardware to manage
 - mutualization → resources optimization
 - software servers sizing (cpu/ram/disk/net)
 - unbind servers from the hardware for higher availability
- Save VMware licences fee

Same level of service ?

- no DRS, no FT, no storage vMotion
- Sort of HA, live migration available



Resource provider motivations :

- Mutualize computing resources with the rest of the IT provisioning
→ use spare resources for computing
- Ensure compatibility with potential users shift

Users motivations :

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

Envisaged alternatives to compute on cloud resources :

- Integrate UGE and Openstack : Unicloud, Vcluster
- Grid middleware that leverages cloud interfaces (like Dirac)
- Direct use of cloud interfaces (EC2/Nova...)

Currently evaluated computing :

- Large Synoptic Survey Telescop (<http://www.lsst.org>)
- Euclid (<http://www.euclid-ec.org/>)
- Atlas with some MC simulation, HammerCloud validation
- Bioaster

Community cloud for whom ?

- IN2P3-CC institutional users (labs/experiments)

Why ?

- Centralize more institutional resources → maximize mutualization, get bigger extra capacity

IN2P3-CC IAAS cloud

Some implementation details

Openstack deployments



Openstack components :

Operational :

- Keystone
- Glance
- Nova
- Horizon
- Cinder
- Ceilometer

Tested :

- Swift

Soon to be :

- Neutron

Clusters :

Infrastructure services
test/dev machines
community cloud
(~400 cores)

Computing
(~500 cores)

One dedicated to eTRIKS project
(~100 cores)



Deployment :

- Scientific Linux 6
- Griddynamics, then EPEL and now RDO
- Configured with Puppet

Test/dev/service/community cluster :

- DELL C6100 and R610 compute nodes (400 cores for test/dev VMs, 100 cores for infrastructure services)
- DELL PE R720xd (30TB Cinder volumes ~ Amazon EBS)
- DELL PE R720xd and GPFS for instances shared storage

Computing cluster :

- 64 DELL M610 (compute nodes, recycled batch WNs)
- 10 ISILON nodes (instances storage)
- Capacity : ~500 VMs (1 core, 20GB disk, 6GB RAM)





Physical

Openstack VM

vSphere VM

Testing and developments systems :

- ~150 VMs (400 capacity) in regard with ~50 physical




: Infrastructure core services

- ~10 VMs (100 capacity) in regard with ~100 physical and ~100 vSphere virtualized

Computing resources :

- ~500 cores in regard with 17000 HT cores (no production yet)



- Nova-network VLAN mode :
 - do you have thousands of available public IPs ?
 - do you support IPv6 everywhere ?
 - don't you care about VLAN isolation ?
- Updates not trivial
- Sparse, not up to date documentation
- Complexity : a lot of features require specific advanced skills (system, networking, storage...)
- Still lacking comfortable features (like vSphere DRS)

- Reliable
- A lot of possible configurations and features
- Huge dynamic community :
 - high development rate
 - clear roadmaps, release cycle
 - bugs are fixed
 - security actively monitored
- Highly scalable and available
- Trivial to patch (and merge upstream)

Impact on datacenter organization

IAAS cloud leads to unified hardware provisioning :

- Use commodity hardware : less different vendors, OSes, configs...
- Capacity planning becomes a global consideration
- Unused resources may be allocated to the most greedy services. Let say computing for HEP sites.
- The more mutualization, the more optimized resources, the more money saved.

Roles evolution, to a certain extent :

- **System administrators** become hardware providers, cloud admins, base image manufacturers.
- **Services administrators** get some of SA former responsibility... for better or for worse.

Openstack (managed by SAs) takes over some other teams traditionnal responsibilities (let say networking/storage team):

- Lot of networking configuration is handled by Openstack (like IPs allocation...)
- Requires a lot of storage skills

Openstack (managed by SAs) takes over some other teams traditional responsibilities (let say networking/storage team):

- Lot of networking configuration is handled by Openstack (like IPs allocation...)
- Requires a lot of storage skills

→ running an Openstack cluster requires spreading the administration amongst teams

Openstack (managed by SAs) takes over some other teams traditional responsibilities (let say networking/storage team):

- Lot of networking configuration is handled by Openstack (like IPs allocation...)
- Requires a lot of storage skills

→ running an Openstack cluster requires spreading the administration amongst teams

... and never leave behind some basic diplomacy skills...


Status, what's next ?

Short term :

- Unicloud / Vcluster evaluation
- Migration to Openstack Neutron

Longer term objectives :

- Migrate everything virtualizable to compute cloud
- Need for lots more IPs, IPv6
- Base hosting offering on cloud technologies
- IN2P3 community cloud offering



(thank you)

Questions ?