



Cloud Computing @ CC-IN2P3

Rencontre LCG-France, SUBATECH Nantes

Alvaro Lopez Garcia
alopez@cc.in2p3.fr

Technical aspects

Virtualization and Cloud

Cloud Infrastructure at CC

Current deployment

Unified Infrastructure

Cloud Federation

Conclusions

Technical aspects

Virtualization and Cloud

Cloud Infrastructure at CC


Current deployment

Unified Infrastructure

Cloud Federation

Conclusions

Virtualization



- Not a new concept, although became a trend from 2000.
 - VMWare ESX 2001.
 - Xen 2003.
 - Qemu 2004.
 - KVM 2006.
- Isolation of multiple machines (guests) on a single host.
 - Different OS guests.
 - Performance penalty unless PVM or HVM.
- Tightly related with cloud computing, but not necessary.

Cloud computing



- Several different concepts are considered "Cloud":
 - Infrastructure as a Service (IaaS).
 - Platform as a Service (PaaS).
 - Software as a Service (SaaS).
 - Network as a Service (NaaS).
 - Whatever as a Service.
- At the CC-IN2P3 we are focusing on IaaS.
 - Provision of machines, leveraging the direct interaction with the hardware.
 - Network control: IP allocation, firewall modification.
 - Tight control of the policies and quotas by admins (network ACLs, CPU, disk, etc)
- Trend in the last few years, lots of projects and products.
 - Commercial products: Amazon, Rackspace, GoGrid, etc.
 - Open Source: Nimbus, Eucalyptus, OpenNebula, Openstack, etc.

Started by NASA and Rackspace.

- Governance by the Openstack Foundation.
- Contributions from HP, SuSE, Canonical, Dell, etc.
- Two different projects: Compute (nova) and Storage (swift).

Solution chosen at CC-IN2P3.


- Open Source, development driven by the community.
- Good development pace.
- Extensive features
- Completely modular.
- Cost.
- Deployment.
- Cloud momentum.
- Openstack is not bounded to an hypervisor
 - Support for KVM, Xen, XenServer, bare metal, LXC, etc.



Openstack nova Architecture



- Openstack Compute (nova) Services
 - API (nova-api).
 - Network manager (nova-network, quantum).
 - Image manager (glance).
 - Identity service (keystone).
 - Dashboard (horizon).
 - Scheduler (nova-scheduler).
 - Volume Manager (nova-volume, cinder).
 - Compute (nova-compute).



Technical aspects

Virtualization and Cloud

Cloud Infrastructure at CC


Current deployment

Unified Infrastructure

Cloud Federation

Conclusions

Hardware



- 16 DELL Poweredge C6100 hosts
 - 2 Xeon 24 Cores X5675 @ 3.07 GHz.
 - 96 GB RAM
 - 2 TB RAID 10 local storage (4 SAS 7.2 krpm)
- Total of 400 cores
- 10Gb NICS with NPAR/SR-IOv
- GPFS Image catalog.
- Private and public IPv4 subnets and VLANs. Full VMs network isolation from CC network.

Software

- Openstack Essex (Folsom will be deployed soon).
- KVM hypervisor.
- Openstack manager. Hosts the following services:
 - API.
 - Network manager.
 - Image catalog (glance).
 - Dashboard (horizon).
 - Identity (keystone).
 - Scheduler.
- Volume manager is on the road.
- Networking: Isolated VLANs per group.

Interfaces



HTTP/REST: Use it programmatically.

- Amazon EC2 (boto).
- Native OS API (novaclient).
- OCCl.

Command Line.

- euca2ools (EC2).
- nova (OS API).

Web Interface.

- Openstack Dashboard

Overview

Select a month to query its usage:

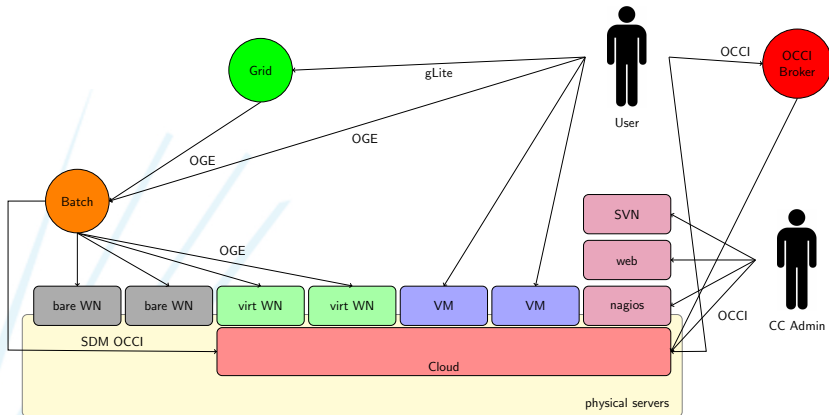
Active Instances: 4 Active Memory: 8GB This Month's VCPU Hours: 274.77 This Month's GB Hours: 8243.20

Usage Summary

Instance Name	VCPUs	Disk	RAM	Uptime
openmaster	1	30	2GB	2 semaines, 6 jours
openw1	1	30	2GB	2 semaines, 6 jours
28c44a1b-631a-4ef4-921e-087149588e42	1	30	2GB	2 semaines, 3 jours
8a790711-62e1-4728-42d9-48b042455c0	1	30	2GB	2 semaines, 3 jours

Deploying 4 items

Unified Infrastrucutre



Plans to extend the cloud testbed and deploy the CC-IN2P3 services on top.

- More flexibility and elasticity.
 - Scale up/down services dynamically according to the instantaneous needs.
 - Dynamic reassignment of resources.
- Zero-time overhead for deployment of new machines.
- Migration of machines.
 - Efficient resource utilization.
 - Hardware interventions without service disruption.
- Added value for users.
 - Instantiation of machines other than Scientific Linux.
 - Use exclusively the resources without sharing.

Technical aspects

Virtualization and Cloud

Cloud Infrastructure at CC

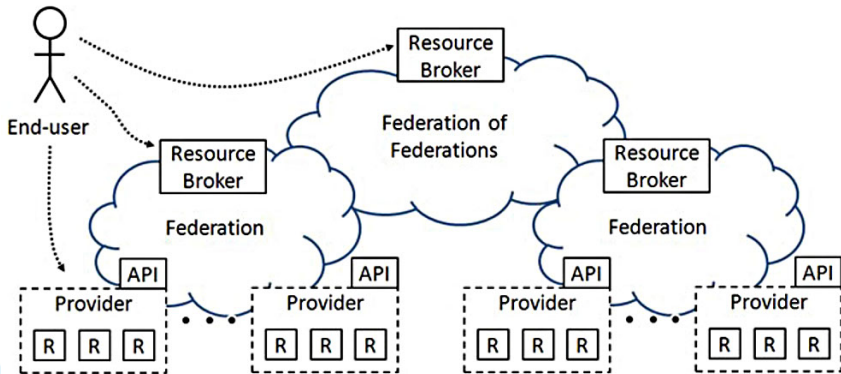
Current deployment

Unified Infrastructure

Cloud Federation

Conclusions

Cloud Federation Outlook



Federation efforts aim to unify the access to the many "scientific clouds" around France and Europe, trying to profit from Grid experience.

- EGI Federated Cloud Task Force.
- France Grilles.
- Collaboration CNRS and CSIC (Spain) to establish a federated cloud.

Different types:

- Strong federation: access through a broker (a-la-grid).
- Loose federation: common services (accounting, authentication, etc.), but no brokering.

A full federation implies some problems.

- Different resource providers, middleware implementations (Opennebula, Openstack, etc.) and hypervisors (Xen, KVM).
- Different APIs with different functionalities.
- Resource advertising of the capabilities.
- Federated authN/authZ.
- Placement policies (data access, image availability, price, energy saving).
- Image catalog, image distribution and overhead.
- Accounting of resources.

Ongoing work in EGI FCTF and the CC-IN2P3.

- Resources advertisement.
- Unified and standard API.
- Federated authN/authZ.
- Placement policies (data access, prices, energy saving...).
- Comprehensive images metadata.
- Aggregated accounting.

Ongoing work in EGI FCTF and the CC-IN2P3.

- Resources advertisement.
 - FCTF: *BDII*.
- Unified and standard API.
 - FCTF: *OCCI*.
- Federated authN/authZ.
 - FCTF: *x509 authentication*. CC-IN2P3: *VOMS authentication*.
- Placement policies (data access, prices, energy saving...).
 - CC-IN2P3 *Compatible One*.
- Comprehensive images metadata.
 - FCTF: *Stratuslab marketplace*.
- Aggregated accounting.
 - FCTF: *APEL*.

Technical aspects

Virtualization and Cloud

Cloud Infrastructure at CC

Current deployment

Unified Infrastructure

Cloud Federation

Conclusions

Evolution for the users



- Self Provisioning of nodes: deployment time decreased to the minute.
- He is able to use a custom environment (system, tools, batch...) and OS not an imposed one (Scientific Linux).
- Software porting from site to site becomes trivial.
- He distributes allocated resources as he wants (processor/memory/disk).
- He sizes his infrastructure dynamically according to his needs (elasticity).
- He may enlarge his infrastructure through public clouds.
- He may gain access to new resources by the way of a simple quota modification.
- With Federation in place he may distribute his resources in several datacenters (failover).

▶ Evolution for the computing center



- Reduced maintenance/administration of user environments.
- Benefits of virtualization: optimized resources.
- May partition and scale the infrastructure dynamically (systems, platforms, users...).
- May provide ressources to IaaS brokers (new revenues).
- Plan and provision extra capacity for only one platform anymore.
- Broaden users communities.

 **This is the end**



Merci!