Cloud platform at CC-IN2P3

May 28th 2015, Julien Carpentier
On behalf Mattieu Puel, Sysadmin team leader

julien.carpentier@cc.in2p3.fr

Outline

- Openstack IAAS Cluster
- Use Cases
  - Testing & preproduction
  - Core services
  - Computing
- Roadmap
- Hosted Projects

Openstack IAAS 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

Preprod cluster :

- 300 CPUs
- 1TB RAM
- 28TB storage

Grand total:

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

+30 TB Cinder volumes

+24 TB Swift S3 storage

TESTING & PREPRODUCTION

For whom : CCIN2P3 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

DEPLOYING CORE SERVICES

For whom : CCIN2P3 Sysadmin team provisionning core services

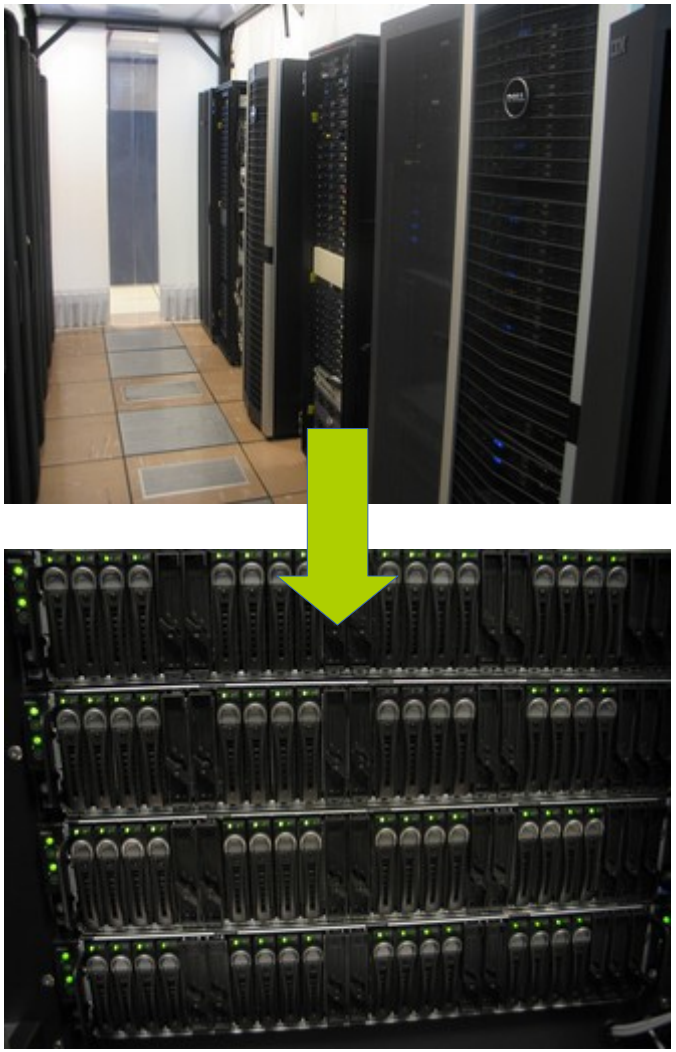

Goals

- 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)

PROS

- Desired sizing of the computing nodes (memory, cpu etc...)
- Control over the execution environment (OS and softwares)
- Isolation
- Hands on scheduling (jobs)

CONS

- More complex setup and operation from the user perspective
- Performance overhead is to be accounted somehow
- Still some technical barriers for high IO jobs

Use cases (Computing)

How ?

- Leveraging cloud enabled jobs framework (Dirac, HTCondor...)
- Using pilot VMs frameworks (Vcycle...)
- Static method : request an amount of specific environment to be spawned and integrated as worker nodes in the batch system
- Reimplemeting a batch system on top of cloud resources

VALIDATED COMPUTING MODELS

Atlas

- First productions since late '14
- WLCG Tier 3 defined on those resources
- MC opportunistic production
- Using HTCondor to interact with the cloud interface

LHCb

- Validated simulation computing in mar '15
- Desire to pledge resources
- Using Vcycle for VM lifecycle management

Under evaluation

- Euclid
- Bioaster

Roadmap

Openstack components :

Operational :

Nova
Horizon
Cinder
Ceilometer
Swift

Still in evaluation :

Neutron

Incoming improvment :

High-Available DB Backend
with MariaDB & Galera cluster

ETRIKS European Project

Budget

23.79M € for 5 years (Oct 2012 to Sept 2017)

Members

10 Pharma, 3 Academics, 1 standards, 2 Commercial suppliers

Goal :

- Provision of a KM Service to support Private/Public Translational Research
- Common platform and single access point for curation and computing

Hardware & Software :

- TranSMART : Java Web Application (KM Platform)

- Servers : Openstack dedicated nodes & DB servers
  → 208 vCPUs / 1.1 TB RAM / 130 TB Block Storage / 114 TB DB Storage

- R : Statistical computing


감사합니다
Natick
Grazie
Danke
Ευχαριστίες
Dalu
Obrigada
Thank You
Köszönöm
Tack
Спасибо
Dank
Gracias
谢谢
Merci
Seé
ありがとう

Questions ?