



Towards a French Academic Cloud

Frédéric Suter

IN2P3 Computing Center



lrfu
cea
saclay

April 26, 2012



Past

- ▶ Then, there was the Grid
- ▶ Some personal assessments

Present

- ▶ Now, there is the Cloud
- ▶ Several views of the same concept

Future

- ▶ A French Academic Cloud
- ▶ First Step and Next Move

10 years ago, there was the grid ...

- ▶ A wonderful promise!
 - ▶ Access to an unprecedented amount of computing and storage resources by **just plugging** your laptop
 - ▶ Analogy with the **power grid**




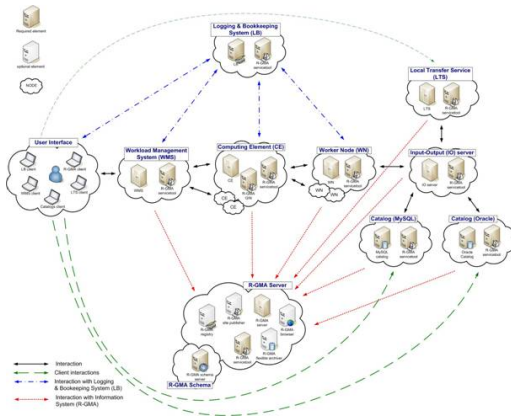
- ▶ Also a crucial need
 - ▶ LHC expected to produce Peta-Bytes of data per year

The (European) Grid Key Points



- ▶ Users
 - ▶ Mostly from the **HEP** community
 - ▶ Also from the **Biomed** community
- ▶ Jobs
 - ▶ Mostly **long and sequential**
 - ▶ Tightly linked to (huge) data
- ▶ Technology
 - ▶ Globus first ...
 - ▶ then **gLite**
 - ▶ Scientific Linux
- ▶ Organization
 - ▶ **Centralized** (around the CERN)
 - ▶ Funded through DataGrid, E-GEE I, II, and III

gLite Deployment and Job Life Cycle



▶ Assessment^(*) after 10 years – Successes



- ▶ The distributed infrastructure is **up and running**
 - ▶ At a European scale
 - ▶ With a 90+% **availability**
- ▶ Tons of jobs running, with tons of data stored and moved!
- ▶ VOs are well **structured**
- ▶ Strong **authentication** and **accounting** mechanisms
 - ▶ X509 Certificates and Nagios
- ▶ Successful transition towards a **distributed** organization
 - ▶ From E-GEE to EGI and its federation of NGIs

(*) This is a totally personal assessment.

▶ Assessment^(*) after 10 years – Fails



- ▶ Software infrastructure far too **complex** and **constrained**
 - ▶ **Discourage** many new potential users
 - ▶ More and more VOs are **bypassing** the job management chain
 - ▶ Generalization of **pilot jobs** (DIRAC, Diane, Glide-ins, ...)
- ▶ Part of the success is due to the specific of the main target community
 - ▶ Physicists are not afraid of computing and coding
- ▶ Bad support of **HPC jobs**
 - ▶ Running an MPI job on the Grid is a hassle
- ▶ Performance quite impossible to model
 - ▶ Too variable
 - ▶ Also prevents reproducibility
- ▶ In France: too **scarce fostering** between **Research in CS** and **Production** (admin and users)
 - ▶ since the DataGrid project!

(*) This is a totally personal assessment.



Past

- ▶ Then, there was the Grid
- ▶ Some personal assessments


Present

- ▶ Now, there is the Cloud
- ▶ Several views of the same concept

Future

- ▶ A French Academic Cloud
- ▶ First Step and Next Move


And Now, Here Comes the Cloud!



- ▶ Replaced Grid (as a buzzword at least) in less than 5 years!
- ▶ Most CS conferences are now rather on Clouds than Grids

But, What is the Cloud?


- ▶ A wonderful promise!
 - ▶ Access to an unprecedented amount of computing and storage resources by **just plugging** your laptop
 - ▶ Everything is hidden by a **cloud**
- ▶ Also a crucial need
 - ▶ Most sciences are now **data-intensive**



Still a large scale distributed computing and storage infrastructure

Main Motivation for Clouds: Money!

► But not the only motivation ...



Source : http://media.amazonwebservices.com/AWS_Cloud_Best_Practices.pdf

▶ A Commercial View



Amazon Web Services

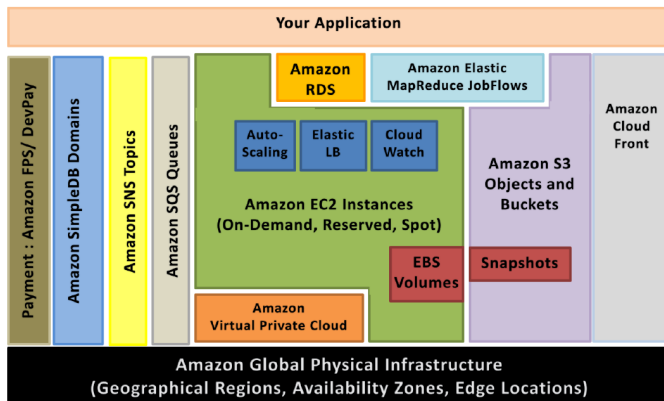
- ▶ Commercial web site amazon.com (1995)
 - ▶ Fluctuating traffic (Christmas burst), 24/7 availability
 - ▶ Strong expertise on data analysis and stock management
- ▶ How to **maximize the outcome** of underused servers?
 - ▶ **Rent them** to individuals!

AWS is THE reference

- ▶ **pay-as-you-go** model and **EC2/S3** interfaces
- ▶ More than 20 products
 - ▶ Dropbox, Slideshare, Netflix, Zinga, ...
- ▶ 5 data centers scattered across the world



General Overview of the AWS Ecosystem



Source : http://media.amazonwebservices.com/AWS_Cloud_Best_Practices.pdf



What Shall We Keep From This View?



- ▶ Cloud computing is a way to **maximize the outcome** of a distributed infrastructure
- ▶ **Diversity** of service offers is a **key to success**
- ▶ Viable **economic model**
- ▶ Importance of **user interfaces**
- ▶ **Elasticity** is an interesting new concept
 - ▶ Offer great production opportunities and research topics

A Governmental View



► Objective

- Build a **trusted** data center to store French **strategic** data
- Namecode **Andromede**
- IaaS Cloud
- Computing, storage, and network offer
- **Development** of an **Open Source** software stack

► Who is it for?

- French administration and companies

► Budget and Operation

- 225 millions euros (75 millions through the **FSN**)
- Operated by **Orange** and **Thales**



- ▶ A concurrent project: **Nuage**
 - ▶ Operated by a consortium of 7 French **SME**
 - ▶ Including two **research teams** on distributed systems
- ▶ **Budget**
 - ▶ 10 millions euros (5 millions through the FSN)
- ▶ Focus on **Green Data Centers**
 - ▶ Original design
 - ▶ Reduced energy consumption and low PUE
- ▶ Claimed strength
 - ▶ Smaller project but easier and faster to set up!



What Shall We Keep From This View?



- ▶ Cloud computing is **strategic** for the French government
 - ▶ At least until May, 6th ...
- ▶ A lot of invested money
- ▶ Strong incentive to **not depend on non-French companies**
 - ▶ Keep a total **control** on where data are stored and analyzed
- ▶ **Energetic** concerns
 - ▶ Through **workload consolidation**

- ▶ The CAPRI project
 - ▶ Answer to the EquipEx call for proposal
- ▶ Objectives
 - ▶ Provide an innovative Data-Scope instrument of unprecedented capabilities
 - ▶ Cover IaaS, PaaS, and SaaS views of a cloud.
- ▶ Approach
 - ▶ User-centric
 - ▶ Build upon some well identified pilot use cases
 - ▶ Close collaboration Applications–Research–Industry
 - ▶ Intrinsically multidisciplinary project


- ▶ The LX Cloud initiative
 - ▶ Migrate the CERN IT to the Cloud
- ▶ Objectives
 - ▶ Service consolidation
 - ▶ Virtualization of batch resources
 - ▶ Both should ease the management
- ▶ 90% of resources to be virtualized by the end of next year

- ▶ EGI Federated Cloud Task Force
 - ▶ Led by Matteo Turilli, Oxford e-Research Centre, University of Oxford
- ▶ Objectives
 - ▶ Write a **blueprint** document
 - ▶ Deploy a **test bed**
 - ▶ Integration of virtualised resources, monitoring and accounting
 - ▶ Investigate **requirements**
 - ▶ Provide **feedback**
 - ▶ Identify **user communities** willing to be early adopters
 - ▶ Identify **indirect issues**
- ▶ **Mandate**: from September 2011 to March 2013

- ▶ Scenario 1: VM Management
 - ▶ *"I want to start a single existing VM image on a remote cloud"*
- ▶ Scenario 2: Managing my own data
 - ▶ *"I want to start a VM instance from an image that I have created"*
 - ▶ *"I want to associate my running VM with a data set in the Cloud"*
 - ▶ *"I want to take snapshots of my running VM for restart purposes"*
- ▶ Scenario 3: Integrating multiple resource providers
 - ▶ *"I want to choose on which resource provider I want to start my single VM"*
 - ▶ *"I need to know about the VMM capabilities the provider offers"*
- ▶ Scenario 4: Accounting across Resource Providers
 - ▶ *"My usage across different resource providers needs to be recorded and reported to multiple aggregators"*
- ▶ Scenario 5: Reliability/Availability of Resource Providers
 - ▶ *"Give me info on the reliability/availability and status of the remote VMs"*
- ▶ Scenario 6: VM/Resource state change notification
 - ▶ *"Told me when the status of the [VM] instance I am running changes"*
- ▶ Scenario 7: AA across Resource Providers
 - ▶ *"I want to use my existing identity, and not re-apply for new credentials"*

EGI TF Roadmap

Roadmap Mar/Sep '12






What Shall We Keep From This View?



- ▶ Cloudification in on the move!
 - ▶ At least at the IaaS level
- ▶ Importance of federation aspects
 - ▶ Need to build upon standards (OCCI, CMDI)
 - ▶ Let the site decide of their own local policies
- ▶ Efforts driven by use cases/scenarios
 - ▶ Do what is needed only
- ▶ Keep the best features of the Grid
 - ▶ Single Sign-on, accounting and monitoring

- ▶ Many Open Source IaaS management systems
 - ▶ OpenStack, OpenNebula, Nimbus, Eucalyptus
 - ▶ Some come from research
 - ▶ Other from industry
- ▶ Two stacks emerge
 - ▶ OpenNebula
 - ▶ Deployed by early adopters
 - ▶ OpenStack
 - ▶ More promising for newcomers
 - ▶ A lot of developing force (a.k.a. money) behind
 - ▶ Strong support of industry
- ▶ All aim at **interoperability**
 - ▶ Through standardized (OCCI, CDMI) or *de facto* (EC2, S3) interfaces





What Shall We Keep From This View?



- ▶ Cloud will probably be **Open Source**
 - ▶ Coherent with governmental trends
- ▶ **Open Stack** is rising!
 - ▶ A feature is missing? Just ask, someone will add it soon!
- ▶ Every stack has its own strengths
 - ▶ Keep them in mind whatever the chosen one!

- ▶ The **Grid'5000** platform
 - ▶ An infrastructure for research on Grid, Clouds and other DCI
 - ▶ Customize, deploy, test, and do over
- ▶ Towards **virtualized** environments
 - ▶ Historical bare metal IaaS mode
 - ▶ Since 2005!
 - ▶ using OAR (batch) and kaDeploy (system image deployment)
 - ▶ Network isolation (KAVLAN)
 - ▶ Virtual machine deployment (Xen, KVM)
 - ▶ Many supported Cloud stacks
 - ▶ OpenNebula, Nimbus, and Open Stack soon

- ▶ Many **hot research topics** investigated
 - ▶ Virtualized resource management
 - ▶ Including scheduling, mapping, and migration
 - ▶ Managing Big Data
 - ▶ Including MapReduce, storage resource management, application deployment
 - ▶ Energy consumption monitoring and resource management
 - ▶ Fault-tolerance and resilience
 - ▶ Sky computing
 - ▶ Using several IaaS management stacks as a single environment


▶ What Shall We Keep From This View?



- ▶ A lot of **top level CS research teams** in France
 - ▶ Many of them involved on topics related to **cloud** and **big data**
- ▶ Will propose a platform dedicated to experiments involving big data
- ▶ Existing **collaborations** with production clouds
 - ▶ Technological and expertise transfers
 - ▶ Dimensioning experiments
 - ▶ Log analysis and injection in a simulation context
- ▶ Grid'5000 provides the **perfect playground** to test new solutions



A User view



THE CLOUD THINGY

A User view – Some Use Cases



- ▶ Have my own nice customized VM with all the software package I need
- ▶ Access to a virtually infinite pool of resources and manage it by myself
 - ▶ Similar to the pilot job approach on the Grid
- ▶ Keep my old fashioned Grid stack on top an IaaS Cloud
 - ▶ The StratusLab approach
- ▶ Do everything through a simple web portal
 - ▶ Required by non-CS specialists
 - ▶ as in the IRT BioAster and IMI eTRIKS projects
- ▶ Access to small subsets of huge data in a cloud fashion



What Shall We Keep From This View?



- ▶ Most of **new** potential users of cloud
 - ▶ Do not know what they expect from the Cloud
 - ▶ Do not (sometimes) even know what they need
 - ▶ Glaring example of the IRT BioAster
- ▶ We cannot assume the same public as for the Grid
 - ▶ Users that see CS as a magical (but frightening) tool
 - ▶ No more skilled physicists
- ▶ Implies to start with some **well defined use cases**



Past

- ▶ Then, there was the Grid
- ▶ Some personal assessments

Present

- ▶ Now, there is the Cloud
- ▶ Several views of the same concept

Future

- ▶ A French Academic Cloud
- ▶ First Step and Next Move

► Definition

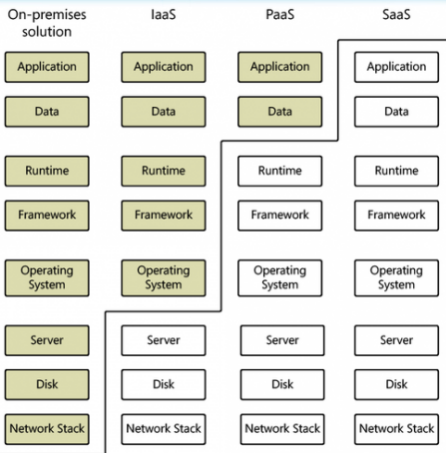
- Deploy a Cloud infrastructure to serve the computing and storage needs of french academic institutes (universities, research institutes)

► Expectations

- Not miss the Cloud (r)evolution
- Control our technical and strategic decisions
- Answer to government expectations related to Big Data challenges
- Be coherent with strategies of partners (CERN, EGI, ...)
 - To what extent?
- Foster collaborations with research in CS and industry
 - Not reproduce the same errors as for the Grid
- Benefit (and make users benefit) of the strong expertise on
 - Data management
 - Federated distributed computing infrastructures

- ▶ Initial test bed
 - ▶ Funded by France Grilles (cont'd in 2012)
 - ▶ Two sites: CC IN2P3 and IRIT, Toulouse
 - ▶ Other potential sites (LAL, LPHC)
- ▶ Technical choices
 - ▶ Thorough study conducted at CC IN2P3
 - ▶ Open Stack holds the rope
 - ▶ Be part of a federation is a main concern (\Rightarrow standards)
- ▶ Use cases
 - ▶ Provided by IRT BioAster and eTRIKS projects
 - ▶ More to come with new projects (Mastodons, Big data, ...)
- ▶ Aims
 - ▶ Run jobs in production
 - ▶ Test solutions coming from research in CS

Which aaS do we target?



- ▶ IaaS is on the way
 - ▶ Virtualized batch
- ▶ Most users would like SaaS
 - ▶ Web portals mainly
- ▶ Will be driven by use cases

Source: "Windows Azure Step by Step" By Roberto Brunetti, Microsoft Press

What's next?



- ▶ Integrate the EGI task force test bed
 - ▶ Once the stack is up and running
 - ▶ Keep an open eye on the choices made by other main partners
 - ▶ Serve the applications of on-going projects
 - ▶ Set up a **group of reflexion** with
 - ▶ People operating the production Grid
 - ▶ Researchers working on cloud and data management
 - ▶ Power users from various communities
- ⇒ Define a strategic roadmap for the next 5 years

*I don't care if my cloud computing architecture is powered by a grid, a mainframe, my neighbour's desktop or an army of monkeys, so long as it's **fast, cheap and secure**.*

Sam Johnston (<http://samj.net/>)

*I don't care if my cloud computing architecture is powered by a grid, a mainframe, my neighbour's desktop or an army of monkeys, so long as it's **fast, cheap and secure**.*

Sam Johnston (<http://samj.net/>)

let's try to avoid the army of monkeys!