

# LAL and GRIF Site Report

Michel Jouvin  
LAL, Orsay  
[jouvin@lal.in2p3.fr](mailto:jouvin@lal.in2p3.fr)

HEPiX, Annecy, May 2014

# Infrastructure: New Era !

- New datacenter delivered on time and on budget in Orsay
  - Shared initiative by 6 labs
  - 100 m<sup>2</sup> / 400 kW, extensible, PUE < 1,3
  - Started in production mid-October 2013
  - See dedicated talk on Wednesday
- LAL moved entirely its previous computing room
  - Completed in 2 months (mid-October to mid-December)
  - Very limited impact on users: 2 x ½ days for major services
  - Managed and done by the IT operations group
    - 6 months of preparation: (almost!) every detail discussed before
    - Very strong involvement of the whole team
    - An occasion to bootstrap a more ITILish process
  - Benefit from a fully monitored infrastructure, including alarms

# Infrastructure: New Era !

- New datacenter delivered on time and on budget in Orsay

- >
- >
- >
- >

- LA

- >
- >
- >



- > Benefit from a fully monitored infrastructure, including alarms

# Hardware (Storage)

- HA Storage: 1 NetApp file FA3220 (70 TB)
  - NFS/CIFS and iSCSI
  - Very reliable
  - Will extend this year but have to plan replacement...
- Performance/Volume: HP blade C7000 with 1 spare server + dense disk JBODs
  - Server swap without reinstallation: just a blade configuration command
  - Very sensitive to power cut: several HW failures
  - HP support close to catastrophic... 10G SFP+ very pricy
  - Looking for an alternative...
- Just received a Dell 3260 with SAS attachment
  - More at a future HEPiX!

# Hardware (Servers)

- Server renewal suffered from the infrastructure problems in the last years...
  - Most servers 6+ years old...
  - A priority for the next years
  - Performance improvement allows to replace 3-4 by 1...
  - Server consolidation raises the need for virtualization...
- 1 rack full of Dell C6100/6200
  - 1/2U server
  - Challenge of cooling (20KW/rack) well addressed by rear-door cooled racks (ATOS)
    - New computing room designed for this kind of racks
  - Very good experience so far... but other GRIF labs with a not so good experience
    - Difference of cooling technology?

# Network

- Network core: 8-year old BD8810 retired, more redundancy
  - > New network core shared with computing room partners: 1 pair of Dell Force 10 4810 (48 10G ports, 4 40G ports, low latency)
    - Modular core switch architecture, extensible
    - Good experience so far
  - > New office concentration : 1 pair of Dell 8132
    - Dual link (LAG) to core switches
- WiFi infrastructure renovated
  - > Captive portal for visitors and internal users (UCOPIA)
    - Internal users can use Active Directory accounts
    - Internal users complain about the need to log into the portal...
    - We are not equipped to manage MAC address registration...
  - > Starting EDUROAM deployment
    - A potential alternative to portal for local users



# OS Changes: Unix

- Only (Scientific) Linux left
- Starting an aggressive campaign to upgrade everything to SL6
  - Still have a few SL4 boxes for some internal services
  - Linked to the move to YUM for package deployment with Quattor
    - See Quattor Update by I. Collier
    - Removing the main “point of pain” in Quattor: maintenance of package list
    - Very minor modification to the site configuration: Quattor backward compatibility is great!
  - Hope to have most of the site upgraded to SL6 by the summer
- Now have a YUM snapshot infrastructure well in place
  - Snapshots managed as hard links to repository mirrors (~0 byte)
  - Move from 1 snapshot to another controlled by Quattor
    - Snapshot managed as immutable: ability to roll back

# OS Changes : Desktops

- Only Windows is centrally managed
  - W8 is the default OS
    - W7 accepted when there is a specific application requirement...
    - Had to face the negative opinion about W8 new interface but we were able to convince users!
    - Deployed by WDS with MS Deployment Toolkit (MDT); a great free tool!
  - Getting rid of WXP: generally involved HW renewal
- National policy about encryption requirements: all desktops/laptops **must** be encrypted
  - Trying to comply with this policy as much as possible...
  - All new computers must be installed with encryption activated
    - Windows (Bitlocker) and Mac (File Vault) are easy, including capacity to configure an admin recovery key
    - Linux harder... requires reinstallation



# Virtualization

- ◉ With 40-core servers, virtualization is a requirement for service machines
  - Also provides easier delegation of some parts of the machine management
- ◉ Started 2 years ago a small private cloud for service virtualization based on StratusLab
  - Not open to any user
  - StratusLab chosen due to LAL involvement in the project
  - Basically working, but some rocky edges
    - StratusLab persistent disk service very valuable for this use case: service data outside of the VM image without the need for a shared FS
    - Lack of dashboard for everyday management
  - Image lifecycle managed with Quattor
    - Very effective, simple through contextualization scripts

# SLAM

- LAL needed a tool to do IP address allocation and have an automatic configuration of related services
  - Mainly DHCP, DNS, Quattor
- A project started 2 years ago with a student
  - No existing tool for this purpose identified: generally more on the machine lifecycle management side, something different...
  - Designed to be a generic tool, extensible to other “backend” service if necessary
  - A web application + a command line tool
    - Web application written using Django framework
  - Have in mind to push it to GitHub and will be happy to see other contributors if any interest
    - Should already been done!
  - IPv4 currently, IPv6 address management to be added soon

# SLAM

**SLAM** Pools ▼ Hôtes ▼ Journal Générer

## Créer un nouvel hôte

Nom

Alias

Catégorie

Nom du pool (optionnel) **ou** adresse (optionnel)

Adresses MAC

☐ Allocation aléatoire  
☐ Ne pas générer d'enregistrements DNS pour cet hôte

### Informations optionnelles

Numéro de série

Numéro d'inventaire

Propriétaire

**Créer**

do IP address allocation and have an  
one of related services

**SLAM** Pools ▼ Hôtes ▼ Journal Générer Langue Déconnexion

## Liste des hôtes

Nom de l'hôte	Alias	Adresses	Adresses MAC
<a href="#">pc-89055</a>	pc-flower	<a href="#">134.158.89.55</a> (pool: <a href="#">lal-pc2</a> )	00:06:5b:bf:66:be
<a href="#">pc-90059</a>	pc-si10 pc-ivan	<a href="#">134.158.90.59</a> (pool: <a href="#">lal-pc2</a> )	00:18:8b:73:ac:62
<a href="#">pc-91044</a>	pc-bao7	<a href="#">134.158.91.44</a> (pool: <a href="#">lal-pc2</a> )	00:10:18:F6:47:B0
<a href="#">pc-90175</a>	pc-sm32	<a href="#">134.158.90.175</a> (pool: <a href="#">lal-pc2</a> )	00:08:74:fe:3d:5c
<a href="#">pc-76190</a>		<a href="#">134.158.76.190</a> (pool: <a href="#">lal-pc1</a> )	b4:b5:2f:c9:7d:eb
<a href="#">pc-vagnucci</a>		<a href="#">134.158.89.131</a> (pool: <a href="#">lal-pc2</a> )	10:60:4b:8c:63:9a
<a href="#">pc-76116</a>	nb-viaud	<a href="#">134.158.76.116</a> (pool: <a href="#">lal-pc1</a> )	00:21:70:B4:AE:91
<a href="#">pc-88012</a>	mba-poeschl	<a href="#">134.158.88.12</a> (pool: <a href="#">lal-pc2</a> )	b8:8d:12:55:97:1b
<a href="#">pc-88020</a>	mbp-allumette	<a href="#">134.158.88.20</a> (pool: <a href="#">lal-pc2</a> )	40:6c:8f:05:c1:88
<a href="#">pc-8840</a>	pc-garault	<a href="#">134.158.88.40</a> (pool: <a href="#">lal-pc2</a> )	10:60:4b:81:b5:56
<a href="#">pc-88041</a>	nb-sm10	<a href="#">134.158.88.41</a> (pool: <a href="#">lal-pc2</a> )	5c:26:0a:7d:b8:7e
<a href="#">pc-88042</a>	pc-bogard	<a href="#">134.158.88.42</a> (pool: <a href="#">lal-pc2</a> )	00:1a:a0:3c:4a:89
<a href="#">pc-caceres</a>		<a href="#">134.158.89.110</a> (pool: <a href="#">lal-pc2</a> )	00:19:b9:02:52:e8
<a href="#">pc-88044</a>	nb-stocchi3	<a href="#">134.158.88.44</a> (pool: <a href="#">lal-pc2</a> )	5c:26:0a:5f:c0:9c
<a href="#">pc-88045</a>	pc-video3	<a href="#">134.158.88.45</a> (pool: <a href="#">lal-pc2</a> )	00:24:21:83:3f:F3
<a href="#">pc-88046</a>		<a href="#">134.158.88.46</a> (pool: <a href="#">lal-pc2</a> )	

done!

address management to be added soon

# GRIF Update

- A grid site in Paris region involving LAL, lrfu and 4 other labs
  - 9000 cores, 4 PB of disks
    - A major site for all 4 LHC experiments
  - GRIF resources representing 2/3 of computing resources at LAL
- Tight budget situation: very limited growth
  - Most spendings for HW renewal
  - LCG France funding 70% of the renewal cost...
  - Starting to offer disk storage to LHCb (T2D)
    - 150 TB this year
- perfSONAR-PS deployed at the 3 main sites with respect to networking
  - Both for LHCONE and the internal 10G network
  - New dashboard (MadDash) found very useful
  - Had a look in the past years to perfSONAR-MDM but gave up

# P2IO VirtualData...

- P2IO: a common initiative all HEP, NP and Astrophysics labs in Orsay to foster synergies between each others
  - Physique des 2 Infinis et des Origines: 8 labs, including LAL and Irfu
  - Several technological platforms: VirtualData = computing
- VirtualData goal: build a computing expertise network around a shared computing platform
  - 130 people involved in computing
    - Development (~75) and Operations (~55)
  - Shared computing platform hosted in 2 shared facilities to enable redundancy when needed
    - Shared facility in Orsay part of this project
    - Second facility being built at Ecole Polytechnique

# ... P2IO VirtualData

- Main challenge today: shared management of these share facilities
  - > Build on GRIF experience: 8-year old successful experience of working together to manage a unique grid site
    - 4 GRIF partners involved in VirtualData
  - > Managing infrastructure is more challenging than computers...
    - Need time to build trust between people not used to work together
    - No lab can do the work for the others...
    - Some concrete work not seen as being rewarding (clean cabling, PDU/rack installation, daily monitoring...)
  - > But a good will and progressing well
    - “Tour de garde” (rota) between labs to do the daily monitoring and problem troubleshooting
    - One group formed from network experts of each lab discussing the network architecture and proposing/implementing solutions



# StratusLab

- Open-source cloud distribution (IaaS)
  - Started in 2009, EU funded between 2010 and 2012
  - LAL always played a major role in the open-source collaboration
    - Cal Loomis has been the EU project director
- A few installations in Europe
  - ATOS adopted it as its cloud platform for participating to the Helix Nebula European project
  - French Bioinformatics network starting a StratusLab cloud dedicated to bioinformatics
- Strong interest at Université Paris Sud for a university-wide cloud
  - Will be build on the existing StratusLab cloud at LAL
  - Funding expected to extend to 1000 cores, 100 TB early 2015
    - Testing at scale, assessment of further needs