LAL and GRIF Site Report

Michel Jouvin
LAL, Orsay
jouvin@lal.in2p3.fr
http://grif.fr

November 2, 2010
HEPiX, Cornell

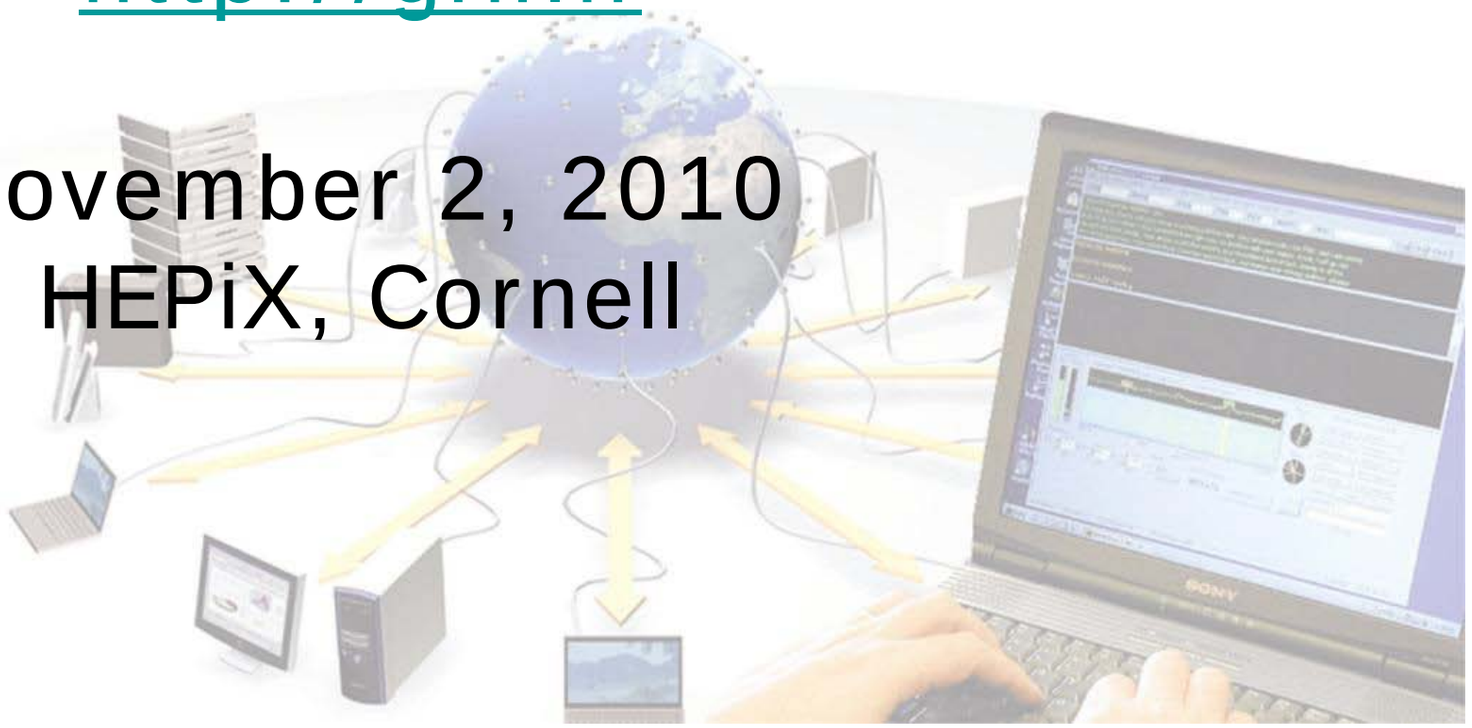

Hardware Changes

- Storage: 2 DDN 6620 (~350 TB usable)
  - High density storage: 60 disks in 4U
    - Up to 2 chained enclosures
  - 2 rdeundant FC raid controllers
  - Very good performance in RAID6
  - Used for grid storage (DPM disk servers)
    - 3 servers per DDN 6620: ~65 TB/server
    - Each server connected with 10 GbE
    - New SAN switches from Brocade
- Network: core switch (BD8810) upgraded
  - Up to 64 10 GbE ports: already 40 installed
  - 10 GbE for disk servers and uplink of compute node racks
- Procurement in progress for Dell C6100
  - 2 Twin in a 2U chassis, X5650, 4 GB/core (2 GB/HT core)

Infrastructure

- 2 small adjacent compute rooms unified
  - Total: ~100 m2, half used
- Testing water-cooled racks (2)
  - Standard racks with “cooling door” from Bull/ATOS
  - Up to 35 kW per rack
  - Open racks
- Plan to extend cooling capacity next year
  - Add 100 kW to existing cooler (150 kW)
- Plan for a new GRIF-shared computing room (COGIS)
  - 400 m2, 1 MW

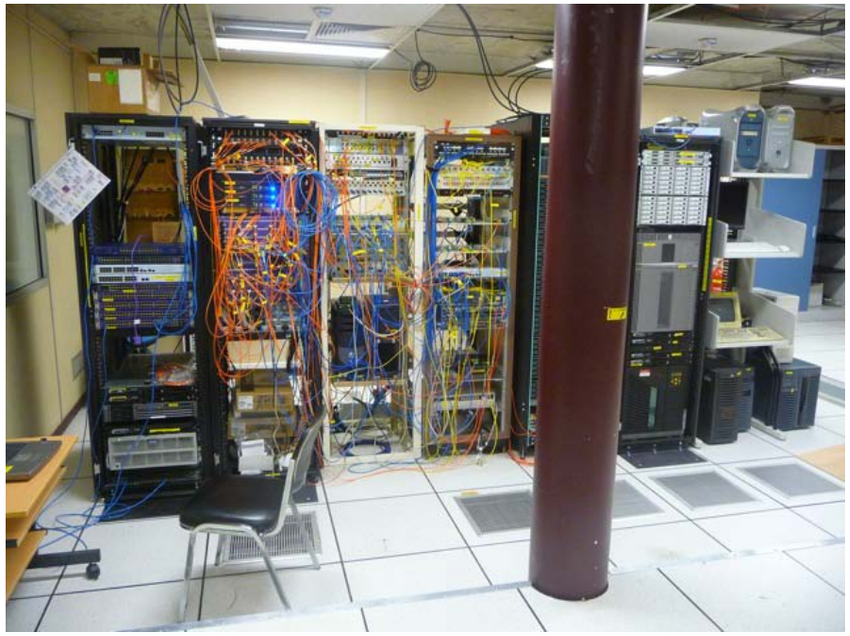

Central File Server

- Started with a 2-node Sun Cluster 2 years ago
  - Working pretty well after initial difficulties
  - Oracle policy makes this solution a dead end: HW no longer sold on its own
- Tru64 cluster (2 nodes) still holding some critical spaces
  - Need to shut it down asap: no more SW support, HW maintenance expensive, bad performances
  - Looking at NFS appliances: NetApp or BlueArc
    - BlueArc technology attractive but not clear they have the clustering capabilities for high availability

Windows

- Active Directory: still running 2003
- Old RIS server moved to a WDS server (W2008)
- Windows 7 installed on all new desktops/laptops
- Antivirus: changed for Microsoft Security Essential
  - End of contract for previous MacAfee

GRIF Status

- Consolidated resources increasing…
  - CPU : 60K HS06; disk : 1.5 PB
  - Spread over 6 locations: 6 CE, 6 SE
- GRIFOPN (10 Gb/s private network) a corner stone for the unified view of the site
- Taking more national responsibilities in the new French NGI (France Grilles)
  - Monitoring
  - Site certification infrastructure
- Looking at CVMFS to replace NFS for SW area
  - RAL/PIC tests looking very promising

GRIF Internal Tools

- Collaborative tools are critical for the (distributed) technical team
  - 20 persons on 6 sites
  - Not everybody full-time = need for “chaotic” participation
- Private chat infrastructure based on Jabber
  - Open-source server : openFire
    - Configured by quattor based on GRIF VO membership
  - Underused but working well
    - Many people prefer emails
- Trac still the master piece for documentation, configuration change tracking…
  - In combination with SVN used by Quattor

OS Changes and Issues

- Linux
  - Most systems running SL5.4+
    - Many systems upgraded from SL4 or earlier SL5
  - Monthly deployment of OS errata if no critical updates
    - Done in whole GRIF, using Quattor
    - Kernel not updated if no critical vulnerabilities
- Disk server freeze under heavy load
  - Only happens with disk servers configured with XFS under SL5
  - Cause not yet clear: problem caused by certain disk drivers
    - In particular on Sun Thumpers
  - ext4 seems not to have the problem: moving to ext4
    - Tests saw similar performances

Cloud Projects

- StratusLab: integrating grids and clouds
  - Implement a grid site on a private cloud
  - EU-funded 2-year project
    - Director: Cal Loomis
  - Producing a toolkit based on OpenNebula
    - First release this week
    - Includes components for configuration management with Quattor
  - LAL participating to the testbed and API development
    - 3 FTEs
- EDGI/DEGISCO: interoperation between production grids and desktop grids
  - 2 EU-funded 2-year projects
  - Bridge between gLite and BOINC/XtremWeb
  - LAL involved in testbed and standardisation
    - 3 FTEs

Quattor

- The key management tool for ensuring GRIF site consistency and allowing a distributed management
  - One unique configuration database, 1000 machines
  - Management of whole GRIF possible from everywhere
  - Non-grid machines at LAL and at LLR, including desktops and Xen-based VMs
  - StratusLab: management of OpenNebula and VM images
- Stronger quattor community despite QUEST project failure at European call last year
  - More and more coordinated work
  - Most of the community trying to share common templates