

LAL and GRIF Site Report

Michel Jouvin

LAL, Orsay

jouvin@lal.in2p3.fr

<http://grif.fr>

May 25, 2009

HEPiX, HPC2N, Umea



Hardware Changes

- 2-node Sun cluster with SAS JBODs for internal services
 - 2x SunFire 4100
 - 2x J4400 (24x 1TB SATA disks) instead of planned 1x J4500 (48 SATA disks)
 - Many HW problems due to unsupported HW sold by Sun...
 - Many restrictions coming from SATA disks : SCSI reservations not supported...
 - Evaluating as a potential replacement for HP Alpha/TruCluster
- 3 new Sun X4550 (Thor = Thumper v2) with 1TB disks for grid storage
 - Not yet installed because of infrastructure problems
 - 10 Gb/s uplinks
 - Running as DPM disk servers (Linux SL)

Infrastructure

- Infrastructure is still the main concern :
 - Harmonics in neutral : installation of an active harmonic filter (from Merlin Gerin) done 4 months late...
 - Cooler at half capacity during 2 months in the winter...
 - Capacity reduced by 50% during 2 months
 - Lost power on cooling 2 consecutive nights in March : reached 60°C in computing room because of lacking thermal security
 - Capacity reduced by 40% during 1 ½ month
- Plan to refurbish computing room late...
 - Need to “unify” 2 adjacent computing rooms
 - Last opportunity to do it is this year
 - Ability to move all active equipments in the most recent room
- Budget situation is more and more difficult
 - Most credit sources dedicated to procurements of HW

Central File Server

- Tru64 cluster (2 nodes) still the corner stone of internal resources but completely overloaded
 - File server (NFS, CIFS) + mail server mainly
 - Very robust cluster file system
 - MySQL and web servers moved to dedicated Linux servers
 - Web server documents still served by cluster
 - Heartbeat: not reliable enough for a file server
- Replacement based on a 2-node Sun Cluster
 - Many problems with the initial configuration due to Sun mistakes in proposed/delivered solution
 - Very difficult to find real experts on Sun Cluster in France...
 - Interested by ZFS features but not supported as a global FS
 - Features Tru64-AdvFS-like, with integrated RAID
 - Each node has to NFS-mount other node file systems...
 - Sun Cluster # to 15-years old Tru64/ASE or Heartbeat...
 - First tests showed good NFS perms : 80 MB/s on a client
 - Going to production next week... 5 months late !

GRIF Status

- Consolidated resources increasing...
 - CPU : 6 MSI2K; 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
 - Assessing impact on perfs with Atlas analysis tests
- Still fighting with Torque/MAUI instabilities and scalability but no real impact on users
 - Spent some time to understand what to monitor and decide the appropriate recovery actions in case of problems
- GRIF operating WMS as a national core service
 - Official WMS for ALICE VO
 - 2 redundant machines
 - 4 generic redundant machines supporting #40 VO's
 - Long fight for stability... but successful since 2 months
 - Team of 6 people participating to management

MPI Jobs & Shared Storage

- Increasing number of MPI jobs (LAL)
 - Mainly from non-LHC VOs
 - Peaks at 300+ 4-process jobs
 - 1 user testing highly // jobs : 128 or 256-process jobs
 - No dedicated HW (Infiniband) : need to work on efficiency
 - Private Gb/s between WNs ? Interested by experiences...
- Reached performance limits for NFS-based shared home directories with MPI jobs requiring writable shared areas
 - Does not affect non-MPI jobs as they don't use home directory
 - Considering experimenting with LUSTRE
 - Already running a test instance of LUSTRE for CARRIOCAS project

GRIF Monitoring

- Several problems in last months underlying insufficient GRIF monitoring
 - Mainly relying on SAM in fact
 - E.g. : inability to quickly detect and set offline a WN acting as a black hole
- Originally based on LEMON
 - Too difficult to maintain with limited support from CERN
 - Alarms requiring (expensive) Oracle for the backend
- Replaced 2 months ago by Nagios + NagiosGraph
 - 1 instance GRIF-wide, organized by group of services
 - Configuration done by quattor, based on machines configured
 - Not using the mainstream templates for Nagios, merge in progress
 - Grid services not yet monitored by Nagios
 - Done by SAM (central grid monitoring service)
 - Central/regional testing will be integrated into our local instance
 - NagiosGraph providing historical data

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
- Started a chat infrastructure based on Jabber
 - Open-source server : openFire
 - Configured by quattor based on GRIF VO membership
 - Plan to configure several chat rooms based on VO groups
 - Many open-source clients for every platform
 - Native client on Mac : Ichat
 - Firewall restrictions : good web-based client
 - Many people still preferring email but reach the limit
 - Some days : 50-100 mails...
 - Not enough interactive : quickly diverge...
- Trac still the master piece for documentation, configuration change tracking...
 - Started to use issue tracker for main actions

OS Changes

- Linux
 - No major change yet... but SL5 knocking at the door
 - Already running a some dedicated SL5 servers (Trac, NFS...)
 - Deploying errata (except for kernel)
 - gLite 3.2 WN (SL5) ready for deployment but need to assess VO readiness
 - 1 test system installed
 - Can be done very quickly thanks to quattor
 - Interactive servers : need to wait for gLite 3.2 UI
- Windows
 - Still a mix of Windows XP and Vista
 - No enthusiasm for Vista migration, not only for good reasons...
 - Laptop users generally happy with Vista new features
 - Trend : wait for Windows v7

Quattor

- The key management tool for ensuring GRIF site consistency and allowing a distributed management
 - One unique configuration database, 700 machines
 - Non-grid machines at LAL and at LLR, including desktops and Xen-based VMs
 - Management of whole GRIF possible from everywhere
- Widened and stronger quattor community
 - Quattor adopted by Morgan&Stanley as the core of their new management infrastructure
 - Already managing 20K machines with quattor
 - Developed a CDB/SCDB alternative : Aquilon and will open-source it
 - 2 workshops/year well attended with new comers
 - Some old contributors left too –(
 - New medium/large sites expressing interest
 - QWG effort more and more successful
 - Seamless integration and easy deployment of gLite
 - GRIF is one major contributor... but several others