

LAL and GRIF Site Report

Michel Jouvin

LAL, Orsay

jouvin@lal.in2p3.fr

<http://grif.fr>

April 5, 2008

HEPiX CERN



- Dominated by commissioning of 250 new 8-core machines
 - Mainly grid worker nodes (200)
 - Computer room expanded last summer with new cooling and power capacity
- Main issue : harmonics in neutral
 - High current : 70A!
 - Cables/breakers temperature : $\sim 45^{\circ}\text{C}$
 - Cause : high number of PS (300) without a UPS
 - Solution under investigation : probably an harmonic filter after the transformer
 - A site already reported a similar problem at a previous HEPiX
 - Interested by sharing experience of electric circuits without UPS
- Old computer room : cooling unit almost dead...
 - Still hosting some critical services : file servers, CE, SE
 - Refurbishing delayed by previous electrical problems

- 10 GbE on Thumpers : Intel/PRO NIC replaced by Neterion, performances acceptable
 - 6 GbE "full-duplex" with iperf (under Linux)
 - Not used for production (DPM) but for CARRIOCAS R& project (LUSTRE testing)
- GRIFOPN being deployed : 10 Gb/s private network between GRIF sites (6)
 - Based on Black Optical Fibres provided by RENATER (french NREN)
 - 1 external connection to CCIN2P3 at 5 Gb/s shared by GRIF sites but dedicated to GRIF traffic
 - Currently 5x1Gb/s shared with other usage
 - Demonstrated ability to saturate 2x1Gb/s (2x100MB/s) during 2 hours with FTS transfers for CMS and ATLAS
 - 2 links already set up last week, GRIF-CCIN2P3 expected this week, other coming by end of June

- Current consolidated resources : 4 MSI2K, 300 TB
 - Spread over 6 locations : 5 CE, 6 SE
- Unified view of storage : all CEs with equal access to SEs as soon as GRIFOPN is ready
 - Reduce number of SE per VO with larger share on each
 - MW not ready for a unified SE but still considered a valid goal
- Main issue : scaling to 2000+ concurrent jobs
 - No problem experienced with NFS server
 - SW area (applications) and home directories (not used during jobs)
 - Current HW : 4-core Opteron, 8 GB of memory, 1 GbE
 - MAUI not stable enough and no support available
 - Currently only 2/3 of cores in production
 - Would prefer to avoid splitting into several clusters/CEs
 - Community support unreliable : mainly sharing of experience

- Tru64 cluster still the corner stone of internal resources
 - File server, web server, mail server, print server, SVN, Trac, MySQL server...
 - Very robust cluster file system
 - 2-machine cluster overloaded, performances are horrible
 - HW is old, too expensive to upgrade, even to add memory
- Move to service-dedicated servers progressing slowly
 - Would like to retain at least automatic failover capabilities
 - First experience with Heartbeat (v1) not very successful
 - 2 web servers sharing service of 20 virtual hosts
 - Files and databases still on the cluster
 - Started a pilot with v2 + Stonit, looks better
 - Next step : redundant service of SAN-based file system
 - MySQL, Web-related file systems
 - Interested by FNAL idea of LVC + virtual machines

- # Library Backup
- Backup based on Networker for 15 years...
 - Server moved off cluster on a dedicated Linux machine
 - 4-core machine with 8 GB of memory, 1 GbE (2 available)
 - No HW redundancy but everything on SAN : easy reinstallation in case of failure
 - Easy migration from Tru64 to Linux : DB compatible
 - Replacement of library STK L40 by a Dell M6020 v2
 - 2 LTO-3 drives and FC I/O blade
 - OEM version of Quantum/ADIC Scalar 500 at 2/3 the price
 - Library and drives dual-pathed (2 independent FC fabrics)
 - Very painful commissioning of new library
 - Mainly minor problems but Dell service incompetent
 - E.g. : at installation, declared library DoA because they were unable to complete of the startup procedure...
 - Seems to work well with very good performances

- Currently based on Tru64 Print Server + LPR client
 - Windows : rely on CERN Printer Wizard
- Hard to integrate recent printer models
 - Painful integration into Printer Wizard
- LPR is not the preferred client for most of the platforms
- Planning to move to a Windows-based server infrastructure
 - Windows client : AD Printing Client/Wizard
 - Linux/Macintosh client : CUPS
 - Pilot setup for a few new printers
 - Still questions about how to provide service resilience
 - Windows cluster ? Any other experience ?

- Windows : migrated installation infrastructure from RIS to WDS
 - Looks much easier to maintain, in particular for driver integration
 - Support for VISTA to come soon
 - Need to change anti-virus : license no longer available for McAfee through university, considering Norton but worried by inability to obtain signature update after end of contract with university (3 y
- Desktop/laptop backup : encouraging use of synchronization features on all platforms
 - Windows : offline folders
 - Linux/Macintosh : rsync
 - On Macintosh, recommending iBackup, an rsync GUI

- EDGES proposal (EU FP7) accepted
 - Goal : R&D for integrating institutional grids (EGEE) and desktop grids (BOINC, XtremWeb)
 - 8 partners in Europe
 - 2 persons to be hired at LAL but difficult to find valuable candidates...
 - Project management, software design and development
- CARRIOCAS : distributed LUSTRE tests
 - Funded R&D project to evaluate 40 Gb/s long distance network
 - Currently 2x10Gb/s, upgrade to 40 Gb/s this summer
 - Works currently focused on understanding network performance