

GRIF Site Report

Michel Jouvin

LAL/Orsay

jouvin@lal.in2p3.fr

Quattor Workshop, Bologna 2008

- 6 geographical sites with a unified SCDB
 - Each site machines grouped in cluster(s)
 - Current total = 600 machines, 25 "clusters"
 - SCDB sites provide a very flexible method of sharing configuration between clusters and/or sites
- Most of the systems are grid systems
 - Management of NFS servers for grid systems
- LAL internal systems
 - Interactive servers : all configured as grid UI
 - Service machines, including SAN/FC configuration
 - Linux desktops

Quattor Configuration

- SCDB, HTTP repositories, AII, QWG templates
 - Site specific templates maintained only for site specific machine types
 - Includes LUSTRE
 - OS hooks integrated in a way similar to gLite
 - config/lal in OS templates
- 1 Quattor server per site
 - Currently used only for AII (DHCP+TFTP)
- 1 master Quattor server
 - All profiles deployed on and served from this server
 - All RPMs served from the central Quattor http server
- Quattor expertise spread over all GRIF sites
 - Every site maintains its site specific templates
- Several contributors to Quattor shared templates, addition to Cal and Michel
 - Guillaume, Frédéric (WMS, Nagios), Christine (LEMON)

- GRIF configuration
 - 300 machines added (300 -> 600)
 - Most machines migrated to gLite 3.1 and SL4.5 x86_64
- **Crossed the upper limit** for deployment with a full recompile : 4mn (limit set to 3)
 - Twice the number of machines managed without increasing deployment time
 - **Main change** : replacement of Quattor server by a dual core Clovertown with 16 GB of memory
 - Dependency checking accounts for 25% : compiler change almost ready to improve this
 - Doesn't include time for compilation before deployment
 - Can be done on a desktop (full compile = ~10mn for a "recent" desktop) or shared server
 - More problematic on a laptop : lower CPU and disk speed
 - May compile 1 selected cluster

- Work restarted recently... basic errata management in place since last autumn
 - Local RPM repositories rsync'ed with SL repositories by cron job (every day)
 - rpmUpdates.pl script used to generate a template with pkg_only for the errata RPMs
 - Template included at the end of every configuration
 - List of errata customizable per node (or cluster, site..) through a PAN variable
- Standing issue : kernel update
 - Need to keep several kernel version on the system rather than replace
 - Version of kernel modules must remain consistent
 - Kernel version part of their name
 - Kernel version must be selectable independently of other parts of the update
 - Initial installation must continue to work
 - RPM doesn't allow to install several versions in 1 transaction

- // compilation of clusters on different servers
 - Compile a cluster on the Quattor server of the site it belongs to
 - Require Quattor servers powerful enough at each site
 - Manual association of clusters and servers needed
- HTTPrep servers distributed on each site
 - Avoid SPOF, increase performances
 - Design : use a cache server on each site
 - 1 RPM content change should involve a file name change
 - Very rare exception seen in the last 5 years (< 1/y) : handle manually
- Management of Xen virtual machines
 - Mainly at LAL for non-grid machines (e.g. : Web serve

- Most HEP grid sites in France (7) using Quattor
 - CPPM (Marseille), GRIF, IPNL (Lyon), IReS (Strasbourg), LAPP (Annecy), LPSC (Grenoble), SUBATECH (Nantes)
 - Clermont/Auvergrid started a new site with Quattor
 - 1 Biomed site in Lyon : IBCP
 - 1 non HEP site (Paris Observatory) considering to use Quattor
 - Between 10 and 100 machines per site
- SCDB + QWG
 - 1 SCDB per site
- Several Quattor “experts” outside LAL
 - Some of them already contributing to QWG templates
- No specific site or mailing list for French users
 - Encourage participation to “global” Quattor
 - Use LCGFR/EGEEFR lists as a space for “discussion in French” to help new users