

QWG Templates Update

Michel Jouvin

LAL, Orsay

jouvin@lal.in2p3.fr

<http://grif.fr>

November 4, 2009

Quattor Workshop, Amsterdam



Releases since London

- gLite 3.1
 - 3.1.0-6: June 25
 - 3.1.0-7: October 28
- gLite 3.2
 - 3.2.0-1: July 10
 - 3.2.0-2: October 28
- Attempt to keep them in sync
 - Services present in both 3.1 and 3.2 should offer the same configuration options
 - Except service change like BDII
 - Probably more gLite 3.2 releases in the future
 - xxx_CONFIG_SITE to be executed first... but not yet implemented consistently
- gLite 3.2 requires gLite 3.1
 - Most service nodes not yet available in 3.2
 - LCG CE will never be ported to 3.2

OS Changes

- SL 5.2 and 5.3 added
 - 64-bit only
 - gLite 3.2 only
- SL 4.7 fixed for gLite support
 - 32-bit and 64-bit
 - gLite 3.1 only
- OS errata support improved
- SVN 1.5/1.6 support
 - Version selected with SUBVERSION_VERSION
 - SVN 1.5/1.6 on SL4: only i386
 - Difficulties to build RPMs on x86_64
 - i386 is currently installed on SL4 64-bit
- File system partitionning fixed for non-LVM config
 - Renumber partitions if necessary to have consecutive numbers and handle properly primary/logical
 - Layout examples for all types of partitionning

gLite Templates (common)

- VO configuration based on CIC Portal (now working)
 - Proper handling of VOMS server certificate upgrade but certificate-less configuration not yet ready
 - And is not supported by WMS so far
 - Support for pool accounts for specific FQANs
 - Account name preserved even though the FQAN list changes
 - Account name based on FQAN
- Pool accounts for specific group/roles
- LCG CE and LRMS on different nodes
- Torque/MAUI configuration based on HW description
 - Eg. Number of cores from CPU templates
- Torque 2.3.6 : new pbs-monitoring cron on WN
 - Handle a bug where a server restart may kill clients
 - Ensure local and NFS file systems are writable, else kill the client

gLite 3.1 Templates

- New machine types
 - VOMS: no using anymore obsolete ncm-glite
 - Hydra
- Missing machine types
 - **CREAM CE: urgently needed for ALICE sites**
 - SCAS server: as part of Argus?
 - Pretty urgent too for LCG sites...
- DPM: xrootd support
 - Previously not working
- Preliminary support for multiple CE
 - Waiting for a site to use it to debug it...

gLite 3.2 Templates

- Machine types supported
 - BDII v5
 - WN
 - UI
 - DPM coming soon
 - Only disk server officially released
 - VOBOX being worked on at GRIF
 - Missing LFC
 - Should be pretty easy, only RPM list to update
- BDII v5: a few issues due to RPM packaging
 - slapd.conf improperly overwritten by new versions of bdii RPM: passwd reset
 - Password configuration not yet handled by ncm-lcgbdii: has to be set manually
 - GIP working files relocated to /var/glite
 - One provider doesn't support it

gLite Templates ToDo

- <https://trac.lal.in2p3.fr/QWG/milestone/QWG-Templates-ToDoList>
- New machine types for gLite 3.1 and 3.2
- Multiple CE support testing
 - Several CE sharing the same WNs
- Review account management using LDAP and a more flexible allocation of UIDs
- Support for SGE : driven by a South-African site just starting

Other Items ToDo

- No real plan to do it yet...
- Rewrite SPMA functions with path auto-escaping
 - Removal of the main loop to find the package in pkg_xxx
 - Should lead to a significant perf improvement as called $O(1000)$ by node
- Redesign/rewrite directory-sync ?
 - Support non-Linux OS, don't rely on a complex 'find'

OS Errata Management...

- Significant improvements in September
 - Present in last gLite releases
 - <https://trac.lal.in2p3.fr/QWG/wiki/DOC/OS/Errata>
- New features
 - Define a default errata level per OS version
 - OS_ERRATA_TEMPLATE nlist: 1 entry per OS version
 - Define an explicit errata level for a specific node
 - NODE_OS_ERRATA_TEMPLATE nlist: 1 entry per (escaped) node
 - undef value disable errata deployment
 - Recommended: set PKG_DEPLOY_OS_ERRATA to true in site/config.tpl

... OS Errata Management

- Kernel upgrade
 - Default kernel for an errata level is defined as part of the errata
 - OS_KERNEL_VERSION_ERRATA nlist: 1 entry per OS version
 - Typically in OS templates: config/os/errata/xxx-init.tpl
 - Templates are installing the right version of kernel modules
- Current drawback: only one kernel version can be installed easily
 - Removal of current kernel during installation, before reboot
 - SPMA improved recently by J. Iven (CERN): kernel and modules not removed if they are active
 - Removed at next reboot
 - Not yet in QWG: need testing
- 3 templates per errata: RPM list, -fix, init
 - RPM list generated with utils/misc/rpmErrata: need to ensure kernel modules are in the RPM list

Profile cloning

- Formerly known as « dummy WN »
 - Currently works only on WN
- Minimalist documentation... hard to find!
 - <https://trac.lal.in2p3.fr/QWG/wiki/Doc/gLite/Experimental/DummySpeedup>
 - Lots of undocumented variables
- Principle:
 - Define a reference node: EXACT_NODE
 - Define nodes that can clone this profile: NODE_REGEXP
 - Define DUMMY_WN=true
 - Can use WN_DUMMY_DISABLED nlist to temporary disable
 - Allow to define DUMMY_WN=true by default
 - When true, most of the profile is copied from the ref node
- Most customization point unavailable, loadpath not defined
 - Would reduce dramatically the performance gain

Release Cycle

- Until now, it has been difficult to have more than 1 release per quarter
- No real impact on sites as the branch is considered the production released
 - Announcement done at every major change in the branch
 - Significant testing before merging trunk changes into a branch
- Release are important for visibility
 - A second release manager would help: private discussions with Christos

QWG Contributions

- Several contributors to templates
 - Several regular contributors, in particular Christos...
 - Mine often hides contributions from others at GRIF
 - <http://grid.ie/svnstats/QWG>
- Need to encourage more contributions
 - Not necessarily for “big things”: HW template, miscellaneous changes, fixes...
 - Ask for a SVN account: philippo@lal.in2p3.fr and jouvin@lal.in2p3.fr
- Areas where contributions are required:
 - Support for new machine types : requires more “experimented people” and **is the current weakness**
 - Configuration of monitoring tool
 - Iptables support/integration : need to review/document what exists and what is missing
 - dCache templates: used to be done by Stijn

Documentation

- <https://trac.lal.in2p3.fr/QWG>
 - Renamed for LCGQWG to QWG
- Probably still too much a personal effort...
 - Feel free to contribute!
- Review of existing/missing part required
 - Producing and reviewing must be done by different people
 - Lack of quick-start pages still remains

Usage Status Changes

- RAL is a new site
 - Impressed that QWG was so T1-ready
 - Thanks for them accepting not everything was perfect
 - Many pb fixes, in particular disk partitionning
- No real news from South-Africa
 - South Africa : 2 sites currently starting with QWG
 - Not sure they really started without...
- Europe:
 - France: 2 new sites
 - Bordeaux University and Paris Observatory