

# QWG Templates Status

Michel Jouvin

LAL/Orsay

[jouvin@lal.in2p3.fr](mailto:jouvin@lal.in2p3.fr)

Quattor Workshop, Bologna 2008

- QWG Components
- Changes
- Errata Management
- Planned Changes
- Release policy
- Contributing to QWG
- Conclusions

# Usage Status

- QWG templates increasing and sites using them b
  - Several sites with 200+ machines (GRIF with 600 mac
- Main known sites :
  - GRIF : 6 sites, 300 machines, grid + non grid usage, 1
  - 8 French sites : 1 SCDB per site, 10-100 machines/site
  - BEGRID : 6 sites, 150 machines, 1 SCDB
  - Irish grid : 18 sites, 400 machines, 1 SCDB
  - UAM : 200 + 500 nodes, 3 clusters, 2 SCDB + 1 CDB
  - Valencia : QWG, SCDB
  - PIC (Barcelona)
  - Germany : Aachen, DESY (?)
  - 1 site in Grecia

- Standard : all templates coming from Quattor core are not QWG specific if maintained by QWG
  - **Quattor core templates unmodified** (quattor/, component/, All)
  - hardware/
  - monitoring/lemon/ and upcoming monitoring/nagios
- OS : one set of template per OS version
  - **Application / grid agnostic** : per application OS-hooks
    - Hooks required by gLite (config/glite) provided in distribution
  - **No dependency between OS versions**
- Grid middleware : one set of template per gLite version
  - **Nothing OS related** : located in OS template hooks
  - Support for mixed OS version/arch in 1 cluster
  - Only 1 gLite version per cluster
- **Site configuration outside QWG templates**
- Rely on panc v7 new features

# Changes : OS templates

- OS versions : SL 4.4 and 4.5 are the main ones
  - Added : SLC 4.5, SL 5.1
    - Not integrated with gLite templates : mainly RPM lists
  - Removed : SL 3.05, SL 4.3
    - Moved in archive/, not maintained
- Namespace layout stabilized : no change
  - Still a few legacy templates in os/ to be moved to conf/
- Kernel version/architecture moved to standard/os
- Kernel version referenced through a variable in kernel version and kernel module names
  - Allow easy and consistent selection of a version different from base version
    - Based on KERNEL\_VERSION\_NUM and KERNEL\_VARIANT variables
    - Kernel RPM name can be defined explicitly (PKG\_KERNEL\_RPM\_NAME)
  - SL 4.4 and 4.5 templates regenerated
  - Kernel RPM no longer part of errata list (to be done)

- gLite 3.0 in maintenance mode, except for WMS
  - WMS support is very late but almost working now
  - Several components added to replace ncm-glite
    - ncm-glite not used anymore except in VOMS server templates
    - Includes **general purpose ncm-mysql**
- gLite 3.1 is the main branch
  - Significantly diverged from 3.0
  - RPM templates restructured to support multiple OS arch
    - New layout is rpms/config.tpl + rpms/*arch*/config.tpl
  - All released node types except VOBOX, VOMS, FTS, AM
    - VOBOX in progress, almost working...
  - Several components upgraded to support gLite flavour
  - Non-shared home directories not working (PBS job ma
- Many new features
  - E.g. Torque and MAUI configuration through variables
  - Pool accounts for VO groups/role (not used by default)

- All gLite updates
  - Generally released to trunk 2-3 days after official release
  - An area for better sharing the effort ?
- VO configuration generated from CIC portal VO capabilities
  - Done by Cédric, being tested at GRIF, almost ready
  - Generates params/ and certs/ templates for all known VOs
    - Backward compatible configuration of pool accounts
  - Implemented as an Ant task : `update.vo.config`
  - No change to VO configuration process
    - Still possible to configure a VO manually
    - Existing changes through `VO_SITE_PARAMS` will be retained
  - More flexibility for local customization
    - Pool account prefix, base uids through specific templates
- Various MPI support improvements
  - 64-bit MPI support coming

- Quattor core templates imported from CVS
- All being upgraded to version 2
  - Use of QWG variables should ease the migration
- Hardware templates : still a few not namespaced
  - Done in the background, not a priority
- Monitoring
  - Significant cleanup in LEMON related templates
    - Namespaced
    - Merged between GRIF, TCD and Brussel
    - Not sure it has been documented...
  - NAGIOS : Fred (IRFU) has a configuration for the server
    - Probably not completely generic
    - I asked him several times to commit them as a starting point... v do it again
    - Who is interested to participate ?



- User management : template to create a set of users on a large number of machines
  - Cross-machine list of users ensuring consistent characteristics for one given user
  - Ability to select on a per-machine basis the subset of users to create
    - Subsets defined as groups
    - Selection criteria may be regexp
  - Management of public SSH keys (ncm-useraccess)
- Support for Xen VMs
  - Done by Stephen
  - Currently only SL 4.5
  - Documented in OS section

- Already there for a long time for gLite
- 1 distinct errata namespace for each gLite update
  - Namespace : update/nn, where nn is gLite update number
  - gLite update to deploy can be selected on a per node or per cluster basis
    - One default update defined in each QWG release
  - 1 template included at end of configuration and doing `pkg_only` (don't add, only replace)
    - `rpmUpdates.pl` produces a template with an entry for the last version of any RPM found in errata RPM repository
  - 1 template for early actions like variable definition
    - `init.tpl`
- 1 additional update : update/unsupported
  - Included in addition to others, if it exists
  - Included after official update : testing of new version

- Variable `PKG_DEPLOY_OS_ERRATA` to control deployment
  - `errata.tpl` part of standard OS templates and updated regularly
  - Configured by `config/os/updates.tpl` : also includes `rpms/update`
- Variable `PKG_OS_ERRATA_TEMPLATE` to use an alternate template name
  - May allow some site « versioning »
- A cron for unattended download of new errata RPM
  - Almost ready, some polishing still required
- Kernel errata : keep already installed kernel on a machine even if no longer in node configuration
  - Require a change in SPMA (in progress)
  - Will be activated by an option in `/etc/spma.conf`
    - List of regexp matching package names
    - Could be used for other specific packages, like java

# Planned changes

- Pre-compiled configuration
  - Dummy WN integration (Stijn) : 1 WN XML profile included into actual WNs
  - VO configuration (using a specific schema)
  - Need integration into SCDB to ensure they are built first
- Reengineering of pool accounts management ?
  - How to add more roles without breaking compatibility ?
  - Need to keep UID/GID consistency across several machines and permissions on existing files/directories
  - Use LDAP or other account/pwd distribution mechanism ?
- Complete MAUI configuration based on variables
  - Fairshare scheduling
  - Short deadline job support

- Compatibility, compatibility, compatibility...
  - My personal obsession...!!!
  - New and changed features should not required any change to an existing configuration if not using it
  - A success for 2 years...
- QWG releases described in LCGQWG site roadmap
  - 1 milestone per release containing a description of changes
  - Next release milestone contains a description of planned changes and work already completed
    - Associated tickets to describe planned/done changes (not intended for support)
- QWG releases announced on LCGQWG mailing list
  - Tagged release also announced on Quattor mailing list
  - Significant improvements in a branch announced as pre-release (or beta)

- 1 QWG branch per gLite version + 1 branch no-grid
  - Made of MW, OS and other standard templates
    - Each branch contains working examples
    - Lack of time to update no-grid : better to use gLite 3.1...
  - Regularly updated with fixes, new features and gLite updates
    - Revisions in branch consist of merge from trunk
    - Try to update branch with gLite/CA updates before 1 week
- QWG (tagged) releases
  - A stable snapshot : always out-of-date...
    - Currently 1 release every 4 months
    - Would site prefer more frequent tagged releases ?
  - Each release tag based on MW version : currently 3.1
    - QWG release number doesn't match gLite update numbers
- 'trunk' is development branch
  - Don't expect it to work at each revision
  - Contributors need to validate against examples
    - `tools/check-compile.sh`

Follow copies  
 Show only adds, moves and deletes  
 Show full log messages

Update

Added   
  Modified   
  Copied or renamed

View changes

		Rev	Chgset	Date	Author	Log Message
<input type="radio"/>	<input checked="" type="radio"/>	<a href="#">@1461</a>	<a href="#">[1461]</a>	03/10/07 10:54:53	jouvin	Merge changes from trunk:1432-1459, mainly : * Update ncm-pbsclient; Add ...
<input checked="" type="radio"/>	<input type="radio"/>	<a href="#">@1460</a>	<a href="#">[1460]</a>	03/10/07 10:47:11	jouvin	Revert previous mod
<input type="radio"/>	<input type="radio"/>	<a href="#">@1446</a>	<a href="#">[1446]</a>	03/06/07 16:14:30	sweirdt	iihe: preparations for new QWG tag v2
<input type="radio"/>	<input type="radio"/>	<a href="#">@1433</a>	<a href="#">[1433]</a>	03/05/07 17:54:53	jouvin	Merge changes from trunk:1428-1432, mainly : * Update ncm-pbsclient; Add ...
<input type="radio"/>	<input type="radio"/>	<a href="#">@1429</a>	<a href="#">[1429]</a>	03/04/07 23:29:53	jouvin	Merge changes from trunk:1426-1428, mainly : \ * Add ncm-accounts as a pre ...
<input type="radio"/>	<input type="radio"/>	<a href="#">@1427</a>	<a href="#">[1427]</a>	03/04/07 23:10:32	jouvin	Merge changes from trunk:1424-1426, mainly : \ * Add ncm-accounts as a pre ...
<input type="radio"/>	<input type="radio"/>	<a href="#">@1425</a>	<a href="#">[1425]</a>	03/04/07 19:14:12	jouvin	Merge changes from trunk:1421-1424, mainly : \ * Miscellaneous fixes for SL ...
<input type="radio"/>	<input type="radio"/>	<a href="#">@1422</a>	<a href="#">[1422]</a>	03/04/07 18:29:14	jouvin	Merge changes from trunk:1418-1421, mainly : * Rename VO param ...
<input type="radio"/>	<input type="radio"/>	<a href="#">@1419</a>	<a href="#">[1419]</a>	03/04/07 16:27:36	jouvin	Merge changes from trunk:1336:1418, mainly : <a href="#">Main changes</a> : * Add ...
<input type="radio"/>	<input type="radio"/>	<a href="#">@1337</a>	<a href="#">[1337]</a>	02/16/07 15:39:53	jouvin	Merge changes from trunk:1332-1336, mainly : * Fix CE_OS_xxx variable ...
<input type="radio"/>	<input type="radio"/>	<a href="#">@1334</a>	<a href="#">[1334]</a>	02/16/07 15:16:29	jouvin	Merge changes from trunk:1318:1332, mainly : * Remove gLite specific ...

## Changeset 1419

**Timestamp:** 03/04/07 16:27:36 (1 week ago)

**Author:** [jouv](#)

**Message:** Merge changes from trunk:1336:1418, mainly :

### Main changes :

- Add gLite update 14 and 15 ([r1370](#))
- Support for a Torque cluster with worker nodes in a private subnet (Torque v2, variable CE\_PRIV\_HOST) ([r1343](#))
- Torque v2 client : directPaths entries are generated from WN\_SHARED\_AREAS ([r1344](#), [r1377](#), [r1415](#))
- VO Configuration : allow VOS\_SITE\_PARAMS[some\_vo] to be a nlist with direct values instead of a template name ([r1375](#))
- VO Configuration : support for locking/unlocking accounts ([r1376](#))
- RPM errata list for SL 3.07 and SL 3.08 ([r1368](#), [r1372](#))

### Miscellaneous changes and fixes :

- Torque v2 client : miscellaneous fixes ([r1345](#), [r1346](#), [r1348](#))

[r1350](#) | [eknoops](#) | 2007-02-20 10:36:22 +0100 (Tue, 20 Feb 2007) | 1 line

- VOMS server certificate updates (FNAL, ESR) ([r1361](#), [r1362](#))
- Fix JAVA\_LOCATION for SL 4.4 ([r1362](#))
- Add missing RPMs for VOBOX required by gLite update 13 ([r1369](#))
- Update VO params for vo.dapnia.cea.fr ([r1371](#))
- Various fixes to vo/functions.tpl ([r1374](#), [r1414](#))
- Define Quattor client RPMs in a template in standard/ rather than per OS and upgrade ccm to support panc v7 ([r1416](#), [r1417](#))

View differences    
Show  lines around each change  
Ignore:  
 Blank lines  
 Case changes  
 White space changes



- Contributions are welcome... and increasing
  - Several regular non LAL contributors
  - Contribution must be done directly in QWG repository
    - Patches sent by email are too much a work to integrate
    - Ask me for an account to get write access to QWG repository
  - A contribution can range from new HW support to more sophisticated changes : no need to be a certified exper
- All contributions MUST be made to the trunk
  - Trunk is a log of work in progress, no guarantee to work at every revision
    - Try to use Trac wiki syntax in your commit messages
  - Each changes are reviewed before being moved to a release branch by the "release manager" (me currently)
    - Expect an email if I have difficulty with your proposed change...
- Tools and procedures documented in 'Developers' Corner' on the wiki
  - Important to follow them to avoid "false differences"

- QWG site : <http://trac.lal.in2p3.fr/LCGQWG>
  - Description of QWG templates (OS and MW)
    - Also SCDB and PAN language
  - 3 areas : documentation, release notes (about release download and installation instructions)
  - Actively updated : in fact often in advance with releases...!!!
  - **Send remarks when you find something wrong or miss**
  - Contributors are welcome... send me your certificate D
- Mailing list : mainly QWG mailing list
  - Announcement of new releases (not beta) also sent to Quattor mailing list
- Support through Savannah ELFms
  - Trac tickets not used for QWG support but for development planning

- QWG templates succeeded to be a flexible, off-the-shelf set of templates for configuring OS and MW
  - Need to maintain more actively non-grid branch
- Enlarging contribution is one of the main challenges for the future
  - SVN leverage the possibility for wide contribution
  - Every one can contribute its local modifications and improvements, not only experts !
  - Providing templates for more OS versions may be shared easily : please commit yours if you run non SL distributions
  - gLite update integration is an other easy-to-share effort
    - How to avoid duplication ?
  - At some point, release manager role may be shared between several people