

QWG Templates Update

Michel Jouvin
LAL, Orsay

jouvin@lal.in2p3.fr

November, 2011

Quattor Workshop, Strasbourg

Outline

- ⦿ Last releases and their new features
- ⦿ Release cycle
- ⦿ New ncm-spma
- ⦿ Pakiti integration
- ⦿ SCDB

Last Releases

- ◎ No new release since last workshop
 - › 1 pending since last Spring but no time to make a formal release
 - › A new release manager is needed...
- ◎ Branches updated actively
 - › Several updates to the branches : ~ once a month until the summer, no update since (but a lot of things in the trunk...)
 - › Can be considered as production quality: doesn't exclude minor problems, fixed quickly
 - › Branch updates announced on the quattor-grid list

Release Cycle

- Until now, it has been difficult to have more than 2 release a year
 - › Some overhead to produce the appropriate change log, release notes...
- No real impact on sites as the branch is considered the production release
 - › Significant testing before merging trunk changes into a branch
- Releases are important for visibility and should ideally be used for production
 - › Current goal: 1 release every 2 sprints (~3 months)
 - › Must share/transfer the release manager role

Main Changes

- ◎ For details, see 3.2 branch log
 - > https://trac.lal.in2p3.fr/Quattor/log/QWG/templates/branches/gLite-3.2?action=stop_on_copy&mode=stop_on_copy&rev=5771&stop_rev=5629&limit=100&verbose=on
- ◎ Most important changes and new features
 - > EMI-1 support started
 - > gLite : ARGUS/glexec
 - > CERNVM FS support
 - > New monitoring templates based on NIKHEF ones

Pakiti integration

- ⦿ Full coverage of all nodes managed by Quattor
 - › Pakiti client installed/configured on every machine as a RPM
 - › Installed as cron reporting to the server configured
 - › Available in last branch versions as part of core templates
- ⦿ Pakiti web interface allows easy identification of nodes without Pakiti or not reporting properly
 - › Either absent from node list or with a last update date too old
- ⦿ Pakiti server configuration included

Pending Issues / Contributors

- Not much progress on most of the pending issues
 - › Lack of time
 - › Most issues tracked in Trac backlog and tickets
- Most issues specific to gLite MW support...
 - › New machine types mainly done by Jerome since 1 year
- ... but a few ones affecting core/OS templates
 - › Structure of OS template namespaces
 - › Handling of package versions
 - › Non-gLite NFS server/client configuration
 - › Package number reduction
- Several new contributors to the trunk

gLite Templates ToDo

- ◉ Part of Quattor backlog
 - > <https://trac.lal.in2p3.fr/Quattor/wiki/Development/Scrum>
 - > Actions in progress in sprint backlog
- ◉ New machine types for gLite 3.2 and EMI-1
 - > 3.2 : Mainly VOMS (soon)
 - > EMI-1: CREAM, WN, UI, DPM, WMS
- ◉ Support YAIM to configure gLite services
 - > Need to find a site interested: NIKHEF
- ◉ New handling of packages based on (improved) default value definitions

ncm-spma v2

- ⦿ Major schema change to better support multiple architectures
 - › Version is now per architecture
- ⦿ Chicken and egg problem: 2-phase upgrade
 - › 1.8: deploy new functions with old schema
 - Able to deal with new schema
 - › Deploy 2.0 when all machines have picked up 1.8
 - Else they won't be able to upgrade...
- ⦿ Should become the default in the next QWG release... but it is late!

ncm-accounts Nightmare

- ◎ ncm-accounts v4 (currently used in QWG) rewrite based on the standard Linux tools failed (was v5)
 - › Too many bugs to workaround in users
 - › Major performance problems
- ◎ Luis wrote ncm-accounts v6 which builds new passwd/group in-memory and overwrites old ones
 - › Very efficient: 10K accounts in 10 mn or less!
 - › Only possible issue: interference with a manual mod done at the same time (no locking)
 - › Unfortunately a couple of bugs prevent using it in production...
 - › 2 weeks ago contact reestablished with Luis who promised to help...

perl-CAF/perl-LC

- QWG still using old versions
- Possible non backward-compatible changes involved in new perl-LC
 - › Most components have been updated already
 - › The main missing piece is ncm-accounts: v4 doesn't work with new perl-LC
- As soon as ncm-accounts v6 is ready, QWG will switch to new versions
 - › Tests already done on test machines: all known issues already fixed

Other Related Tools

- ◉ No real plan to do it yet... lack of manpower
- ◉ A script to check a RPM list is satisfied by an existing (XML) profile
 - › RPM list retrieve from YUM grouplist with all dependencies
 - › Check that all required packages are present with a version \geq to YUM requirements
 - › Different purpose than checkdeps
- ◉ Redesign/rewrite directory-sync ?
 - › Support non-Linux OS, don't rely on a complex 'find'
 - › A python script created but doing nothing...

Documentation

- ◉ <https://trac.lal.in2p3.fr/QWG>
- ◉ Probably still too much a personal effort...
 - > A few contributions in the last months
 - > Feel free to contribute!
- ◉ Review of existing/missing part required done by Andrea 1 year ago
- ◉ Main action done: merge of contents of former MediaWiki and Trac with some reorganization
 - > No real feedback since...

Open Questions

- ⦿ Have clearly identified QWG subsets
 - › Almost the case already: OS, gLite, Monitoring, All...
 - › How to package these subsets?
- ⦿ No-grid templates effort didn't succeed
 - › Work in progress, mainly by Christos, but no time to complete it in the last 2 years
 - Web server, smtp server...
 - Mainly SL4-based?
 - › Do we really need such a release or properly identified (packaged?) subsets are enough?

SCDB

- ◎ Main change: VOConfigTask modified to support new VO card format (Jerome)
 - > Not yet released in a production version but easy to get from trunk
 - > Used to generate VO-related templates from CIC portal
 - > No known issues

Backup Slides

New Handling of Packages

- ⦿ Current implementation has drawbacks
 - › Need to reapply OS updates every time an OS-provided RPM is added to the configuration
 - › Lot of pkg_repl for the same packages: very inefficient
- ⦿ New approach based on default versions
 - › Define default version of each possible package in a OS/gLite version
 - › Define the updated default versions for each update/errata level
 - › Very efficient as it involves only a simple nlist entry replacement
- ⦿ For each service, define the list of packages to use without an explicit version
 - › Pkg added in the configuration without a version
 - › Version resolved at the end of the configuration process
 - Just before resolving repositories

Number of Packages

- ⦿ Reducing the number of packages has been on the todo list for quite a long time and still no progress
 - SPMA/rpm issue because of too many open files
 - Security issue with unused vulnerable packages installed
 - Promising work in progress by James: status?
- ⦿ Short term: pkg_del() everything added to the configuration but not useful and reported as vulnerable by Pakiti
- ⦿ Medium term: rebuild RPM lists based on core group only