QWG Templates Update

Michel Jouvin
LAL, Orsay
jouvin@lal.in2p3.fr
http://grif.fr

October 12, 2010
Quattor Workshop, RAL

Outline

- Last releases and their new features
- Pakiti integration
- New handling of package versions
- Reducing number of packages
- QWG development process

Usage Status Changes

- No news from site in Israel
- Romania: 1 site started but need to establish contact
- Contact lost in Africa
- One site in Beijing pretty active
  - A meeting in the future?!

Last Releases

- Releases announces at last workshop done in July…
  - gLite 3.1.0-8: support for multiple CEs per cluster, CREAM 1.5
  - gLite 3.2.0-3: support for multiple CEs per cluster, CREAM 1.6
- gLite 3.2 is the main release
  - Some service configuration actively maintained only for 3.2 (e.g. CREAM CE)
  - Most thing still in sync
    - When they are identical
- gLite 3.2 still requires gLite 3.1
  - Several service nodes not yet available in 3.2: WMS, MyProxy
  - LCG CE will never be ported to 3.2
    - May be deprecated in the coming year
- Branches updated more actively

Main Changes

- Most of them described in my QWG Update during Thessaloniki workshop
  - http://indico.cern.ch/materialDisplay.py?contribId=5&sessionId=1&materialId=slides&confId=84432
- Other important changes/additions
  - SL 5.5 added recently added to the branches
    - 64-bit only
    - gLite 3.2 only: already tested for several machine types
  - Templates for OS errata produced after each major vulnerability
    - Should be done every month as GRIF producing them
    - Can also be done by any site at any time
- gLite 3.1 in maintenance mode… but benefit from most gLite 3.2 developments
  - CREAM CE available in 3.1 but the version to use is 3.2

CREAM CE Support

- 3.2 recommended and only version maintained
- No new machine type: machine-types/ce
  - Use *CE_HOSTS_CREAM* to determine CE variant
- New CREAM-specific variables for CREAM-specific features
  - Documented on the wiki
  - Eg. sandbox configuration
- Support NFS-sharing of CE sandbox area
- Documentation updated
  - https://trac.lal.in2p3.fr/Quattor/wiki/Doc/gLite/TemplateCustomization#CEConfiguration

Multiple CE Support...

- Was necessary for CREAM CE support
  - WNs generally shared with LCG CE
- Main configuration variable: *CE_HOSTS* (list)
  - Replaces (and defaults to) *CE_HOST* (single value)
- Configures a shared gridmapdir between CEs
  - *GRIDMAPDIR_SHARED_SERVER* allows to specifiy the server. Defaults: 1st CE in *CE_HOSTS*
  - Default sharing protocol: NFS
  - Be sure to configure the same VO list on all Ces (to be fixed)
- Torque/MAUI: configures GIP plugins in cache mode
  - Plugins run on the Torque/MAUI server and produce a file with the relevant information
  - GIP plugins on CEs do a `cat` of the cache file
    - By default located in dteam SW area

... Multiple CE Support

- Torque/MAUI server could in principle not be a CE but currently untested and a few oddities expected with publishing into BDII
  - GIP configured only on a CE
- Good contact with CREAM developpers who send us early information about new releases...
  - Interested as the official process is very long...

OS Errata Management

- Used several times in many sites in the last 6 months and working pretty well
  - New ‘protectkernel’ option is SPMA is a must
  - A chicken and egg problem with openAFS in some circumstances
- Templates for errata easily produces with rpmErrata.pl script
  - No need to wait for GRIF to produce them!
  - Templates produces are completely generic
- See presentation at last workshop for the details
- Probably the main reason for using QWG OS templates even if not using the gLite templates
  - Trying to reimplement this in a site-specific context is not worthwhile, better to spend time using QWG OS templates!

Profile Cloning

- Formerly known as « dummy WN »
  - Currently works only on WN
- Reworked with new variable names starting with *PROFILE_CLONING_*
  - Old variables *(DUMMY_)* still accepted (backward compatibility)
- A complete documentation is available
  - https://trac.lal.in2p3.fr/QWG/wiki/Doc/gLite/WNCloning
  - All variables documented
  - FAQ documenting main known issues
- Most customization points unavailable, loadpath not defined
  - Would reduce dramatically the performance gain

gLite Templates ToDo

- https://trac.lal.in2p3.fr/QWG/milestone/QWG-Templates-ToDoList
- New machine types for gLite 3.2
  - Waiting for their official release
  - FTS: need RAL to do it as it is the only one to run it
- Review account management using LDAP and a more flexible allocation of UIDs
  - Current allocation of UIDs is a nightmare
  - A temporary approach could be to parse a XML file to retrieve allocated UIDs in update.vo.config
- Integration of new LRMS: LSF, SLURM, SGE
  - Need a site using them
  - Plans with SGE in South Africa didn't materialize

Other Items ToDo

- No real plan to do it yet… lack of manpower
- 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 personnal effort…
  - A few contributions in the last months
  - Feel free to contribute!
- Review of existing/missing part required done by Andrea
  - Action from monthly meetings
  - No attempt to implement any change yet: need some sort of roadmap for documentation too
  - Merge between former MediaWiki content and Trac still to be done
    - E.g. 2 sections for developpers
    - More generally “Quattor Toolkit” section content need to be reshuffled

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
- 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 still to be done

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 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
  - SPMA/rpm issue because of too many open files
  - Security issue with unused vulnerable packages installed
- 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

Enabling Wider Usage

- Have clearly identified subsets
  - Almost the case already: OS, gLite, Monitoring, AII…
  - How to package these subsets?
- gLite: allow the use of YAIM rather than QWG for the grid service configuration part
  - Use QWG to configure base OS
  - Use QWG to configure grid service RPM list
  - Swith to yaim.tpl instead of service.tpl
  - Need a site interested
- No-grid templates
  - Work in progress, mainly by Christos
    - Web server, smtp server…
    - Mainly SL4-based?
  - Not able to make any release…

QWG Contributions

- New contributors to templates
  - CREAM CE done by J. Pansanel (Strasbourg)
  - 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

Development Process…

- Very few discussions about the main priorities for QWG
  - Most of them established by me… based on the feedback I receive!
  - Contributions from others are mainly their local developments to solve problems faced at site
- No real team to implement current todo list and ideas for future enhancement
  - Development ideas generally presented at workshops
  - Most of them also have a ticket in Trac
- QWG would benefit from a more collective and structured handling of the todo list
  - Probably possible today as we have more contributors
  - May help others to join

... Development Process

- Agile-like (agile-light) process?
  - Product backlog: we have it!
  - Sprint: next development cycle
    - Typically less than a month, could be 2 in our case
    - An initial meeting to decide which part of the backlog we process
    - Normally a short daily meeting, could be a weekly meeting in our case, to review progress.
  - Sprint release: could be implemented as a branch update
- StratusLab using this process
  - MEB is an evangelist...
  - May get some feedback on how to usefully adapt the process to our needs
- (short) Meetings dedicated to QWG, distinct from monthly meetings
  - Must be frequent enough to create some dynamics

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 released
  - Announcement done at every major change in the branch
  - Significant testing before merging trunk changes into a branch
  - Branch development pretty active since last workshop
- Release are important for visibility
  - Time-based released?
  - A real process behind QWG development may help to release what has been done
  - A more collective involvement may also help to share the release manager role