Quattor WG Status Report

Michel Jouvin
LAL/Orsay
jouvin@lal.in2p3.fr

GDB, Prague
April 4th, 2007

- History : Why QWG ?
  - QWG mandate
- Current Usage
- Achievements
  - Templates
  - SCDB
  - HTTPrep
  - Documentation
- Future
- Conclusion

- Machine configuration described in term of final sta
  - No description of how to reach the state (actions)
- Abstract, service oriented description of machine configuration
  - Written in PAN language and compiled
    - PAN is a declarative-like language
  - Organized in “templates” that can be heavily factorized
  - Configuration consistency checked at compile time
    - Information schema and validation functions
- Configuration stored in a versioned central databas
  - Allow very easy rollback of any change.
- Configuration implemented through “components”
  - Components are plugins running on client : 100+ availab
  - Extensible framework : easy to add new components
- Manage both installation and configuration changes

Why Quattor for LCG ?

- EGEE/LCG MW installation tool : YAIM
  - In fact arrived after Quattor first templates for grid…
  - Shell scripts configured through environment variables
  - Initial installation only, reconfiguration difficult
    - Not adapted to the need to frequently deploy patches
  - Machine oriented rather than service oriented
    - Assume a machine dedicated to a service
- Quattor : service oriented management tool
  - Quattor developped by EDG WP4
  - Installation AND reconfiguration
  - PAN templates are easy to organize by service
  - Easy mix of service on a machine
    - As long as MW supports the mix (port conflict…)
    - No side effect of one service configuration on another service
  - Configuration consistency of common low level service
    - Accounts, cron, NFS mounts…

Why Quattor Working Group

- Official working group of LCG GDB started in 2004
  - Mandate : ease and coordinate use of Quattor in grid context
    - http://lcg.web.cern.ch/LCG/Boards/GDB/quattor-mandate.html
  - Not a task force with significant manpower... Never been more than spare time from a few people !
  - Chairman : Michel Jouvin took over responsibility from C Loomis mid 2006
- Main responsibility : PAN templates and NCM components for EGEE/LCG MW configuration and deployment
- Contribution to maintenance and development of Quattor
  - SCDB (Subversion based CDB) as an alternative CDB implementation
- Probably never reported to GDB but very active...

Usage of Quattor

- Quattor usage increased over last years
  - 40+ sites, mainly in Europe
    - Several group of sites managed from a unique database
  - Some countries “quattorized”
    - France, Spain, Ireland, Belgium
  - Still very few usage outside LCG : Ireland
    - Quattor sites are not dedicated to HEP
- Wide range of site size
  - CERN : ~5000 nodes
  - Several T1s : NIKHEF, CNAF, PIC
  - Many T2s and T3s, including large ones (GRIF, DESY) and small ones (<50 nodes)
- Main usage is management of grid resources...
  - Starting point to use Quattor
- But growing non grid usage at these sites
  - Site internal servers, desktops, virtual machines (XEN)

Usage Status / QWG Templates

- Usage increasing : some sites with 200+ nodes
  - GRIF : 5 sites, 200 machines, grid + non grid usage, 1 S
    - +150 machines to arrive in June
  - 6 french T2/T3s : 1 SCDB per site, 10-50 nodes/site
  - BEGRID : 3 prod sites, 200 nodes, 1 SCDB (+ 1 testing s
  - Irish grid : 15+ sites, 350 nodes, 1 CVS SCDB
  - UAM : 700 nodes, 3 clusters, 3 SCDB
  - Valencia : QWG, SCDB, 175 nodes
  - Aachen : 1 SCDB, 35 WNs, 7 dCache SEs, 130 UIs
  - Protvino (Moscow region) : still LCG 2.7.0…
- 1 site migrated off QWG templates : NIKHEF (2006)
  - First generation of templates not generic enough
- Several large sites actively considering switching
  - DESY, PIC
- QWG sites generally switched to SCDB

Machine Profile Template

- Goal : allow site customizations from a central template without editing standard templates
- Example 1 : WN

```
object template profile_grid22;

include machine-types/wn;

include repository_common;
```

- Example 2 : CE

```
object template profile_grid10;

#variable CE_STATUS = 'Closed';
#variable CE_STATUS = 'Draining';

variable CE_CONFIG_SITE = "pro_ce_torque_grif";
variable BDII_SUBSITE = "LAL";

include machine-types/ce;
```

Machine profile template

- 1 file (template) containing all site customization
  - Can be refined in a specific (set of) machine profile
- Example

```
variable CE_HOST ?= "grid10."+SITE_DOMAIN;

variable SE_HOSTS ?=
  list('grid05.'+SITE_DOMAIN,
       'grid03.'+SITE_DOMAIN);
variable SE_TYPES ?= nlist(SE_HOSTS[0], "SE_dpm",
                           SE_HOSTS[1], 'SE_dpm');
variable CE_CLOSE_SE_LIST ?= SE_HOSTS[0];

variable LFC_HOST   ?= "grid14."+SITE_DOMAIN;
variable RB_HOST    ?= "grid09."+SITE_DOMAIN;
variable PX_HOST    ?= "grid02."+SITE_DOMAIN;
variable BDII_HOST ?= "grid01."+SITE_DOMAIN;
variable MON_HOST   ?= "grid08."+SITE_DOMAIN;

variable VOS ?= list('alice',
                     'atlas',
                     'auger',
                     'biomed',
                     'calice',
                     'cms',
);
```

Implementation : templates

- 2 main sets of templates : OS and gLite
  - OS : 99% generated from RH/SL comps.xml
  - gLite : rpm list generated from gLite description
- gLite templates are really generic and pretty flexibl
  - Complete re-engineering of 1st generation (LCG templat
  - Now in use in different contexts
  - Upgrading template releases very easy
    - Site config clearly separated from standard templates (OS and MW
  - Cover all supported services
    - dCache available, gWMS/gCE not very well tested yet
- Explicit release cycle with a roadmap : 10 in 9 mon
  - gLite updates delivered as part of new template releases
    - First delivered a week after official release in pre-releases
  - Ability to select gLite update to install on a per node bas and independently of template releases
  - Backward compatibility is an (personal) obsession

Introduction : Templates

- Added value by QWG templates compared to YAIM
  - SL4 64-bit support on WNs since spring 2006 (LCG 2.7.0
  - Fix for DPM publication script using virtual GIDs
  - NFS configuration for shared VO areas or home dirs
    - Support many conf scheme, including dedicated NFS server
  - Very simple VO configuration : VO added to a list...
  - BDII off CE since the beginning
  - Torque/MAUI v2 for several years
    - Including configuration examples for MPI/Short Deadline Jobs support
  - Torque submit filter customization
  - Gsissh accessible public UI
  - Non dedicated UIs for deployment on desktops or internal servers
  - Ability to close/drain a CE or selected WNs
  - Ability to restart all or selected WN LRMS clients
  - Default/close SE per VO

- SCDB : plug and play replacement for original CDB
  - Mandate : identify changes in the Quattor tools which wi simplify the configuration
  - Secondary goal : focus Quattor development on core : PAN compiler + components
    - Rely on standard products for other parts
- Subversion based instead of CVS based
  - No specific layer to maintain (CDB SOAP layer)
  - File system layout of templates matching Quattor view
- Client part (+panc) written in Java : administration possible on every platform (Unix, Windows, MacOS)
  - Administration using “ant” tasks : compile, deploy
    - 1000 LOC written in Java
  - GUI available : Eclipse
    - Nothing specific needed to manage Quattor
    - PAN syntax highlighting
  - Take advantage of CPU available on admins desktops

- HTTPrep : plug and play replacement for SWrep
  - Mandate : identify changes in the Quattor tools which wi simplify the configuration
  - Secondary goal : focus Quattor development on core : PAN compiler + components
    - Rely on standard products for other parts
- Rely on a plain HTTP server
  - http is the only access protocol
  - Software repositories implemented as directories on the server
    - Require access to the server for updating RPMs
  - Access controlled by HTTPD + file system permissions
- Administration : update of repository templates reflecting repository contents
  - Done via an “ant” task : update.rep.templates
  - Can be done from any machine
- Not a requirement for using SCDB

- 1 of the main deficiency of Quattor at QWG creation
- New wiki-based LCG QWG site started March 06
  - http://trac.lal.in2p3.fr/LCGQWG
  - Everybody is welcome to contribute
  - Covers QWG templates, SCDB, new PAN compiler
  - Contains roadmap of releases for each product
- Documentation now pretty complete
  - Includes installation guides and release notes
  - Used by many new sites : improved accuracy !
  - Try to keep it in sync (in fact in advance) with new features
    - Documentation must be written for a new feature to be released
  - Rework of main Quattor web site is in progress
- Release contents documented and accessible through the roadmap

- Quattor Workshop every 6 months since March 06
  - Mandate : Coordinate with the Quattor developers to ensure a smooth evolution of configuration components
  - Bring together users and developpers: 20-30 participant
  - March 06 (LAL), Oct 06 (DESY), March 07 (Dublin), Oct 07 (Madrid)
  - Improved a lot synergy between developpers and avoided risk of split (e.g. SCDB vs. CDB)
- Participation to several tutorials and other dissemination initiatives
  - 1 day tutorial during WLCG T2 workshop (June 06)
- Enlarge contributions : now ~5 people
  - Everything in a Subversion directory
  - Separate areas for contributions and “stable” branches
  - Write access given on request

Future Plans

- Increase the number of supported grid services
  - Polishing and testing of gWMS/gCE
  - Support for AMGA, HYDRA
- Complete cleaning of non gLite templates
  - Move to namespace for easier administration
- Integration of monitoring configuration
  - Release generic Lemon templates
  - Provide template examples for configuring Nagios
- Support for virtual machines
  - Share hardware between VMs dedicated to different serv
- More active support for non grid usage
  - Growing interest from Quattor sites to manage non grid systems, including desktops and virtual machines
- Dissemination inside and outside LCG and grids
  - Talk planned to next LISA conference

- Abandon of gLite configuration system
  - Core component : description of all configuration options a service in a XML file
    - Maintenance of XML file was the responsibility of service developers
  - Allowed entirely generated templates for all pure gLite services based on XML configuration
  - Very bad decision : no more official source of config opti
    - Move the responsibility to GD who doesn't have the knowledge and manpower
    - YAIM is the only and incomplete source… Back to LCG !
    - YAIM could benefit from this official source about config options
    - May GDB ask about reconsidering this decision to drop options description and ask MW developers to maintain it ?
- New service oriented YAIM as an alternative ?
  - YAIM implementes a limited subset of configuration optio
  - YAIM is an install tool : no fine-grained control of service
  - If not, may raise the need for new components…
    - Hopefully most of what is needed is already there

Conclusion

- QWG did a great work and fulfill its mandate
  - Providing generic templates and improving tools
  - Building a Quattor community to share the maintenance and development effort
    - New contributors on a regular basis : ~5
- QWG effort is critical to leverage Quattor value
  - Allows out-of-the-box usage of Quattor for grid
  - Hides new/small sites the complexity of setting up and maintaining a site
- Quattor is now a community-backed product
  - Main developers and expertise outside CERN
  - Remains a critical product for CERN : never been a problem

- New sites are welcomed...!!!
  - Community ready to help/support them
  - Lot of experience with setting up a new site or migrating to QWG templates and/or SCDB
- Still a lot to achieve...
- QWG asks for mandate renewal
  - Better coordination and increased visibility of the effort
  - Mandate adjustment ? Not sure worth the effort...
    - Could add dissemination and support effort

Documentation

- QWG site : http://trac.lal.in2p3.fr/LCGQWG
  - Description of QWG templates (OS and MW)
  - Description of SCDB
  - Description of PAN language
  - Actively updated
- Quattor main site : http://quattor.org
  - Mainly basic architecture and NCM components
  - Less actively updated…
  - Major rework planned for more attractive and up to date information
- Tutorials (pretty old materials…) :
  - English : http://quattor.web.cern.ch/quattor/documentation/tutorials/