

# SCDB Update

Michel Jouvin

LAL, Orsay

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

<http://grif.fr>

November 4, 2009

Quattor Workshop, Bruxelles



# Outline

- Last SCDB releases
- SCDB 2.3 and its new features
- SCDB 3.0 and beyond

# SCDB

- SCDB is a set of tools to administer a site with Quattor
  - Some tools common with other CDBs: panc, checkdeps
    - Imported from SF SVN repository
  - Ant interface to configuration database
  - Some general-purpose utilities (scripts) in utils/
  - Source of server-side scripts in src/
- Interdependencies mainly between ant config file and externals/
  - In particular panc, scdb-ant-utils, svnkit
  - svnkit version must be synchronised with SVN client version used on the same machines
  - quattor.build.xml and scdb-ant-utils must be synchronised when task options are modified

# SCDB Releases

- 2.3: originally planned end of March but released mid-July...
  - Followed by 2.3.1 (cluster groups) and 2.3.2 (fixed update.vo.config)
  - Last version to support svnkit 1.1 (SVN client 1.4)
- 3.0: currently in beta
  - Main focus: SVN 1.5/1.6 support, including new deployment scripts

# SCDB 2.3 Main Changes

- Panc 8.2.8: add support for new debug(), fix dependencies with profile cloning
- scdb-ant-tools 7.4.5 :
  - Fixed update.vo.config to generate VO config parameters from CIC portal
  - Fixes for profiles-info.xml generation: new profiles not added during first deploy
  - Ability to ignore templates in dependencies evaluation
- checkdeps.py: fixes, imported from SF SVN
- Quatview updates
- Support for sudo rather than ssh between SVN and Quattor merged into post-commit
  - Used to be a separate script in 2.2
- Support for cluster groups
- Support for quattor.build.properties in the repository
- src/utils/ renamed utils/ : a symlink still exists

# Cluster groups

- Allow to group clusters in several subsets
  - They used to be all under `cfg/clusters/`
  - A cluster group is a directory under `cfg/clusters/` containing the clusters
  - Same cluster name can exist in several groups
- Disabled by default, enabled by *clusters.groups.enable=true*
  - Normally defined in `quattor.build.properties`
  - <https://trac.lal.in2p3.fr/QWG/wiki/Doc/SCDB/Usage#ClusterGroups>
- Handled transparently by `ant compile/deploy`
- Group-wide defaults: `cluster.group.build.properties`
  - In the group directory... but probably useless currently
- Possibility to select a group: *-Dgroup.name=group*
  - Not supported in 2.3

# Ant-Tools new Features

- Features are activated with *-Doption=...*
- Compile debugging: `pan.debug.include=ns/templ`
  - Value must match template namespaces
  - Wildcard allow (regexp syntax)
  - Also `pan.debug.exclude` to exclude some templates from the previous selection
  - Multiple options allowed
  - Interface to a feature implemented by `panc`
- Dependency evaluation: ignore some templates
  - Option `pan.dep.ignore=templ_name`
  - `Templ_name` is the file name relative to `cfg/`, including `.tpl`
  - Regexp allowed
  - Interface to a feature implemented by `panc`

# checkdeps

- Allow to check RPM dependencies before deployment
  - Use /software/repositories to configure YUM repositories
  - Uses /software/packages from the XML profile
  - Use YUM API to assess the dependencies or propose RPMs for missing dependencies if any present
    - Require YUM 3.2.19+ (SL5, impossible on SL4) on both server and client sides
- Presented in Amsterdam last year
  - <https://trac.lal.in2p3.fr/QWG/wiki/Meetings/20081027#checkdeps.py-S.Childs>
  - Not SCDB specific: sources maintained in SF SVN
  - In SCDB: *utils/checkdeps/checkdeps*
- Typical usage:

```
utils/checkdeps/checkdeps [-d 3] -I build/xml/ -x profile_ipnvobox.xml
```



# quatview

- Web application to list machine characteristics from their templates
  - XML profiles loaded into a SQL db (SQLite?)
  - Probably not specific to SCDB but maintained in SCDB utils/
  - To be moved to SF if not SCDB dependent...
- Presented in Amsterdam last year
  - <https://trac.lal.in2p3.fr/QWG/wiki/Meetings/20081027#QuatView-T.SuerinkNIKHEF>
  - Originally from NIKHEF, contributions from TCD (S. Childs)
  - In Amsterdam, mentioned that backend could be merged with CDB2SQL...
    - Any update?

# Deployment Scripts

- SCDB deployment handled by 2 scripts:
  - post-commit: must run on the SVN server (SVN hook)
  - build-tag: launched by post-commit to recompile and deploy
  - Both can run on the same machine
- New Python-based version (\*.py)
  - Used to be bash (post-commit) + perl (perl-tag.pl)
  - Use pysvn API for accessing SVN: no direct use of SVN
  - Unified configuration file: /etc/quattor-deploy.conf
  - Return messages in case of error in a format compatible with SVN client (valid XML)
    - No longer 'BUILD SUCCESS' in case of an error
  - Almost ready for // compile on several deployment servers
    - Mainly implementation of binding between clusters and servers
- Require pysvn and Python 2.4
  - Cannot run on SL4, will be the default in SCDB 3.0

# HTTPrep

- No change

# SCDB 3.0

- SVN 1.5/1.6 support on client and server sides
  - Will be a requirement on client side due to major svnkit API changes
  - On server side, will require new deployment scripts
    - SL4 unsupported
    - Solaris should work (or any OS with pysvn available)
- Parallel deployment on several servers for better performances in large configurations
  - Clusters or group of clusters associated with a deployment server through build properties
    - Typically true on distributed sites but may work on a single site too
  - Hook script will launch build-tag on all deployment servers
  - A deployment server will compile and deploy only the clusters associated with it
- Status: SVN 1.5/6 support ready but client not yet in production at GRIF

# Future

- See <https://trac.lal.in2p3.fr/QWG/milestone/SCDB-ToDoList>
- 3 main items:
  - Ability to precompile some parts of the configuration used in many nodes, like VO information
    - Based on profile cloning for WN ideas
  - Replacement of database-like templates by real DBs + an ant task to generate the templates from the DBs
    - May improve performance as a node may be impacted only by a change affecting itself if the template contains only its information
  - Ability to validate before deployment without compiling all profiles
    - Attempt to improve performance of validation phase done one a desktop in large configurations
- Just ideas, no date...
  - Anybody interested to contribute?