

SCDB Update

Michel Jouvin

LAL, Orsay

jouvin@lal.in2p3.fr

<http://grif.fr>

March 17, 2010

Quattor Workshop, Thessaloniki



Outline

- Last SCDB releases
- Git and SCDB

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

- 3.0.0 released in January
 - SVN 1.6 (1.5) support
 - New deployment scripts (written in Python)
 - Parallel compilation not yet available
 - Development version
- 2.3.3
 - Fixes for update.vo.config (VOConfigTask)
 - New deployment scripts optionally available
 - No more developments

SCDB 3.0 Main Changes

- Panc 8.2.11
- SVN 1.6 (1.5) support
 - scdb-ant-tools 8.0.0 : same features are 7.4.?
 - Svnkit 1.3.0
 - Need to replace by 1.2.x for SVN 1.5 support
- Support for cluster groups: option *group.name* renamed *cluster.group* for consistency

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: *`-Dcluster.group=group`*

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 in 2008
 - <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] -l build/xml/ -x profile_ipnvobox.xml
```

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

Open Questions

- Build Ant tasks with Java 1.6 rather than 1.5?
- What about removal of src/utils symlink?

Future

- See <https://trac.lal.in2p3.fr/QWG/milestone/SCDB-ToDoList>
- Main items:
 - // compilation of clusters
 - 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 willing to contribute?

Git with SCDB

- Git is a distributed VCS: no central repository
 - Everybody has its own repository *clone*
 - *Push* and *pull* between repositories (implies *merges*)
 - Very different workflow compared to SVN
- Replacing SVN by Git would be very difficult
 - SCDB is based on the concept of a central repository used for deployment
 - SCDB relies on SVN to ensure that central repository and admin working copy are in sync before deployment
 - Very difficult to do with Git and probably not very efficient
- Git can be used as a SVN client: *git svn* command
 - A feature of all major DVCS (eg. Mercurial)
- First documentation available at <https://trac.lal.in2p3.fr/QWG/wiki/Doc/SCDB/Git>

Why Git as a SVN Client?

- Full repository history available in the « working copy »
 - « working copy » is a Git clone of SVN repository
 - Size of git repository + working copy < SVN working copy
- Easier to work on several things in //
 - Can use Git branches with easy Git merges and intermediate commits not published into SVN
 - Git branches are not exported to SVN
 - Ability to rewrite history is great!
 - Allow to squash several commits into 1 before pushing to SVN
- *git svn* command instead of *svn update*
 - *svn update* = *git svn rebase*
 - *svn commit* = *git svn dcommit*
- -: Git is more complex to use than SVN...

What is Missing in SCDB?

- In fact nothing to replace SVN client but it is not possible to deploy from Git
 - ant deploy relies on SVN to check that repository and working copy are in sync
 - WC must be a valid SVN WC
 - Currently requires a SVN WC in addition to Git for deploying
- No JAVA API currently available for Git
 - One being developed for Eclipse integration (JGit) but still very preliminary
- No API for git-svn (this is only available at a command)
- Possible approach: a wrapper doing git-svn rebase+dcommit and creating the deployment tag
 - Could also investigate a Git hook for *git tag* but not clear if it can trigger/run git-svn

Other Git Issues

- Git integration with Eclipse is still very preliminary and basic
 - Should improve in the coming year...
 - It's possible to use standard Git in // for git and git-svn operations
 - +: Provides a bash shell on Windows...
- Git (and other DVCS) has no matching features for SVN:externals
 - A problem at least for GRIF where external/ directory is managed through SVN:externals
 - A trick using Git submodules is available (and documented on QWG)
 - Relies on some symlinks: a problem on Windows...