

SCDB Update

Michel Jouvin

LAL, Orsay

jouvin@lal.in2p3.fr

<http://grif.fr>

October 11, 2010
Quattor Workshop, RAL



Outline

- New releases
- SCDB 3.1 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

- No new release since last workshop (Thessaloniki)
 - No pending issue I'm aware of...
 - Except 1 with rpmErrata : last version detection because of changes between '.' and '_' in version number...
- Several new features/improvements introduced in 2.3/3.0
 - Cluster groups
 - Dependency calculation optimisations
 - Support for last versions of svnkit (1.2/1.3)
 - New deployment scripts : more robust, improved debugging messages, improved configuration
- Includes a two core tools
 - checkdeps: a few known issues... but basically working
 - Quatview

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`

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, is the default in SCDB 3.0

SCDB 3.1

- 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
- checkdeps improvements
- Rewrite rpmErrata/rpmUpdates using the RPM Python API
- To monitor progress, see <https://trac.lal.in2p3.fr/Quattor/milestone/SCDB-3.1.0>
 - No active work currently

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?