

SCDB Update

Michel Jouvin

LAL, Orsay

jouvin@lal.in2p3.fr

<http://grif.fr>

Octobre 28, 2008

Quattor Workshop, Amsterdam



- About to be released... branch contains an almost final version
 - Last release is 2.1.5 released mid-February
- Main changes
 - panc v8.2 (8.2.2 in final release)
 - scdb-ant-tools 7.3.1 : update.vo.config to generate VO configuration parameters from CIC portal
 - Ant, svnkit, Saxon updated to last version
 - Svnkit still SVN 1.4 compatible version (1.1.x)
 - SVN 1.5 breaks deployment hook script
 - checkdeps.py: check RPM dependency issues before deployment
 - Based on recent YUM, version available only in SL5 (or FC?)
 - Support for sudo rather than ssh between SVN and Quattor
 - Still a separate version of hook script, post-commit.sudo
- Documentation : major rewrite of SCDB initial installation guide

- An optional SCDB component...
 - Implemented as an ant task, update.rep.templates
 - No change in SCDB v2.2.0
- Support for RPM server cache in reverse mode currently under testing at GRIF (LLR)
 - Reverse cache: transparent for the client, client uses cache URL instead of server URL
 - Replace host name by a variable in repository templates
 - Variable defined in a cluster, site, or profile template
- A web-based RPM upload interface currently being developed at GRIF
 - Allow to add new RPMs to existing repositories without interactive access to the RPM server
 - Not yet clear how generic it would be... Any interest ?



- Current deployment model leads to 3 different issues with a large number of nodes (several 100)
 - Compile twice: on sysadmin machine + on Quattor server
 - Compile + notification on Quattor server
 - Currently all machines unconditionally notified, 1mn for 600 nodes
 - Compile on sysadmin machine may be very long on a desktop/laptop (~1/2h for 600 nodes)
 - If compile on sysadmin machine is too long, risk of a change committed before deploy checks
- Possible solutions:
 - Pre-compile some parts of the configuration, like VO config
 - Extension of WN dummy trick
 - // compile on several Quattor server responsible of a subset of the clusters each
 - Use SVN locks to get an exclusive right to deploy
 - Use tmpfs on Quattor server: first tests disappointing...



- Current SCDB based on a simplistic model that ensures consistency of the whole configuration
 - Nothing can be deployed until successful compilation by sysadmin of all of its mods merged will all already committed mods
 - Deployment based on a SVN copy (under tags/...) of the validated trunk revision (that must be the last one)
 - All machines affected by a change notified
- Problem to scale to a large number of nodes due to requirement to compile on a sysadmin machine first
 - May a desktop/laptop
 - Integrated submission to a compile server "à la CDB" ? User turnaround may be a problem
 - Based on a scdb/panc web service with load balancing ?
- Stage deployment implemented through clusters
 - Include path defined per cluster
 - A node is easily moved from one cluster to another

- Too much full rebuild because of addition of new RPM
 - Caused by update of templates for RPM repositories used on all nodes
 - E.g. components
- RPM-less components: profile rebuilt only if the component is used on the node
 - Component embedded in the profile
 - Impact on compile time and profile size (writing profiles time consuming)
 - Bug in ncm-cdispd (Savannah ???)
- Ignore dependencies on RPM repository templates
 - Easy to implement, no impact on compile time and profile size
 - Issue with handling of moved RPMs: undetected at compile time, profile not updated
 - An optional feature with a way to force rebuild on deployment



- Rewrite hook script in Python using SVN API
 - Currently a shell script parsing svnlook output: fragile, broken by SVN 1.5
 - Easier support for // compile on different Quattor server
 - Python has a good API to schedule defferent threads
- Add support for compiling only one profile (or a set profiles) during compile phase
 - Currently only per cluster but a cluster can be large...
- Rename src/utils to utils ?
 - Increase visibility of these useful utilities
 - src dedicated to providing scripts used by Quattor server convenience
- Organize clusters by sites (already done by Stephen)
- QWG External distribution point: was not intended to be used in production by other sites but...
 - Seems to used by some sites (Ireland?), any feedback ?