

SCDB Update

Michel Jouvin

LAL, Orsay

jouvin@lal.in2p3.fr

<http://grif.fr>

March 12, 2009

Quattor Workshop, London



- About to be released... branch contains an almost final version
 - Last release is 2.2 released last autumn
- Main changes
 - panc v8.2.7 : debug() + -Dpan.debug.include/exclude
 - scdb-ant-tools 7.3.? : fixed/improved update.vo.config to generate VO config parameters from CIC portal
 - No change to Ant, svnkit, Saxon updated to last version
 - Svnkit still SVN 1.4 compatible version (1.1.x), 1.5 breaks deployment hook script
 - checkdeps.py: a few fixes, see Toos session
 - Support for sudo rather than ssh between SVN and Quat merged into main post-commit
 - Ability to exclude RPM repository templates from dependency calculation ?
 - Need to assess performance impact on dependency calculation
 - src/utils/ renamed utils/ : a symlink still exists

- An optional SCDB component...
 - Implemented as an ant task, update.rep.templates
- Support for RPM server cache in reverse mode currently deployed at all GRIF sites
 - Reverse cache: transparent for the client, client uses cache URL instead of server URL
 - No change to SCDB or generic templates



- 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

- 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
 - - : API not really documented, need to reverse engineer examples
- Document as the standard configuration running deployment script as a non-root user
 - Only All requires root privileges
 - Any site already using this config ?
- Add support for compiling only one profile (or a set profiles) during compile phase
 - Currently only per cluster but a cluster can be large...
- Problems generation OS templates for pure RHEL5
 - Layout quite differerent compared to v4



- Not a lot of manpower available to implement all the ideas...
- Would like to have at least one SCDB release at each workshop...
- Highest priority:
 - Rewrite SVN hook script and implement // deployment on several servers
 - Have checkdeps.py fully operational and documented
 - Cleanup problems with generation of templates for pure RHEL5
 - Clusters organized by sites
 - Already committed by Stephen ?

