

Using the SAM framework for the CMS specific tests



Andrea Sciabà



System Analysis WG Meeting
15 November, 2007



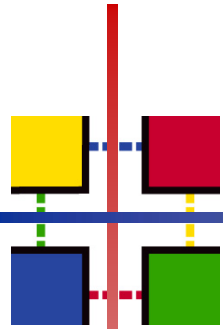
Outline

- SAM for CMS
- CMS SAM tests
- Using SAM with OSG sites
- SAM and VOMS
- SAM and OSG
- Issues
- Plans



SAM for CMS

- Why SAM?
 - SAM is explicitly developed to run periodic sanity checks on Grid (and experiment) services
- How can it be used?
 - Relying on ops test results
 - The easiest option, done for years
 - Running some standard tests under the CMS VO
 - e.g. to spot problems occurring only with VOs other than ops
 - Running custom CMS tests
 - The most effective option



Using ops tests as critical tests

- CMS uses since a long time some ops tests as critical tests
 - Job submission
 - CA certs version
 - csh test
- The failure of any of this tests is definitely a serious problem!



Using CMS custom tests in SAM

- A CMS instance of the SAM client is installed at CERN
 - Tests are submitted every two hours to “real” CMS sites
- The SAM framework allows to easily plug in new tests for existing sensors
 - Added to “testjob” sensor \Rightarrow run on the worker node

Test name	What it does
basic	Checks that the CMS software area is defined and exists, and the CMS site local configuration file is correct
swinst	Checks that the required versions of the CMS software are correctly installed
Monte Carlo	Checks that the stage out of a file from the WN to the local SE is working correctly
Squid	Discovers from the local site configuration file the name of the Squid server and makes a simple query through it
FroNtier	Reads calibration data using CMSSW via the local Squid server



SAM and VOMS roles

- Different tests may need different VOMS roles
 - The `/cms/Role=lcgadmin` role is preferred because
 - It allows to write in the experiment software area
 - It has a higher priority at sites
 - However the `/cms/Role=production` role is needed for the "Monte Carlo" test
 - To take advantage of any write access privileges granted only to that role
 - Solution
 - It is necessary to submit two jobs for every CE instead of one



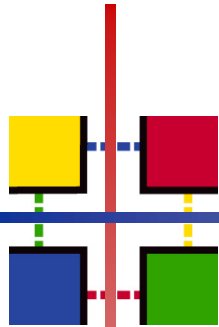
EGEE and OSG

- The job submission is done using the LCG Resource Broker for both EGEE and OSG
 - For EGEE sites it must work by definition
 - For OSG sites it requires some effort
 - The site must be in the central EGEE BDII to be in the SAM database: **OK**
 - The CA certs and CRLs must be kept up to date: **OK**
 - The lcgadmin and production roles must be supported: **OK**
 - The middleware installed in the OSG WN's must be "friendly" to the LCG job wrapper: **OK**
- The SAM tests run nicely on OSG!
 - After an initial phase where lots of problems were found and fixed, now job submission problems are rather infrequent



Description of the SRM tests (I)

- SRM-v1-get-pfn-from-tfc
 - Given the SE name, looks in the TMDB for the corresponding lfn-to-pfn rule for the test LFN
 - `/store/unmerged/SAM/testSRMv1_070628_081219`
 - Returns a warning if in TMDB transfers go to an SRM different from the input node
 - Returns an error if it could not map to a PFN
- SRM-v1-put
 - Copies with `srmcp` a test file from the UI to the PFN
 - **Retries are handled by the script, not by `srmcp`**
 - Returns an error if `srmcp` fails
 - Returns a warning if the pfn-from-tfc test could not map to a PFN
 - **It's not SRM's test if the catalog has not the right information!**



Description of the SRM tests (II)

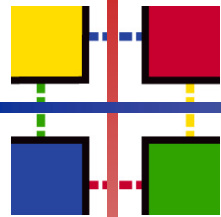
- SRM-v1-get-metadata
 - Uses `srm-get-metadata` on the PFN to retrieve size and checksum
 - Gives an error if `srm-get-metadata` fails or the size or the checksum differ from the original file
- SRM-v1-get
 - Copies with `srmcp` the PFN to the UI
 - Gives an error if `srmcp` fails or the copied file differs from the original file
- SRM-v1-advisory-delete
 - Uses `srm-advisory-delete` to delete the PFN
 - Gives an error if `srm-advisory-delete` fails
 - NOTE: for CASTOR the method is dummy, so test files will grow in number; manual cleaning is required once in a while



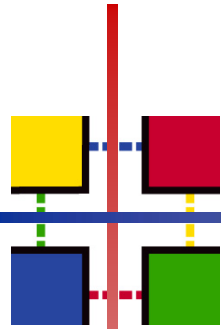
Critical tests for SE/SRM

- Not clear the distinction between the SE and the SRM sensor
 - A legacy of the past...
- CMS runs SRM tests and no SE test
 - It used to run the lcg-cr test
- Since last Monday, there are no critical tests for the SE
- Planned to make the SRM-v1-put test critical for CMS
 - With care, as sites tend to be sensitive about critical tests! They don't want to look bad in the availability calculation...

Availability calculation in GridView



- Bug found in the algorithm
 - <http://savannah.cern.ch/bugs/?31233>
 - The service instance status is UNKNOWN if no test is critical \Rightarrow it should be UP
 - The service status stops being computed from the moment when no test is critical any more
- Will be fixed very soon



FNAL availability

- The FNAL availability in GridView is flawed because the site (in GLUE sense) USCMS-FNAL-WC1 used by GridView has not the CEs
 - The CEs are at another sites: one at uscms-fnal-wc1-ce and one at uscms-fnal-wc1-ce2
 - FNAL will always look better than it is, if the CEs are ignored!!
 - GridView should be able to aggregate more "GLUE" sites in a "GridView" site



Other issues

- None of relevance in the framework
- Tests change from time to time, generally improving to avoid "false alarms"
- After increasing the timeouts from 10 to 20 minutes, the time needed to submit all tests increased to ~1.5 hours, dangerously close to the 2 hours period of the cron job
 - Most of the time is taken by the SRM tests
 - Will upgrade ASAP to the latest version of the SAM client, that prevents a sensor to be run if there is still another instance running



Plans

- Add a test for the "analysis"
 - Try to read a small dataset like the JobRobot
- Closely monitor the SRM-v1-put before making it critical