



Enabling Grids for E-science

CREAM: current status and next steps

**EGEE-JRA1 All-hands meeting
Amsterdam, February 20-22, 2008**

Massimo Sgaravatto – INFN Padova
On behalf of the EGEE Padova team

www.eu-egee.org



- **Last summer CREAM passed the acceptance tests defined by the project**
 - Reliability and performance tests
 - Test results:
 - *>8 days unattended running*
 - *~90K jobs submitted via gLite WMS*
 - *No errors due to CREAM*
 - *No performance degradation observed*
- **The TCG then decided to increase the effort on CREAM to make it production ready**
- **SA3 and TCG defined a checklist, with a set of activities to be completed for making CREAM ready for certification process** (<https://twiki.cern.ch/twiki/bin/view/EGEE/CECheckList>)
 - installation
 - configuration
 - documentation
 - functionality
 - operations

- **In particular the following issues were identified:**
 - Proxy renewal
 - Worked properly under light load
 - Problems even with moderate load conditions
 - Proxy renewal was not used in the CREAM acceptance tests
 - Scalability issues identified also for other operations
 - In particular lease
 - *Mechanism to avoid “zombies”*
 - *Basically each job submitted has an associate lease time, which is periodically renewed by ICE*
 - *Should the lease time expire before the termination of the job, job is killed by CREAM*
 - Lease was disabled in the CREAM acceptance tests
 - Some performance issues

- **CREAM (and ICE) architecture and code had to be revised to address these issues**
- **In particular**
 - New approach for CREAM back-end
 - Now based on relational DB (MySQL)
 - Revision and optimization of the CREAM interface (WSDL)
 - Some operations (e.g. proxy renewal and lease) completely redesigned and reimplemented
 - Old interface preserved for testing and debugging purposes
 - *To be able to start testing with old clients while the new ones were being implemented*
 - Other revisions of architecture/code
- **Major changes**

- **These changes are going to address the raised issues and improve the reliability and performance of the system, but they took much more than expected**
 - Besides CREAM, also ICE had to be modified
 - Changes needed in CEMon (cream-job sensor) as well
 - Changes needed in CREAM CLI
 - Lot of code changed → lot of re-testing needed
 - We used a testing environment at CNAF, but it had to be switched off because of electrical power problems at CNAF
 - Major problem for our tests
 - Just “fixed”: WNs physically moved from CNAF to Padova
 - Probably planned and implemented too many improvements
 - But very likely no room for further major changes (e.g. interface changes) in the future

- **Porting to SL4 and VDT 1.6**
 - CREAM acceptance tests done on SL3 and VDT 1.2
- **Porting to ETICS**
 - During the acceptance tests we were relying on the old gLite build system
 - Took some time ...
- **YAIM based installation procedure**
 - We had a INFN-GRID YAIM based installation procedure during the acceptance tests
 - It was necessary to port it to the official gLite yaim 4
 - Some merge/integration with the LCG CE installation procedure done as well
 - Many software components used by both CREAM CE and LCG CE
- **Documentation**

- **Batch system support**
 - The interaction with the batch system is fully managed by BLAH
 - Support for Torque/PBS and LSF in place since the beginning
 - Submissions to these batch systems via CREAM verified
 - BLAH BLparser reimplemented (also to facilitate the porting to new batch systems)
 - Basically referring to the batch system status/history commands instead of parsing the batch system log files
 - A first implementation of this new BLAH BLParser with the relevant "plugin" supporting Condor has been done
 - The relevant changes in the CREAM software have been implemented
 - PIC SA3 persons gave us access to their Condor environment to test and debug this BLAH and CREAM integration
 - Some problems found and already fixed
 - PIC people had to address a problem with environment of jobs (took a while)
 - Yesterday we were able to submit and successfully run via CREAM
 - Still tests to do

- **Foreseen implementations done**
- **Problems still to be addressed**
 - ICE crash not fully understood
 - Further stress tests needed with proxy renewal (on-going)
 - Many problems and very late to have a WMS with proxy renewal service working properly
 - *Problems identified in the “email” field in the subject of the certificate used for tests*
 - *Looks like this causes interoperability problems when different openssl versions are considered*
 - Several problems that we had to address
 - *Last one was a problem reported by BLAH due likely to a bug in Globus*
- **What’s next**
 - Stress and scalability tests done by developers (on-going)
 - Tests to be done by independent NA4 users
 - Release for certification
 - We assume that VOMS 1.8 certified at that time

- **Address all problems that will be raised during certification process**
- **Submission to CREAM by Condor**
 - Some work already done with the “old” CREAM
- **Allow use of CREAM even without requiring the deployment of the BLParse**
 - Even with lower performance
- **Better integration between CREAM and LB**
 - CREAM able to log information to LB
 - Right now this is done just by the job wrapper (as for the LCG-CE)
 - Enhance LB events with further information
 - Use of LB tools to monitor CREAM jobs
 - Also for the non WMS-jobs (i.e. the ones submitted directly to CREAM)
 - Discussions already started in the IT-CZ Rome meeting
- **New development model for CREAM and WMS job wrapper**
 - CREAM and WMS job wrapper have many common parts
 - Not good and dangerous to have duplicated code

- **High availability/scalable CE**
 - CREAM CE front end and pool of CREAM machines doing the work
 - Main needed functionality introduced with the revised CREAM implementation
 - Multiple CREAM CEs sharing the same DB
 - E.g. a job can be submitted to a CREAM CE, and can then be cancelled on another CE
 - Still some issues to address
- **CREAM used also to access a relational DB**
 - Requested by some GDSE people
 - With the revision of CREAM architecture and code, CREAM is now a general purpose command executor
 - Default command executor: job management
 - So it is just a matter to implement and plug an executor to access a RDBMS
- **WM-ICE integration ?**
 - ICE as thread of WM ?
- **Authorization**

- **From MJRA1.7**

- “CREAM CE uses two authorization frameworks: gJAF for authentication and authorization decisions in java code and LCAS/LCMAPS within glexec.
- Recommendation: The gJAF framework should be abandoned and replaced with a simple authentication check of the certificate and a simple call-out mechanism to the new site authorization service.
- Reason: The use of two authorization frameworks in the same service (i.e. the CE) is not justified and may lead to inconsistent authorization decisions.
- Comments:
 - gJAF will no longer be supported in EGEE-III.
 - If it turns out that a richer functionality than a minimal authorization call-out is needed at the CE, then the Globus authorization framework should be considered as a solution. It is independent of the Globus code and its continued maintenance seems to be better guaranteed than gJAF.”

- **It looks like we are the only one using gJAF, even if it was supposed to be the official EGEE authZ mechanism for Java**
- **It is true that using two authorization frameworks in the CE may lead to inconsistent authorization decisions**
 - We will see if/how the new site AuthZ service can address this issue
- **We are not willing to use the Globus authorization framework (suggested as interim solution ?)**
 - How much should we spend to integrate it ?
 - Does it really meet all our requirements ?
 - Who is going to maintain it ? Globus ? Ourselves ?
- **We prefer to take the current gJAF code, “import” it in the CE code and maintain it**

- <http://grid.pd.infn.it/cream>