



Enabling Grids for E-science

CREAM and ICE

Alvise Dorigo & Massimo Sgaravatto – INFN Padova
On behalf of the EGEE JRA1 Padova Group
JRA1 All-Hands Meeting
Helsinki, June 18th–20th 2007

www.eu-egee.org



- Memory leak fixed in CREAM
- With WMS DAGless now the nodes of job collections can be submitted also to CREAM CEs (no modification in ICE code)
- Split of JobRegister/JobStart in ICE improved the reliability but reduced a little bit the performance
- Support for jobDir in ICE (not fully tested yet)

- Completed and fully tested new ICE feature: in order to get job's status it now subscribes to CEMons on behalf of the owner of the jobs (not using anymore the WMS's host credentials). Not needed any more to have the WMS host's DN on the grid-mapfile of the CE
- Fixed some memory leak in CEMonitor C++ client API (used by ICE to interact with CEMon). This is included in the last TAG, deployed on the `prod-ce-02.pd.infn.it`
- Memory leak in a Gridsite function used for delegation. Applied a patch in CREAM client code (suggested by A.McNab) waiting for a new Gridsite version with the proper fix. The patch mitigates the problem but doesn't fix all the memory leaks. This is included in the last TAG, deployed on the `prod-ce-02.pd.infn.it`

- **3 tests of Glite-CE (WMS using JC+LM+Condor) and 3 tests of CREAM-CE(WMS using ICE):**
 - Test1: Submission of 1000 jobs by 1 user
 - Test2: Submission of 2000 jobs by 1 users
 - Test3: Submission of 2000 jobs by 2 users (1000/user)
- **Used an UI @CNAF for both tests**
- **The Glite-CE testbed was “less powerful” than one used for CREAM-CE in Padova; also LRMS and WNs were different.**

Test results presented at the review (2/3)

Tests of ICE/CREAM have been made without the split of Register/Start that would provide better performance but less robustness

Test Procedure:

- Turn off JC/ICE
- Submit all the jobs
- Turn on JC/ICE

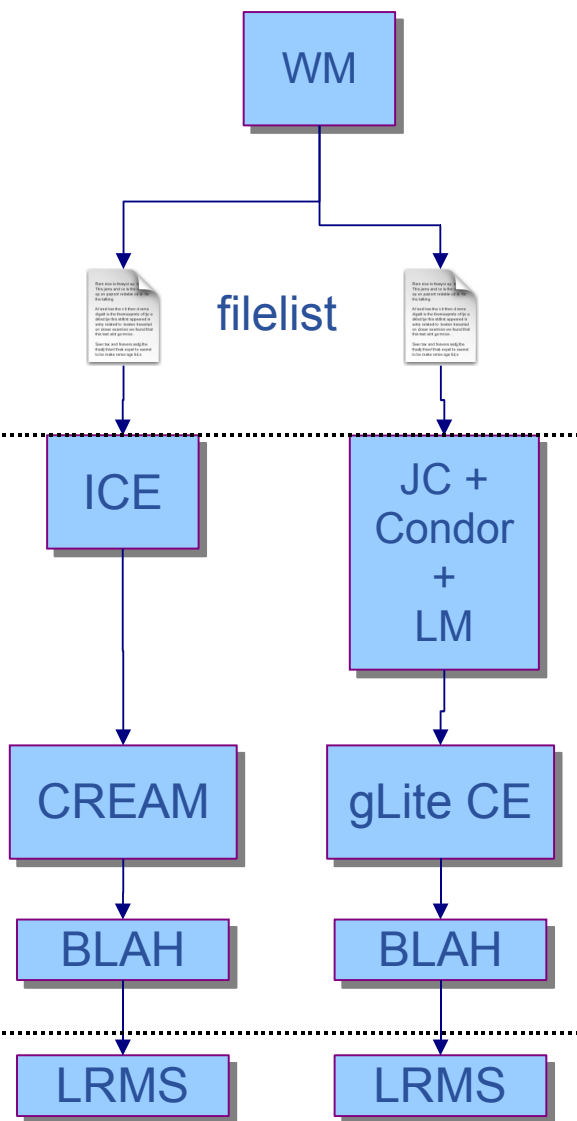
Tstart
(first "Dequed" LB event logged by JC/ICE)

↑

overall time to submit to the LRMS

↓

Tstop
(last submission to the LRMS logged by BLAH in the accounting log file)



Measured:

- efficiency
- throughput

- Results:**

Test	Done(Success)	Done(exit code != 0)	Aborted	Waiting	WMS JobID not returned	efficiency	throughput (jobs/min)
glite CE - Test1	993	0	3	4	0	99.70%	17,44
glite CE – Test2	1835	0	155	10	0	92.21%	10,74
glite CE – Test3	1973	2	22	2	1	98.80%	19,45
CREAM CE - Test1	1000	0	0	0	0	100,00%	47,39
CREAM CE – Test2	2000	0	0	0	0	100,00%	42,37
CREAM CE – Test3	2000	0	0	0	0	100,00%	41,17

- Details about aborts:**

- gliteCE test1: 3 due to File not available. Cannot read JobWrapper output, both from condor and from Maradona
- gliteCE test2: 145 due to “removal retries exceeded”, 10 due to “File not available. Cannot read JobWrapper output, both from condor and from Maradona”
- glite CE test3: 7 due to “Cannot take token, Standard output does not contain useful data. Cannot read JobWrapper output, both from Condor and from Maradona.”; 14 due to “File not available. Cannot read JobWrapper output, both from Condor and from Maradona”; 1 due to "Job got an error while in the CondorG queue"

- **Tried to install a CREAM CE on SLC4 using RPMs built via ETICS**
 - So VDT 1.6
- **Installed Java 1.5 and Tomcat5 via JPackage, following the instructions reported in Steve Traylen's web page (<https://twiki.cern.ch/twiki/bin/view/EGEE/GLite31JPackage>) as discussed in the EMT mailing list**
- **Some non-blocking problems:**
 - Problems with dependencies of some security RPMs (bug #25490)
 - Some modifications needed on Steve T.'s procedure (under discussions with him)

- ... but we were able to submit on such a SLC4+VDT 1.6 CREAM CE
- **What is missing ?**
 - gridftp voms-lcas-lcmads enabled
 - Needed only for direct submissions to CREAM CE and when files have to be staged from UI node
 - So not needed when submitting via WMS-ICE
 - Discussing with Oscar Koeroo how to use his gt4-lcas-lcmads-interface component
 - More tests to see if there are some other problems

- **WMS and CREAM installations in the preview testbed not really tested by users**
- **Only some tests performed by Andrea Sciaba` several times ago**
- **We are using the CREAM CE in Padova now for internal tests**

- **ICE enabled WMS deployed on the experimental services at CNAF**
 - The ones used by Simone Campana, Enzo Miccio, Andrea Sciaba` for their tests
- **Right now this WMS sees a single CREAM CE at Padova**
 - Two other CREAM CEs being deployed at CNAF and Bari
- **Simone Campana started performing some tests with ICE&CREAM**

- **1st test**
 - Submission every hour of a collection of 150 nodes ("hello world") targeted to the CREAM CE
 - 9000 jobs in total
 - Deep and shallow resubmissions disabled
 - 1 single failure
- **2nd test**
 - Submission of collections of 100 jobs targeted to the CREAM CE
 - 10000 jobs in total
 - Jobs are "sleep 2 hours"
 - Deep and shallow resubmissions disabled
 - Major problems with lease (already adressed, the fix is in design phase)

- **3rd test**
 - Same as test2 but with lease stuff disabled (we have to fix it)
 - Test still running (LSF configured to run only 50 jobs simultaneously). So far:
 - *9398 jobs scheduled*
 - *51 running*
 - *539 done*
 - *12 aborted (all because failures in delegation)*
 - It took about 420 minutes to submit all the jobs to the CE (~ **24 jobs/min**)

- **Installation instructions for manual CREAM installation available**
- **There is also a semi-automatic procedure for installation and configuration based on glite-yaim**
- **A wiki page provides installation and configuration instructions:**

`http://igrelease.forge.cnaf.infn.it/doku.php?id=doc:guides:install-cream`
- **When we have sorted out the existing problems (see previous slides) the yaim procedure will support also the installations on SLC4**

- **Testing the new jobDir**
- **In order to reduce the server's load we're designing in ICE a new system of lease update and a new system to update jobs's status that reduces the net traffic;**
- **Reducing ICE's memory usage and making faster the access to internal metadata using BerkeleyDB and a [de]serialization facility better than classad's [un]parser (probably the boost's serialize APIs)**
- **Make the ICE's logging compliant as much as possible with recent recommendations (see last Ales's email)**