CREAM and ICE

Massimo Sgaravatto – INFN Padova
On behalf of the EGEE JRA1 Padova Group
JRA-1 All-Hands Meeting
Catania, March 7-9, 2007

Information Society and Media

Updates since last All-Hands meeting

- Several bug fixes and enhancements have been made in CREAM and also in other software components used in CREAM
  - Bug #20357 (race condition in BLAH) fixed
    - This caused problems in concurrent submissions (in particular when done by different users)
    - This also was the cause of the failures because of “gridftp problems” reported last time
  - The failures seen as “glexec problems” reported last time were understood and fixed: it was a bug within CREAM
  - Submission time to LRMS (schedule time) improved a lot
- Several bug fixes and enhancements also in ICE
  - *putProxy* and *JobRegister* in the same call to fasten job submission to CREAM

BES related activities

- **First implementation of BES support done in CREAM**
  - Shown at SC'06 (Tampa-FLORIDA) in a interoperability demo with other computational services
- **CREAM-BES developments done in collaboration with the OMII-EU project**
- **The idea is to enhance CREAM with an additional BES-compliant WSDL interface**
- **The BES interface will coexist with the current one (i.e., the same WSDL will provide two different port-types)**
- **Issues**
  - BES does not provide any security mechanism
    - Interim solution: Basic Authentication Profile with Username Token
    - OMII-EU JRA3/Security group will provide better solutions
  - The BES specification is still evolving
    - The specification should be released for public comments this week

Porting to SLC4

- **CREAM tested so far only on SLC3**
- **Not tested yet on slc4_ia32 also because not all needed RPMs are available in the ETICS repository**
  - The following gLite modules don't available in ETICS rep. (as of March 8, 2007)
    - *glite-security-lcas*
    - *I am told that the problem has been fixed*
  - Can't find the following RPMs on the ETICS repository:
    - *edg-gridftpd*
      - Needed only for direct submissions to CREAM from UI, and only when needed to stage jobs from UI node (e.g. not needed for submissions through WMS/ICE)
      - To be used also within WMS
      - Being discussed at the EMT its inclusion in ETICS
    - *fetch-crl*
      - Needed for all Grid nodes

CREAM and ICE on the Preview testbed

- Same testbed composition reported last time
  - 1 UI @ INFN-CNAF
  - 1 WMS (ICE enabled) @ INFN-CNAF
  - 1 BDII @ INFN-CNAF
  - Still only one CREAM CE @ INFN-Padova
    - cream-01.pd.infn.it:8443/cream-lsf-grid01
    - 4 WNs
- NIKHEF made available another CE (+3 WNs), where to deploy CREAM + Glexec on WN
  - Unfortunately a 64bits OS was installed and so it was not possible to install the software
  - OS just reinstalled, so we can proceed with CREAM installation
- This Preview Testbed is open to users willing to test CREAM and ICE

Update on test results

- Focus on tests on submission to CREAM CE via gLite WMS (ICE enabled)
- Done also tests of direct submission to CREAM CE using the command line CREAM UI

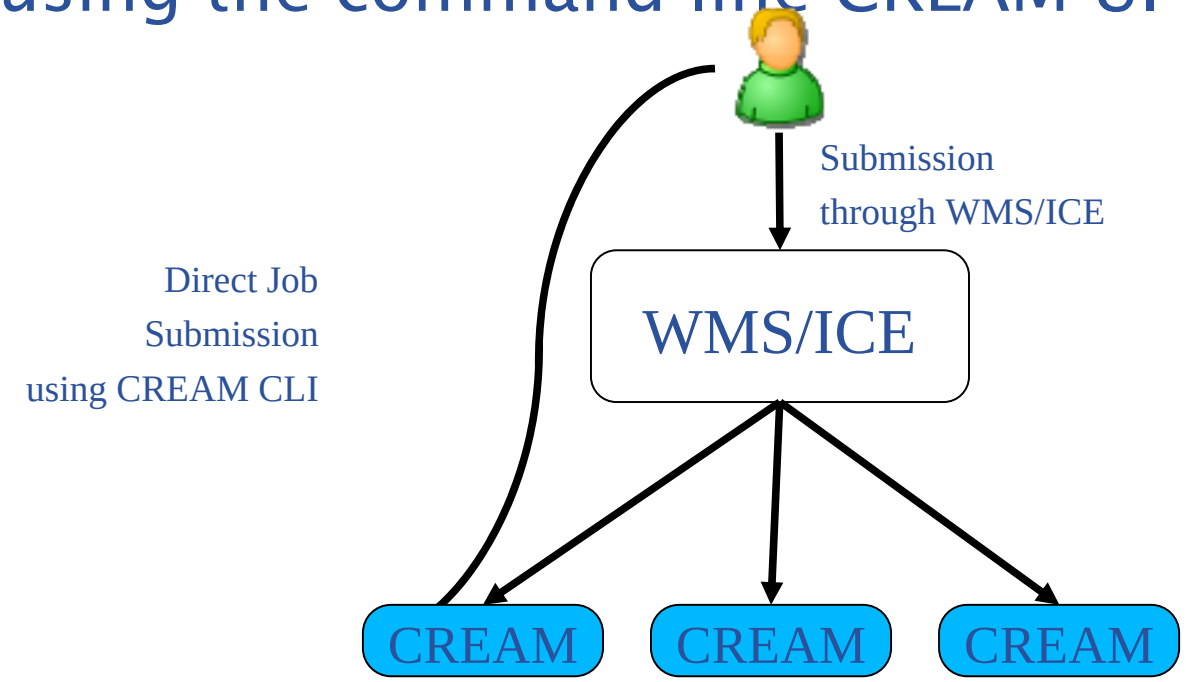

Direct job submission tests

- **Stress tests:**
  - Submission of an increasing number of jobs from UI @ CNAF (`pre-ui-01.cnaf.infn.it`) to CREAM CE @ Padova (CEId `cream-01.pd.infn.it:8443/cream-lsf-grid01` with 4 worker nodes)
    - Submission of **100** jobs from 1 thread
    - Submission of **250** jobs from 1 thread
    - Submission of **500** jobs from 1 thread
    - Submission of **1000** jobs from 1 thread
    - Submission of 2**000** jobs from 1 thread
  - Tests have been made using a pre-delegated proxy
  - Measured values:
    - The number of failed jobs (taking into account the reported failure reasons)
    - The time taken to submit each job to the CREAM CE (i.e. the time needed to get back the CREAM JobID)
    - The time needed to submit the job to the LRMS via BLAH (i.e. the time needed to get the BLAH jobid)

Direct job submission tests: results

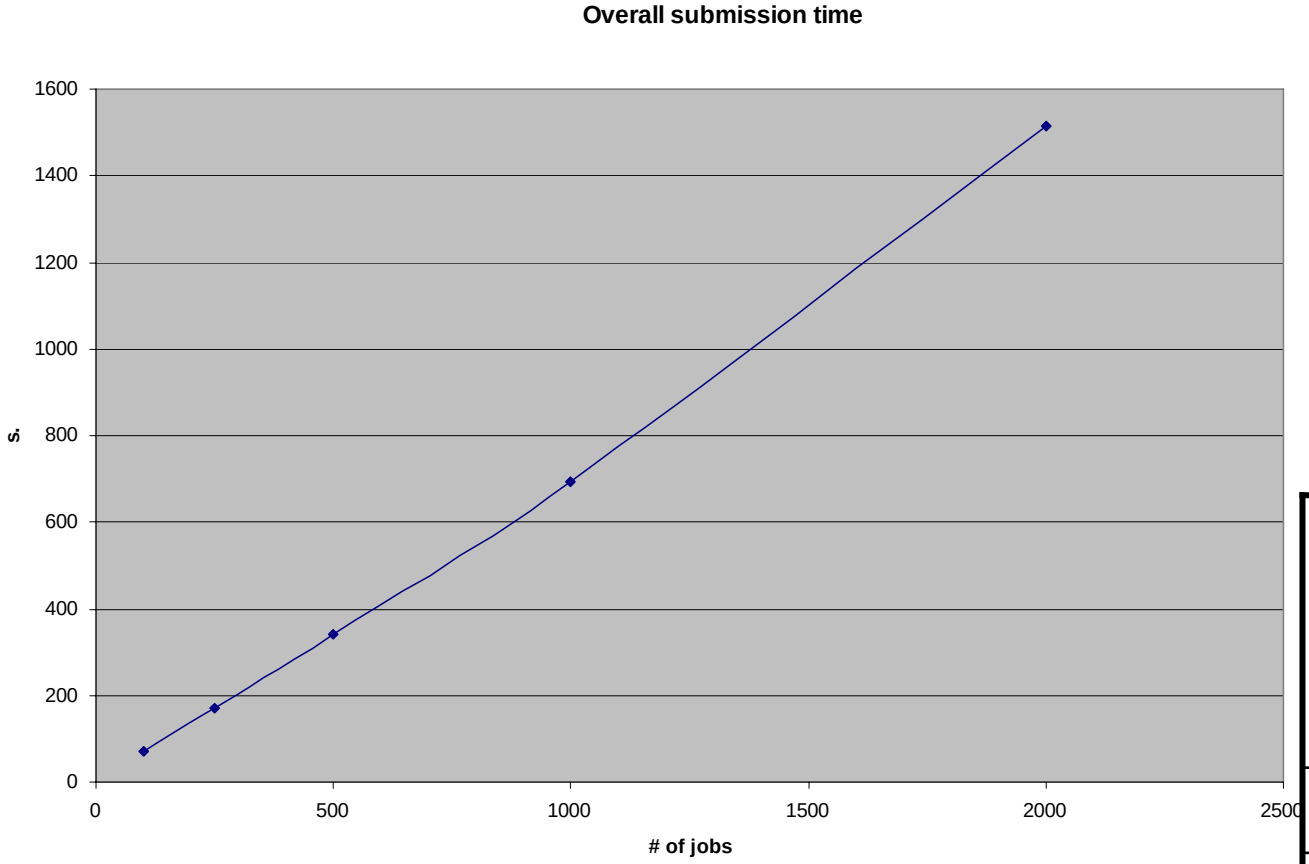

Time needed to submit to LRMS

# of jobs	Average schedule time (s.)
100	3.152
250	2.968
500	3.002
1000	3.087
2000	3.087

Failures in these tests: 0

Overall submission time can be improved submitting from multiple threads

Submission through WMS

- **Tested**
  - CREAM CE vs. gLite CE and vs. LCG CE
  - ICE vs. JC+Condor+LM
- **Same JDL used in all scenarios**
  - Shallow and deep resubmissions disabled
- **Testbed configuration**
  - WMS, BDII, UI @ INFN-CNAF (used in all tests)
  - A single CREAM CE @ INFN-PADOVA configured with 50 threads (`cream-01.pd.infn.it:8443/cream-lsf-grid01`)
  - A single gLite 3.0 CE @ INFN-PADOVA (`cert-04.pd.infn.it:2119/blah-pbs-long`)
  - A single LCG CE @ INFN-PADOVA (`cert-12.pd.infn.it:2119/jobmanager-lcglsf-cert`)
- **What has been measured**
  - Efficiency (reporting the number of failed jobs along with the failure reasons)
  - For both JC+Condor+LM and ICE: for each job, the time needed for the submission to the LRMS and the corresponding throughput
  - Only for ICE: the time needed for the submission to the CREAM CE and the corresponding throughput

```
Job JDL:
Executable = "test.sh";
StdOutput = "std.out";;
InputSandbox = {"gsiftp://grid005.pd.infn.it/Preview/test.sh"}
OutputSandbox = "out.out";;
OutputSandboxDestURI = {"gsiftp://grid005.pd.infn.itPreview/O
RetryCount = 0;
ShallowRetryCount = 0;
```

```
#!/bin/sh
echo "I am running on `hostname`"
echo "I am running as `whoami`"
sleep 600
```

Submission through WMS

- **How the tests have been performed**
  - ICE/JC turned OFF
  - Submission of 1000 jobs to the WMS in order to fill the ICE/JC input filelist
  - When all the requests have been inserted in the ICE/JCinput filelist, ICE/JC turned ON, *so it can start to satisfy the submission requests*
- **How the measurements have been performed**
  - **Tstart** = LB timestamp of first ICE/JC dequeued event (i.e. request removed from the filelist, i.e. ICE/JC started its work)
  - **Tstop1** = LB timestamp of the last “Transferred OK to CE” event (when measuring throuput to submit to CE for ICE scenario)
    - Not straightforward to distinguish submission to CE vs submission to LRMS in the JC+Condor+LM scenario
  - **Tstop2** = timestamp of last submission event in the DGAS accounting log file (when measuring throughput to submit to LRMS for both ICE and JC+Condor+LM scenarios)
- **Throughput to submit to CE = # jobs / (Tstop1 - Tstart)**
  - Measured only for ICE scenario
- **Throughput to submit to LRMS = # jobs / (Tstop2 - Tstart)**
  - Measured for all scenarios

Test scenario

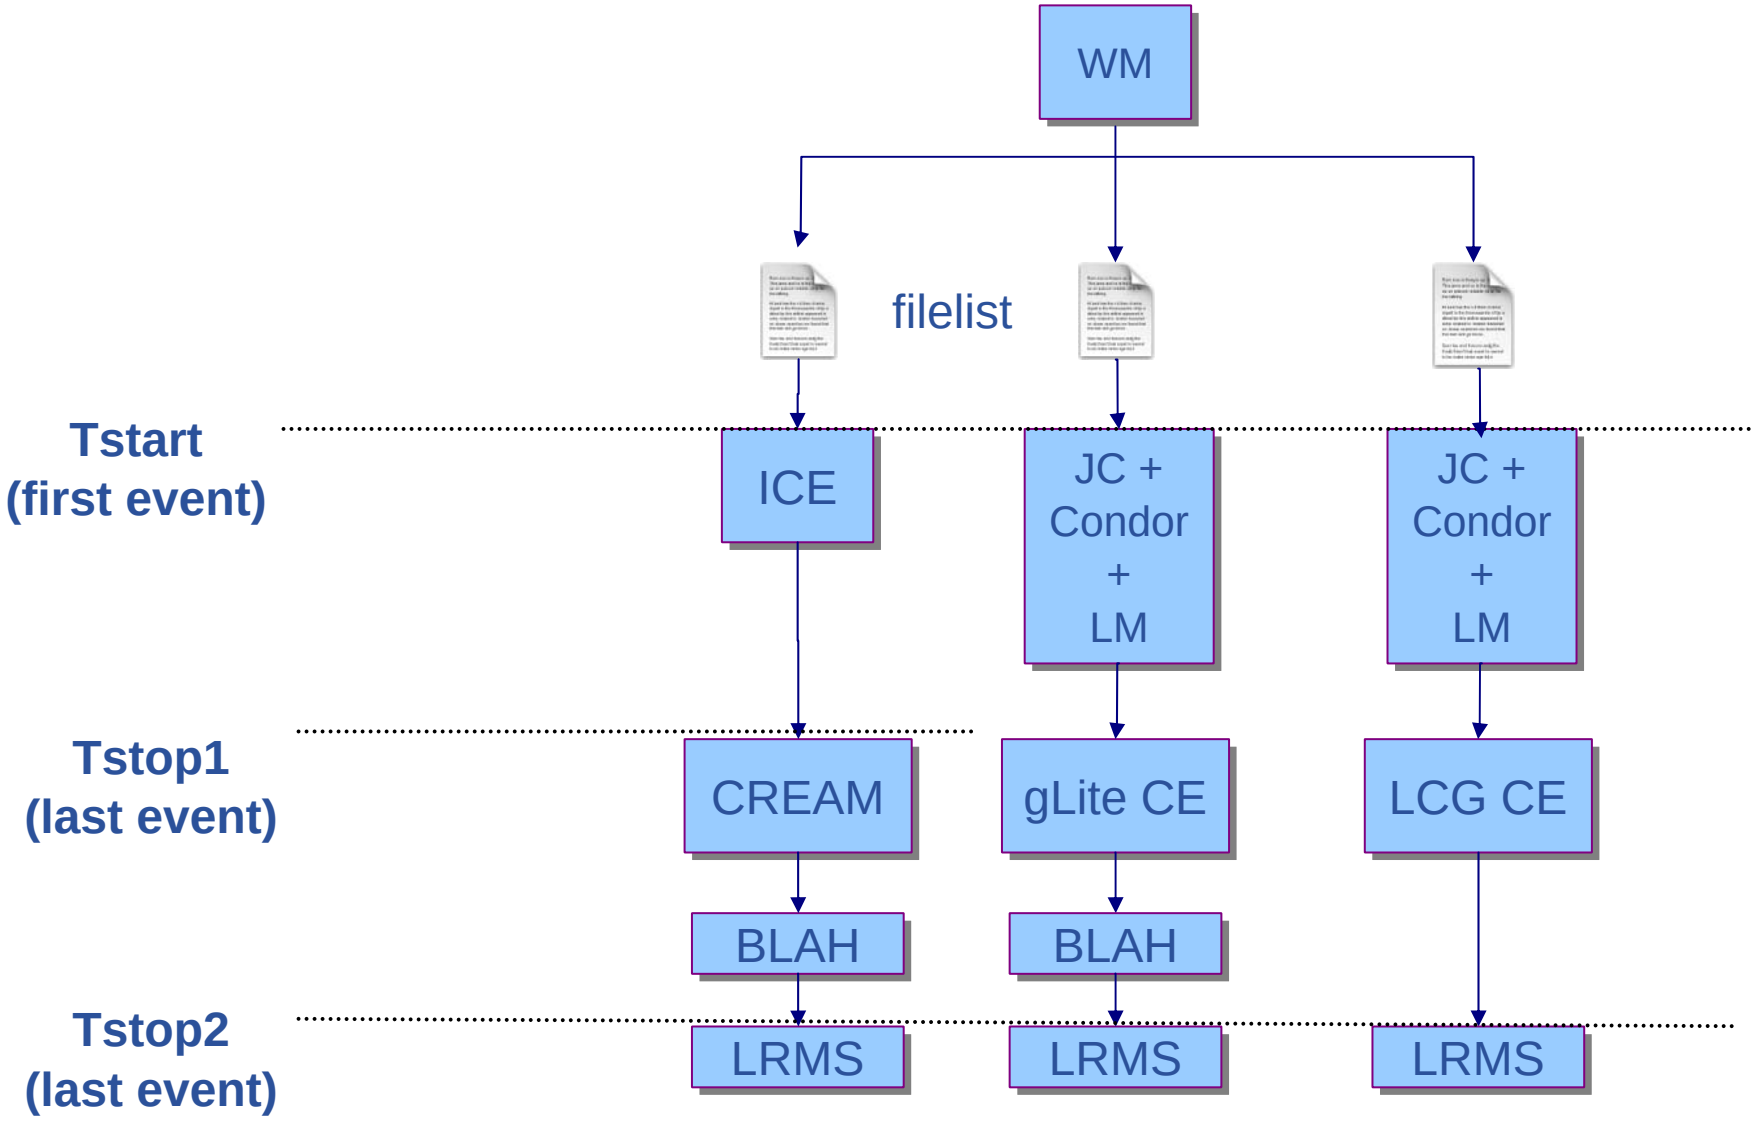

ICE & CREAM Test Results

- Submission of 1000 jobs by 4 users

# ICE threads	% success considering only jobs managed by ICE	% success considering all jobs	Throughput to CE (jobs/min)	Throughput to LRMS (jobs/min)
5	100 %	99.6 %	38.63	38.55
10	100 %	99.5 %	37.73	37.54
15	100 %	99.4 %	38.45	38.40

- All failures happened at submission to WMProxy (Grid-jobid not returned to user, because problems registering job in LB)

ICE & CREAM Test Results

- **Tests submitting to ICE&CREAM performed also by Andrea Sciaba` on the Preview Testbed**
- **He submitted 4491 jobs**
- **4 failed at WMproxy level because “Unable to untar ISB file”**
- **All other (4487) jobs managed by ICE and then submitted to CREAM were successfully executed**

- **PS: Andrea was not able to submit using his new cert signed by the new CERN CA, because of bug #23534**
  - glexec and/or LCMAPS and/or VOMS problem ?
  - Under investigation by Oscar

gLite CE - JC+Condor+LM test results

Try No.	% success considering only jobs managed by JC+Condor+LM	% success considering all jobs	Throughput to LRMS (jobs/min)
1	94.6	94.3	2.4
2	93.9	93.9	2.4
3	96.0	96.0	2.4

Same test results reported last time: no updates

- 1000 jobs submitted by 1 user
- Low throughput because of Condor bug #21529

  As far as I understand now fixed but new Condor not yet deployed
- Failure reasons:
  - 3: Submission to Condor failed
  - 106: Cannot read JobWrapper output, both from Condor and from Maradona
  - 40 Job got an error while in the CondorG queue
  - 2 Removal retries exceeded
  - 7 jobs in waiting

LCG CE - JC+Condor+LM test results

- Submission of 1000 jobs by 1 user

Try No.	% success	Throughput to LRMS (jobs/min)
1	100 %	13.52
2	99.9 %	13.85

Measured
37.54 – 38.55 jobs/min
with ICE-CREAM
(see previous slides)

- 1 single failure because of “Submission to Condor failed”

Work in progress

- **Changes in ICE and CEMon so that subscriptions (to get notifications about job status changes) are done by users and not by WMS**
  - So not needed anymore to have WMS host DN in CE's grid-mapfile
    - Big problem for deployment
  - Under internal tests
- **CREAM "automatic" installation via YAIM**
  - Basically done
    - http://igrelease.forge.cnaf.infn.it/doku.php?id=doc:guides:install-cream
  - Under test
- **Bug fixes and support**
  - Not only for EGEE but also for GridCC and OMII-Europe

Next steps

- **Continue testing and debugging of ICE and CREAM**
  - Necessary to perform testing of ICE in a larger scale, in particular considering more than a single CREAM CE
  - Preliminary testing plan being discussed within INFN:
    - CREAM to be deployed in 4 sites (probably Padova, CNAF, Torino, Bari) with at least 5 WNs per CE
- **Integration of ICE with DAGless WMS**
  - So nodes of bulk jobs can be submitted to CREAM CEs
  - Not necessary to modify code in CREAM and/or ICE
- **More information:**

  CREAM web site: http://grid.pd.infn.it/cream