

Roadmap for WMS—CREAM—LRMS parameter passing

*D. H. van Dok, Nikhef
GDB 10 February 2010*

www.eu-egee.org



Parameter passing to the batch system (LRMS) was discussed during the December GDB by Douglas McNab (ScotGRID). To implement this effectively it will have to be supported by the middleware and YAIM without additional effort of the site administrators.

All of this is tracked on [https:](https://twiki.cern.ch/twiki/bin/view/EGEE/ParameterPassing)

[//twiki.cern.ch/twiki/bin/view/EGEE/ParameterPassing](https://twiki.cern.ch/twiki/bin/view/EGEE/ParameterPassing).

The idea is to turn requirements in the JDL like this:

```
CERequirements = "GlueHostMainMemoryRAMSize > 100 &&  
other.GlueCEMaxWallClockTime >= 400 &&  
other.GlueCEMaxCPUTime >= 300 ";
```

...

...into this:

```
# PBS job wrapper generated by pbs_submit.sh
...
# PBS directives:
#PBS -S /bin/bash
#PBS -o /dev/null
#PBS -e /dev/null
#PBS -l mem=100mb
#PBS -l walltime=24000
#PBS -l cput=18000
...
```

This is the initial set of parameters to aim at. Essentially selecting on memory, wall clock and/or cpu time has obvious applications and benefits for both users and sites.

MainMemoryRAMSize

MaxWallClockTime

MaxObtainableWallClockTime

MaxCPUTime

MaxObtainableCPUTime

Especially for MPI, the following parameters will also be supported

SMPGranularity

WholeNodes

The proof-of-concept, packaged as RPM can be found here:

<http://www.nikhef.nl/~dennisvd/blah-pbs-hook/>

For trial purposes this was installed on CREAM CE
stremsel.nikhef.nl.

- Direct submission to CREAM
- submission through WMS

Direct submission to CREAM is OK.

- CPU, memory and walltime requirements are picked up and enforced

- overrunning the CPU time results in an

Status = [DONE-FAILED]

ExitCode = [W]

FailureReason = [pbs_reason=271]

- overrunning the wall time leads to

Status = [CANCELLED]

ExitCode = []

- Asking too many resources results in a job abortion and a BLAH error, but this is fixed in newer BLAHs (not yet in production).

The BLAH bug mentioned before results in leftovers in /tmp:

```
-rw-r--r-- 1 tomcat tomcat 37 Mar 23 16:46 ce-req-file-1269359209171688  
-rw-r--r-- 1 tomcat tomcat 37 Mar 23 16:49 ce-req-file-1269359388531765  
-rw-r--r-- 1 tomcat tomcat 37 Mar 23 16:52 ce-req-file-1269359569560461  
-rw-r--r-- 1 tomcat tomcat 37 Mar 23 16:55 ce-req-file-1269359750346065  
-rw-r--r-- 1 tomcat tomcat 37 Mar 23 16:58 ce-req-file-1269359930309989
```

Which are the requirements translated to shell code.

Submission through WMS is not yet completely worked out, but it should be better with the next version. At least memory and cpu time requirements are passed along. But it is still not what you'd expect: a job asking 300 minutes CPU time got matched to the 'short' queue.

This is the outcome:

```
===== glite-wms-job-status Success =====
```

```
BOOKKEEPING INFORMATION:
```

```
Status info for the Job : https://lb.testbed:9000/022BYKBGR_q1Kpvj9AjoLw
```

```
Current Status:      Done (Success)
```

```
Logged Reason(s):
```

- BLAH error: submission command failed (exit code = 1) (stdout:)
 (stderr:rm: cannot remove '/tmp/ce-req-file-1269393003236504':
 Operation not permitted-qsub: Job exceeds queue resource limits
 MSG=cannot satisfy queue max cput requirement-) N/A (jobId =
 CREAM231105474)

- Job Terminated Successfully

```
Exit code:          0
```

```
Status Reason:      Job Terminated Successfully
```

```
Destination:        creamce.gina.sara.nl:8443/cream-pbs-short
```

```
Submitted:          Wed Mar 24 02:09:57 2010 CET
```

```
=====
```

- inclusion in upcoming TORQUE_utils metapackage patch #3650 and similarly LSF_utils, SGE_utils, CONDOR_utils. Certification process still needs to start.
- YAIM variable to toggle turning the script on (open question, should it be default **ON** or **OFF**?)
- WMS bugfix for bug #42288 needs to be deployed

Do we need a (early) trial run?

- selected sites
- selected users

What happens with Glue 2.0?

- <https://twiki.cern.ch/twiki/bin/view/EGEE/ParameterPassing>
- <http://grid.pd.infn.it/cream/field.php?n=Main.ForwardOfRequirementsToTheBatchSystem>