

# EGEE Strategy for Replacing the LCG Compute Element with the CREAM Compute Element in the Production Service

EGEE Technical Management Board (TMB)

21<sup>st</sup> January 2009

## Introduction

In the third quarter of 2007, a decision was made that the CREAM CE should be developed to the point that it is ready for full scale production usage and that once this had happened, the LCG CE should be retired from service. In October of 2008, the CREAM CE was first made available in the EGEE production infrastructure software repositories and so now is an appropriate time to consider the transition plan for transitioning from the LCG CE to the CREAM CE and retiring the LCG CE.

## Criteria for Transition to Start

Before the LCG CE can be replaced by the CREAM CE, a number of criteria need to be met by the CREAM CE:

1. The CREAM CE should provide at least equivalent functionality and performance as the LCG CE.
  - a. Preferably without the ability for users to fork processes on the CE
  - b. There is no list of functionality for the LCG CE (and no-one willing to create the list), therefore the easiest and pragmatic approach is to try out the CREAM CE in production and see what's missing.
2. Condor-G submission to CREAM must be available in production with no significant bugs.
  - a. In this context, "significant" will be decided by the grid release teams (integration, certification & operations), with input from VOs.
3. The ICE enabled WMS must be in production with no significant bugs.
  - a. In this context, "significant" will be decided by the grid release teams (integration, certification & operations), with input from VOs.
4. The ICE / CREAM job submission chain should be able to meet all performance criteria and otherwise perform at least as well as the WMS / LCG CE submission chain.
5. The CREAM CE must use an acceptable (to sites) proxy renewal mechanism.
6. An adequate set of monitoring probes (SAM/Nagios) must be available for the CREAM CE.
7. There is a clear plan, with agreed implementation timelines, for migration of the CREAM CE away from gJAF.
8. The following batch systems must be integrated by default:
  - a. LSF
  - b. PBS-Torque/Maui
  - c. Sun Grid Engine
  - d. Condor

The process for integrating other batch systems must be fully documented

9. Performance criteria [*taken from EGEE TMB document "EGEE Strategy for the Compute Element"*]
  - a. Performance
    - i. At least 5000 simultaneous jobs per CE node

- ii. Unlimited number of user/role/submission node combinations from many VO's (at least 50), up to the limit of the number of jobs supported on a CE node
- b. Reliability
- i. Job failure rates in normal operations due to the CE < 0.1%
  - ii. Job failures due to restart of CE services or reboot < 0.1%
  - iii. 1 month unattended running without significant performance degradation
  - iv. Graceful failure or self-limiting behavior when the CE load reaches its maximum (e.g. if a CE node can support only 5000 jobs it must not crash or become unresponsive with more than that)

**N.B.:** The mechanism for passing parameters to BLAH has been removed from the list of criteria for transition, although it is still expected to be available with the production CREAM CE.

### Tasks

- Get the CREAM CE out into the production service as soon as possible.
  - The ROC managers have been asked to ensure that at least one CREAM CE is installed per region.
  - Have also asked the WLCG Management Board to ask tier-1 sites + 10 tier-2 sites to install CREAM CEs.
  - Need to encourage VOs to use these CREAM CEs (the Alice VO is already doing this)
- Meanwhile:
  - Push for Condor-G and WMS-ICE job submission solutions.
  - Push for the acceptable proxy renewal mechanism.
  - Make sure that the monitoring probes are in place.
  - Test that the publishing of accounting data is working correctly.
    - This should be fine as the LCG CE uses the same log format as the CREAM CE (BLAH).
- Once the transition criteria are met (t=0):
  - Switch all existing CREAM CEs to “production” status so they are matched by the ICE-WMS.
 

**Time: immediately**

    - Need to think about double counting of resources here, although it probably won't be any worse than already exists in production.
  - Announce the end of life of the LCG CE, including:
    - Standard updates (new functionality + bug fixes) to stop.
 

**Time: t + 6 months**

      - N.B.: This may be influenced by plans for SL5.
    - General support to stop.
 

**Time: t + 6 months**
    - Security updates to stop (and so all support withdrawn)
 

**Time: t + 12 months**
  - Ask sites running >1 CE to switch CEs over one at a time.
  - To prevent many sites going into scheduled down-time for several days each, the best strategy for sites running only 1 CE will be to:
    - Install, configure and test a CREAM CE without making it available in the information system (therefore not generally visible in production)

- Once happy, put it into production in parallel with their already existing LCG CE
- Once the CREAM CE is shown to be behaving properly, drain the LCG CE and retire it from production
- However, for any sites running the batch system on the same physical node as an LCG CE, this migration will be more tricky and may have to involve some site down-time.

### **Risks**

- 1) At no time can we break interoperability with OSG.
- 2) The upgrading of grid middleware to SL5 may affect the timelines as there is no effort to port the LCG CE to SL5.