



## WLCG Service Interventions

### Introduction

Since early May 2007, we have maintained a daily and weekly view<sup>1</sup> of all scheduled and unscheduled interventions at WLCG Tier0 + Tier1 sites. At the highest level, there is at least one intervention every weekday over the sum of these sites and a disturbing fraction of these are unscheduled.

Given that CMS plan a 30-day challenge – CSA07 – due to start September 10<sup>th</sup>, how will this relatively high frequency of interventions affect the exercise? Can we avoid the ‘super-human efforts’ that were required to run the service challenges but which are not sustainable in the long term? Can the situation be improved prior to 2008 / 2009 data taking?

It is clear that a number of the interventions are due to preparations for the experiments’ Full Dress Rehearsals and the deployment of the so-called ‘residual services’. A number of others affect services that are redundant, so that the service seen by the users is at worst degraded (e.g. one out of several CEs or RBs is down for some hours).

If we exclude such interventions, we are still left with a number of critical issues that have a major impact on the overall services. Perhaps unsurprisingly, these include:

- Power, cooling and network problems, responsible in the worse cases for complete site downtime for several hours or even days;
- Problems with storage and related services, e.g. CASTOR or dCache (by definition);
- Problems with back-end database services or the interaction between the application layer and the database, affecting services such as LFC, FTS, VOMS, SAM etc.

These three categories need to be addressed separately, but all with priority.

### Infrastructure Issues

It is proposed that each site analyse the infrastructure failures – including those external to the site, such as LCG OPN links between the Tier0 and the site in question – to understand whether their infrastructure upgrade plans can be expected to result in a measurable improvement in site availability on the timescales of September this year (CSA07) and early 2008 (site readiness for LHC data taking in 2008). Such a report should be submitted to the Overview Board for its next meeting.

### Storage Services

The instability of storage services has been an issue for several years and has a big impact on many of the key services that must be provided by Tier0 and Tier1 sites, according to the MoU agreement. Given the current deployment schedule for SRM 2.2, it cannot be considered likely that stability will improve significantly on the timescales mentioned above – indeed production SRM 2.2 deployment may well introduce further instability for a number of months, including to the FTS service, which may have to adapt to any changed or unforeseen behaviour in these new implementations. A key element of the FDRs will therefore be to ensure that there is



21 June 2007, Jamie Shiers

sufficient redundancy in the services, i.e. in terms of buffer space and catch-up capacity. In the longer term, significantly better stability could result in a reduction in costs, i.e. through lower (but never zero) catch-up capacity (the converse also being true) and manpower, by avoiding the need for special effort during challenge periods.

## Database Interaction

There are well known techniques for making multi-tiered applications more resilient. A plan to address not only the core Grid services but also those critical to the experiments is in the process of being drawn up. It typically requires:

- Redundant middle-tier deployment via DNS load-balanced services;
- Redundant database deployment, e.g. on high-availability clusters;
- Modifications to the application – which can be non-trivial in the case of complex systems – to handle the necessary failover / DB reconnection etc. correctly.

## Summary

The weekly view of service interventions at WLCG sites has proven very useful in planning the service in the short – medium term. It has also allowed a simple analysis of the key causes for such interventions and helped prepare a plan to improve service stability. Provided that the plan is consistently executed, measurable improvement should be seen within a few months.

---

<sup>i</sup> See <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGServiceLog>.