



## LCG Services Report March – May 2008

1 June 2008

Jamie Shiers

The primary focus for the past 3 months has been the preparation for – and execution of – the 2<sup>nd</sup> phase of CCRC'08. This phase took place from Monday May 5<sup>th</sup> until Friday May 30<sup>th</sup> inclusive. The preparation included face-to-face meetings in March and April, as well as a week-long WLCG Collaboration workshop, also in April (21 – 25 at CERN). With respect to the February challenge, only very minor updates to the middleware and storage-ware were made. These are detailed below. The challenge was run in the same overall manner as in February, with a number of small optimizations and full 2008 workflows / loads. In addition, full 2008 hardware resources were expected to have been in place (due in production by 1<sup>st</sup> April 2008). In some cases, this schedule was not respected (details below). Whilst this was not a problem for this specific challenge, if hardware resources are required and deployed late in future years, this can be expected to have a significant negative impact on the ability to support the data taking and production for the LHC experiments for the year in question.

Notes from the daily operations meetings, held at 15:00 Geneva time, were sent to the [wlcg-ccrc08@cern.ch](mailto:wlcg-ccrc08@cern.ch) mailing list shortly after the end of the meeting. This had a noticeable positive impact on the readership of the minutes. In the final week of the challenge it was decided to hold the operations meeting also on Mondays, leaving the Joint EGEE/OSG/WLCG operations meeting, held the same day at 16:00 Geneva time, to focus on longer term issues.

Although on the first day of the May challenge a number of problems were reported, these and most other problems encountered during the month were resolved very rapidly. Thus, not only were fewer problems reported on a daily basis than in February, but they were also resolved more quickly. In this respect, we did better than our target – which was to perform as in February, despite the higher workload. One site (RAL) suffered a micro-power cut that resulted in a downtime of over 8 hours for the most complex services (i.e. storage) to be resumed. Regrettably, the DB serving the GOCDDB suffered a corruption during this incident and took several days to recover. The full details are given in the DB section below. A cooling problem hit NIKHEF, requiring the power-off of worker nodes for an extended weekend. Finally, on the last day of the challenge a power cut on the CERN site took out the physics services, which again required several hours before they were fully resumed. These events remind us how exposed we remain to such problems – it is unlikely that HEP budgets will allow us to be fully protected from power outages, and thus we run at risk. Cooling similarly requires significant investments and remains a key infrastructure challenge for large scale centres.

The operator alarm mailing lists for CERN were setup and configured correctly. They were only required on one occasion and for a single VO (ATLAS). Unfortunately, the procedures were shown to be not yet fully in place – the problem triggering the use of the list was solved independently. These problems have since been solved, yet a frequency of once per month (once also in February) – good news in itself – suggests that the solution implemented must be simple. A proposal for equivalent lists at the Tier1

sites was made, but not yet consistently implemented. The experience from the Tier0 setup should be taken into account, together with the service requirements on Tier1s.

The need for emergency site contact phone numbers also came up once and with a single site. Due to staff turnover, the procedure did not work correctly, highlighting again the need for simplicity – as well as periodic testing.

At the time of writing, CNAF are promising 2008 resources in the July timescale, with improvements in deployed capacity expected over the coming weeks also at BNL, FNAL and IN2P3. BNL added additional disk space during the May run, as did NL-T1, which was able to make available 120/900TB of the 2007/8 pledges. NDGF foresee adding disk capacity incrementally “on-demand”.

### Middleware

The middleware versions that sites were required to upgrade to for the May run of CCRC’08 are given in the table below. A further release to DPM – 1.6.10 – was available in production, but it was felt that there had been insufficient time for testing to recommend sites to upgrade prior to the challenge. NIKHEF did upgrade to this version during the challenge, as a result of running into a problem (lifetime of resized spaces was reset to < 2 days) fixed in the more recent version.

M/W component	Patch #	Status
LCG CE	Patch #1752	Released gLite 3.1 Update 20
FTS (T0)	Patch #1740	Released gLite 3.0 Update 42
FTS (T1)	Patch #1671	Released gLite 3.0 Update 41
gFAL/lcg_utils	Patch #1738	Released gLite 3.1 Update 20
DPM 1.6.7-4	Patch #1706	Released gLite 3.1 Update 18

### Storage-ware

The following table shows the SRM / backend versions and patch-levels that were recommended for use at all sites during the May run of CCRC’08. These versions were finalized at the face-to-face meetings at the beginning of April, giving adequate time for sites to upgrade.

Implementation	Version / Patch-Level
CASTOR + CASTOR SRM	v1.3-21 (SRM); 2.1.6-12 (CASTOR)
dCache	1.8.0-15, patch levels 1, 2 or 3 (preferred – patches are cumulative).
DPM	See above
StoRM	1.3.20

### DB Services

The Database services continued to run smoothly during the run-up to the May run of CCRC’08, as well as during the run itself. No significant change in load was seen on the Tier0 database services, which raises the question of whether the access patterns seen

are indeed typical of those to be expected during 2008 data taking and production. Despite some concerns with the new multi-core hardware acquired for startup, the migration to this hardware was performed smoothly prior to the beginning of the May run, using the well-established “Data Guard” technique. Not only does this significantly reduce the service downtime with respect to more classical approaches, but in this case also enabled the previous hardware to be retained in standby mode to the new. This would have allowed the service to have been resumed on the old hardware within a short period (some 30’ during working hours), should the hardware problems have turned out to have been worse than expected.

Unfortunately, there was insufficient time for adequate application testing of the latest Oracle release – 10.2.0.4 – and this is now foreseen for June 2008. This version includes numerous bug fixes, as well as the latest quarterly security patch set and experience earlier in the year re-confirmed the importance of being on the latest release if important patches are to be available within a short time (required for several of the services labeled by the experiments as “very critical”).

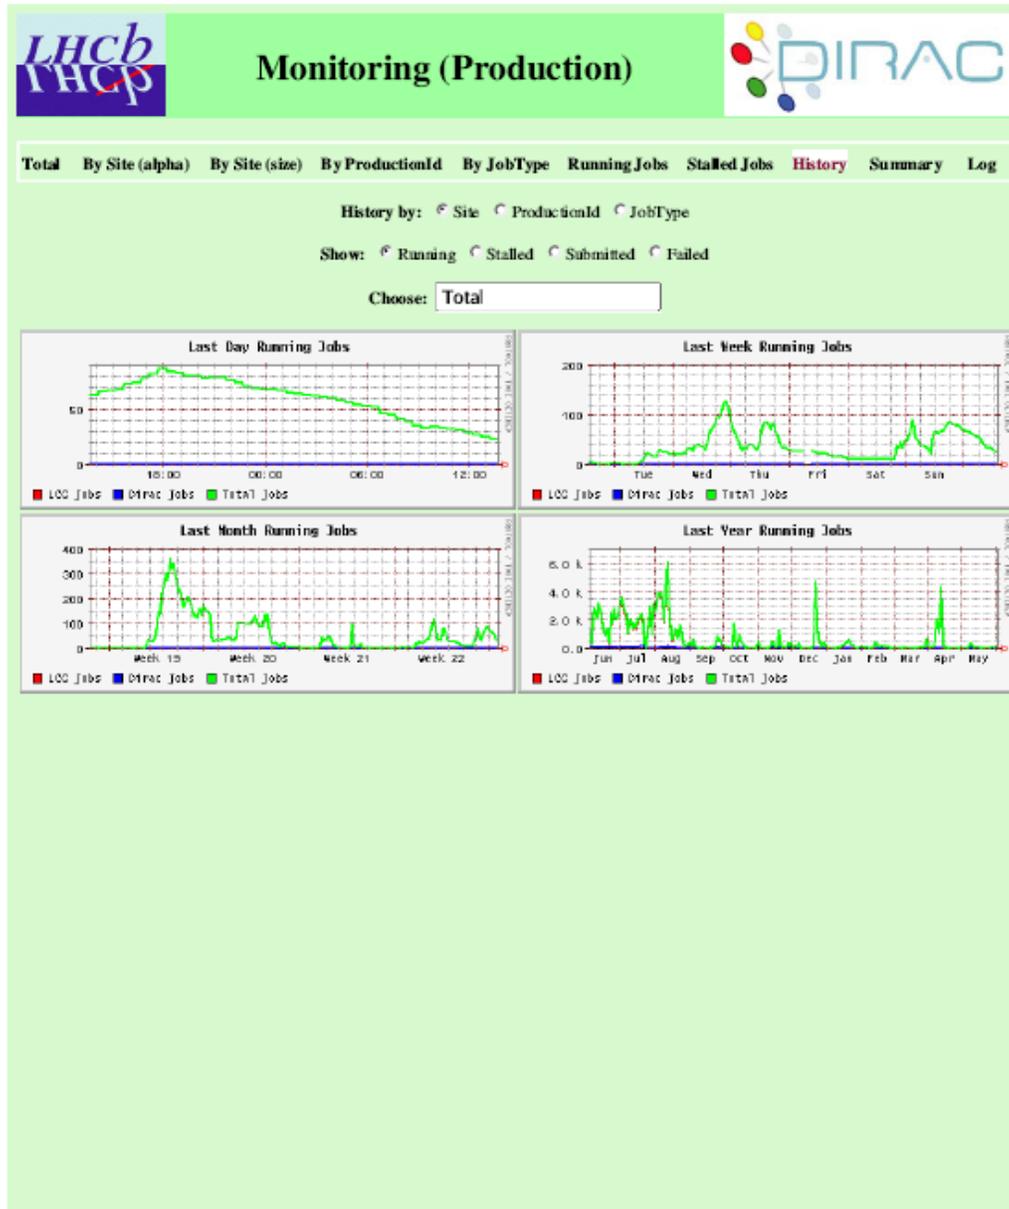
Some scalability and stability issues were seen throughout May with the DB services behind the CASTOR SRM. Pending application-level fixes, some parameter settings are foreseen at the database level which would render the service more deterministic and stable, but require extended handling of error conditions / timeouts by the experiments data management frameworks. Again, these changes are expected to take place in early June 2008, coupled to targeted load tests by the experiments to verify that the expected improvements in behaviour are indeed seen.

Database services at the Tier1 sites continued to run smoothly, but again the expected ramp-up in load, for example from ATLAS conditions due to re-processing activities, was not seen.

### **Monitoring, Logging & Reporting**

Following numerous presentations at the Grid Deployment and Management Boards in late 2007 / 2008, a clear model for the hour-by-hour and day-by-day running of experiments’ productions was confirmed. Basically, this consists of experiment “shifters” who use the Dashboards, experiment-specific tests executed through the SAM framework (together with additional monitoring, e.g. PhEDEX, Mona Lisa etc.) to follow the progress of the work and spot any problems. Using the elog-books and agreed support channels (GGUS etc.), issues are reported and generally resolved rather rapidly. This is a simple mechanism that works satisfactorily. Whilst with experience the amount of human “eye-balling” of screens and increased automation (and robustness!) can be expected, this model works sufficiently well to support 2008 data taking. Examples of key monitoring tools are given below.





## Readiness for LHC Data Taking

Based on the experience with the two phases of CCRC'08, together with on-going and previous challenges (dress rehearsals, data challenges, service challenges etc.) we can now – with justification – claim that we are ready to *face* LHC data taking. This does not mean that the service runs all of the time without problems (it does *most* of the time), but implies that issues – when they occur – are both at an acceptably low level and are solved within a reasonable time frame. Improvements in both of these aspects were clearly observed between the February and May runs of CCRC'08. Furthermore, we have now shown over a period of at least years that we are able to find solutions – albeit stop-gaps on occasion – to even the most daunting problems. Not all aspects of every computing model were fully demonstrated even in May – some “late surprises” that may

come only with the full load and access patterns of real data taking and processing should be expected. However, the above experience suggests that we will be able to work through any issues encountered, finding a short-term solution as well as longer term strategies, within a relatively short period.

### **Outlook for the remainder of 2008 and beyond**

The LHC experiments have been instructed to be ready for beam(s) at any time from the second half of July. Whilst the exact commissioning schedule of such a complex machine cannot realistically be defined in advance, we must assume that we are already in data taking mode. Only very minor and well motivated changes can now be foreseen (typically at the level of configuration / parameter settings and aimed at improving robustness). These must be scheduled taking into account not only their individual impact but also the overall service – it is not an “opportunity” for each individual service to make last minute changes. Additional (bug) fixes and / or well planned and needed upgrades (e.g. security patches) will be required throughout the remainder of the year and will need to be scheduled taking the accelerator and first data processing timelines into account.

Even though an (annual) accelerator shutdown can be expected, re-processing of the previous year’s data, together with preparation for the next, will continue at full tilt. A growing list of changes that can be expected for 2009 is already appearing and managing the deployment of these concurrently with a full production service will be one of the key challenges ahead.