



LCG Services Report 2007 Q4

29 February 2008

Jamie Shiers

The main focus for the LCG services in this quarter continued to be preparation of the needed services for the experiments' dress rehearsals and the Common Computing Readiness Challenge (CCRC'08). Two phases are foreseen for the latter – from February 4th – February 29th and May 5th – May 30th. An agreed set of middleware and storage-ware versions for use was published and all sites urged to upgrade to these versions. Closely coupled to the preparation of the challenge has been deployment of SRM v2.2-capable storage managers at WLCG sites: the target was to have these in production at CERN and the Tier1s by the end of 2007, with upgrades at Tier2s continuing through January and February 2008. The first part of this milestone was met – although configuration of the space management features of SRM v2.2 continued into February. CMS, for example, reported a significant improvement during the month, with almost all sites passing their certification tests by the end of the month.

Preparation has been more thorough than for any previous challenge, but it has still been difficult to stick to the agreed schedule and bug fix releases have been made right up until the end of January. In order to improve on this for the May phase, it is proposed that a set of versions be agreed at the April face-to-face meetings, based on software that is already released or close to release (e.g. in PPS, in certification).

A [CCRC'08 wiki](#) has been setup, linked directly from the LCG [home page](#). In addition to the monthly face-to-face and weekly planning conference calls, daily “operations” meetings have been started since early 2008. Amongst other details, such as the experiments' requirements and the various metrics that were used in February to determine our degree of success, the wiki contains a list of [known issues](#), with their consequences and possible workarounds.

The February run of CCRC'08 is considered to have been a success. Whilst it did not exercise every aspect of the service, it nevertheless provided a very important focus and allowed us to debug many aspects of service delivery. The daily meetings proved a good means of ensuring information flow and problem follow-up and the use of the elog book for this purpose proved valuable. The mechanisms by which experiments can initiate expert call-out (via the central operator) need further publicity and ATLAS have requested that such mechanisms be extended also to Tier1 sites.

Targets for intervention and resolution at the Tier0 were agreed (see table), with corresponding targets also for Tier1s and Tier2s. If these targets are not met a post-mortem is automatically triggered.

Only two issues triggered a post-mortem during February – a prolonged site downtime due to power outage – after which some core services took many hours to recover – and data export problems from CERN. The former problem highlighted the need to refresh – and also use – emergency site contact phone numbers, whereas the latter would have been addressed by expert call-out as in the table below.

Concurrent data exports from all experiments occurred for several days, with a new record for peak exports of some 2.3GB/s. This is roughly 1GB/s more than the average sustained during Service Challenge 4, running as *dteam*.

Time Interval	Issue	Target
End 2008	Consistent use of all WLCG Service Standards	100%
30'	Operator response to alarm / call to x5011	99%
1 hour	Operator response to alarm / call to x5011	100%
4 hours	Expert intervention in response to above	95%
8 hours	Problem resolved	90%
24 hours	Problem resolved	99%


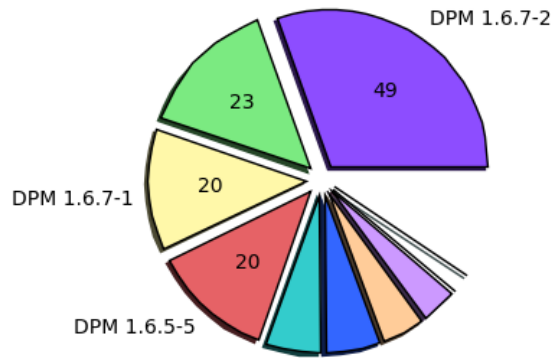
SRM v2.2

As reported in the last quarterly report, a production deployment schedule was agreed at the end of September 2007. In October 2007, production deployment of the storage services supporting SRM v2.2 started. By the end of December 2007, the Tier0 and all Tier1s had completed the deployment as per the agreed schedule. Numerous Tier2 sites had also upgraded, although the schedule for such sites is somewhat more relaxed (the last Tier2s – those that source dCache through OSG – are scheduled to have upgraded by the end of February 2008.) As a result of considerable effort by all involved, all major remaining problems in the implementations were fixed during this period. A few are still unresolved and it will be a key priority for the coming quarter to find and agree fixes – or most likely work-arounds – that are consistent with the timescales of the May run of CCRC'08 and pp data taking. The experiments have defined their needs in terms of site configuration, but this will be an iterative process in which we learn from production usage in February and most likely further tune the configurations prior to May.

The production deployment and configuration of SRM v2.2 services at the majority of WLCG sites can be considered an important outcome of CCRC'08. A snap-shot of the status of deployment is given in the figure below, with up-to-date monitoring available at <https://twiki.cern.ch/twiki/bin/view/LCG/GSSDSRMv2DeploymentMonitoring>.

Current status of deployed SRMv2.2 endpoints (Sum: 161 Count)

24 Hours from 2008-02-28 09:25 to 2008-02-29 09:25 UTC



Middleware

Apart from additional bulk methods for DPM/LFC – made available in Q3 but only moved to production after the ‘green light’ from the experiments – no significant new functionality was made available during this period, with the exception of the inclusion of AMGA as part of the middleware stack. (AMGA has however already been used by the experiments and other communities.)

More importantly, the move to SL4 and VDT-1.6 has been implemented for most of the services. The WMS and LB are in certification, while the FTS is being prepared for certification.

Whilst this did not add new functionality it is nevertheless crucial to permit new resources – which typically only run under SL4 – to be exploited. In addition the move to 64bit versions started and resulted in 64bit versions for the Data Management components.

Numerous bug fixes were also made, related to the introduction of SRM v2.2 and in relation to preparation for CCRC’08. Following the “Critical Services” lists from the experiments and a WLCG “Service Reliability” workshop, an action list has been drawn up for hardening further still both WLCG and experiment services. The target is to demonstrate “measured improvement” by the time of the WLCG Collaboration workshop in April and to have met the overall targets for all services by the time of CHEP 2009.

DB Services

The Tier0 Physics Database Services have been going through a preparatory phase for a major migration to new hardware, foreseen to be deployed in production by the end of Q1 2008, ready for the May run of CCRC’08 and consequently 2008 data taking. In preparation, all the experiments’ integration and test Oracle Real Application Clusters have been migrated to 64bit OS and Oracle server version, to allow the

experiments and Grid community to perform extensive tests of their applications prior to their deployment in production on the same platform.

A significant feature of these migrations is that they have been performed using Oracle Data Guard with a minimal service downtime, independently of the database size. The backup and monitoring repositories have been moved to a high-availability cluster configuration and improved unavailability tracing has been established, providing an overview of scheduled and unscheduled interventions.

During the preparation phase of CCRC'08 – i.e. during January 2008 – an Oracle streams bug, triggered by table compression switched on for a different application hosted on the same cluster – prevented distribution of conditions information from the ATLAS online to offline environment for a period of one week. Even though this problem was escalated with maximum priority – and indeed a fix already existed, but only for 64bit versions of Oracle – this incident should be considered a sanity check in terms of problem resolution. It is important to note that this service is critical to ATLAS and an interruption of this nature during data taking would effectively stop ATLAS production.

Outlook for Q1/Q2 2008

After many years of preparation, we look forward to first data taking from pp collisions in the LHC in the summer of this year. It is clearly essential that we demonstrate the full readiness of the LHC Computing Environment and all experiments and the sites that support them are working on this with top priority. The primary goal of CCRC'08 is to demonstrate all production use cases at all 2008 rates and capacity for all experiments and at all sites simultaneously. Progress will be followed on a daily basis with regular reporting to the management bodies of the WLCG based on three sets of metrics: those defined by the experiments themselves, the “critical services” and availability requirements for these and finally the services and targets defined in the LCG Memorandum of Understanding. There is no doubt that we will find some issues and problems during this period. We must, however, remain focused and pragmatic and address these problems systematically, using techniques that have been proven successful on numerous occasions during the data and service challenges of recent years. In-depth reviews of the achievements in CCRC'08 are scheduled for March (first results), April (WLCG Collaboration workshop) and June (final post-mortem of the two phases of CCRC'08). The 2nd – and final – run of CCRC'08 leads into full production services for 2008 data taking and beyond.

Action Items

A relatively short-list of actions arose from the February run of CCRC'08 and these are listed below:

1. CERN-IT: publicize the on-call services now in place, their scope and operation.
2. CERN-IT: update and publicize the mechanisms by which named contacts within the experiments may trigger on-call / export call-out.
3. Tier1s: update the emergency contact phone numbers and implement and publicize emergency contact procedures similar to those offered by the Tier0.
4. Experiments: provide contact names to follow-up on “Critical Service” issues (procedures, contacts etc.) and tests for the various “Functional blocks” aka “MoU Targets”