



December 18th 2007

WLCG CCRC'08 Status

Jamie.Shiers@cern.ch

Since being proposed at the WLCG Collaboration workshop in Victoria, BC, in September, and a “kick-off” face-to-face meeting at the pre-GDB meeting in October, planning meetings have been held weekly (phone conference) with monthly F2F meetings. In addition, a great deal of work has been done within the experiments on defining their FDRs and the functionality to be tested as part of CCRC'08. Together, this has led to significant progress in defining the details of the two challenges foreseen (February 4th – 29th and May 5th – 30th).

Very little time remains prior to the February run and some compromises are to be expected – for example, additional functionality over what is currently available in production cannot be expected; the window for additional bug fix releases is very limited – additional fixes will be ‘fast tracked’, e.g. through use of pilot services (e.g. FTS, LFC) and/or the Applications Area as appropriate.

Much of the discussion has focussed on storage / data management – in particular on the setup and configuration that is needed at participating sites for production SRM v2.2 services. Other areas have not been discussed in such detail and “late surprises” are not to be excluded.

SRM v2.2 Issues

The details for ATLAS and LHCb have been provided and sites are in the process of performing the necessary configuration. A ‘[dash status](#)’ has been developed, showing the status of the specific end-points at the various sites. This shows that a significant amount of work is still required in the little remaining time – working production end-points are required no later than mid-January to allow pre-challenge ramp-up to begin.

The situation regarding ALICE and CMS is less clear. However, it turns out that – in particular for (multi-VO) dCache sites, supporting space tokens for some experiments and not for others is not a realistic option.

It is therefore proposed that both ALICE and CMS use a single space token (see also the discussion below), e.g. ALICE_DEFAULT.

As agreed at the December GDB, a document describing the implications of the current behaviour with respect to space tokens for Get and BringOnline operations is being prepared. After iteration with the storage and site experts, it is proposed that this be presented in January (GDB or CCRC'08 F2F) – but this does not alter the decision to go with what we have for February, as the need for production experience before further discussions is still valid.

Recuperation of Tape Space

The default assumption is that data created during the challenge will be deleted thirty (30) days after the end of the challenge. This is essential to recuperate the tape space. Mixing permanent data with temporary challenge data on the same tapes is to be avoided, as tapes will be recycled in bulk operations and not one by one.



December 18th 2007

Any data that is to be retained must use well identified storage tokens (and back-end infrastructure) accordingly.

This has been discussed at several CCRC meetings and has been successfully tested with ATLAS data at CERN, as well as with CMS during CSA07 (with somewhat less success – some 300 tapes have to have important data copied off them before they can be fully recycled).

Critical Services

Following the WLCG Service Reliability workshop, an action plan has been established covering the services classified as critical – both in terms of the CCRC’08 challenges and with the goal of obtaining measured improvement in service quality by the April WLCG Collaboration workshop. Whilst this is still very much work in progress, a prototype Gridmap-style “Service Map” has been created and it is proposed that this be used to monitor progress in this area.

Site Concerns

Following the presentations by the experiments at the December F2F meeting, additional questions were raised by some of the sites. These were presented to the ATLAS Tier1 jamboree and a revised version has been produced (see MB agenda page). Experiments are asked to provide the requested information no later than the January F2F meeting (30’ slot per experiment) and if possible by e-mail beforehand to the CCRC’08 [mailing list](#).

Tracking the Challenge

Whilst the various tools for monitoring the state of the services and the experiment views (dashboards) are essentially in place, the way in which operational aspects of the challenge are followed up on a daily basis still have to be finalised. As most experiments have a ‘morning meeting’ – in addition to the central operations meeting – this slot is excluded. Most likely, a short-term meeting either early afternoon or end of the day will be setup, so that important issues seen by the experiments can be followed up in a timely manner.

January F2F Meeting

This meeting is scheduled for Thursday January 10th 2008 from 09:00 in B160 1-009. The draft agenda can be found [here](#).

Tiny Files

The record so far for small (rawdata) files is 22 bytes. This is not a single case – it applies to a very large number of files – and is disastrous for the tape system and the total bandwidth to persistent storage. It is unlikely that the target data rates to tape at the Tier0 can be satisfied unless these persistent issues with small files (<< the tape label!) are resolved. There is also a non-negligible overhead related to FTS setup (current transfers of all files show an average of 30’’ total – 10’’ at source and 20’’ at sink, with a long tail out to a few minutes) for the transfer of such files and it is also likely to be an issue for Tier1 sites.



December 18th 2007

Summary

To reiterate a previous message – we need to focus on things that are known to work. Other possibilities may be desirable in the future – possibly on the timescale of the May challenge – but the list of outstanding actions for February is such that adding new issues (be they bug fixes or feature requests) when there is already an existing solution only stacks the odds against us. It is not the purpose of this summary document to list all technical details and the associated options. However, the January F2F meeting has a slot during which the detailed technical choices need to be re-confirmed, based on the experience gained up to that date.