CCRC'08 Recommendations

Executive Summary

1. Where they do not already exist, regional Tier2 coordinators (plus backup(s)) should be established, responsible for the bi-directional communication between their community and the overall WLCG project. They should attend (physically or virtually) the main deployment and operations meetings (GDB, weekly operations), as well as the on-going CCRC'08 meetings (and whatever replaces them). They should be informed of the decisions of the MB and communicate these to their communities. By analogy with the existing list for WLCG Tier1 contacts, the proposed list is wlcg-tier2-contacts@cern.ch;
2. DB coordinators should be established similarly (wlcg-db-contacts@cern.ch);
3. A short-term (time scale: now – May 2008) "Storage Solutions" group should be established, to propose solutions to problems found during (and around) the February run of CCRC'08. The proposed solutions must be implementable on a time-scale that is consistent with the preparations for the May run of CCRC'08 as well as 2008 pp data taking. This group should contain 1 – 2 technical / operations experts from the experiments as well as representatives (technical and managerial) from the storage implementations and WLCG.

Introduction

By definition, CCRC'08 is by far the biggest challenge – in terms of scope and scale – that we have ever organised. In many respects, the planning is more detailed and further advanced. There are, nevertheless, a number of areas of concern, which give rise to the actions proposed above. We discuss the two main areas – communication and technical issues regarding the services – in more detail below.

Communication

In a project of this size, communication is clearly a challenge. Despite many attempts to improve information flow, there are still misunderstandings, misinterpretations and simply lack of sufficient detail or structure in the innumerable wikis, web-pages and other sources of information. For this reason, we propose that coordinators are appointed in the areas of Tier2s and databases with clear responsibilities for insuring information exchange, as listed above.

Technical Issues

It should come as no real surprise that the main technical problems that face us in the final run-up to CCRC'08 are related to the so-called "residual services" that have been deployed over the last year or so. It is important that we use the February exercise to obtain as much production experience with these services – to shake them down and if necessary fix or work-around any issues that emerge. We cannot afford to put off testing against well-defined metrics to a later date – there is no opportunity beyond May other than full production data taking. This has the following consequences:

1. Metrics must be defined for the Oracle conditions and related services that will be tested in February. This includes the impact of and recovery from prolonged down-time of elements of the streaming service (recently down for ~1 week from ATLAS online to offline), with the identification and testing of

possible solutions to “de-block” production activities, should they be stalled but such problems. The necessary monitoring and alarms must identified and put in place. Metrics for the conditions service – akin to those defined initial for DBL3 and later extended to HEPDB in the LEP era – should also be established. (Conditions data must be made available within one hour onsite and a few hours at external sites; data integrity metrics etc.);

2. Data management and in particular the use of new features of SRM v2.2, such as space management, are already known to have some issues. These must be monitored closely with work-arounds and / or solutions implementable and deployable on the timescales of May and 2008 pp data taking must be identified.