

10th January 2007

Summary of User Board Discussions on dCache – Castor2 Migration.

SUMMARY

The proposal to migrate from dCache to Castor was discussed at the User Board meeting held on 4th January 2007. Comments were also received via email. The proposal was considered reasonable except for the following issues:

1. The Castor version needs to support all mass storage options (e.g. disk1tape0) before it can be considered a full-production service for the LHC experiments.
2. Although the deadline of 30 June 2007 was considered appropriate for ending write access to dCache, some concern was expressed at the proposal to end dCache read access from 30 June. It was requested that the Tier 1 centre consider extending read access by a few months.
3. A request for clarification on tape access using the tape/vtp software to be circulated.
4. A request for clarification of the policy concerning disk cache on batch worker nodes to be circulated.

EXPERIMENT SUMMARIES

ATLAS (Roger Jones): Reservations about the migration until a production version of Castor is proven to work (with disk1tape0). Until the end of December, dCache had been used, but some tests had been done with Castor.

CMS (Dave Newbold): Already using Castor (e.g. for CSA06), but not tested all the different storage classes. Need testing of simultaneous use by experiments.

LHCb (Raja Nandakumar): Currently using dCache for production (DC06). A lot of testing needed of full chain including reconstruction and stripping needed before migration to Castor can proceed. The end of read-only access to dCache after 30 June may be an issue (supported by other experiments, what are the manpower needs to do this?).

BaBar (Fergus Wilson): Migration to Castor does not affect BaBar.

DO (Gavin Davies): No comment.

MINOS (Nick West/Simon Peeters): Castor is an unknown to MINOS (no allocation). Problems continue with dCache (reported to help desk) and is often unusable due to multiple hangs. CMS comment that is why they gave up on dCache and went to Castor, which should be more reliable. Request for disk space (0.5 - 1TB) to aid Castor migration for Minos.

SNO (Nick West/Mike Kordosky): The experiment should finish Tier 1 work by the end of 2007 and would therefore like to be able to continue until the end of the year with their current configuration (1.3TB of NFS disk and DataStore software). Question raised over the future of tape/vtp software as SNO has written a lot of data using this and require this interface for continued access.

MICE (Paul Kyberd): Prefer (and encouraged) to use Castor as a new experiment at the Tier-1.

H1 (Dave Sankey): Don't believe that Castor migration will pose difficulties as GRID use of dCache is only as temporary cache.

ZEUS (James Ferrando): Wish to start with dCache in 2007 and then migrate to Castor before June deadline. Reason is that software already works at DESY using dCache. Support the request for extending read-only access to dCache.

UKQCD (Alan Irving): Currently just using NFS disk at Tier 1. Ramping up for GridPP3 era where will adopt software compliant with SRM. No problem with June deadline as will go away and modify software and come back afterwards.

In the subsequent discussion, the policy on use of disk-cache on batch worker nodes was also raised. A definitive answer was not available, but it was suggested that a ticket be raised requesting an answer.