From: Ron Trompert [ron@sara.nl]
Sent: Monday, May 08, 2006 18:52
To: Jamie Shiers
Cc: sanden@sara.nl; Jurriaan Saathof
Subject: Report of SARA for disk2disk and disk2tape tests

Hi Jamie,

Here is some extra input for SARA for the review tomorrow. I hope you find this usefull.

Best regards,

Ron

Disk2disk test

For this test we have used 5 dcache pool nodes and one admin node.
The first day april 3rd we have experienced the same old network problem again which was solved that day. On april 4th and 5th we have had some trouble with a pool node which is is probably kernel related. This problem was controlled by shutting down the dcache gridftp door on that node and leave only the pool running. From april 5th to the end of the disk2disk test everything went without any incident on our side and we managed to do 154 MB/s over a 15-day period. The problems that were encountered were external like the FTS database problems and Castor DNS aliasses dropping out.

Disk2tape test

For this test we have used 6 tapedrives and the dcache pools of the disk2disk test.
For the disk2tape test we were considerably less successful unfortunately. Over a week we managed to do about 27 MB/s. We have investigated the performance bottleneck and we found that the DMF/CXFS server that controlls the copies from disk to tape had an I/O-related memory problem which was largely caused by a nonoptimal configuration.
It is possible to improve and tune this better but that has the consequence that we needed to reboot de DMF/CXFS server and two supercomputers that depend on it. We did not want to do this. We intend in the near future to separate the LHC and supercomputer storage environments so we are better able to tune the LHC storage infrastructure without causing problems for our supercomputer users. We also plan to test disk->tape and tape->disk internally to optimise things and see if we can achieve the required rates. In addition, we have also optimised our dcache in several ways, for example, by incorporating prestaging which will improve the rate we get reading from tape.