From: Jos van Wezel
To: Jamie Shiers; Harry Renshall;
Subject: SIR storage FZK/GridKa
Date: Thursday, July 02, 2009 14:14:20

Dear Jamie,

not sure where this fits but I propose you consider it
as SIR for GridKa

regards

Jos

--

GridKa uses TSM as tape backend to dCache. The months before STEP09
we found that it would be necessary to pool tape drives for access
by different users. At GridKa CMS, ALICE, LHCb and ATLAS all use a
single TSM system and share all tape drives. Because the dCache setup
for CMS is so much different than that of ATLAS, they have almost
all dcache pools configured as pools with tape access, CMS activity
is able to allocate all available drives. TSM is not able
to limit drives to specific operations or users.

A separate tape library manager which has been tested and used since
early 2008 can prevent drive starvation. Managed access to several CMS
and ATLAS datasets was successfull. We decided to update to a newer
version because minor issues, e.g. support for one of our tape libraries,
had been fixed and we could then use the LTO4 drives deemed necessary
for the required rates. STEP09 was also a test for us in this respect.

The update resulted in random failures accessing the drives. IBM engineers
suspected SAN problems, since we had added 12 LTO4 drives, and some error
indicators indeed pointed in that direction. One week before STEP09 we
still had no reliable tape access and decided to setup an second emergency
SAN for the LTO4 drives in order to exclude SAN errors. However only
a few nodes could be connected to the emergency SAN. These nodes had
spare disk space and 10GE networking which would allow them to operate
as stagers to service all other pools.
At the same time one of our engineers found a new configuration in the
library manager software which directed drive control to a special IBM
driver. Switching this off almost immediately restored proper behaviour
but we could not be sure this was not the result of the second SAN

being used.

Few days before start of STEP09 we had to decide if we could participate or not. The emergency SAN and stager pools would not allow full planned rates but seemed stable otherwise. GridKa communicated our status and we asked CMS not to use tape initially, leaving the tape system for ATLAS only use.

First week of STEP09 plagued with HW problems (disks, library, tapedrives), frequent out of mem on pools, core dumping TSM software, etc.

Second week: opted to use 2 more dedicated stager hosts after it became clear that the load cannot be handled by the existing 4 nodes which were already running dcache pools. Reasons were lots of p2p traffic also in between pools on the same host. Stability proved much better on dedicated stagers. Getting stable rates in between 100MB and 150MB/s. Some files not staged because some tapes flagged read errors.

The hasty addition of the dedicated stagers caused a configuration error which resulted in cache trashing on the CMS stager. Files got deleted to make room for newly arrived files before they were copied to the requesting disk pools. Tape reading was stable at 50-100MB/s for CMS

ATLAS files except for 2, all arrived on disk albeit too late to be counted for STEP09.

It is not clear what the cause is for the trashing and this is being investigated. According to dcache.org it is proven to work. Several tests with changed settings to find the cause are underway in cooperation with ATLAS and CMS.

Tape recall tests performed the last days together with ATLAS proved the system is able to sustain nominal rates.