

SAM (REX/Ops) Operations Meeting 2007.11.13

Present: Adam (typist), Randolph, Angela, Robert, Krzysztof, Anil, Art, SteveS, Dave, Dehong, Parag

For reports submitted to the meeting...

The report text is in this font

Comments made during the meeting are in this font inline with the report

Releases:

No new releases this week. Steve is building a new DB server. SAM (client) v8.2.3 (with the new SAM DH Get Lite that locks files). Something in SAMgrid needs to be modified so that it is used. What is that?

Randolph put together DB Server v8.4.0 to include the new Request system states. [including hold_pending, hold_partial]. There's specific support to handle unqualified hold states in the DB and user unspecified hold state changes (the system can figure out which hold to choose). Tested in unit tests.

Still waiting on free time to do the retirement tests. Randolph will do this.

Must make a new sam_products for Art (see 20071106 minutes).

DØ:

Issues/Events :

Nov. 6: Oracle security patches have been successfully applied to d0ofprd1 and d0ofint1 between 7:30 and 8:30am FNAL time. Steve reconnected manually the SAM db servers. Afterwards, all DB's were well.

Nov. 6/7: After enstore was paused for configuration changes, shifters reported that it took some time to come back. Mover (D61BLTO) still needed a hard disk replacement.

Nov. 8: One tape in NOACCESS (PRJ795) but now it seems OK.

Nov. 10 and 11: A shifter reported the encp history page had not been updated properly. enstore-admin seemed to need to fix this problem more than 1 time (?). It seems normal now.

Unresolved tickets :

IT# 3130 : "samTransfer problem at Manchester" --- the problem seems to

have disappeared but just not sure why ... Anybody is investigating or ?

No, one is investigating. Glad it works now. Must have been a problem at Manchester.

IT# 3132: jobs to spgrid.if.usp.br failed and this seemed to be due to INACTIVE SRM disks.

*This is fixed - was a broken version of sam bootstrap. Steve will resolve this ticket. [sam_bootstrap 8.2.0 is bad - fix was hacked in place. Need to get the fix into a new version of sam_bootstrap] [**]*

IT# 3133 : mc jobs to ce00.phy.ph.ic.ac.uk fail (ce00 or gw39 ?)

Seems to be a site configuration problem. Procedure is to send e-mail to the lcg-integration list to request the D0 LCG collaborators to look into this. Usually Andrew/Gabriele/Parag does this.

IT# 3137: Problems with file delivery from d0rsam01 .

There was a misconfiguration that made d0srv071 useless for Mike. This was fixed. Now there's another 2T disk for Mike. Robert wants to remove the work groups, since that makes it hard to know what fraction of the disk is used. Don't use mc99 (Joel?). Just use the dzero group? Production uses d0production. Robert wants to use d0srv072 to balance d0srv071 and stop using d0rsam01 (which is old and needs to be retired). samgrid jobs should also use sam dh get lite (v8.2.3)

IT#3138 : D. Wicke has problem in setting up AliceNext and Parag said that it is because the globus gatekeeper is not able to get its credentials correctly.

Trying to reconstruct lost configuration files.

CDF:

here follows the SAM shift report for the week 11.6 to 11.12.
In general a pretty quiet week, most notable issues:

Nov. 6:

movers 9940B11 and 9940B25 both offline, 9940B25 with a stuck tape (IAC391): issue solved.

Nov. 7:

CDF-LTO3 library manager dead (since ~2am), restarted by George Szmuksta: issue solved.

Nov. 8:

volume IAD503 is marked NO ACCES and encp's are failing on cdf. Tape cleared and back available for use. (No access due to errors occurring on the enstore side (George Szmuksta): issue solved.

Nov. 11:

dcache pools fcdpdata138-1/2/3/4 OFFLINE and not pingable, helpdesk ticket open. Machine fcdpdata138 rebooted and now up and running (Etta): issue solved.

*CDF SAMTV is not accessible from outside of Fermilab? -- No, not the problem -- Angela looked into this - the "current" link seems to be broken. The log file from SAMTV has "unexpected stuff" in it. [**] For some reason the current link is not set correctly. Not running the MIS station, so can't upgrade to SAMHDTV easily.*

dCache - pools are crashing and are restarted automatically. But there are lingering active movers. An active mover is a transfer that is reported by the system -- should be visible in two monitoring pages (doors and movers). These are only visible in the movers (not the doors). These are cleaned up every three days - this is not happening. This prevents setting correct limits on a pool as to how many transfers a pool can handle. Other experiments don't have this problem since they set their pool limits sky high. Alex will report this problem to DESY. Alex has an idea for what to flag so that pools don't crash (but needs testing).

MINOS:

General -

Condor queues are being used on the Minos Cluster interactive systems

Trying glideinWMS for better control, and access to Farm resources.

Enstore -

DCache -

BlueArc -

GPFARM mounts of /minos/data and /minos/scratch have been made and tested. Now need AFS/NFS translated or remounted mount of products and releases.

SAM -

IT 3128. sam_products for sam_station v6_0_5_22_srm
v4_29 describes sam_station v6_0_3_9p1 -q GCC-3.1
IT 1894 Need 'sam locate' in C++ API

Oracle -

OPEN ISSUES

IT 1146: cannot test existence of sam tape location
 IT 1894: sam locate via C++ api, for simple client needs
 IT 1979: Compiler warnings from sam_cpp_api
 IT 2071: quadratic project startup time growth
 IT 3128. sam_products for sam_station v6_0_5_22_srm

PLAN -

Migration to 9940B tape can proceed anytime

DB:

Oracle released the critical patch update on Oct 16.
 We will be scheduling this patch on dev/int and prd databases.

D0 :

- Scheduling the Oracle's critical patch on d0ofint1/d0ofprd1 on Nov 06 Tuesday - Done

- SAM cut v6_10 on d0ofdev1 - To add composite index on lower(file_name) and retired_date -done

- v6_4 index on create_date of data_files against d0ofdev1 done

- Service Request with Oracle on Full table scan Vs Index scan on data_files_param_values.

Development is working on Bug Fix. Bug# 5601578 No time estimate yet for bug fix.
 SAR# 5837876.993

Alerts:

- Tablespace [TS_CFPD_03312000] is [81.71 percent] full

Action : Added 1200MB

- Message=Tablespace [CONSUMED_FILES_P01] is [80 percent] full

Action: Added 2GB

- d0dbsrv2 status (failover node)

The d0dbsrv2 fail over node is having disk or motherboard issues.

The sysadmin has contacted D1. This machine is not under coverage and cost will be time

and materials. \$140 for initial diagnosis and Steve will follow up with the final estimate.

Approval will be needed from D0 for the final cost. In the meantime we've been in

contact Steve White. He has acquired d0dbsrv3 from SAM and will deploy it as the fail over until the other box is resolved. Depending on the cost, the preference is to return to the fail over machine d0dbsrv2.

CDF :

We need a date for Oracle Qrtly fcdfora4~cdfopr2. 1HR
 We need a date for O/S Qrtly fcdfora4~cdfopr2 2HRS

Alerts:

MINOS:

- Need to schedule the Oracle's qrtly patch on minosprd

dCache/enstore:

enstore/dcache report

Scheduled Downtimes:

- o SL8500 at GCC - Thu, Nov 15 ???
 - if we can arrange it with Sun, we will take a brief downtime on the SL8500 to switch the firmware image to a the most recent version, already loaded. This affects mostly CDF and CMS.
- o STK silos on FCC mezzanine - Thu, Nov 15 ???
 - restart ACSLS in order to clear out some stuck mount requests

Operations:

- o d0en
 - ADIC replaced the Robot 2 arm. Later, a pneumatic hose broke. We scheduled all the movers on side 2 down, allowing production to continue using side 1. The hose was replaced the next day.
 - Primary enstore servers (5 systems) were upgraded to new h/w and new raids. Downtime was considerably longer than expected, but d0en was back in production by about 10 pm. Some issues, mostly related to monitoring, were fixed during the rest of the week. A few are still being worked on.
- o cdfen

- CDF-LTO3 library manager problem appears to have been fixed.
- o cdfdca (from Vladimirs's report)
 - A number of tape transfers were in "Wait" state. It turned out that the tape requested by them had been damaged in the tape drive. It has been copied and is back in service.
- o stken
 - back to its full complement of 15 LTO3 drives.
 - stkensrv3 (backups) was rebooted to clear network problems.
- o fndca
 - A number of files (belonging to LQCD) on stkendca10a were larger than 8 GB, and the installed version of encp was an older version that lacked the option necessary to handle this in the dcache. So the encp's failed and the resulting error messages accumulated in the log files until they filled the root partition, causing the system to hang.
 - The volatile pools have been restarted with some changes in their configuration to prevent stores from queueing.