

SAM (REX/Ops) Operations Meeting 2007.10.09

Present: Adam (typist), Randolph, Art, SteveS, Andrew, SteveW, Angela, Robert, Parag, Krzysztof, Doug (virtually), Dave

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:

Nothing new. Stability is good!. SAM DH GetLite is in CVS and is ready to be released. Angela is working on an improvement to SAM-At-A-Glance so that dropped servers don't merely disappear from SAAG (and thus get missed by the shifters).

Testing of Retired Date at CDF delayed by problems with CDF integration DB (see CDF section below).

Robert may make release this week with SAM DH Get Lite (see last week's report for importance). Angela hopes to have SAAG improvements done by Friday.

DØ:

From Kin:

Events :

--- Oct. 3/4: After the Tue. downtime, SAMDbServer.user_prd and SAMDbServer.web_prd_old couldn't be restarted by our SAM shifter. Robert eventually fixed it, probably because some directory links were screwed ?

No, directory links were a red herring - the v5 DB servers were not restarted after the DB shutdown, including the tomcat server used by the new Dataset def editor. For some reason their entries were commented out in the server list file (oh - SteveW did this when he shut them down). The shifter should have turned the DB servers back on, but perhaps the e-mail alerting them never got out. Will work on this better next time. Fortunately not many v5 users. v7 DB server do not require shutdown and restart.

-- Oct. 8: 8 LTO movers were in status 'bad'. As a possible consequence there were >100 pending read requests on CAB.

After the shifter communicated with enstore-admin, all but 2 movers were back in operation.

Some enstore problem (see enstore section below)

-- Shifters have frequently seen/complained about SamTV having irregular/long updating periods.

Known problem since very heavy usage of cabsrv1. SteveS has made the new

SamHDTV available. Shifters should start using it.

Only 2 new unresolved issues/events :

IT# 3068: "Stuck projects" --- this is a ticket which just came in the last 30 minutes or so.

The enstore problem again.

IT# 3067: "high batch job failure rate on cab2"

From Joel/Mike about SAMGrid jobs. Don't know the solution? Probably Fermigrid or OSG issue.

Older unresolved tickets

IT #3042: "Bug in Phased dataset jobs with notification in Samgrid JDF"

Parag has work around

IT #3038: While trying to merge files at GridKa, J. Snow can't get "UK files" (from dodge.hep.man.ac.uk:/data/sam/durable)

IT# 3022/3030 : xmlldb (in GridKa) corrupted....

IT# 3023: `sam.getRequestInformation(afterDate=reqdategt)` has taken a long time.

See last week - Randolph is working on a new version of this command.

CDF:

From Doug:

09/28 ~18:00~21:00 problem with one of the CDF-9940B library arms;
resolved by SSA

10/01 ~04:00~08:15 dCache problem with staging files in the B Pools group; resolved by moving the B Pools to Test Pools and converting equivalent number of recently upgraded ReadOnly Pools to B Pools (to be addressed during the 10/03 downtime)

10/03 ~7:00 dCache and sam nodes were brought down before 07:00 , this e.g. brought sam naming service down and caused two nodes: cdfsam12 and cdfsam22 to develop problems (several dCache nodes had to be revived by hand) many file systems not mounted after the reboots;

"Solution: jonest@fnal.gov sent this solution:
The raids are now correctly mounted on the servers.
we shutdown the PC and rebooted the atabeast
and then powered up the server."

10/03 pnfs/dcache/oracle upgrade

downtime summary

- .) offline db - oracle upgraded to v [10.2.0.3](#)
- .) sam - servers restarted after oracle upgrade
- .) enstore - pnfs node migrated to new hardware
- .) dcache - dcache code upgraded to version
production-1-7-0-45 (=patch level 45)
(many memory management & file restore/copy issues addressed)
- ..) java upgraded to v6 (64bit where OS allowed it)
- ..) hyperthreeding turned on on remaning group of nodes
- ..) regrouped nodes as per dcache team recommendations
- ..) changed limits on the maximum number of clients
reading from the pools to decrease load on the machines

10/04 ~14:00 a group of dcache nodes had to be reconfigured &
restarted due to unmounted file systems at the time of the upgrade,
see above, (cafs were down, so almost no impact)

10/05 network related problems ~7:00:~11:00 affecting dcache/enstore
and many other services including desktops (failure B0 offline network)

cdfsoft missing from sam upload nodes; pnfs mount
missing from some cdfsam* nodes

10/05 - Prior to releasing the GroupCAF to CDF users -
SAM CAF interface testing was performed.
Jobs complained about missing SAM IOR. Original CAFexe restored.
Cafexe testing will occur on test CAF 9-Oct and 10-Oct.

10/06 caf issue with the "local copy" left on from caf tests;
CAF configuration file had test settings. Fixed by CAF team on 10/06.

10/07 ~22:00 station cdf-fncdfsrv0 restarted by RI per RLC to cleanup old projects

10/08 - SAM MC upload failed due to renamed mounts to CDF software. (fixed)

*dCache downtime was performed as scheduled, but perhaps not all problems of
previous version are solved. Still watching system. CAFs were down till the weekend,*

so low load for the past few days and right now.

CDF integration database lost the data_file_locations table. Reason unknown. Randolph reconstructed it (including the volume id column). Should be workable now. Krzysztof confirms. Robert has tried it too and it worked for him. However, the integration database has not been refreshed since Feb 4 2006. So DBAs should do the space studies and a refresh.

Krzysztof did ups start sam_bootstrap for a DB server and for some reason got two DB servers started. One was crashing and burning since was repeated in the nameservice. And got 100,000s of trace and log files. Deleted them and killed the bad DB server. Impressed that Linux could handle this. But the directory table is huge. Need a downtime to fix this. SteveW: Perhaps add code to the DB server run wrapper so that if have too many crashes it won't bother to restart.

MINOS:

From Art:

Minos status 2007 Oct 09

General -

Power glitch Wednesday - sam/mysql servers rebooted, recovered cleanly.
Took till Monday to get Ganglia monitoring back.

Enstore -

DCache -

SRM hung up Wednesday morning, probably due to the power glitch.
Restarted Thursday evening.

BlueArc -

Deployed SataBeast/BlueArc as /minos/scratch and /minos/data

Starting user configuration and testing.

SAM -

Will start declaration of Monte Carlo data in production this week.

IT 1894 Need 'sam locate' in C++ API
Users really need this feature.

Oracle -

minosora3 Oracle upgraded to [10.2.0.3](#), no impact.
Request production upgrade before approx 23 Oct beam startup.

OPEN ISSUES

IT 1146: cannot test existence of sam tape location (noted by shifter)
Here is the problem -- the SAM database keeps a list of tape locations. CDF and DØ use autodestd which keeps this list up to date automatically. MINOS doesn't use autodestd and so if a tape location does not exist in the DB, it must be added by hand by an admin user. There is no way to check for the existence of a location ahead of time except to try and check for failure. Think of a simple way to add a check.

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 (from Nov 1 2005)

PLAN -

Migration to 9940B tape can proceed anytime - when ?

Update Minos dbservers and clients to v8

Insert very LOUD fire alarm here.

DB:

No report???

dCache/enstore:

From Dave:

Scheduled Downtimes:

- o all - Thu, Oct 11 - networking
 - Brief network outage around 6:00 am to reboot FCC1 core switch
 - Public enstore/dcache and all robot communication are affected
 - all enstore libraries will be paused

- o cdfen, stken - Thu, Oct 18 - networking
 - Brief network outage around 6:15 am to reboot FCC core switch
 - CDF, DO, CMS workgroups will be isolated
 - cdfen and stken libraries will be paused

These two downtimes will require draining of the CAF? If so, this is a huge inconvenience. Check on this.

- o d0en - Tue, Nov 6
 - ADIC will replace Robot 2 arm; AML/2 will be down all day
 - upgrade of entire suite of enstore servers to SDE rack
 - enstore s/w updates

Operations:

o d0en

- Scheduled downtime went well:
 - * enstore s/w updated
 - * Library managers and media changers did not get restarted during the update. The libraries were later paused, and the services restarted.
 - * ADIC made adjustments to the robots, which seems to have fixed the misread barcodes.
- A different problem developed Thursday night, with many mount failures on the Robot 2 side, extending thru the weekend.
- ADIC came Monday morning, diagnosed worn bearings, and will replace the entire arm at the next opportunity.
- The samlto2 LM stopped responding Monday afternoon, in turn causing Inquisitor to hang. Both were restarted.

(LM = library manager; inquisitor writes the monitoring web pages)

- We have received a list of 1273 LTO2 tapes to recycle.

o cdfen, cdfdca

- Unplanned power outage Wed morning delayed downtime work (we had to bring the systems back up to take them down)
- Scheduled downtime went fairly well:
 - * pnfs server (cdfensrv1) has a new home
 - * enstore s/w updated
 - * dcache s/w updated to latest production version
 - * Not all dcache disks were mounted after shutdown, leading to problems during restart.
 - * the raid on cdfensrv3 (enstore backups, dcache monitoring) had apparently not recovered well from the power outage; fixed by power cycling it again, which required rebooting cdfensrv3, which was done without interfering with enstore operations.

o stken

- nothing to report

o fndca

- monitor node (fndca2a) needed to be rebooted Tuesday morning, and again Wednesday morning. SRM remained down until Thursday evening.

Went down this morning too. Seems to go down every Tuesday morning. Hmmm.

- SRMwatch is now available (see dcache status page)
- Two write pools on stkendca9a are stopped due to some sort of disk problems. This happened before, some weeks ago, but no

cause was found. We will investigate again.

- Investigation of problems with RMAN backups is making progress. There were changes in common ftp code during Jul 19 downtime, which changed (corrected) the handling of loss of the control link. The perl script which runs the transfers was OK with the previous behavior, under which it had been developed, but not with the current behavior.