

## *SAM (REX/Ops) Operations Meeting 2008.04.01*

Present: Adam (typist), Krzysztof, Angela, Eric, Art, Robert, SteveW, Randolph (driving), Dave

Meeting Length: 45 minutes

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:**

*8.4.4 DB server was released with new Omniorb and new version of Python. Release notes from Omniorb say they've fixed a file descriptor leak, and we had been noticing a slow leak of socket resources. Will watch if this problem really has gone away. New version of Python is latest 2.4.5 (bug fix release). Deployed at D0 right now. Not at CDF yet, nor MINOS.*

### **DØ:**

*Worth noting:*

*--- scheduled d0ora2 downtime on Apr. 1 (08:00 to ~10:00 am)*

*Updating firmware in SAN controller.*

Had configuration differences on the forwarding nodes, OUHEP00 was overloaded, restarted, and killed production at FNAL for no good reason. Andrew claims he's fixed this (not sure if it was deployed).

Close to 32,000 sub-directory limit. Parag is trying a solution for a permanent fix (moving some job info to different directories and putting that dir name in the JDL so Condor could find it).

*Unresolved issues (all seem to be OSG/samgrid-related) :*

*#3490: "OSG: samgrid jobs at CIT\_CMS\_T2 leaving large logfiles" : problem from samgrid running at CalTech ?*

*Need to look at this one*

*#3485: "GridKa FSS full of files" : Joel complained that the time between authentication and stager being contacted is ~12 hours.*

*Joel said this spontaneously fixed itself*

*#3479-3481: "Increased failure rate on OSG ([Fwd: Re: Lots of exit status 1 on d0cabosg2])" : Any ideas ?*

[https://plone3.fnal.gov/SAMGrid/tracking/3481/pcng\\_issue\\_view](https://plone3.fnal.gov/SAMGrid/tracking/3481/pcng_issue_view)

*One worker node became a job black hole. FEF fixed.*

*#3474 : "station\_prd2 stager died at osg-ouhep: status 134" : Is this the same as 3479-81 ??*

*Same as overloaded OUHEP00 machine (files appear faster than they disappear).*

*New version of station 6.0.5.24\_srm on clued0 - fixes the max project limit problem (someone hit the limit yesterday, so this works now). Has been on clued0 for a week, considering moving to CAB. There are other nice fixes as well (better behavior when one empty cache disk - doesn't become a black hole for files, also can disable disks even when in high use).*

## **CDF:**

*Tue Mar 25: One tape that was marked bad just needed to have the NOTALLOWED flag cleared.*

*fcdfdata321 errors needed a dataset registration.*

*Wed Mar 26: Enstore maintenance and major SAM upgrade*

*Thu Mar 27: Enstore returns but everything else is still down*

*Mon Mar31: SAM came back for a bit, but still lots of problems and it remains in the hands of experts.*

*One dead mover: LTO3\_05.mover (powerup fails)*

*Current status: All SAM services seem to be running on new machines, except for SAM-At-A-Glance. [CAF is still not up, since not all of their worker nodes are using the new fcdfcode3 instead of the old fcdfcode1]. Part of the upgrade included moving the SAM nameservice - this is causing monitoring to fail. Because not everything is using the alias.*

*Yesterday, when trying to start services - services were writing to a log area that had a quota limit (by mistake). FEF removed the quota. Hitting the quota limit caused a strange seemingly unrelated error message. Stuff was running by 4pm.*

*Problem now is finding the non-updated nameserver references (e.g. SAAG). And none of the SAM parts of DH at a glance are working. CSLtoEnstore monitoring fails at the SAM DB server heartbeat. SAMTV is also broken. Running blind right now. [SAM-At-A-Glance int works fine - note that nothing changed for int]*

*Offsite push: K and A are creating a new sam\_client\_products with the nameserver*

*alias. If they download that, it will fix the nameservice problem. Remote people will have to update their sam\_config by hand, since their config is different than ours -- have to do a "merge" by hand. Thomas has already done Gridka. Need to do Italy, Manchester, ...*

*Margaret has been putting some of these things into JIRA.*

*CSL is giving max retries errors.*

*Randolph's files uploaded monitoring is also broken.*

*Top priority is to restore monitoring, then help with remote sites and CSL.*

*Stay alert for CAF coming up.*

## **MINOS:**

### General

Minos Collaboration Meeting - Sussex 3-6 March

AFS - Timeouts continue

Enstore -

LTO-4 copies continuing, over 5 TB so far

DCache -

SAM -

Need to upgrade station to v6\_0\_5\_23\_srm - April

Oracle -

## OPEN ISSUES

IT 1146: cannot test existence of sam tape location

IT 2071: quadratic project startup time growth

PLAN -

Migration to 9940B tape can proceed anytime

## **DB:**

No report today. Downtime for D0 today.

## Enstore/dCache:

### Scheduled Downtimes:

o d0en - Tue, Apr 1, 8 am to 10 am

- enstore s/w update, particularly movers and library managers. A memory leak in ftt (used by the mover) has been fixed, and the LM stability has been improved due to experience with CMS.
- OS security patches

### Operations:

o cdfen

- Last week's upgrade went well, at least for enstore. We were finished at 1 pm on Thur, only an hour behind schedule.
- postgres updated, pnfs databases dumped and reloaded
- ACSLS updates were skipped, will be done another time.

o cdfdca

- stopped and restarted

o stken

- stkensrv3 rebooted to clear a problem with its raid array. this node handles backups, monitoring, and is none-critical.

o fndca

- A pool node had to be rebooted when the system disk filled up due to unchecked growth of log files.

*Stan announced at CD Ops meeting a possible downtime stken on 4/17 -- possibly for PNFS. Check in next week.*