

SAM (REX/Ops) Operations Meeting 2008.02.12

Present: Adam (typist), Steve W, Randolph, Jerry, Margaret, Krzysztof, Art, Robert, Rick St. D (virtually), John Yoh, Steve S, 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:

DB server v8.4.3 is in production at D0, in integration at CDF. Steve W wants this DB server in production for everyone, since it will take care of the cache file issues (move them to the history table). For D0, the move of old cached files to history is complete. Krzysztof says that there will be a downtime on 3/18 and that's when the DB server can be upgraded (Steve W will be FUR that week). After moving to the new DB server, run a script to move current records.

New sam_cpp_api released v8.4.0 -- should be viewed as a test release. Now uses Omniorb instead of Orbacus (mysterious hang on exit is fixed). New features: handling of asking for max files of zero is now consistent with everything else (zero=infinity), sam locate is in (!!), there is now a timeout for waiting for file deliveries (is an argument to the get next file function with a default of forever). Added routine to put parameters into metadata directly from C++. This version *should* be source compatible with old code, but not extensively tested. Incompatibilities are unintentional -- Report problems back to Robert. Robert will try to put this version into the D0 build system. For CDF - ask Thomas first (though this is not his duty anymore). The library size is now less than 50% of the old library (probably due to Omniorb).

DØ:

--- User "Herb Greenlee" reported/complained a couple times about SAM HDTV broken ~Feb. 7-9, 2008 (as in #3365).

Steve is in communication with Herb.

--- Oracle security patch to be applied to d0ofprd1, d0ofint1 and d0oflump databases on Feb. 12, 2008.

Unresolved issues:

#3372: Remove d0rsam01 from the forwarding node mapping list --- waiting for Robert to say it's OK to do so.

Waiting for check of premature file deletion fix.

#3369: Premature file deletions from sam cache: a test version of sam_client installed and the situation is being observed ?

Steve S has deployed the new code, but can't check due to cabsrv2 downtime.

#3368: Problems with station context not found -- Andrew said:"The real reason is likely due crash of the context_service itself when trace file rolled over 2G. The other problem is that it did not restart by itself. That needs to be looked at." ???

Robert disagrees - further investigation

#3363: Status of the location of `dodge.hep.man.ac.uk:/data1/sam/durable` ??

New cert request was rejected twice with no feedback. Steve S opened a helpdesk ticket. Awaiting reply.

#3351: Adding OSG site --- waiting for Tom to confirm whether he's satisfactory as Joel said that he could run a job already.

We think its working - need verification from Tom.

#3344: MC jobs held --- Adam said that the OSG station will be upgraded to a "non-leak" station Tuesday (Feb. 12).

Robert is deploying the latest station today.

Resolved:

User complained that jobs were timing out --- user had stored a very large number of files in the DB but had set the size to 1 MB, though the size was really 0.5 GB. This confused the station. The meta-data code in the DB server does not check the size if you supply it (if you don't supply it, it will fill in the size). Look into adding size verification in the DB server. [!!] Robert corrected the DB for these files.

Broken tape: not worth recovering since all but ten files are in cache. The 216 files have been saved to a new tape.

Some other SAMGrid issues were resolved too yesterday.

LCG Wuppertal forwarding node is fixed and is supposedly up!

samgfw03 motherboard swap won't happen till Friday at the earliest.

CDF:

- o Known and continuing problems
- o `fcdfdata320` offline

- o Issues added after last meeting

- Issue 3555 : when user requests 0 files, jobs don't run properly since the consumer process thinks it is done. A user is likely to do this with a buggy script that sets the number of files in an environment variable. If the C++ interface realizes it is using SAMDH then it should exit and clearly state that zero files were requested.

- Issue 3557: Project Master Polling Failure; Project master gets a failure in polling a consumer process. The user program aborts prematurely.

- o Pool draining

- r-fcdfdata138-1 moved by Angela to testpool to train, then returned to regular service

- o TNS operation timeout

- had one but recovered

- o Bad movers and Tapes (fixed without shifter intervention, Thank you Enstore group)

- LT03-04,23,27,30 tape IAD 419

- LT03-05,26

- LT03_29, IAD988

- o Enstore Downtime as noted by David (and seemed to go smoothly, thanks, again)

- o Alarm due to CRC checking error; Quite a long discussion about levels of checking and enstore/dcache/sam. Defer to experts for final discription in Dave Berg's report.

- o User Problem: Chris Neu could not run a top stntuple dataset. Turned out a snapshot had not been taken. He took a snapshot.

- Two comments:

1. Taking a snapshot of a dataset after defining it should be part of the procedure since the dataset definition can be changed until the snapshot is taken. This is something to watch for since a user may take a standard definition and append _0 or A or something and if two users happen to do this at the same time, it can lead to confusion as to who gets there first.

2. If a sam dataset is used to submit to caf, is it true that the dataset is verified and that works only if a snapshot is taken? Then if a project is started does the snapshot get taken by SAM? How does the first snapshot get taken? This is related to a later query by Ray Culbertson concerning who takes snapshots.

- o Diskpool node was down. Was reported to dcache incorrectly by me, then sent to cdf diskpool but no action was taken until a CDF user (Lina Galtieri) could not access data. Then it was corrected.

- o Upload issues:

- User (Craig Group) forgot to define dataset, was contacted and corrected problem and thanked us.
- User (Sasha Golossanov) had a couple of failures in loading pskt3w. First, no such file or directory mentioned. Rob Harr had not seen this before but instructed Sasha that after a failure, files were put back in a safe area for retry. Secondly, there was an inability to get metadata for the same file. This problem remains.

o User Problem: User unable to access bottom stntuple. The problem was that the stntuple uses SAM as the catalog but not to deliver URLs. That it forms in its code by querying the SAM database directly for the file location and builds the dcache URL. The user had a very very old copy of the code. I recognized this problem in June 2006 and Ray Culbertson patched the code. For future reference, the stntuple user sees files listed with something like:

```
dcap://cdfdca3.fnal.gov:25159/pnfs/fnal.gov/usr/cdfen/filesets/NB/NB00/NB0003/
NB0003.5/xi035993.0196bhdi
```

```
xbhdii.002 1307495424.000
61973 219539 3066869 219539 3798018 0.000 0.000
```

```
dcap://cdfdca3.fnal.gov:25159/pnfs/fnal.gov/usr//cdfcaf-backbone:dcap://
cdfdca3.fnal.gov:25164/pnfs/fn/xi035993.0196bhdi
xbhdii.002 1307495424.000 61973 219539 3066869 219539 3798018 0.000 0.000
```

The first line is fine, it is the second line, with cdfcaf-backbone that is the problem.

The code is Stntuple/loop/TStnCatalogServer.cc and was fixed in 2006/06/22. It is difficult to track stntuple as the tag name is kept the same and the code is changed in CVS. But this was rev 1.34 of the file. One should check user code for this.

A user has to have very old code and I am very surprised that this code works at all for many other reasons! This may never come up again.

o Pools down: r-cdfdata135-1 through 4 (turned out that admins said the node was not pingable and hardware folks worked on it)

r-cdfdata111-1 through 4

This goes to dcache-auto so I only note this. Angela sent mail. Is that needed? DCache group magically seems to fix this. (Thankyou)

o Angela is getting a new maillist deployed for the sam test jobs

o CSL error came up but it seemed a communication issue. But discovered the name of the node is not that on SAG, have to use b0gateway or b0doorway to log in and I am not in k5login. What should we be doing on shift with CSL?

K: Only CSL experts are in the .k5login ; There will be a new production series - some files from the "0j" group will be declared bad (don't need to be retired or removed since correct files will have new names). These files will be replaced by the "0k" production run. Randolph created datasets for production and standard ntuples for 0k and files will be re-processed in the next few days (period 14). Are any of the "0j" files in the station cache already?

Randolph: DBSmon (tracks DB usage by user) had a remote logging server failure (the report of the top five users using the DB wasn't populated) - was restarted (first time in a year something had to be restarted). Randolph will document the restart procedure. No evidence of the failure was left behind (bummer). DBSmon runs under "cdfdb".

MINOS:

General -

AFS - Timeouts continue.

Ran a Minos loon framework job under Parrot, on a node having no Minos software installed.

Glideins for analysis jobs now allow selection of GPFARM nodes having AFS.

Enstore -

DCache -

Writes blocked by CDMS writing 30K 10 MB files to the write pools, Monday afternoon. May take all week to clear.

SAM -

Working to use the new C++ API with sam locate. (!!!)

IT 3128. sam_products for sam_station v6_0_5_22_srm

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's quarterly Security patches are expected to be release on Jan 15, 2008
 CSS-DSG group will be scheduling downtimes for Dev/int and production databases
 in coming weeks.

D0 :

- Scheduling deployment of security patch on d0ofint1/d0ofprd1 on 2/12/08
 - Scheduling Schema cut on d0ofprd1 during March downtime - To add composite index on lower(file_name) and retired_date
 - Drop the unique index on lower(file_name)
 - Recreated as non-unique on file_name.
 - dropped the unique constraint on file_name and retired_date
- Alert(s):

Message=Tablespace [TS_CFPD_03312000] is [80.1 percent] full

Action: Added 1.6GB

- Thu Feb 7 06:01:01 CST 2008
 Information for: d0ofprd1

Owner Object Type Extents Used Extents Allocated Extents Left

 SAMPRD FI_WG_FK_I INDEX 196 200 4

Action: Increased the max extents to 500.

CDF :

- Need to schedule the patch on cdfofpr2 and v6_10 SAM Schema cut
Will probably the 3/18 downtime.

Alert(s):

MINOS:

- Need to schedule the patch on minosprd as well as v6_10 SAM schema cut

Randolph: Need a number pairs global temporary table has been requested - Anil will review after his vacation/FUR.

Enstore/dCache:

Scheduled Downtimes:

- o D0 - Tue, Feb 12, 8 am to noon
 - enstore s/w updates
 - cfengine upgrade
 - yum updates
- o STK - Thu, Feb 21, TBD
 - selective enstore s/w updates (esp. media changer)
- o CDF - Tue, Mar 18, all day
 - upgrade of entire suite of enstore servers to SDE rack

Operations:

- o d0en
 - LTO2 to LTO4 migration has begun.
 - A problem with the AML/2 media changer caused many LTO2 drives to go offline. Developers updated the code on Thursday and it has been OK since.
 - Robot 1 in the AML/2 was down for about 2 hours Monday evening. For part of that time, enstore remained operational using only Robot 2. The problem was fixed by shutting down the AML/2 and restarting it. Quantum will check the log files.
- o stken
 - LTO4 drives are available.
 - About half the old SLF 3.x movers have been updated to SLF 4.4
- o fndca
 - nothing to report
- o cdfen
 - Last week's s/w update went well. Unfortunately, the update had a bug that caused it not to fix the problem for which it was

intended; at least it doesn't appear to have made it any worse.

o cdfdca

- A file could not be restored from tape because of a CRC mismatch between layer 2 (dcache) and layer 4 (enstore). The latter was in fact correct, and the former had somehow become corrupted.

Randolph: Desirable to have access to enstore DB to determine if file is on a tape that is unavailable (request from Rick St. D - wants to know if there will be problems getting files upon job submission).