

SAM (REX/Ops) Operations Meeting 2007.12.11

Present: Adam (typist), SteveW, Art, Doug, Randolph, Angela, Anil, Krzysztof, Dehong, Dave, Andrew, Robert, SteveS, Valeria (remotely, but connection problems).

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:

Anil will put new data_files retired_date indices into production so that tests can be run.

Robert looked at retry logic given the d0ora2 problems. Robert found that if d0ora2 is up but the db is down, the wrong exception is thrown (an internal exception is thrown by the DB server -- though this is apparently deliberate according to the code comments). Robert changed it to throw a DB_SQL exception. Robert added new retry logic to the client to retry on that exception -- this new client requires the new DB server. Testing is tricky (how do we test -- can test by giving a non-existent DB name -- or try during a DB patch update). While doing this, improved error message when d0ora2 is down. New code to dump the request system transition table for testing and users FYI. Now waiting on Steve's fix to the request system python API.

The SAM C++ API that we build is against SL3 which, for D0 and MINOS, runs on SL4 too. But for some reason, at CDF, the API library fails. Robert will try to build with SL4 libraries. We are using a very old version of Orbacus -- can't upgrade since the newest version is commercial and they've dropped some functionality that the sam station uses. The CDF C++ API had orbacus built for RH 7.1. The "0b" version of the C++ API has orbacus built for SL3. CDF will try this version first, and if it fails Robert will do a SL4 build.

DØ:

Kin Yip to sam-design

Events worth stating :

- ~Dec. 8/9: d0ora2 went down 2 times leading to a CPU being replaced after noon on Dec. 10, 2008.
... and somehow somebody needed to remember to restart those v5 DB servers.*
- Dec. 5: d0ora1 had CPU replaced.*
- Dec. 4 (~10 am) : Database upgrade (to oracle v10.2.0.3) and system patching has been completed. d0oflump is available to users now. (Diana)*

Unresolved issues:

#3236: grid job completed while batch jobs continue to run
Parag is still looking into it

#3235: unl transfer problem to d0srv071/d0srv072: due to wrong host certificate

Installation of 072 had a copy of the host cert for 071. Oops. For now, 072 has been declared invalid. Getting cert. (Robert has a request in for a host cert).

#3234: It is suggested that the option "Loopcheck= NoCRC,NoBody" should be added to the sam-oncall mailing list to avoid duplicate notification. Is this being done ?

From Randolph - Adam will look at this.

#3233: Open Science Grid: Failures from samgrid.indiana.edu at Nebraska . Can we help them ?

New cache at Indiana was missing a product. Disk is declared invalid. Thom at Indiana has to do the install correctly.

#3226: "SAM not responding" : initially d0ora2 problem and right now seems to be due to "cabsrv1"

DB down

#3222: "Station 'osg-ouhep' not responding" not responding .. waiting for the user to reply to close the ticket.

Closed

#3220: "MC production (@stitch.oscer.ou.edu and spgrid.if.usp.br) fails, merge succeeds" ... Please help.

Parag is looking at this.

#3219: "Some jobs not arriving at the LCG sites". Also checking with the user whether this is still a problem.

Parag is looking at this.

Arrrgh! D0ora2 is down again!!!!

CDF:

Valeria Bartsch to Adam, sam-design

Hi all,

this is the CDF SHIFT REPORT:

Wed Dec 5 02:37:12
downtime from the 4.th of December is over.

Thu Dec 6 01:52:53

ENSTORE:
CDF-9940B bad mover 9940B27
CDF-LTO3 bad mover LTO3_29, LTO3_31

Thu Dec 6 11:27:30

Alex wrote:
kerberized dcap door logs are full of messages when job with wrong credentials (=expired cred.) tries to access door.
Please grep dcache logs to get host IPs, there are few of them.
There records are like follows :
[enstore@fcdfdcach5 dcache-log]\$ fgrep "Tunnel verification failed"
~enstore/dcache-log/dcapK02-fcdfdcach5Domain.log
[enstore@fcdfdcach5 dcache-log]\$ host 131.225.240.20
20.240.225.131.in-addr.arpa domain name pointer fcdflnx8.fnal.gov.

Fri Dec 7 03:44:28

In order to enable new sam project summary and recovery features related to the new CAF-SAM interface being deployed on CAFs we had made sam v8_1_6 current on the central clusters.
This version, among others, provides more information in the output of the sam get project summary command.
(\$SAM_DIR/RELEASE_NOTES provide more details)
Krzysztof (for DH)

over the weekend:
problems seen for the FermiCAF that jobs running on SL4 nodes were not able to transfer their input sandbox to the worker nodes. The problem is being investigated.

Angela discovered a dCache group of popular files and followed the procedure to spread out the files.

Doug states that using the C++ API, which was brought in by the SAM-caf interface in the CAF executable. The executable was built on SL4, it worked on SL3, but does not work on SL4 (see above). Strange. See above for action items. Doug notes that many jobs on FermiGrid were lost by this problem.

MINOS:

Minos status 2007 Dec 11

General -

AFS - Timeouts of some servers are continuing

Enstore -

DCache -

BlueArc -

AFS/NFS translated/remounted mount is not working yet.
Proceeding with AFS installation on limited farm nodes.

SAM -

Monte Carlo input and output are fully declared,
in keepup mode since 1 Dec.

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:

D0 :

- D0ora2 down 12/08/07 10:13AM - 12:05PM
12/09/07 8:33PM- 09:51PM

SystemAdmin is working with Sun. Seems like one of the CPU got faulted.

Sun Engineer is planning to come on-site and run the diagnostics on Monday
12/10/07

d0ora2 is down again -- error is reported on the same CPU that was replaced!!

- CPU was replaced on d0ora1 on Wed 12/05/07 Bad CPU was cause of the crash.
- On 12/04/07 SAM cut v6_10 on d0ofdev1/dofint1 - 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
- On 12/05/07 SAM Cut on d0ofdev1/d0ofint1 - To add history tables for CACHED_FILES and CACHED_FILE_PROJECT_USAGE tables.
- v6_4 index on create_date of data_files against d0ofdev1 done

Alerts:

Message=Tablespace [PARTITIONED_IDX_01] is [80 percent] full

Action: Added 8GB

CDF :

- On 12/04/07 SAM cut v6_10 on cdfofdev/cdfofint.
- Done - Dec 04 8:00Am-12:00PM
DB patches fcdfora4~cdfofpr2
O/S Qrtly fcdfora4~cdfofpr2
Also OBP upgraded to version 4.22.34

Alerts:

MINOS:

- On 12/05/07 SAM cut v6_10 on minosdev/minosint
- Scheduling the Oracle's qrtly patch on minosprd in second week of Dec 20 (Decache may be down)
dbas will pick a day, since every day is bad for MINOS.

dCache/enstore:

Scheduled Downtimes:

Operations:

- o SL8500s at GCC
 - The ACSLS patch to recognized LTO4 drives was installed.
 - 17 LTO4 drives are installed at GCC and connected to movers, which have had OSeS installed and are being commissioned.

- o cdfdca (from Alex's report)
 - Pools were restarted to pick up configuration changes.
 - Head node was patched to fix exception in request controller.

- o fndca (from Vladimir's report)
 - A disk in the third raid array on stkendca13a went bad; a service call has been placed.

- o doen
 - Burn rate and slot usage plots have been fixed.

- o stken
 - All the stken LTO3 movers went offline. The primary was paged, and found the system disk on stkensrv4 full. The LTO3 library manager debug log file was the culprit.

KLG: When will LTO4's be commissioned? Running now on the test stand. Making duplicates of old D0 tapes to the LTO4's.