

SAM (REX/Ops) Operations Meeting 2007.10.30

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

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:

New db_server_base_cx is released (sam-db-server-package 8.3.1). Robert hasn't done the new Omniorb yet (minor bug fixes only -- go into next release).

v8.3.1 has nameservice re-registration. No Oracle client change here. We're running the latest Oracle client (last checked about a month ago). This version is not yet deployed! Waiting on the retired date tests. New CDF MC procedure will benefit from retired date (e.g. test the MC files AFTER they've been put into SAM - so may have to retire files that are bad).

DØ:

Kin Yip to sam-design

Incidents:

--- Oct. 25, 2007: At around 8:45 d0ora2/d0ofprd1 crashed. Also SAM was not usable, and several webpages like "SAM at a glance", where not available. SAMDbServers/SAMOptimizers all seemed to be down (red).

Experts rebooted d0ora2 and eventually all went back to "green".

Still no clue as to why the machine crashed. Machine went off network at 8:09. Robert noticed that there was a DB server DB connection that was in a SELECT call when the DB went down. When the DB came back up, the DB server did not notice and became stuck. This stopped the fnal-cabsrv1 station until the DB server was restarted. This problem must be a bug in the Oracle client or cx_oracle (note that Steve has released a db_server_base_cx with a newer version of cx_oracle [see above] but there's nothing in the release notes regarding this problem).

--- Oct. 27, 2007 : Shifter noted that cinvestav-station was not responding and mailed its admin. It was put back to work

--- Oct. 29, 2007 : slowness in clued0 station noted.

Due to fcp failing (all the time on jetsam). Unfortunately, tests did not catch this problem. Angela is looking into it.

Unresolved issues :

#3101: "Willem complained that sam/LCG didn't submit all jobs". Can some experts comment please ?

SteveS looked into it - two of his three jobs he complained about did complete. One is still in Held state. Willem has a new ticket that we'll have to look at. Parag: possible problem, lcg client sometimes chokes on submitting jobs to lcg - then the SAMGrid code times out. Workaround for now is to resubmit the jobs. Not sure why Willem thinks that some jobs had failed but are indeed working. User education is necessary here. What is the sign that people should look for to determine that their jobs are indeed really stuck?

#3098: "What is the sam dimension constraint to restrict data files to those recorded during a store?"

Is there such a thing ? Anybody knows the answer

We don't believe that there is any meta-data that specifies if a file was really taken during a store (since runs that are taken after the store still have the last store number in their meta-data -- this is an online feature). The data quality group goes through these runs after the fact, though, so non-store runs can be filtered. There used to be a flag manually set by the DAQ shifter to say if a run is within a good store - but the input was unreliable and is ignored.

Question : When will we switch to use SAM HDTV ? From last week's minute, I thought it should be available by now ?? At least <http://d0db-dev.fnal.gov/sam/backdoor/> now still points to the old SAM TV.

I will update the sam tv links.

First Tuesday downtime is a week from today. No specific SAM work (but SAM will be down due to other activities that day). Lots of activities for SAMGrid - upgrading forwarding and queuing nodes. SAMGrid submissions will not work that day.

CDF:

BARRY J BLUMENFELD to Adam, sam-design, genser

Hi,

This is a summary of my SAM shift. Problems included bad power supplies, several bad tapes and some problematic LTO3 movers.

Tue Oct 23 09:45:36

Bad Mover/Bad Tape IAB070

and a few hours later another Bad mover/bad tape LTO3_01

Resolution from David Berg about the bad movers:

These two movers were in error state after mounts failed during cap

activity in the SL8500 at GCC. The two noaccess volumes were home. I cleared the volumes and restarted the movers.

Tue Oct 23 12:14:44
fcdldata318 was offline

Resolution: A machine having its operating system re-installed was grabbing the ip address of fcdldata318 - that has been corrected.
Needs its dcache re-installed after system wipe. Should happen today.

However, still troubles.
Resolution: dcache did not start, I'm restarting dcache on this node. Alex.

Tue Oct 23 16:53:10
User had an apparent sam problem: I am trying to pull some files over using "sam get dataset --fileList".

Resolution from Krzysztof:
this file is not cached yet:
dc_check (after setup dcap) in place of dccp will say that.

Wed Oct 24 12:27:21
Several mover problems:
LTO3_26.mover alive : OFFLINE cdfenmvr126a 2007-Oct-24 12:24:43
LTO3_27.mover alive : OFFLINE cdfenmvr127a 2007-Oct-24 12:24:43
LTO3_28.mover dead

eventual situation:
LTO3_27 and 28 are back in service,
but LTO3_23,24,25,26 are marked suspect.
This is still the situation at end of shift.

Thu Oct 25 08:26:26
Bad Tape:IA6978 181.47GB (NOTALLOWED 1024-1728 none 0619-1350)
CDF-9940B

Thu Oct 25 15:30:35
LTO3_19.mover problem alive : ERROR - ('1', "mount IAB633 failed: ('ERROR', 13, ['ACSSA> mount IAB633 0,2,1,3', '2007-10-25 15:15:39 16615 Volume VOF309 reactivated.

Thu Oct 25 15:34:13
A user had an apparent dcache problem
Resolution: it was due to incorrect forwarding of the Kerberos ticket.

Fri Oct 26 07:42:40

fcdfdata099.fnal.gov is not pingable.

Resolution:Machine brought back up. Two of the three power supplies have failed.

Sun Oct 28 08:35:53

Bad Mover/Bad Tape

LTO3_19.mover alive : ERROR - ('1', "mount IAB434 failed: ('ERROR', 13, ['ACSSA> mount IAB434

0,2,1,3', 'Mount: Mount failed, Unreadable label.', 'ACSSA> logoff', "], ", 'MOUNT 13: mount

IAB434 0,2,1,3 => 0,Mount: Mount failed, Unreadable label.')")

Mon Oct 29 15:49:27

LTO3_26.mover alive : ERROR - ('MOUNTFAILED', 'max_consecutive_failures (3) reached')

cdfenmvr126a 2007-Oct-29 15:47:48

Cheers,

Barry

Krzysztof: Still no dcap client that works correctly under scanning conditions (see last week for symptoms). All relevant parties are aware of the problem. If no progress on dcap soon, then may request scanning exemption for CAF. Or block the scanning nodes from talking to CAF (must be done within security rules). dcap client improvements normally come from DESY. Art: Wants to know how to test this - since *any* person running a dcap client would be hit by this problem if scanned. This problem seems to be a major dcap client vulnerability.

SteveW: What is the plan for testing retired date? Adam wants to see a project run on a replaced file to make sure the station delivers the replaced file.

Thanks to Anil and co. for refreshing integration from production DB!!

MINOS:

Arthur Kreymer to sam-design, minos_sam_admin

Minos status 2007 Oct 29

General -

Deployed Condor on Minos Cluster interactive systems last week.

A typo was corrected in the config files, jobs are running.

First users are trying it out.

We are using Condor 6.8.6, which has restart bug fixes

which are not present in the Condor 6.8.5 running on the farm.

Using Condor on the interactive cluster (about 50 cores). Runs like a batch system (runs as yourself in your home area). This system is meant to recover idle cycles. Who is supporting this? No one right now. Run II Sys installed it. Fermigrid support told them what to do and gave them config files. Support is being worked out.

Enstore -

DCache -

BlueArc -

Continuing to load datasets, no problems so far.

Continuing to discuss GPFARM mounts, still need AFS readonly/executeable. Can't easily go the route of tar-balling sandboxes since not enough people to do that work.

SAM -

Installed sam_bootstrap v8_1_1 as current in dev/int/prd

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

PLAN -

Migration to 9940B tape can proceed anytime - when ?

DB:

Anil Kumar to sam-design

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

D0 :

- d0ora2 crashed on Thu 10/25 at 8:30AM Power Cycled the box. It was up by 9:30AM

Cause is unknown yet. Nothing interesting in any logs, so no further investigation. Yikes!

- Scheduling the Oracle's critical patch on d0ofint1/d0ofprd1 on Nov 06 Tuesday
- 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 [OPERATIONAL_IDX_01] is [80 percent] full

Action : Added more space. Now it is 72% full.

- Message=Tablespace [TS_CF_03312000] is [80.07 percent] full

Action : Added 1 GB

-Owner	Object	Type	Extents Used	Extents Allocated	Extents Left
SAMPRD	FL_PK	INDEX	196	200	4

Action: Increased the max extents to 400.

CDF :

- Refresh SAM schema @d0ofint1 on 10/24 Wed. - Done
It took 12 hours for refresh.And 10-12 hours to size the datafiles and export of samprd.

Thanks!!

- fcdfora1 had bad disks. Bad Disk was swapped out on Monday. Note: fcdfora1 box doesn't have any vendor support. *This hosts CDF integration and development. Disk was removed.*

Alerts:

MINOS :

Robert: There is in fact a new Oracle client that goes with Oracle v11. But there are no release notes available. Anil: Oracle v11 was released in July, but no plans for upgrade yet. Maybe next year if no pressing need.

enstore/dCache:

David Berg to sam-design, enstore-admin

enstore/dcache report

Scheduled Downtimes:

- o SL8500 at GCC - Thu, Nov 1
 - if all the barcodes have been replaced by this time, 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.

Should be a short downtime, since new firmware is already loaded.

- o d0en - Tue, Nov 6, 7 am to 5 pm
 - 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

Replacement part was shipped from Germany - but was damaged in shipping (!!). Hopefully the re-replacement part will be here by next week.

Operations:

- o d0en
 - recycled 1270 LTO-2 volumes from a list of 1273. There were 1866 "unknown" files that had to be marked deleted. The last 3 volumes have files that are not marked deleted, but do not appear to exist in pnfs.
- o cdfen
 - CDF-LTO3 library manager hung and needed to be restarted, twice.
 - working on bringing new LTO3 drive online for CDF
 - barcodes have been replaced on all CDF LTO3 volumes with the older, "small" barcodes.
- o stken
 - barcode replacement continues on CMS LTO3 volumes with the older, "small" barcodes. About 500 remain.
- o fndca

- a couple of failed disks in raids; the arrays rebuilt,

Thanks,
David