

SAM (REX/Ops) Operations Meeting 2007.09.25

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

Releases:

DB server release 8.3.0 changes to bring the mini-DB up to date with retired date stuff and new cx_oracle. Randolph is testing at CDF with a project.

No new client stuff.

List requests command (IT-3029)- fails in old release but works in a new release - but don't know why it is failing in the old release now. Can't figure out why this is so. IDL mismatch? Error is not as such - problem is that a string is terminated early (before the null terminator). No idea why yet. Server is the same for both. The client has changed - but changes should not effect this. And removing the changes does not fix the problem. Claim from users is that it used to work with the old client, but then suddenly stopped. Problem really is in the software (problem occurs on different machines). Perhaps CVS tags are wrong, but that happens automatically.

Since the new one works, perhaps not put too much effort into making the old one work. But want to know why this fails. Hard to make a unit test out of this because lots of data need to be in the database. Mystery! But at least the new one works!

Robert is working on DH Get Lite - make sure the file is not released until the transfer is complete (solves a problem that Mike Diesburg has doing production). Under test now.

DØ:

Kin's report is in this font. Notes from the meeting are in this font.

Kin Yip <kinyip@bnl.gov> Mon, Sep 24, 2007 at 5:16 PM
To: sam-design@fnal.gov

Quite a lof of unresolved issues/events :

IT# 3019: Willem continued to have problems with "merging jobs" (even though Andrew thought he'd resolved this issue).

IT-3019 is dCache failure in UK. Andrew has a procedure to uncache the file from Willem's station and will get the file from somewhere else. Would prefer that SRM would give some status that file is no longer available (doesn't do that).

IT# 3022 : xmldb (in GridKa) corrupted.... waiting for J. Snow to confirm whether he's followed the instructions

IT# 3030 : xmldb error at GridKa seems to be a follow-up on #3022.

IT-3022/3030: There are instructions for what to do in 3022. Parag cannot connect to Gridka's XMLDB from FNAL (firewall? service down at GridKa?). Need help from GridKa. We no longer have a contact at GridKa.

IT# 3023: sam.getRequestInformation(afterDate=reqdategt) has taken a long time. An "index" on the effective date seems to be the solution.

Waiting for Anil to implement this ??

IT-3023 is being worked out. The index did not solve the problem.

IT# 3025: fnal-cabsrv1 crashed ~Sept. 18 (?) and red bullet on SAAG page. Andrew said he'd look at it but he's given a conclusion yet (?).

fnal-cabsrv1 crash: Andrew is looking into a bug. Station crashed again because log file hit 2 Gig. Perhaps put in code to open a new log file when it gets too big. There are lots of projects (~400).

IT# 3028: Willem complained "query.py" failed to get the next MC request.

IT-3028 - Problem for Joel? We don't do queue.py.

IT #3029: "corba failure in sam/v8_1_5 but not in sam/v8_2_*" with the command sam.getRequestInformation(afterDate)

Robert and Steve Sherwood haven't figured this out.

See Releases section above.

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)

We think same problem as IT-3019 (see above)

CDF:

ByeongRok's report is in this font. *Notes from the meeting are in this font.*

ByeongRok Ko <brko@phy.duke.edu> Mon, Sep 24, 2007 at 5:13 PM
To: sam-design@fnal.gov

Overall calm shift.

Tue: Sep 18

CDF-LTO3 library manager is back.

email from George Szmuksta

The SL8500 is operational again. The CD-LTO3 and the CDF-LTO3 library

managers have been released.

It appears that the problem is a known firmware issue. The firmware cannot be updated until all of the tapes in the robot have been relabeled.

Tape labels are now being outsourced. New labels are much better. Have to replace all labels in the robot.

Wed: Sep 19

fcdpdata320 is offline

==>Etta Burns fixed it

Fri: Sep 21

email from George Szmuksta

The network interface on cdfensrv1 (PNFS) had to be reset.

--> enstore ball is red

email from George Szmuksta

This was due to the problem on cdfensrv1. The ball will disappear shortly.

--> PnfsManager is offline

--> All dcache doors are not responding.

email from Alex Kulyavtsev

There were problems with dcache pnfs manager domain.

We restarted pnfs manager and all dcache doors

After resetting the network interface on cdfensrv1 (PNFS), everything is back

When pnfs manager goes offline, all jobs get unhappy. Hope that new hardware is more reliable.

Kudos to SAM shifter for reporting pnfs outage immediately.

Mon: Sep 24

fcdldata150 is offline,

email From : Alex Kulyavtsev

I took node 150 for reconfiguration.

I will take out node 138 as well.

Krzysztof: fcdldata150 was taken offline because system was not staging files. Alex worked to try to get it unstuck. What worked was to move all of the stage pools to different nodes. New requests could not be put in because of memory growth on the read-only pools. Fortunately, fcdldata150 was being drained for a reconfiguration, so it was easy to take it down.

Angela: new version of fcp server that does multiple queues. Running on fcdfprj1. Needs testing.

MINOS:

Art's report is in this font. Notes from the meeting are in this font.

Arthur Kreymer <kreymer@fnal.gov> Mon, Sep 24, 2007 at 10:00 PM

To: sam-design@fnal.gov

Cc: minos_sam_admin@fnal.gov

Minos status 2007 Sep 25

General -

Preparing for Minos Collaboration meeting Sep 27-30

Enstore -

Would like to move our 'vault' second raw data copy to LTO-3 whenever convenient. About 4.5 TB .
Discussed this with Sasha M.

DCache -

20 Sep DCache upgrades went smoothly for us.
No file handle problems so far, too early to tell ?

Many monitoring web pages are still stale (Jul 31)
(login lists, door plots, pool contents, etc)

Slow ROOT DCache reads in root >= 5.12
We still need to rerun tests on larger test files.

BlueArc - Preparing for deployment of 30+ TB SataBeast/BlueArc
as /minos/scratch and /minos/data
Encouraging test of 1/2/3 disk failures before deployment.

SAM -

One of the Minos users surprised us by wanting to run
parallel jobs in the classic SAM fashion,
not caring about the order of file delivery.
Liz quickly updated the SAM framework to allow this.

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

Oracle -

Minosora1 firmware was upgraded. 20 Sep
during the DCache maintenance.

OPEN ISSUES

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 observers and clients to v8

Offsite user could not get to the database browser because their site had blocked the DBB port. DBB does not run on port 80 because of security issues at FNAL.

DB:

Anil's report is in this font. *Notes from the meeting are in this font.*

Anil Kumar <akumar@fnal.gov> Tue, Sep 25, 2007 at 8:02 AM

To: sam-design@fnal.gov

D0 :

- upgrade of d0ofdev1 to 10.2.0.3 for Sep 18 Tuesday - done

- Scheduling upgrade of d0ofint1/d0ofprd1 databases to 10.2.0.3 Oct 2, 2007

Downtime at D0 on 10/2 to upgrade DB to Oracle 10.2.0.3 (since GCC will be down on 10/2).

- SAM cut v6_10 on d0ofdev1 - To add composite index on lower (file_name) and retired_date -done

- SAM cut of v6_11 on d0ofdev1/d0ofint1/d0ofprd1 - Add index on effective_date on work_request_status_histories table.

- 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 :

- There are memory error alerts on d0ora1. As per Steve K.
"These are corrected memory problem. I talked this over with D1. We will replace the DIMM if corrected again."

CDF :

- Upgrade to 10.2.0.3 done for cdfofdev and cdfofint on Thursday 09/13/2007
CPU was replaced on Tuesday 09/18/2007
- SAM cut of v6_11 on d0ofdev1/d0ofint1/d0ofprd1 - Add index on effective_date on work_request_status_histories table.
- Scheduleing upgrade of cdfofpr2 databases to 10.2.0.3 Oct 3, 2007.
Downtime for CDF to upgrade DB to Oracle 10.2.0.3 on 10/3 for lab wide power outage. Needs confirmation from Rick.

MINOS :

- 09/20/07 Thursday, planning to upgrade bios and firmware on minosora1 -
done
BIOS and firmware had already been upgraded on minosora3
- SAM cut of v6_11 on d0ofdev1/d0ofint1/d0ofprd1 - Add index on effective_date on work_request_status_histories table.

Enstore/DCache:

Dave's report is in this font. *Notes from the meeting are in this font.*

David Berg <berg@fnal.gov>
To: sam-design@fnal.gov
Cc: enstore-admin@fnal.gov
enstore/dcache report

Mon, Sep 24, 2007 at 9:02 PM

Scheduled Downtimes:

- o d0en - Tue, Oct 2
 - enstore s/w updates

- o cdfen - Wed, Oct 3
 - dcache s/w updates, incl. fix to "too many open files"
 - upgrade cdfensrv1 (pnfs server), PC and raid
 - enstore s/w updates

Operations:

- o d0en
 - nothing to report

- o cdfen
 - cdfensrv1 (pnfs server) dropped network
 - handbot failure in SL8500 at GCC

Krzysztof:: Is the arm problem due to firmware (is the arm making a painful movement?)? Dave: For example, handbot can become misaligned. New firmware will try to detect and correct this problem.

- o cdfdca
 - Four files reported with CRC (adler32) errors.
 - Two pools stopped restoring files from tape after enstore libraries were paused; restores to other pools got to "Waiting state" as well. Fix will be applied during upgrade.
 - pnfs manager and all doors were restarted after pnfs server came back on the network.
 - Several pools were restarted.

- o stken

- downtime went smoothly, all goals accomplished:
 - enstore s/w updated
 - 9940B drive firmware updated
 - Nexsan raid firmware updated
 - stkserv1 (pnfs server) raid disk replaced, array rebuilt

o fndca

- dcache code successfully upgraded. It includes fixes to deal with issues like pool crashing due to "too many open files"; memory and resource retention issues; improvements in scheduling of restores from tape (more relevant to cdf).
- few pools were restarted to fix "too many open files" problem; we have not seen this again since the upgrade.

4000 DLTs for SDSS need to be removed because tower is going out of warranty. SDSS is copying files and this process may affect DØ's usage of the ADIC.

3rd SL8500 will be moved into FCC2 south of the ADIC. May be delivered today or tomorrow or soon.