

# NDGF All Hands 2020-2

ENDIT daemons status update

# ENDIT overview

- Two parts
  - ENDIT dcache plugin
    - A HSM provider plugin for dCache, ie Java
    - krishnaveni expected to take the role as maintainer
    - <https://github.com/neicnordic/dcache-endit-provider/>
  - ENDIT daemons
    - Perl scripts that do the actual transfer of data to/from IBM Spectrum Protect (TSM)
    - Mostly hacked on by nikke nowadays
    - <https://github.com/neicnordic/endit>
- Documentation
  - Repository README:s aims to be generic documentation
  - [https://wiki.neic.no/wiki/DCache\\_TSM\\_interface](https://wiki.neic.no/wiki/DCache_TSM_interface) has more details and NDGF-specific bits

# ENDIT daemons overview

- Main daemons
  - tsmarchiver.pl
  - tsmretriever.pl
  - tsmdeleter.pl
- Hint file generators (choose 1)
  - tsm\_getvolumecontent.pl
    - “Server side” - via TSM admin interface (dsmadm) requiring separate credentials
  - tsmtapehints.pl
    - “Client side” - via dsmc (same as main daemons)

# ENDIT daemons - ordering of files on tape

- We have never been quite satisfied with the get/retrieve performance, lots of time is spent doing LOCATE to wind tape to correct position.
  - Surprisingly much, as the group of files written tends to be read with similar grouping
- While investigating the actual order of files on tape it was discovered that tsmarchiver archived them in readdir() order, which might or might not be the same order as the files were written to the pool.
- Changing this to sort according to mtime and testing, we discovered that the actual order of files written to tape matched an alphabetical sort...
  - Regression in TSM client, documented that files should be processed in file order
  - First IBM response was that they would just update the docco to current behavior
  - After a rather stern reply asking why it's suddenly OK to change the docco/design instead of fixing the actual issue it apparently was assigned an APAR and thus scheduled to be fixed

# IBM Authorized Program Analysis Report (APAR)

- The means by which a problem with an IBM program is documented, tracked, and corrected.
- Essentially identifies and documents a specific bug, usually with the intention for development to fix it.
- You'll likely never hear an IBM:er refer to an issue as a bug, but rather an APAR.
- Just remember that APAR is IBM speak for program bug, and you'll be OK :-)

# IBM TSM client APAR IT33143

- OPERATIONS THAT USE THE -FILELIST OPTION DO NOT PROCESS FILES IN THE ORDER THEY APPEAR IN THE FILE LIST
- Problem Conclusion:
  - “Upon the fix, the backup or archive operations which use -filelist option will send objects in the order specified by the -filelist. That will be in accordance with the current -filelist option documentation.”
- <https://www.ibm.com/support/pages/apar/IT33143>
- Has much higher impact than ENDIT bug, due to always sorting the file list
- Haven't investigated when they broke it
- Scheduled to be fixed in version 8.1.11, ETA later this year

# ENDIT daemons - ordering of files on tape (2)

- Likely affects HPC2N most, due to faster tape drives and thus having set the migration start threshold higher to get decent chunking there is a greater probability of mixing unrelated groups of files.
- Fix will not affect written data.
- Requires update of TSM client to get rid of sorting done there.
- Will tag ENDIT daemons 1.2.0 from master once TSM client with fix is released.

# IBM TSM client APAR IT33682

- DSMC COMMAND MAY EXPERIENCE A CONNECTION DELAY WHILE LOGGED WITH NON-ROOT USER
- Problem Conclusion:
  - “Backup-archive and administrative client will not experience a connection delay anymore.”
- <https://www.ibm.com/support/pages/apar/IT33682>
- Some cert-handling “feature” busy-loops trying to open the root-owned cert store in write mode after having read the parts it needs...
  - Retries for approx 2m45s, so a lot of retries...
- Affects ENDIT daemon operations if set up in the recommended way (ie. not running as root) following the IBM documentation on non-root usage
- Scheduled to be fixed in version 8.1.11, ETA later this year



The End