

Data management - storage maintenance

Volodymyr Yurchenko, CERN

[<volodymyr.yurchenko@cern.ch>](mailto:volodymyr.yurchenko@cern.ch)

Grid Operations

Since April 2018 we have been running Grid maintenance operations by a command-line recovery tool that is able to:

- Recover lost files on a storage
- Clean up orphaned entries in the central file catalogue
- Find dark data and sync SE
- Replicate/move SE content
- Storage decommission
- Restarting data transfers

This tool handles the most common recovery tasks in a uniform way and provides detailed statistics on each aspect of operation.

Input Data Format

A text list of GUIDs/PFNs, e.g.:

96ADE76C-12DE-11DE-AB90-001E0B4E0440

CF76E0B8-180A-11DE-AF50-001E0B4E0440

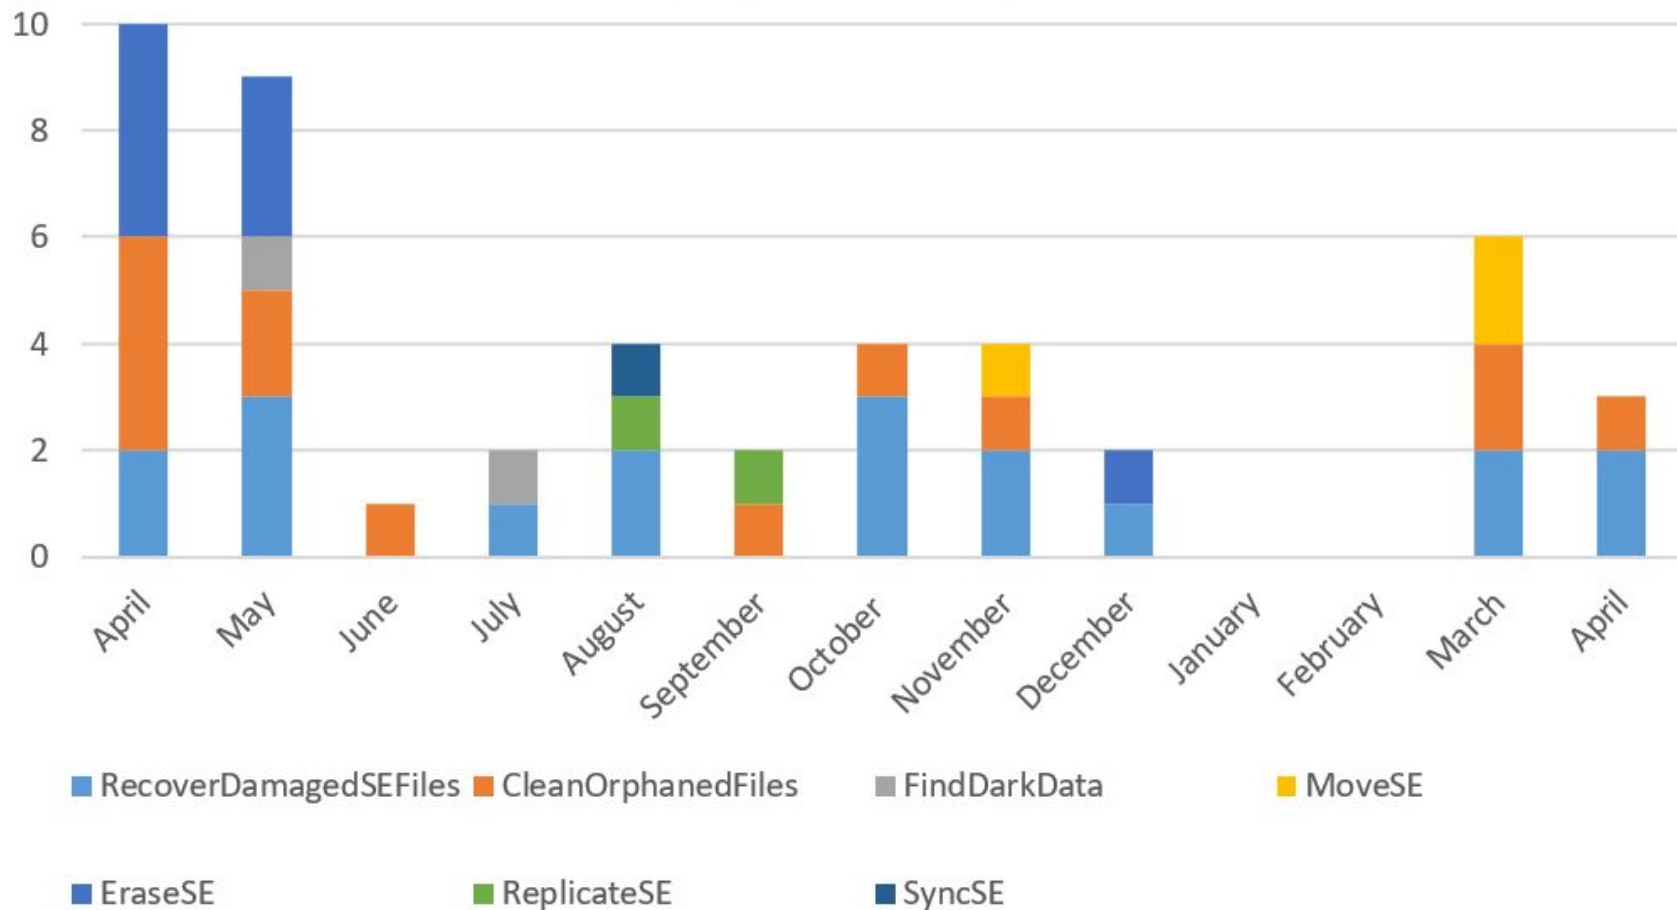
E98585F2-155F-11DE-AF71-001E0B4E0440

root://epgsr5.ph.bham.ac.uk:1094//10/17439/96ade76c-12de-11de-ab90-001e0b4e0440

root://epgsr5.ph.bham.ac.uk:1094//11/09198/cf76e0b8-180a-11de-af50-001e0b4e0440


root://epgsr5.ph.bham.ac.uk:1094//14/05089/e98585f2-155f-11de-af71-001e0b4e0440

Recovery Operations By Months




In average we receive 3.8 requests/month from storage admins

Data movement



We have moved 2157 TB during last 12 months

Catalogue Clean Up



More than 28.8 million GUID entries (and similar number of LFN entries) were removed from the central DB

Time Estimates

100 000 files are scheduled for transfer in **~30 min**

100 000 orphaned entries cleaned up in **~14-15 hours**

100 000 files (4.7 TB) transferred in **12 days** – *heavily depends on the target storage network throughput, 3rd party copy enabled, internal configuration, etc*

Report Issues

Directly to Costin <grigoras@cern.ch> and
Volodymyr <volodymyr.yurchenko@cern.ch>



Mailing list alice-lcg-task-force@cern.ch



Jira <https://alice.its.cern.ch/jira/projects/GO>



TrackRefs Removal

Have to repack old MC ESD archives to remove the above mentioned file from the ZIP archive (so read / unpack / rm / pack / upload):

Total number of found files: **68 561 375**

Total number of processed: 10 062 181 (**14,68%**)


Total size of found files: **3660 TB**

So far reclaimed: **475 TB** (12,98%)

Failed to reclaim: 175.4 TB (4,79%)


Part of TrackRefs in Productions

found volume



■ 2007 ■ 2008 ■ 2009 ■ 2010 ■ 2011 ■ 2012
■ 2013 ■ 2014 ■ 2015 ■ 2016 ■ 2017 ■ 2018


reclaimed volume



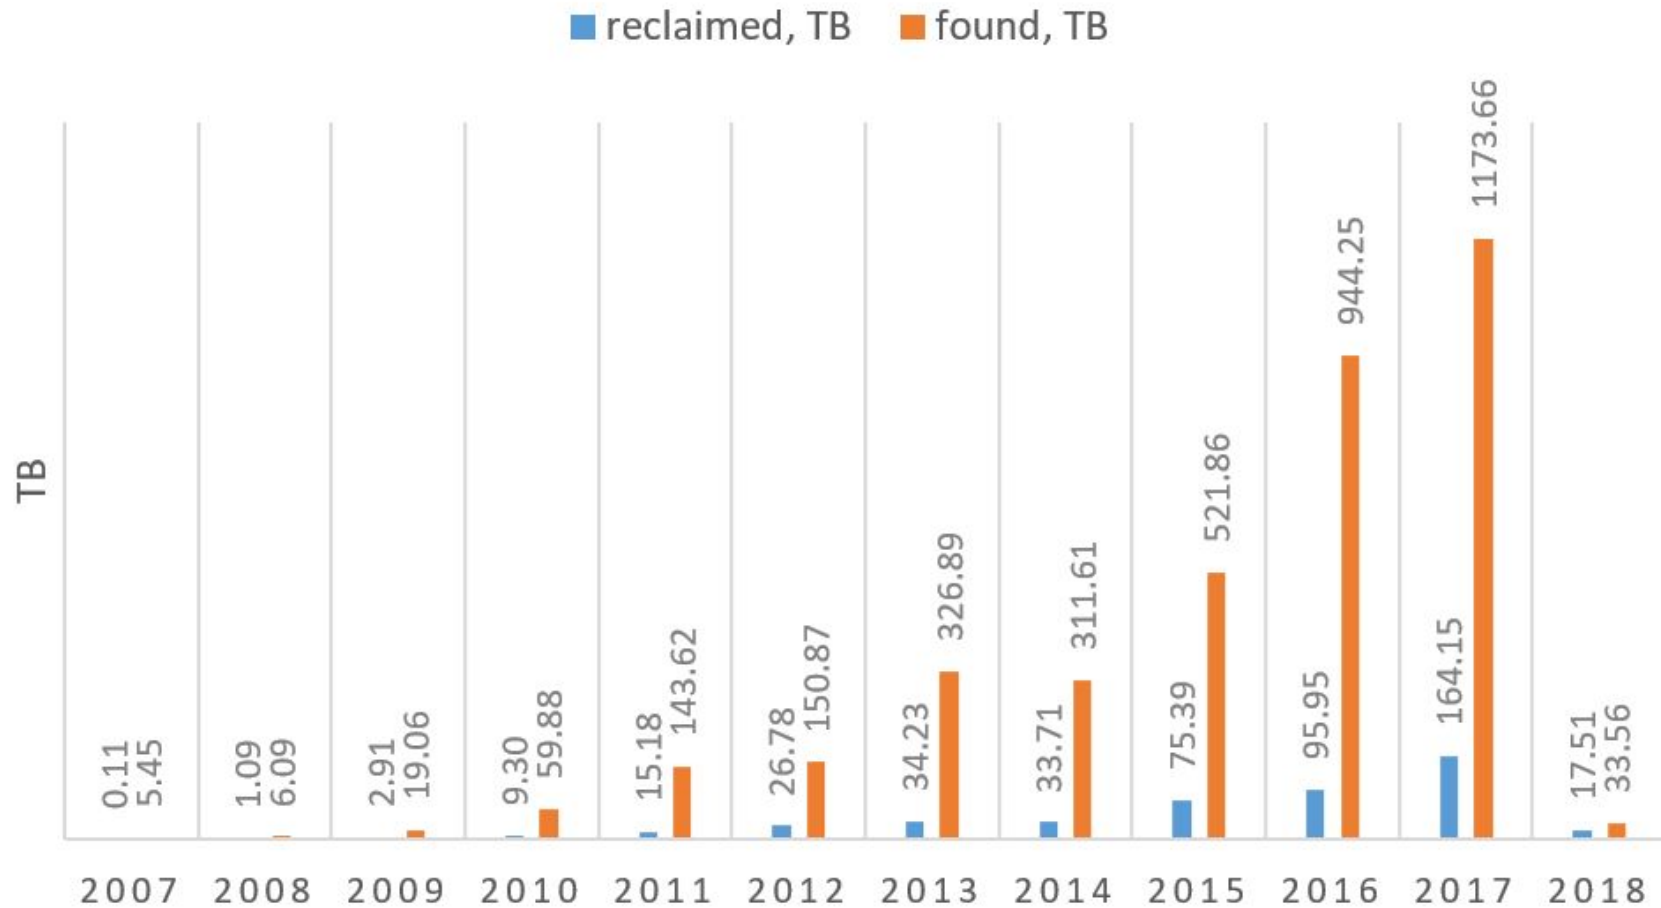
■ 2007 ■ 2008 ■ 2009 ■ 2010 ■ 2011 ■ 2012
■ 2013 ■ 2014 ■ 2015 ■ 2016 ■ 2017 ■ 2018

PROCESSED FILES

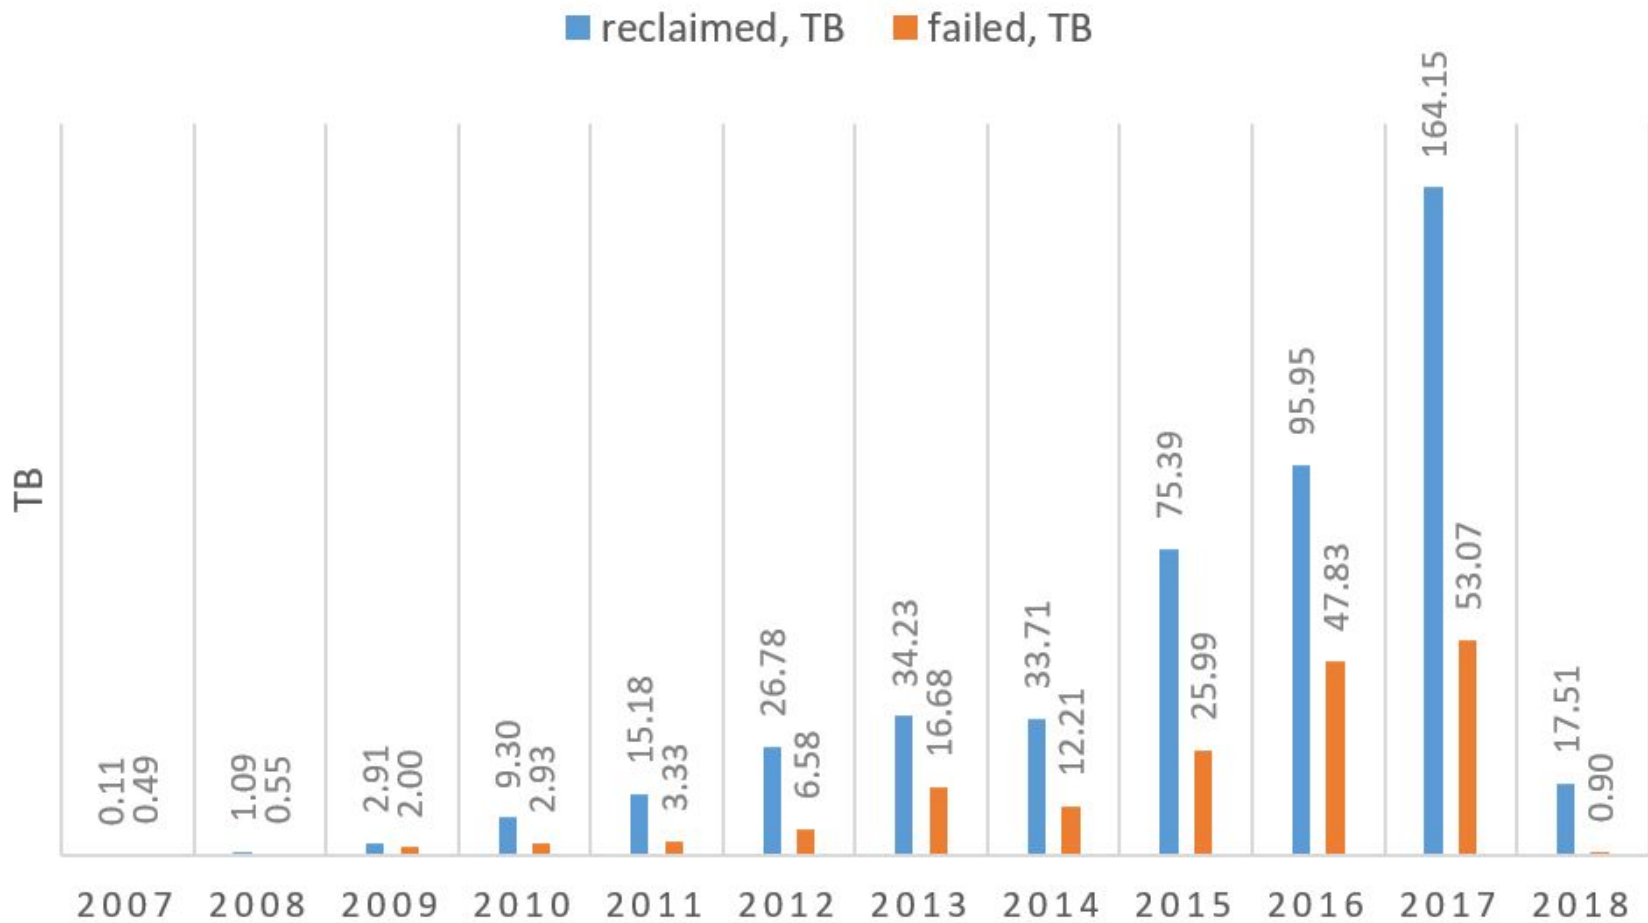
■ processed files ■ found files



RECLAIMED SPACE



RECLAIMED VS FAILED



TrackRefs: Summary

- One masterjob (~2400-2500 subjobs, **45000** files, 1 resubmission round) takes **24:30** to be executed
- Provided we run 10 masterjobs concurrently 24/7, we will be able to finish TrackRefs removal in **133** days
- In some cases the jobs don't receive authentication tokens, we are polishing the procedure

XRootD Reminder

- XRootD version - at least 4.1, latest is the best
- [alicexrdplugins](#): system auto updatable installation via RPM packages
- If you use [xrd-installer](#), please *do* upgrade

| Version ▾ | EOS Version | add | get | rm | 3rd |
|---------------|-------------|-----|-----|----|---------|
| Xrootd v3.2.6 | | | | | Test... |
| Xrootd v3.2.6 | | | | | Test... |
| Xrootd v3.3.4 | | | | | Test... |
| Xrootd 3.3.6 | 0.3.268 | | | | Test... |
| Xrootd 3.3.6 | 0.3.35 | | | | Test... |

3rd Party Copy (EOS)

Configuration parameters to be added to all FSTs:

- SLC6:
 /etc/sysconfig/eos
 export EOS_FST_NO_SSS_ENFORCEMENT=1
 service eos restart
- CC7:
 /etc/sysconfig/eos_env
 EOS_FST_NO_SSS_ENFORCEMENT=1
 service eos@fst restart

<http://alimonitor.cern.ch/stats?page=SE/table>

| Version | EOS Version | add | get | rm | 3rd ▲ |
|---------------|-------------|---------|---------|---------|---------|
| Xrootd v4.8.5 | | | | | Test... |
| Xrootd 3.3.6 | 0.3.35 | | | | Test... |
| Xrootd v4.9.1 | | Test... | Test... | Last... | Test... |
| Xrootd v3.3.4 | | | | | Test... |
| Xrootd v4.8.1 | | | | | Test... |
| | | Test... | Test... | Last... | Test... |
| Xrootd v4.5.0 | | | | | Test... |
| Xrootd v4.8.5 | | | | | Test... |
| Xrootd v4.8.5 | | Test... | | Last... | Test... |
| Xrootd v3.2.6 | | | | | Test... |
| Xrootd v4.9.0 | 4.4.32 | Test... | Test... | Last... | Test... |
| Xrootd v4.8.4 | 4.4.23 | | | Test... | Test... |
| Xrootd v4.9.1 | | | | | Test... |
| Xrootd v4.8.4 | | Test... | Test... | Last... | Test... |
| Xrootd 3.3.6 | 0.3.268 | | | | Test... |
| Xrootd v4.4.1 | | | | | Test... |
| Xrootd v4.9.1 | | | | | Test... |
| Xrootd v3.2.6 | | | | | Test... |
| | | | | | Test... |
| | | | | | Test... |
| Xrootd v4.0.4 | | | | | Test... |
| Xrootd v4.8.1 | | | | | Test... |
| | | | | | Test... |
| Xrootd v4.6.1 | | | | | Test... |
| Xrootd v4.8.4 | | | | | Test... |
| Xrootd v4.0.0 | | | | | Test... |
| | | | | | Test... |

3rd Party Copy (dCache)

- 3rd party copy works for X.509-based authentication
- ALICE uses token based authentication, passing both source and target access envelopes for the respective operations
- To enable TPC for ALICE dCache XRootD client needs a separate plugin
- Ticket was opened in dCache tracking system on 19 Sep 2018, but ALICE TPC patch wasn't the first target to implement
- If you run dCache, please remind the developers of this ticket

Thanks! Questions?