

The GridKa tape storage: latest improvements and current production setup

Haykuhi Musheghyan, GridKa Team



GridKa Tape


Tape Library

- 1 Oracle SL8500 Library
- 35 T10k-D Drives
- ~10,000 cartridges

dCache file servers

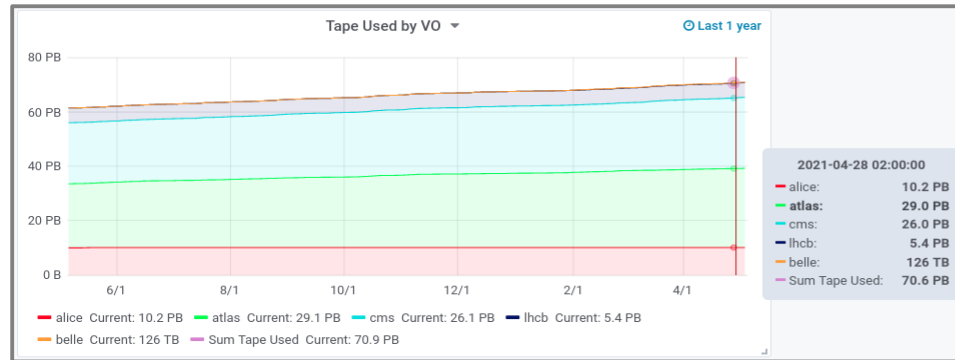
- Separate instances for different experiments
- dCache pools are directories in a shared GPFS
- Primary and fallback pools for stage/write pools

KIT operates more tape libraries outside of Tier-1 context.



GridKa Tape Production Setup Today

- dCache (ATLAS/BelleII/CMS/LHCb) & xrootd
- GPFS
- TSM
- TSS-client: queuing and scheduling software developed at GridKa based on TSM API
- Endit-TSS: intermediate software between dCache Endit-Provider plugin and TSS-client
 - In production since Jan 2020 (glue scripts (before Jan 2020))



Back to 2018 & Summary of Various Tests

After the 1st Data Carousel Test, various perf. tests were done on the GridKa tape to find possible bottlenecks.

- The number of concurrent requests can be increased from 2,000 to 30,000.
 - New limit (30,000) due to limitations on the TSS side
 - ~50% improvement in overall tape recall rate
- TSM allows duplicated files.
 - Removing them significantly reduces the number of tape mounts
- Writing and then reading large files (>10GB) directly affects the tape recall/migrate rate.
 - Increases the overall tape recall rate ~3,0 times
 - Writing large files are expected from VOs

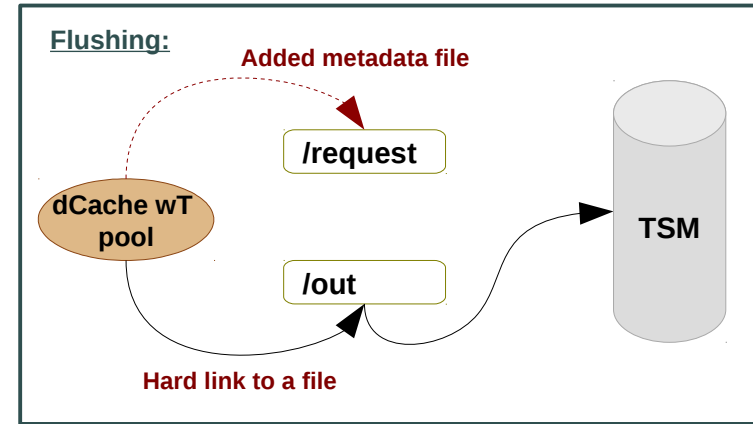
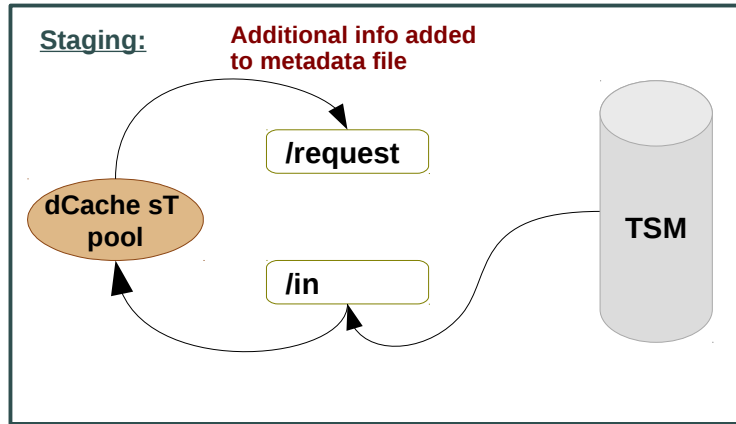
With these tests, we were able to improve the overall tape performance from a minimum of 50% to a maximum of 200%.

More details in the [CHEP article](#).

Bringing Test Results into Production...

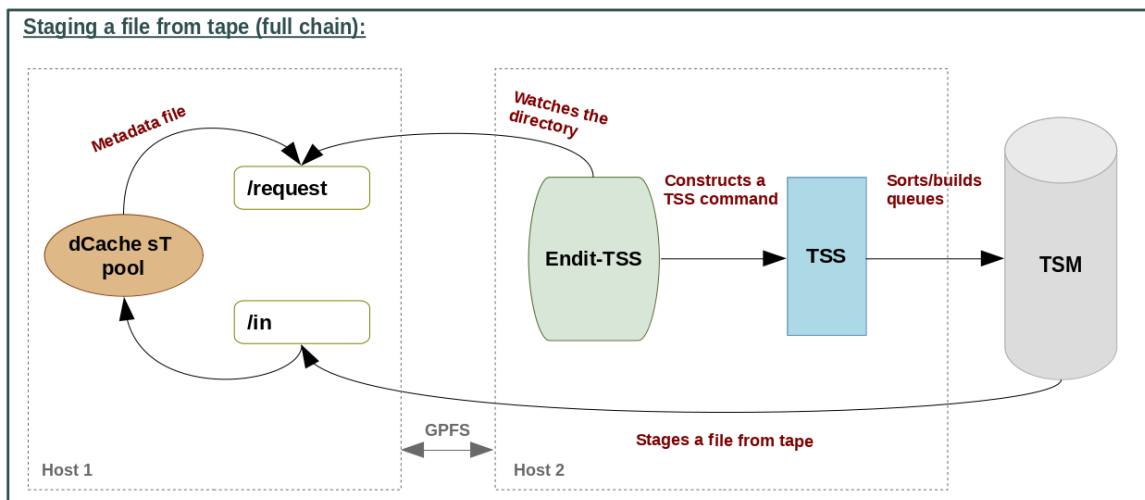
Using the dCache Endit-Provider Plugin

- An efficient dCache interface to TSM ([DCache_TSM_interface](#))
 - Connects the dCache pools with TSM
- Used by NDGF-T1
 - GridKa modified a little bit the original source code of the Endit-Provider plugin
 - Limit of 2,000 for maximum concurrent requests is gone!



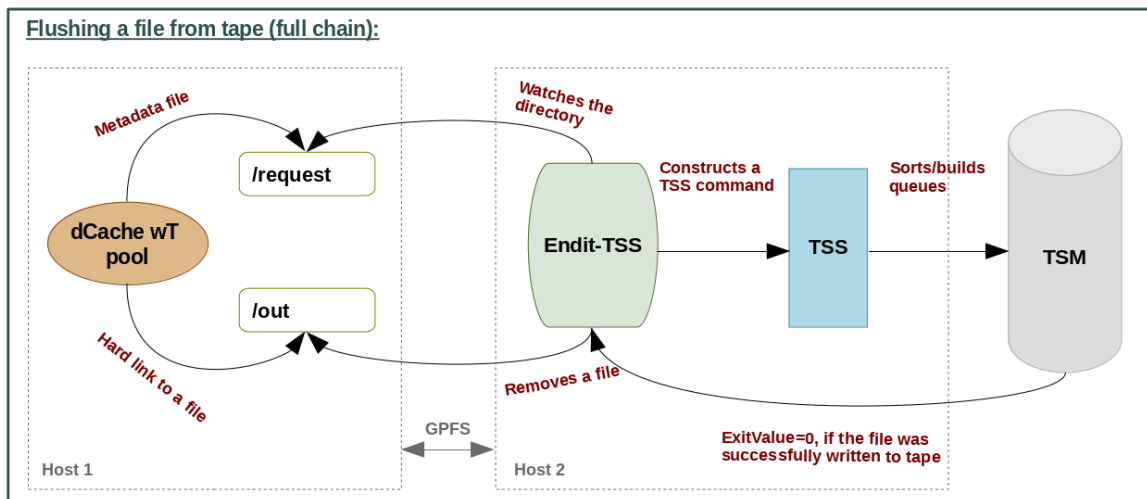
Staging from Tape: Full Chain (GridKa-usecase)

- Moving out from the limit of 2,000 (MaxRequests)
- New limit (30,000) due to limitations on the TSS-client side
- 2 different hosts for the dCache pool & Endit-Provider; Endit-TSS & TSS
- Staging: in production since January 2020



Flushing to Tape: Full Chain (GridKa-usecase)

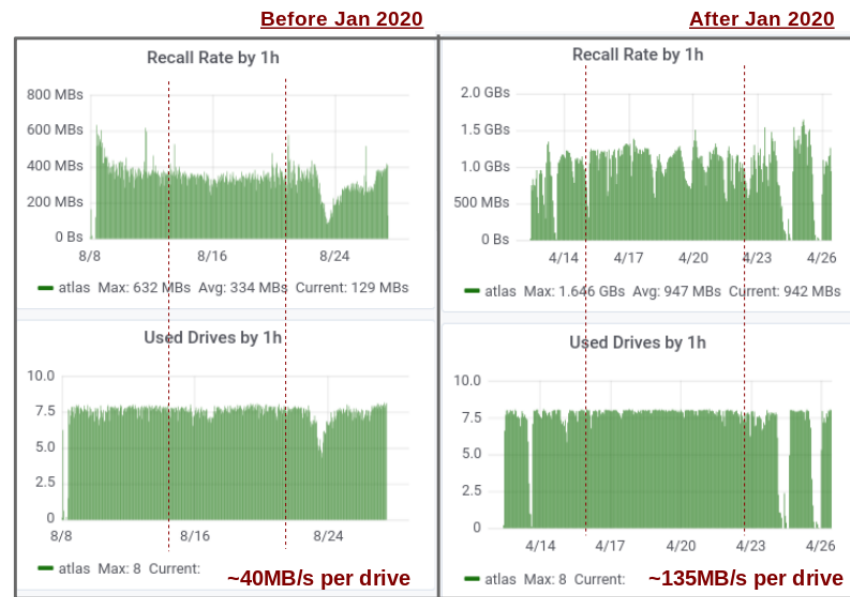
- Moving out from the limit of 2,000 (MaxRequests)
- New limit (30,000) due to limitations on the TSS-client side
- Flushing: in production since January 2021 for BelleII only



Latest Production Results & Conclusions

Latest improvements include:

- detection and elimination of bottlenecks,
- use of the dCache Endit-Provider plugin,
- new software Endit-TSS for efficient recalling from tape,
- new hardware and faster network.



The new setup works pretty well and the tape recall rate is improved by more than factor of 3 per tape drive.

Thank you