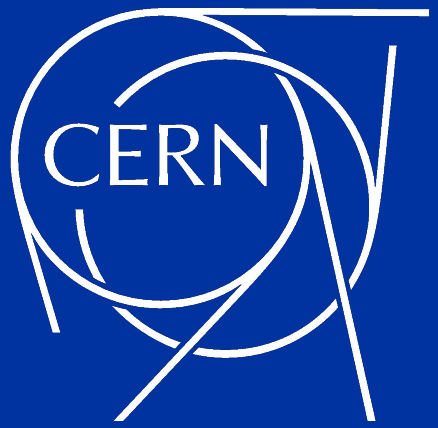

CTA Production experience

Julien Leduc

17/3/21 - HEPIX

EOS+CTA Architecture

- **EOS+CTA is a pure tape system.**
- **Disk cache duty consolidated in main EOS instance.**
- **Operating tape drive at full speed full time efficiently requires a SSD based buffer: EOSCTA**

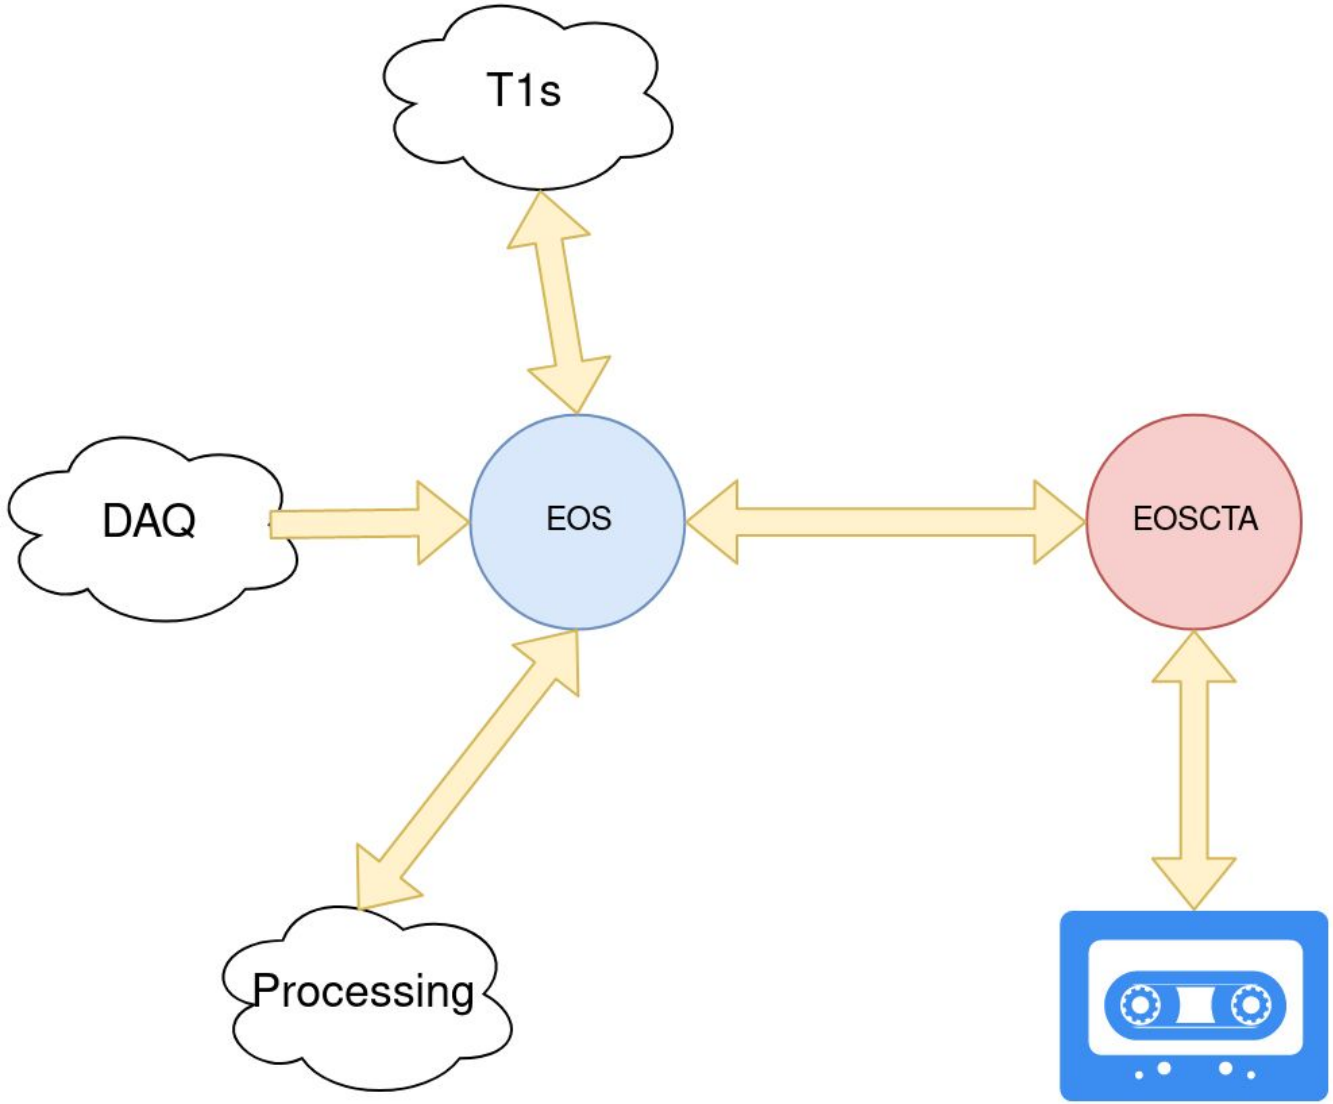

EOS+CTA Architecture

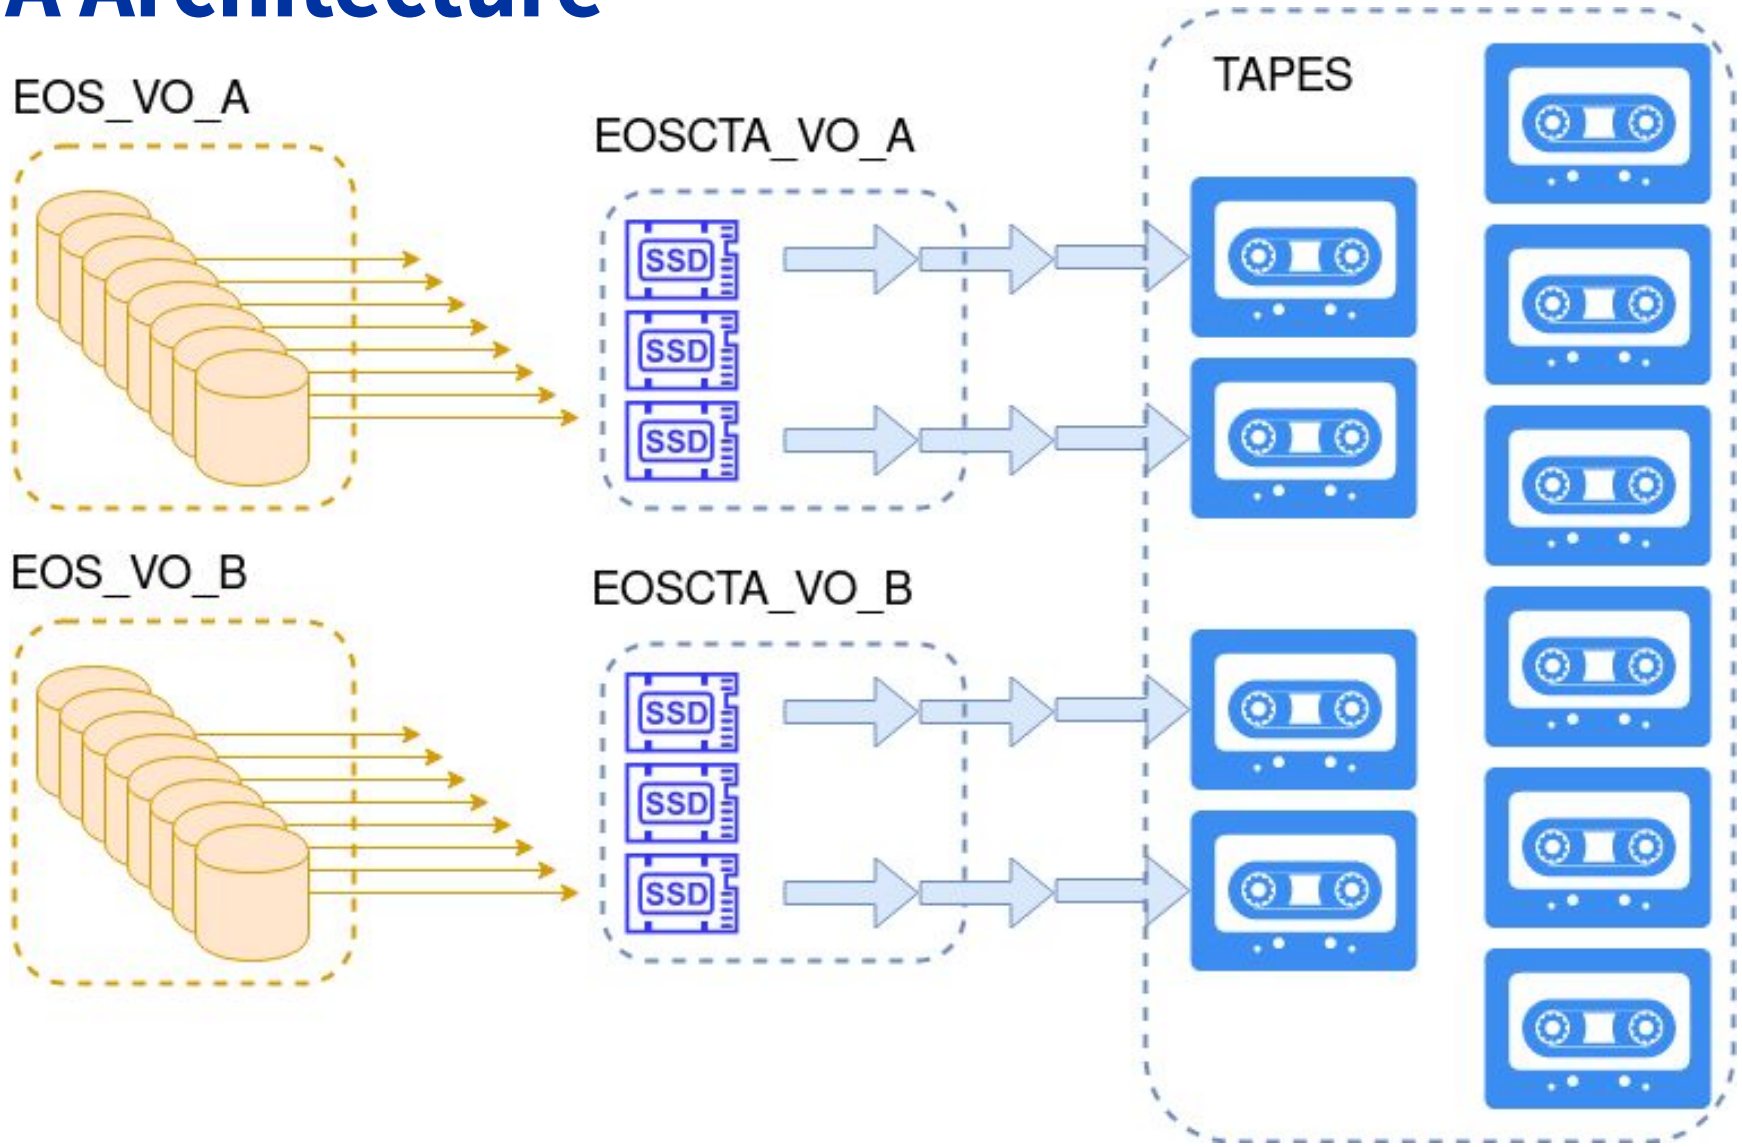

EOS+CTA Architecture

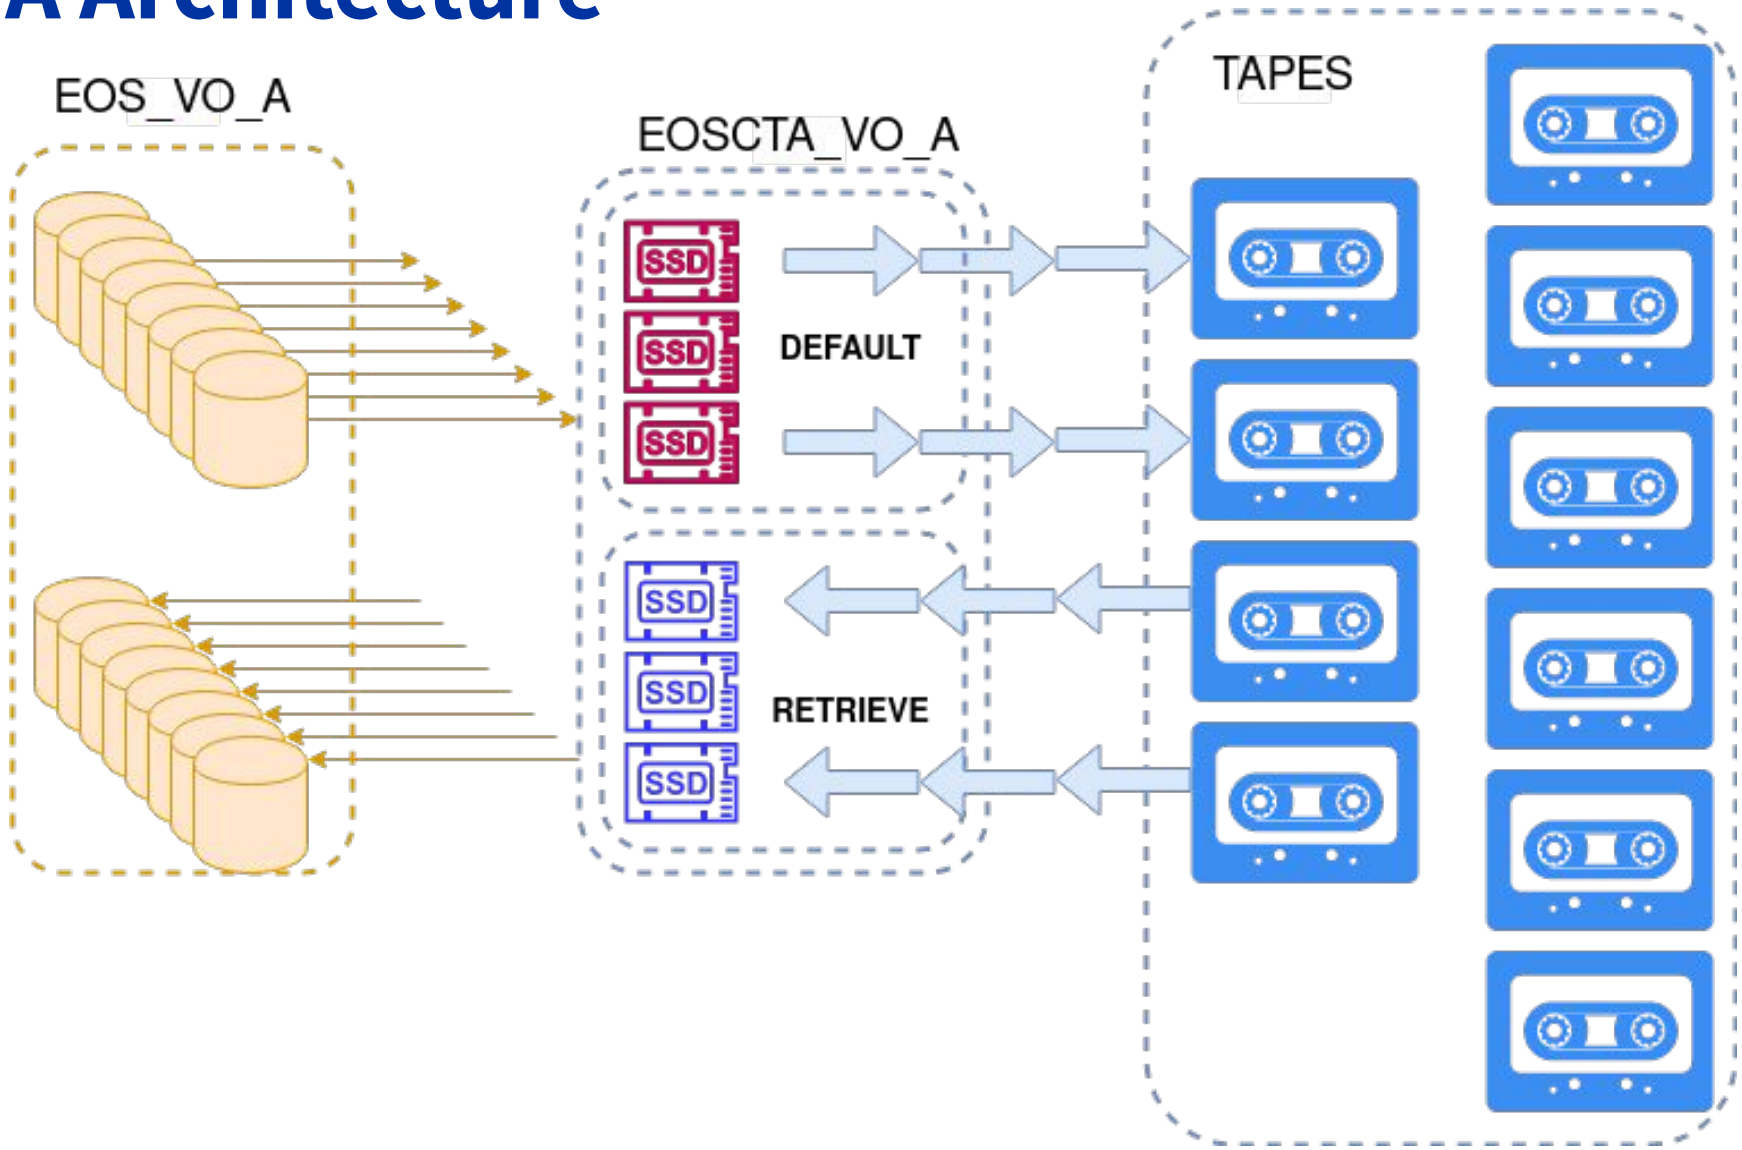

EOS+CTA Space Properties

- **No file redundancy in tape buffer**
- All files in *default* space are on their way to tape:
  - `d1:t0` and disk *default* replica deleted when successfully written to tape
- All files in *retrieve* space are on their way to EOS disk coming from tape:
  - `d1:t1` and disk *retrieve* replica deleted when successfully transferred out
  - disk replica deleted after 24 hours by the diskserver Garbage Collector

SSD spaces are (mostly) empty when everything is fine

Fig 1: Default space during write to EOSCTA public instance

EOS+CTA Architecture *spinners* Addon (ALICE)

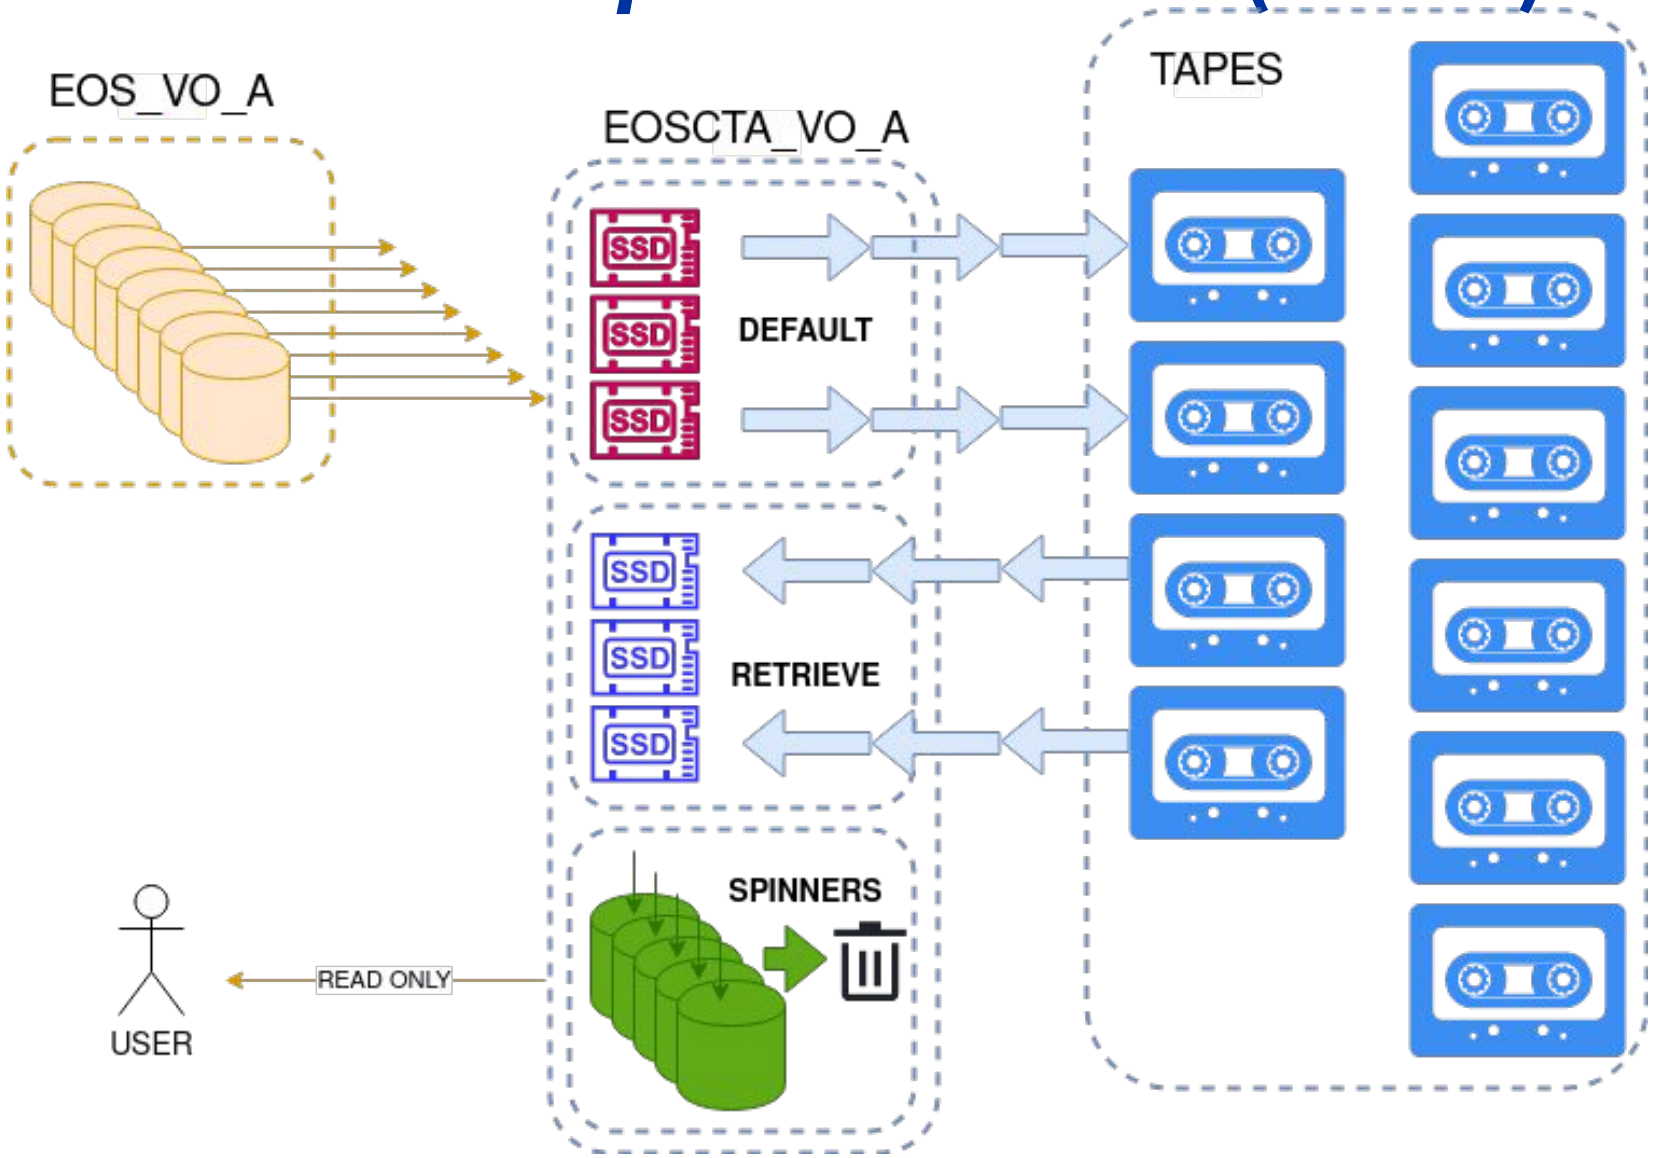

EOS+CTA Timeline

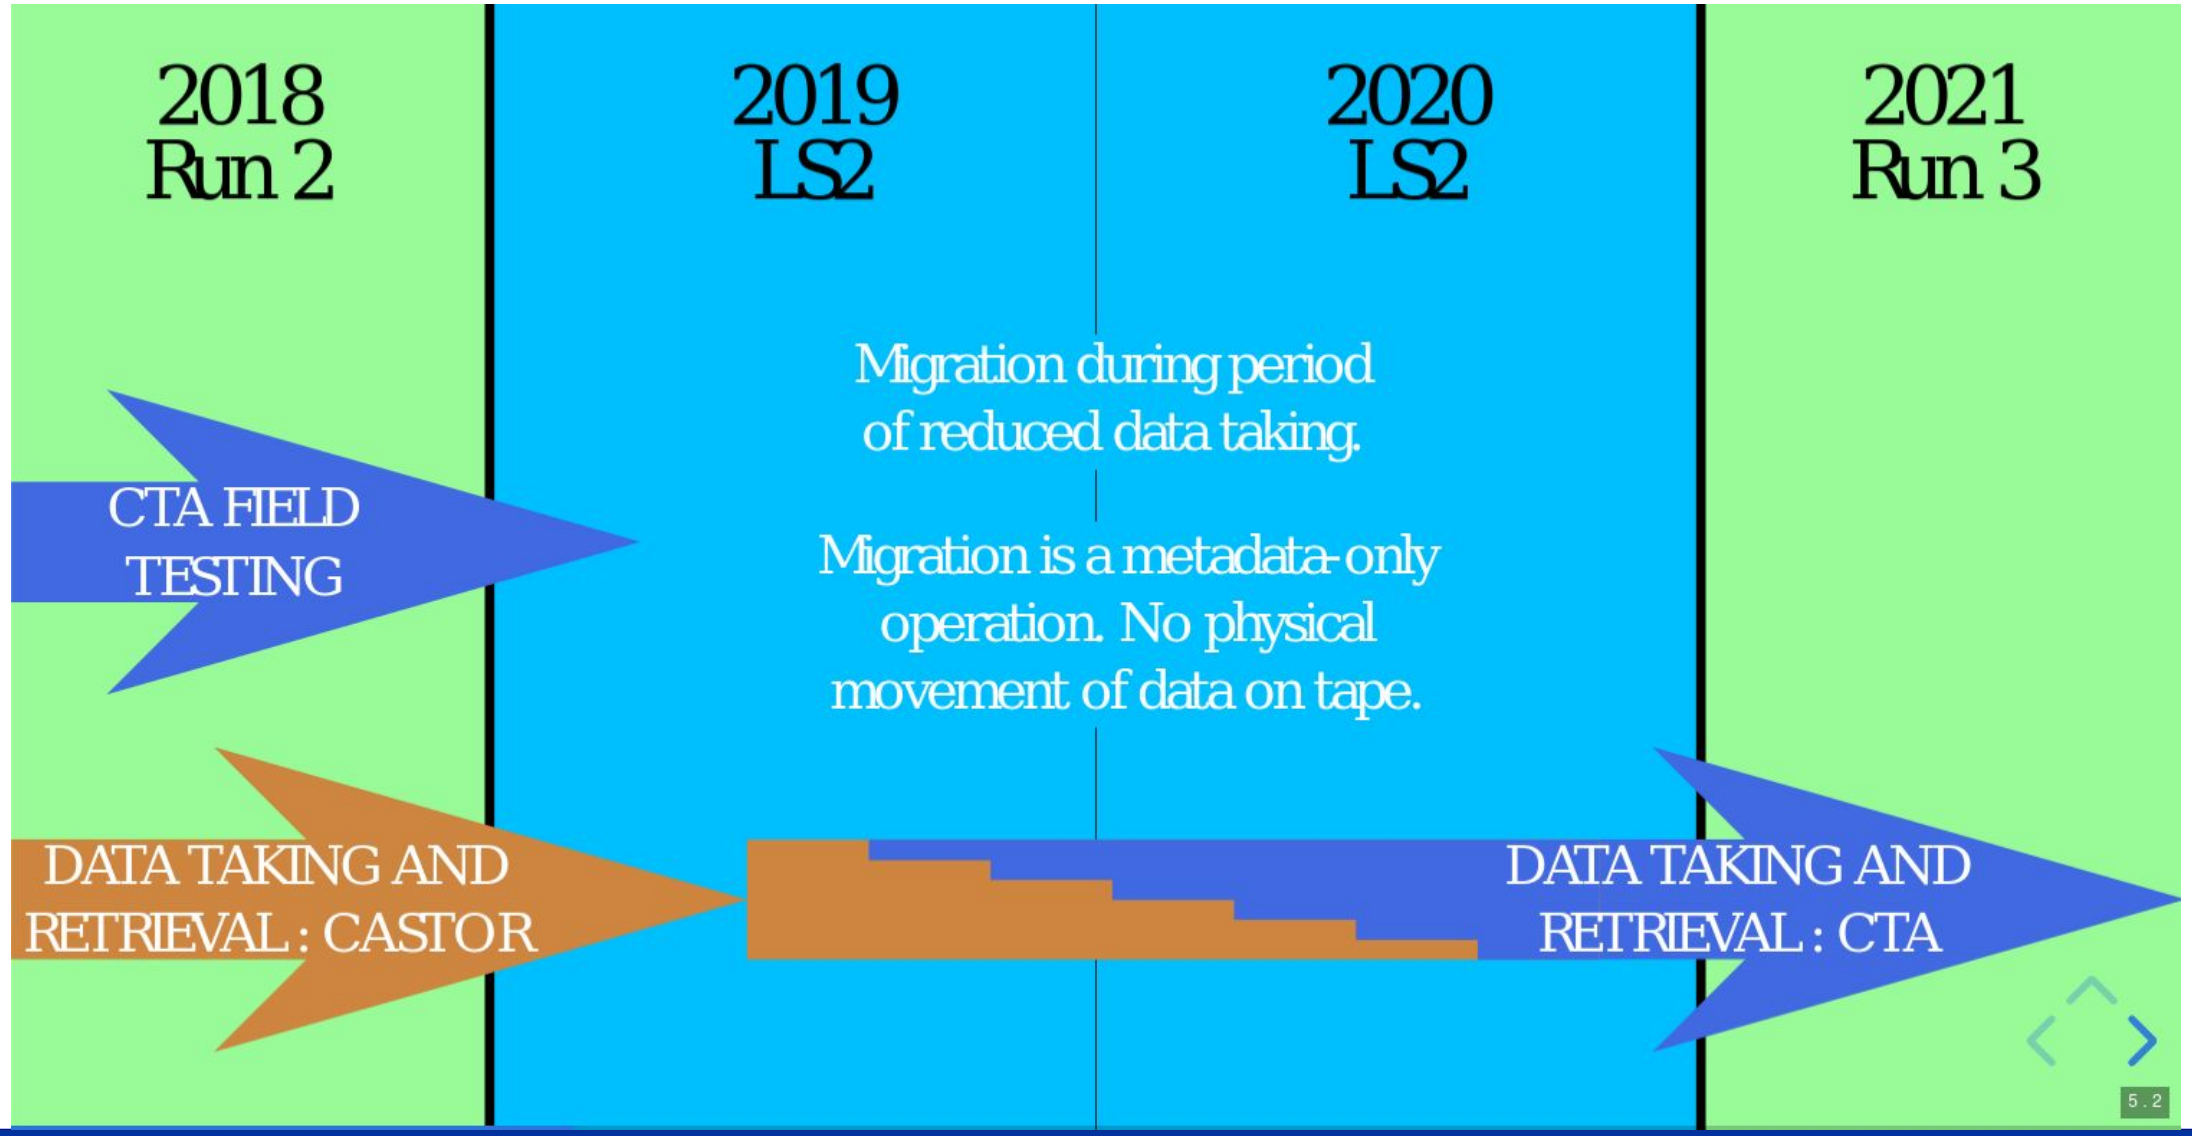

EOS+CTA Timeline

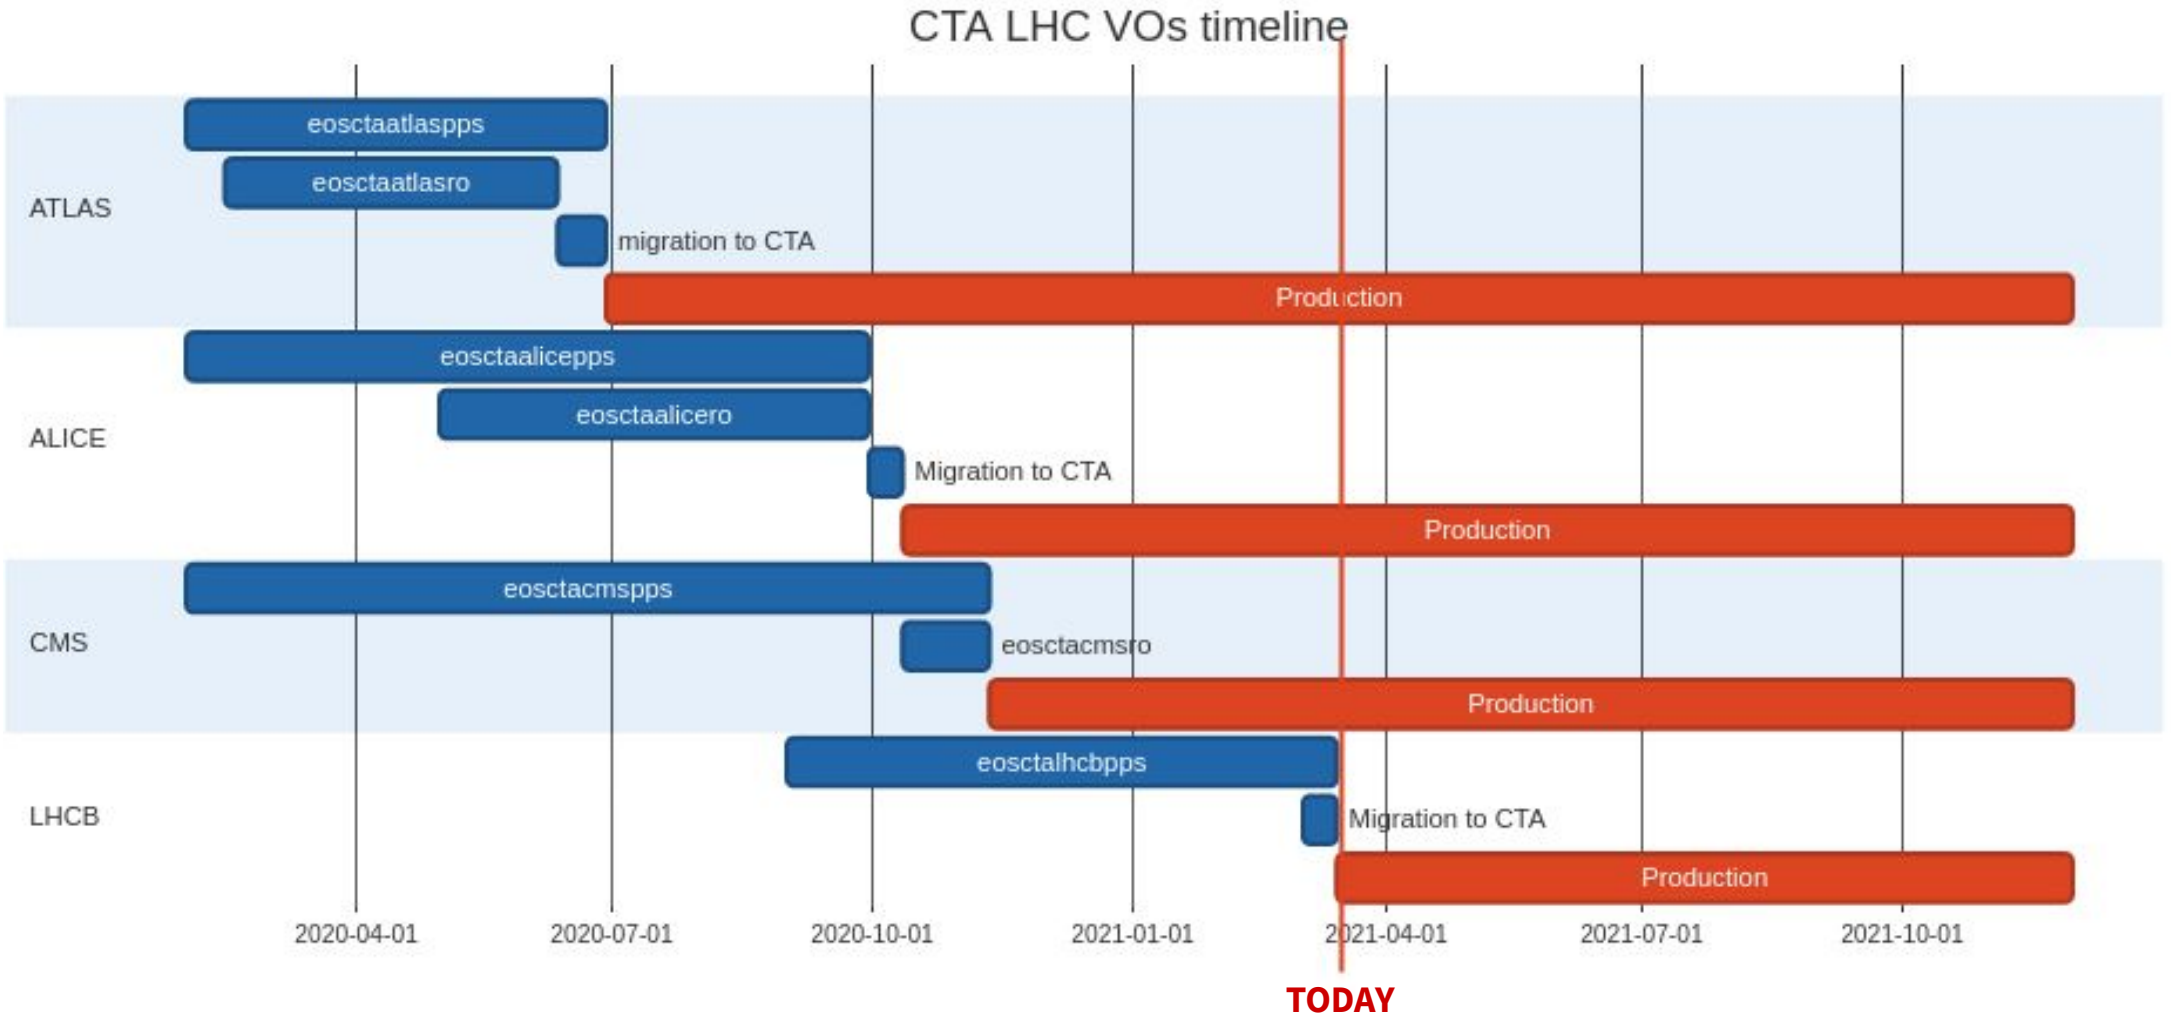

EOS+CTA Timeline

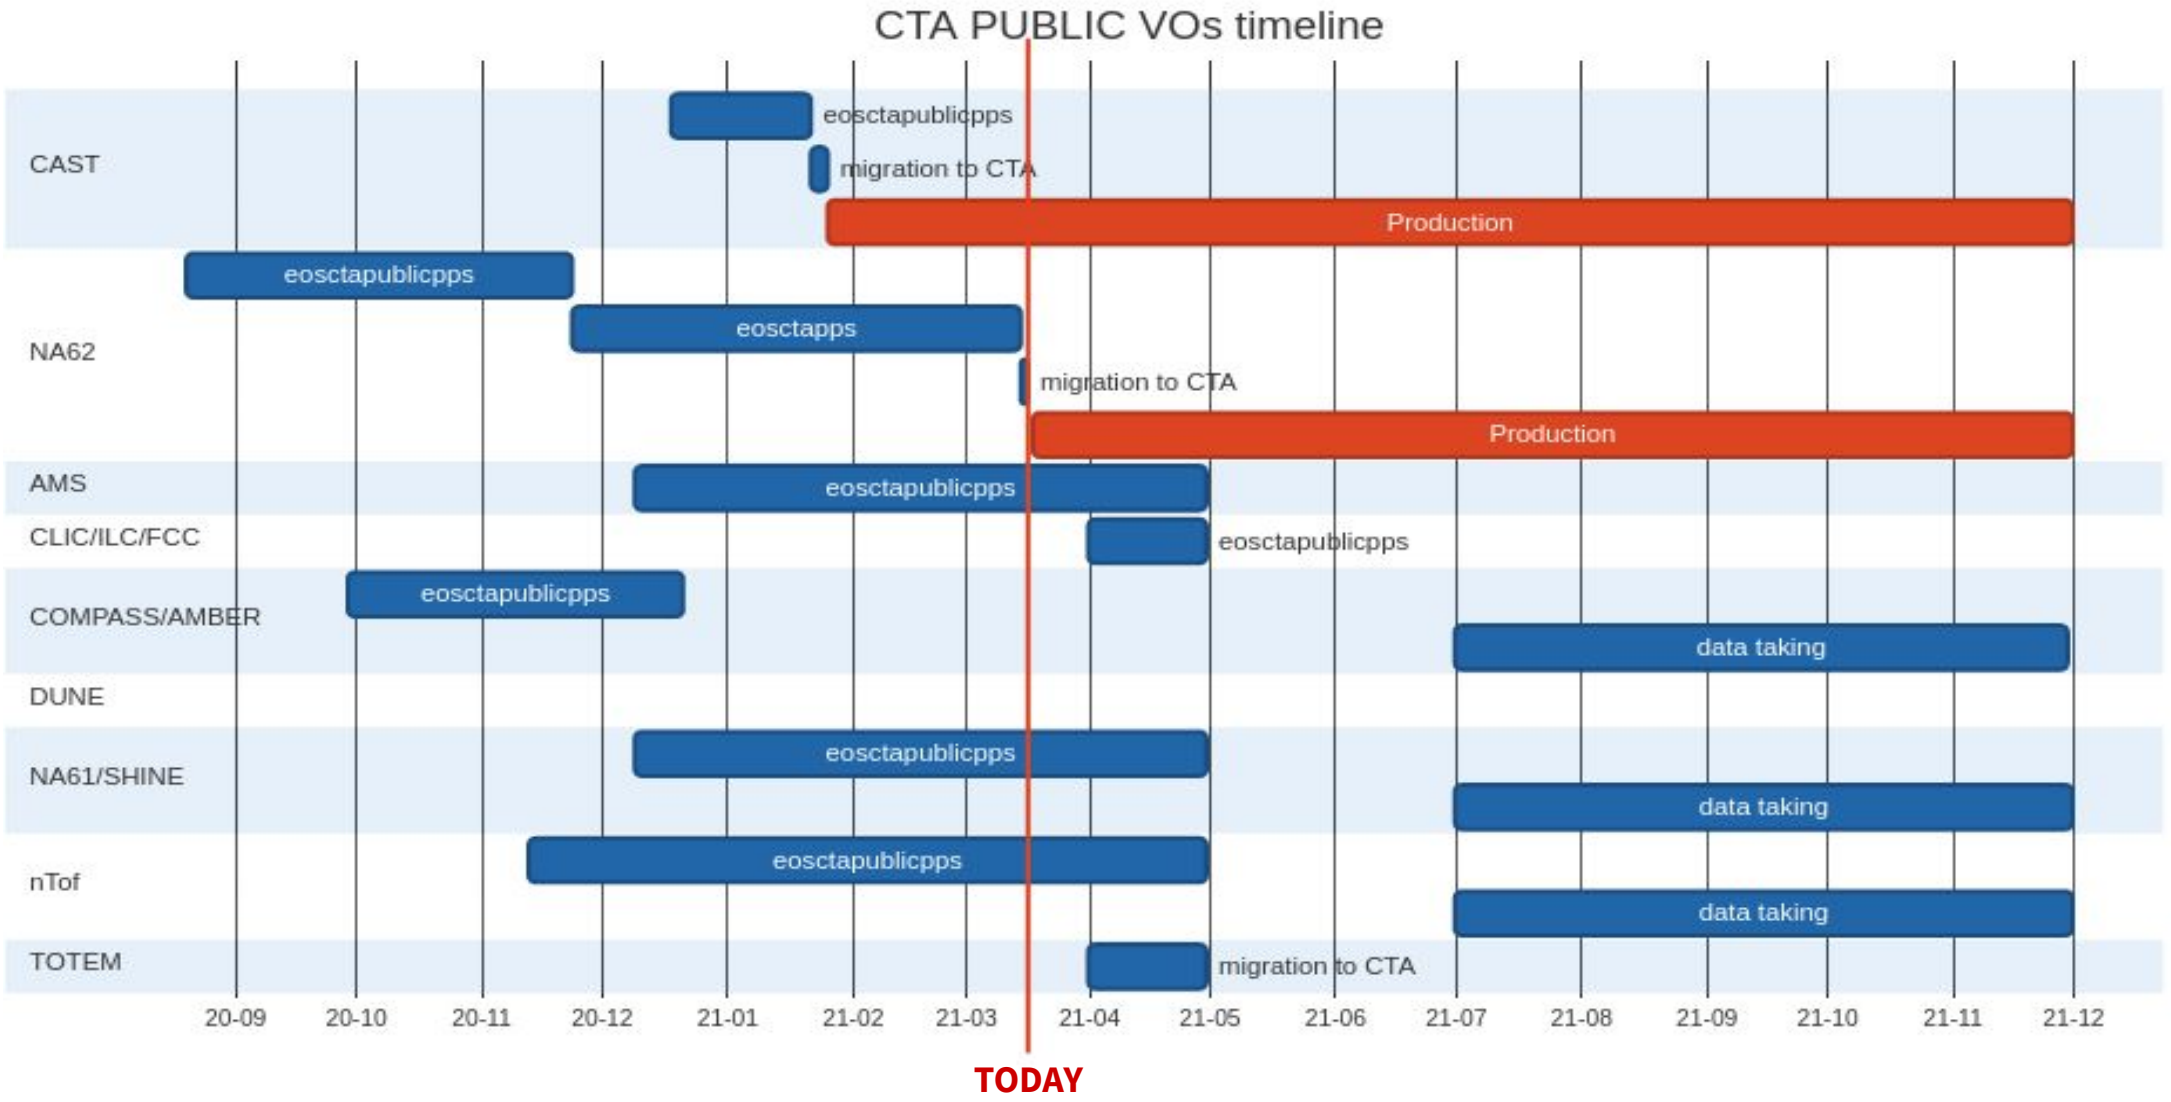

Physics data in CASTOR and CTA

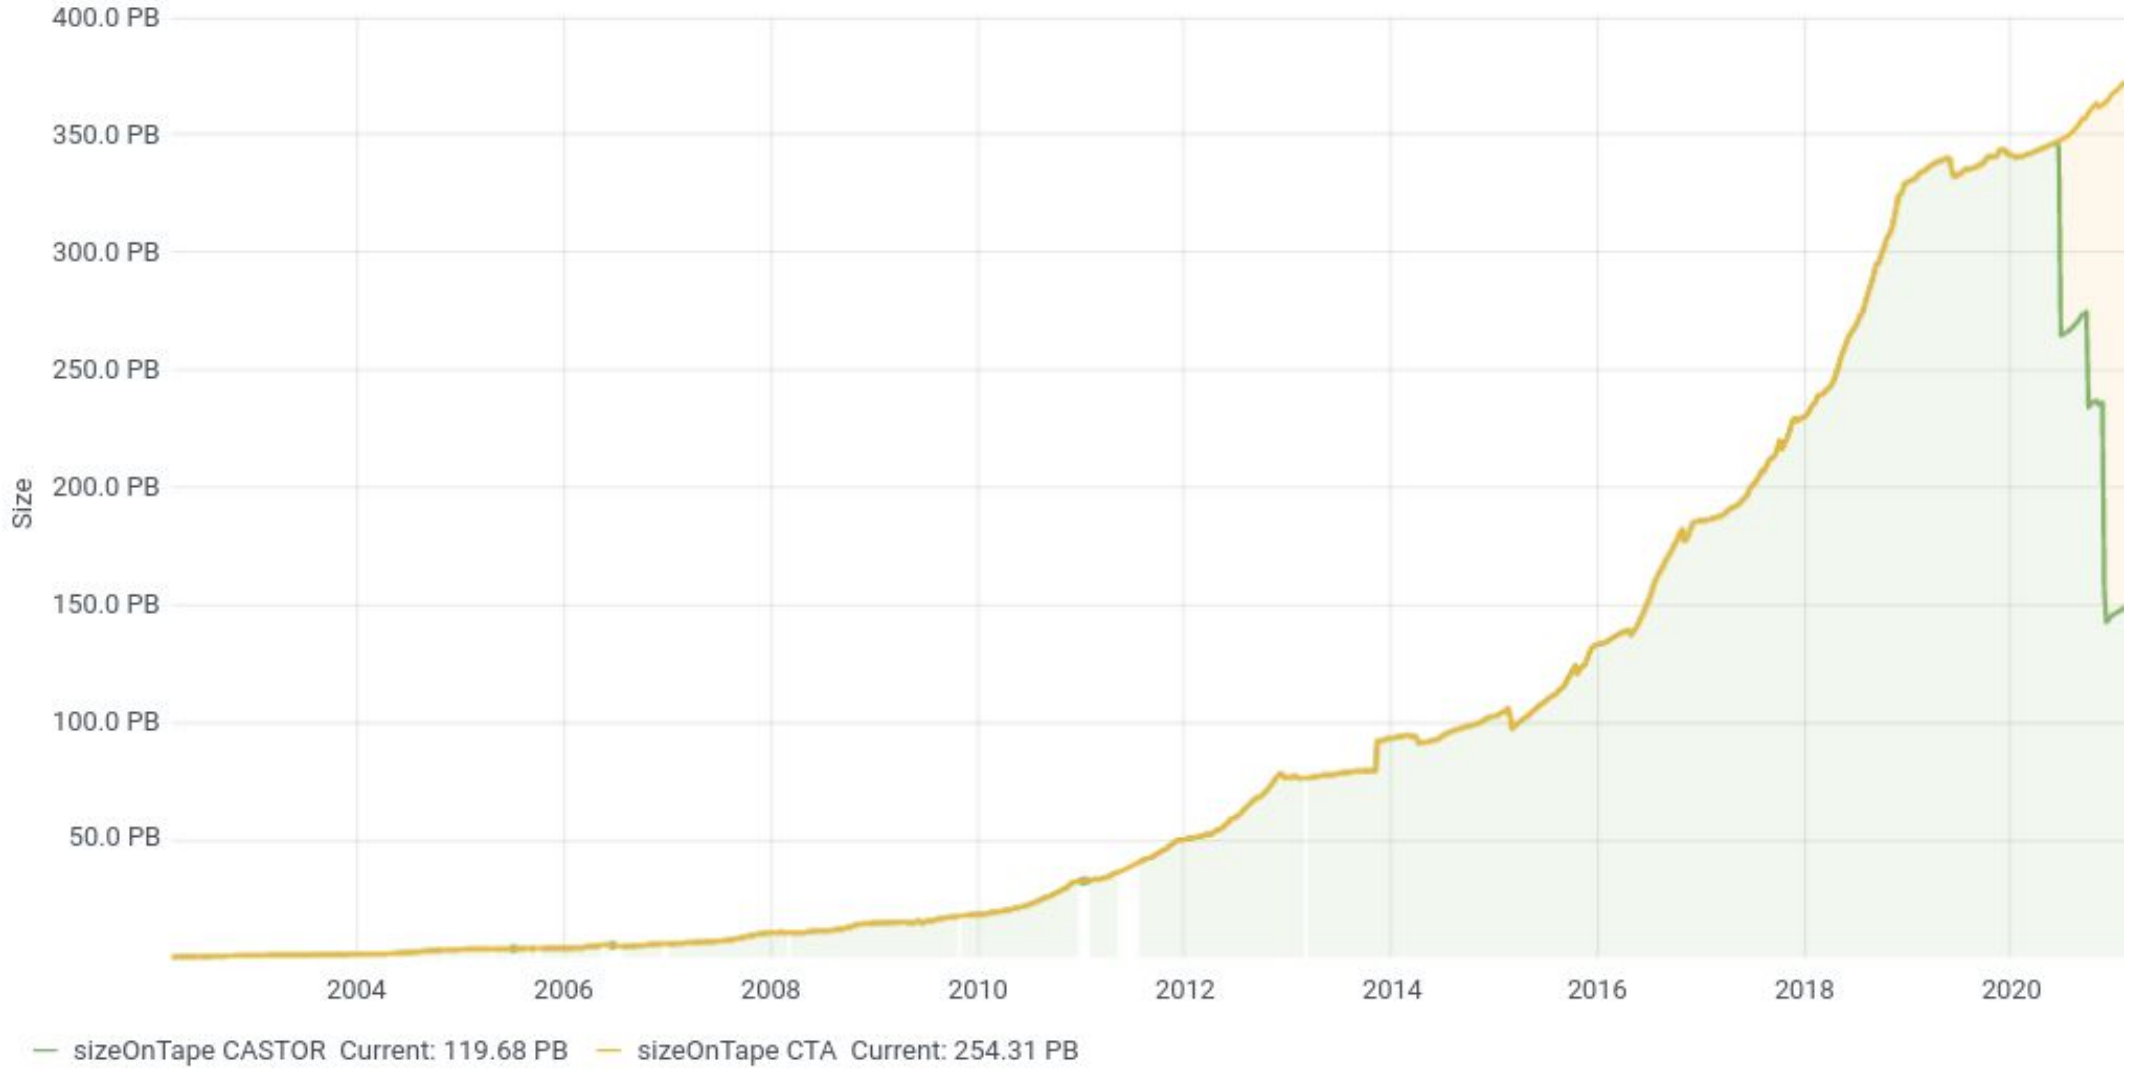

CTA Production Cumulated READ

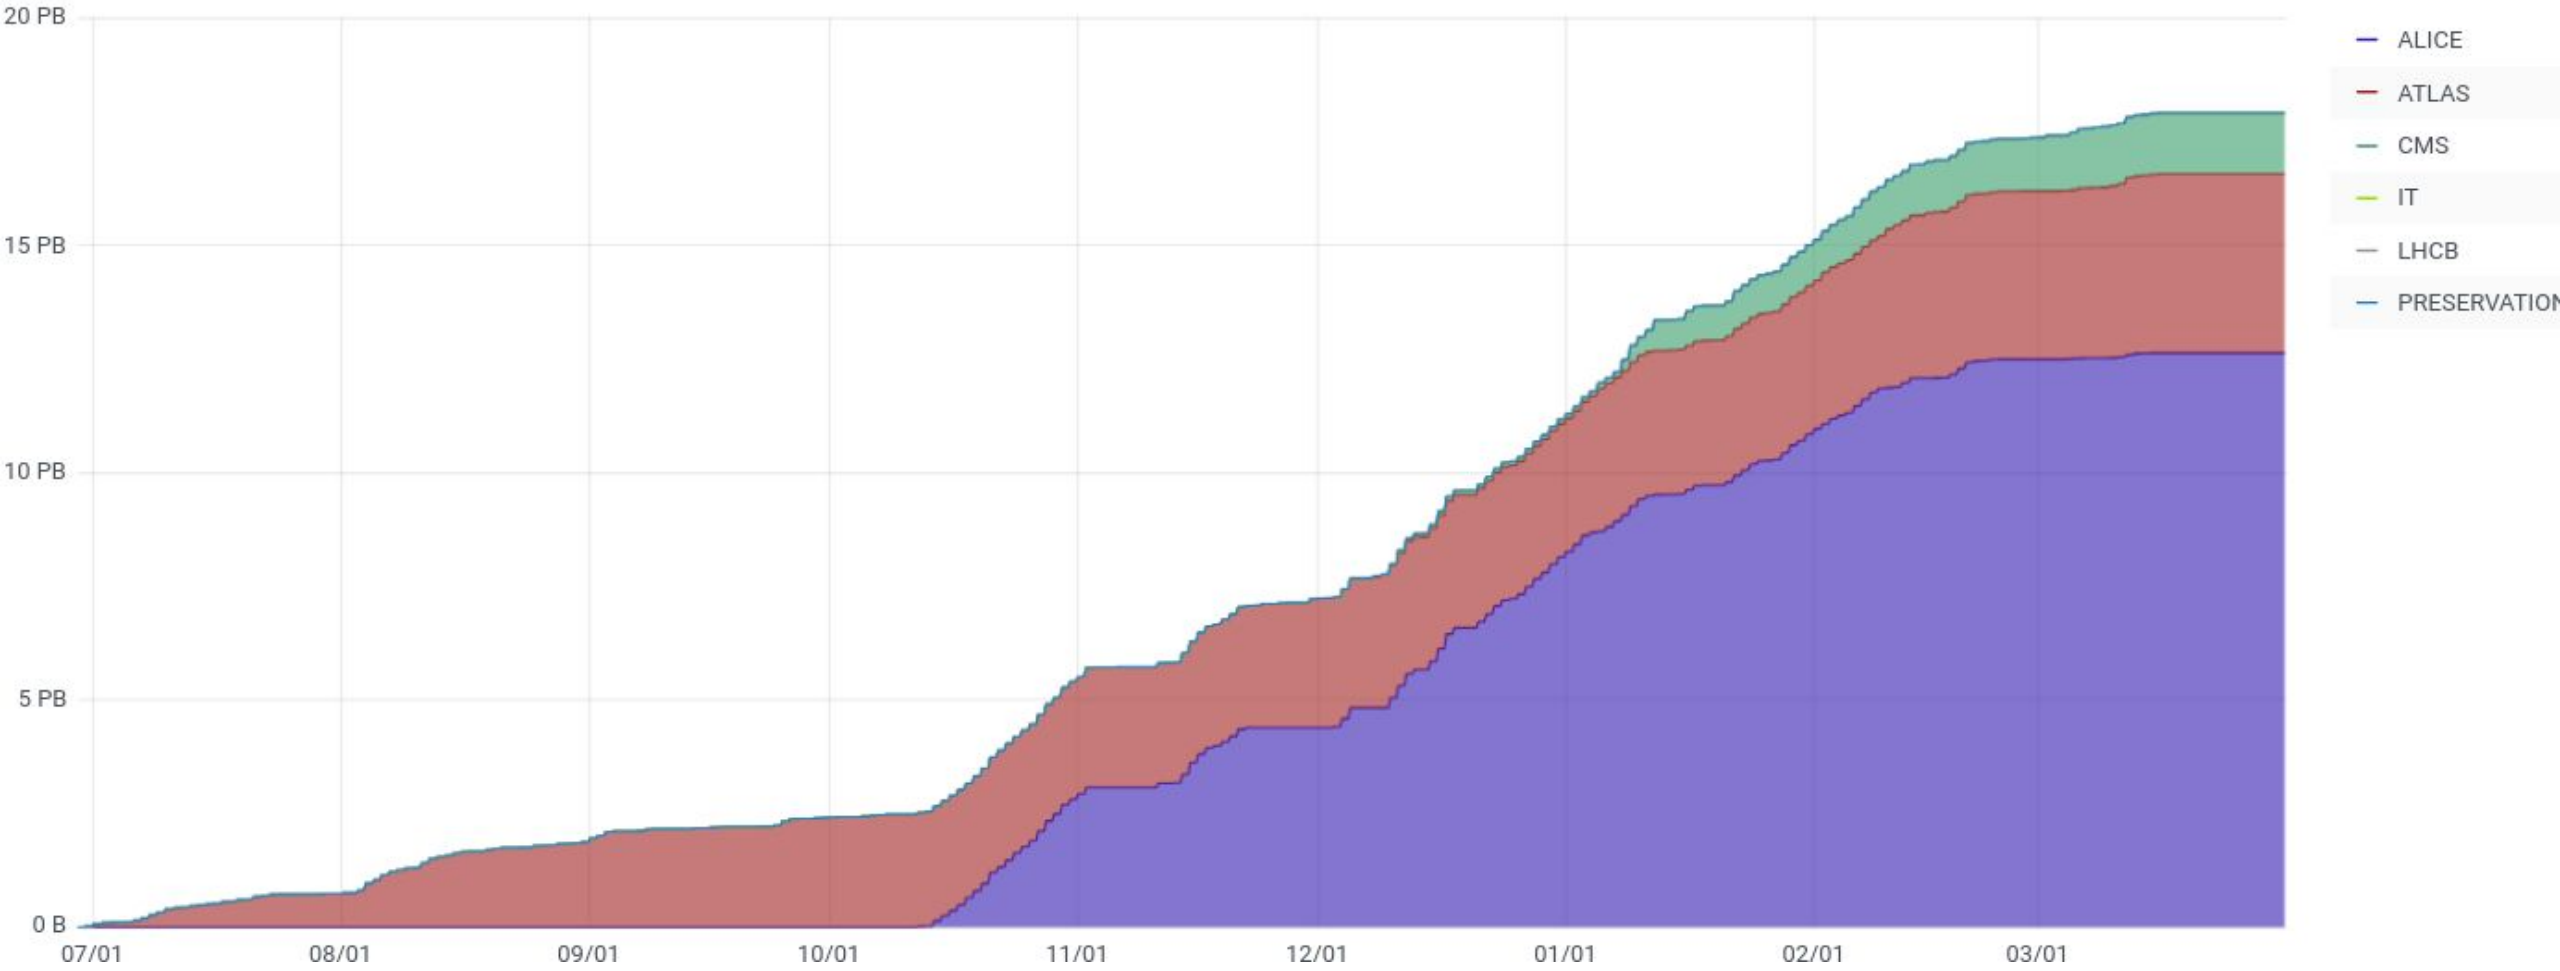

CTA Production Cumulated WRITE

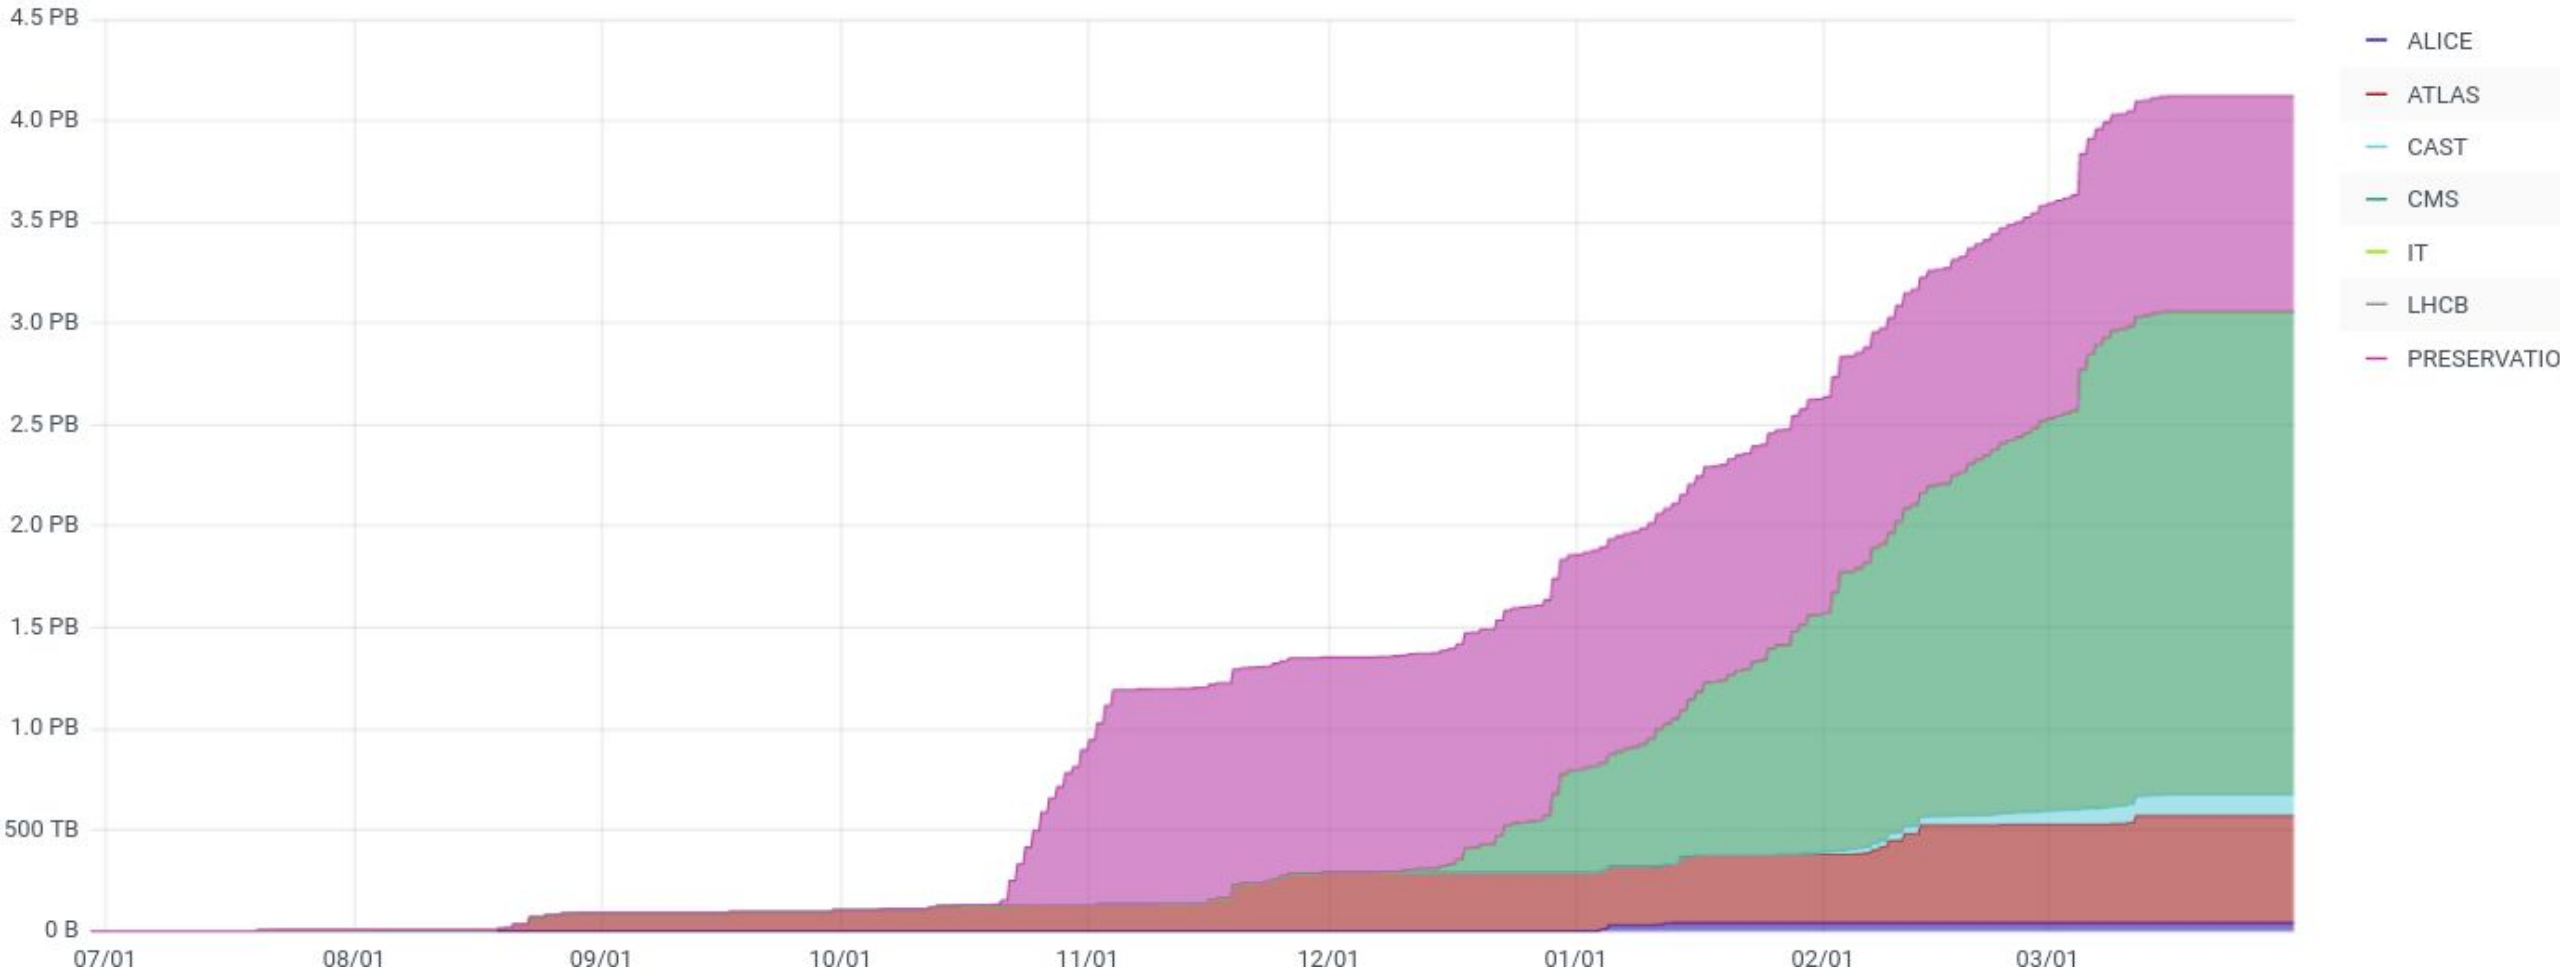

EOSCTA infrastructure for Run3

- Run3 constraints:
  - >60GB/s of bandwidth
  - 8 hours of cache

- 64 buffer servers installed:
  - 200GB of RAM, 500GB-1TB NVMe (OS + logs)
  - 16x2TB SSDs, 25Gb/s each
    - total: 2PB at 200GB/s simplex

- 100Gb/s Router uplinks (no stacking)
  - ~⅔ blocking factor

Current operations mode

- Current production at bandwidth nominal speed

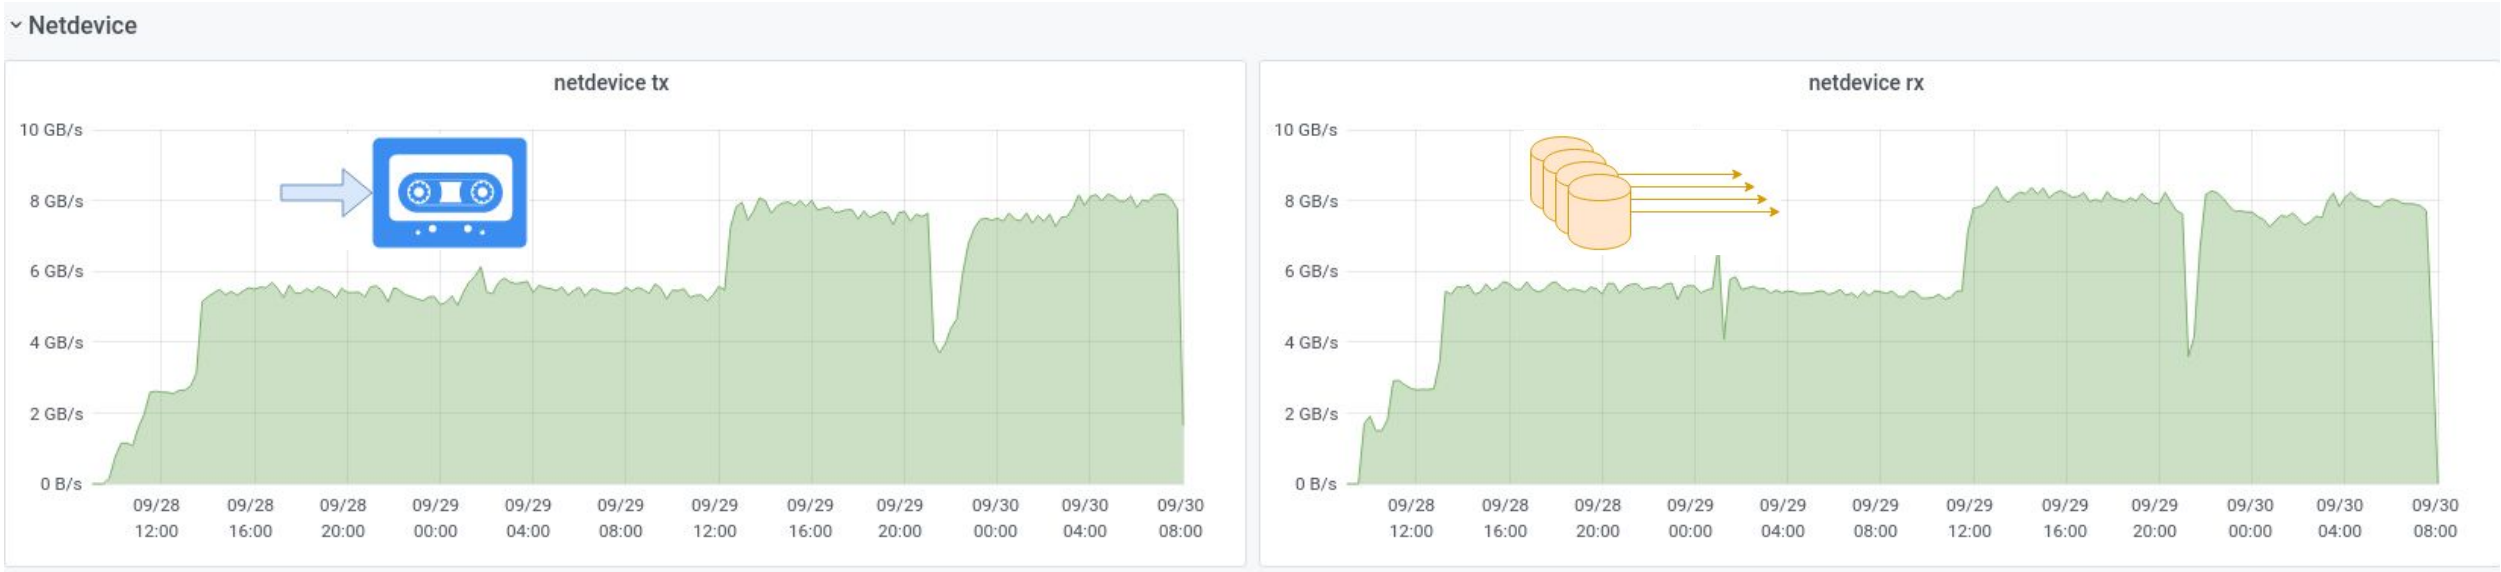

3 buffer servers saturating the tape infrastructure during 1PB write test

- 90TB of buffer only (230TB at Run3 nominal capacity)

FTS CheckOnTape feature deployment

- All files in *default* space are on their way to tape:
  - in *archiving* state in FTS until it is on tape
- If file is on tape before *archive_timeout* expires and no tape transfer error
  - successful FTS transfer
  - failure otherwise

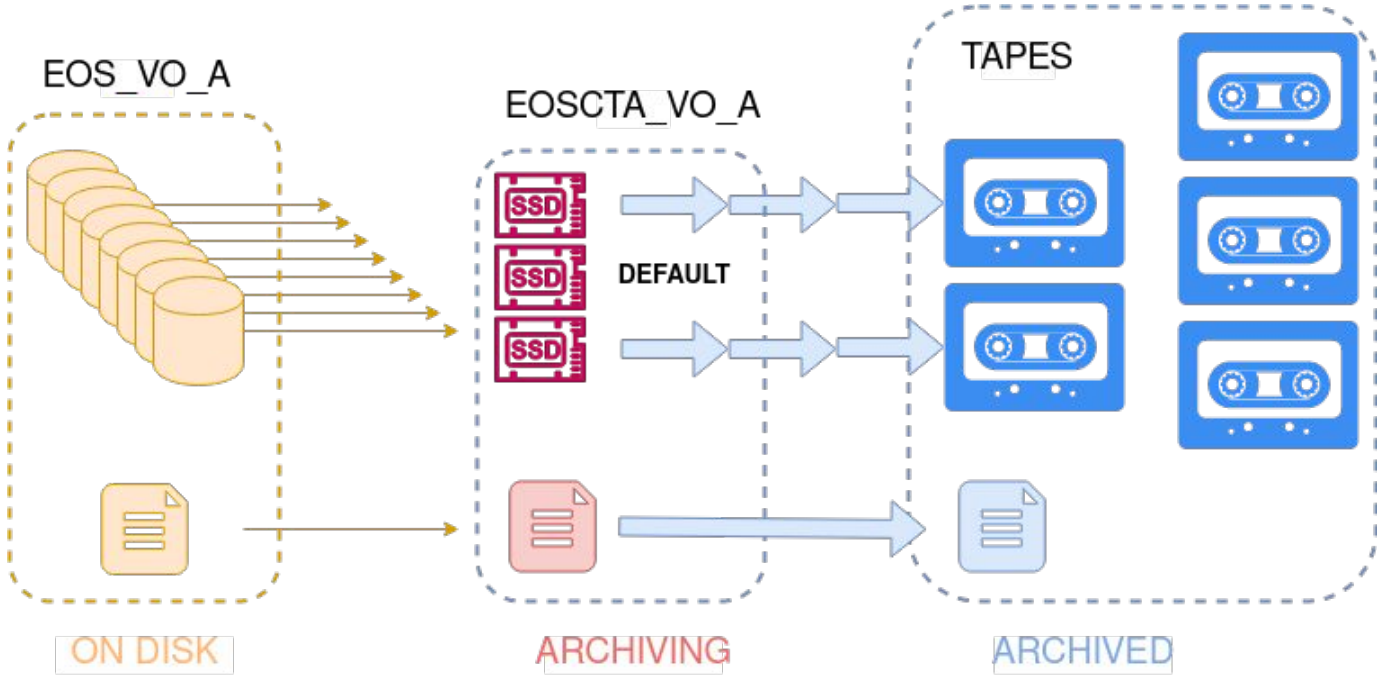

Heavily tested with FTS developers and the experiments
Protects files against disk replica loss, bitflips, truncation...

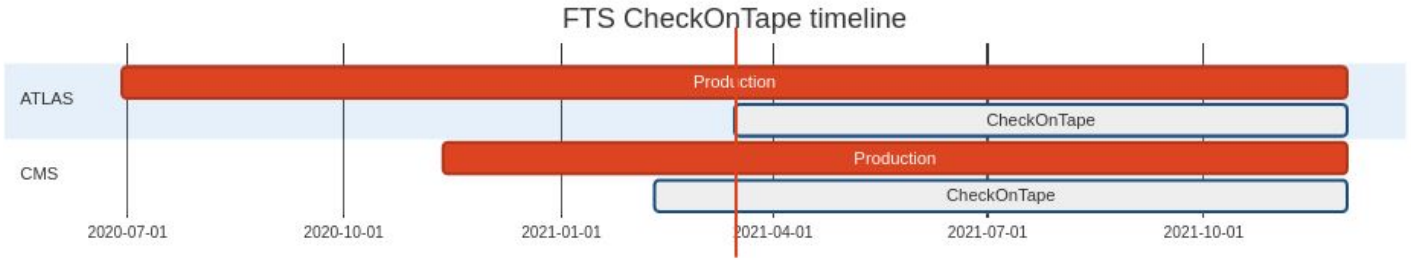

Conclusion

- EOSCTA Run3 tape buffer capacity installed
- **All LHC experiment have been migrated to CTA**
  - Ongoing Public VOs migration
- FTS CheckOnTape now enables **extra safe transfers with no compromise on performance**
- CTA Run3 tape infrastructure on its way (see V. Bahyl presentation)
- DAQ tests and Run3 data challenges in the pipeline

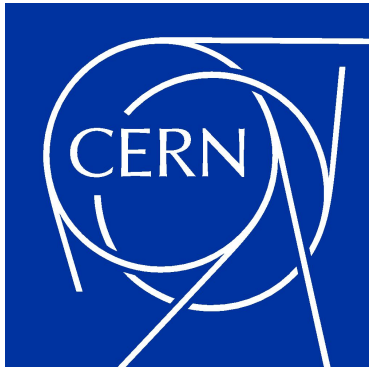

home.cern