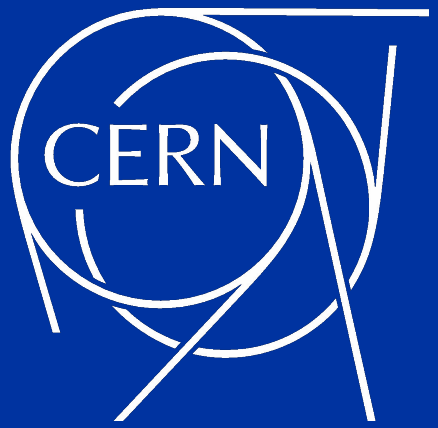

CTA operations best practices

Julien Leduc

2024-03-18 - CTA workshop

T0 Tape: CTA Service Level Agreements

- **Store LHC experiments data for Run3**
  - Absorb incoming data at experiment DAQ rate
    - At least 10GB/s of tape archival throughput per LHC experiment
    - Loss of 1 tape library should not impact LHC archival rate
    - Arbitrate best effort bandwidth over 10GB/s
  - Provide enough staging performance for data replication campaigns to T1s
- **Support non LHC experiments and non physics use cases**
  - Small and Medium Experiments (SMEs)
  - IT backup use cases
  - Preservation use cases (LEP, Babar,...)

4x10GB/s + 8GB/s + 10GB/s (library redundancy) = 6 x libraries with ~30 tape drives each

T0 tape: Run1 to Run3 evolution

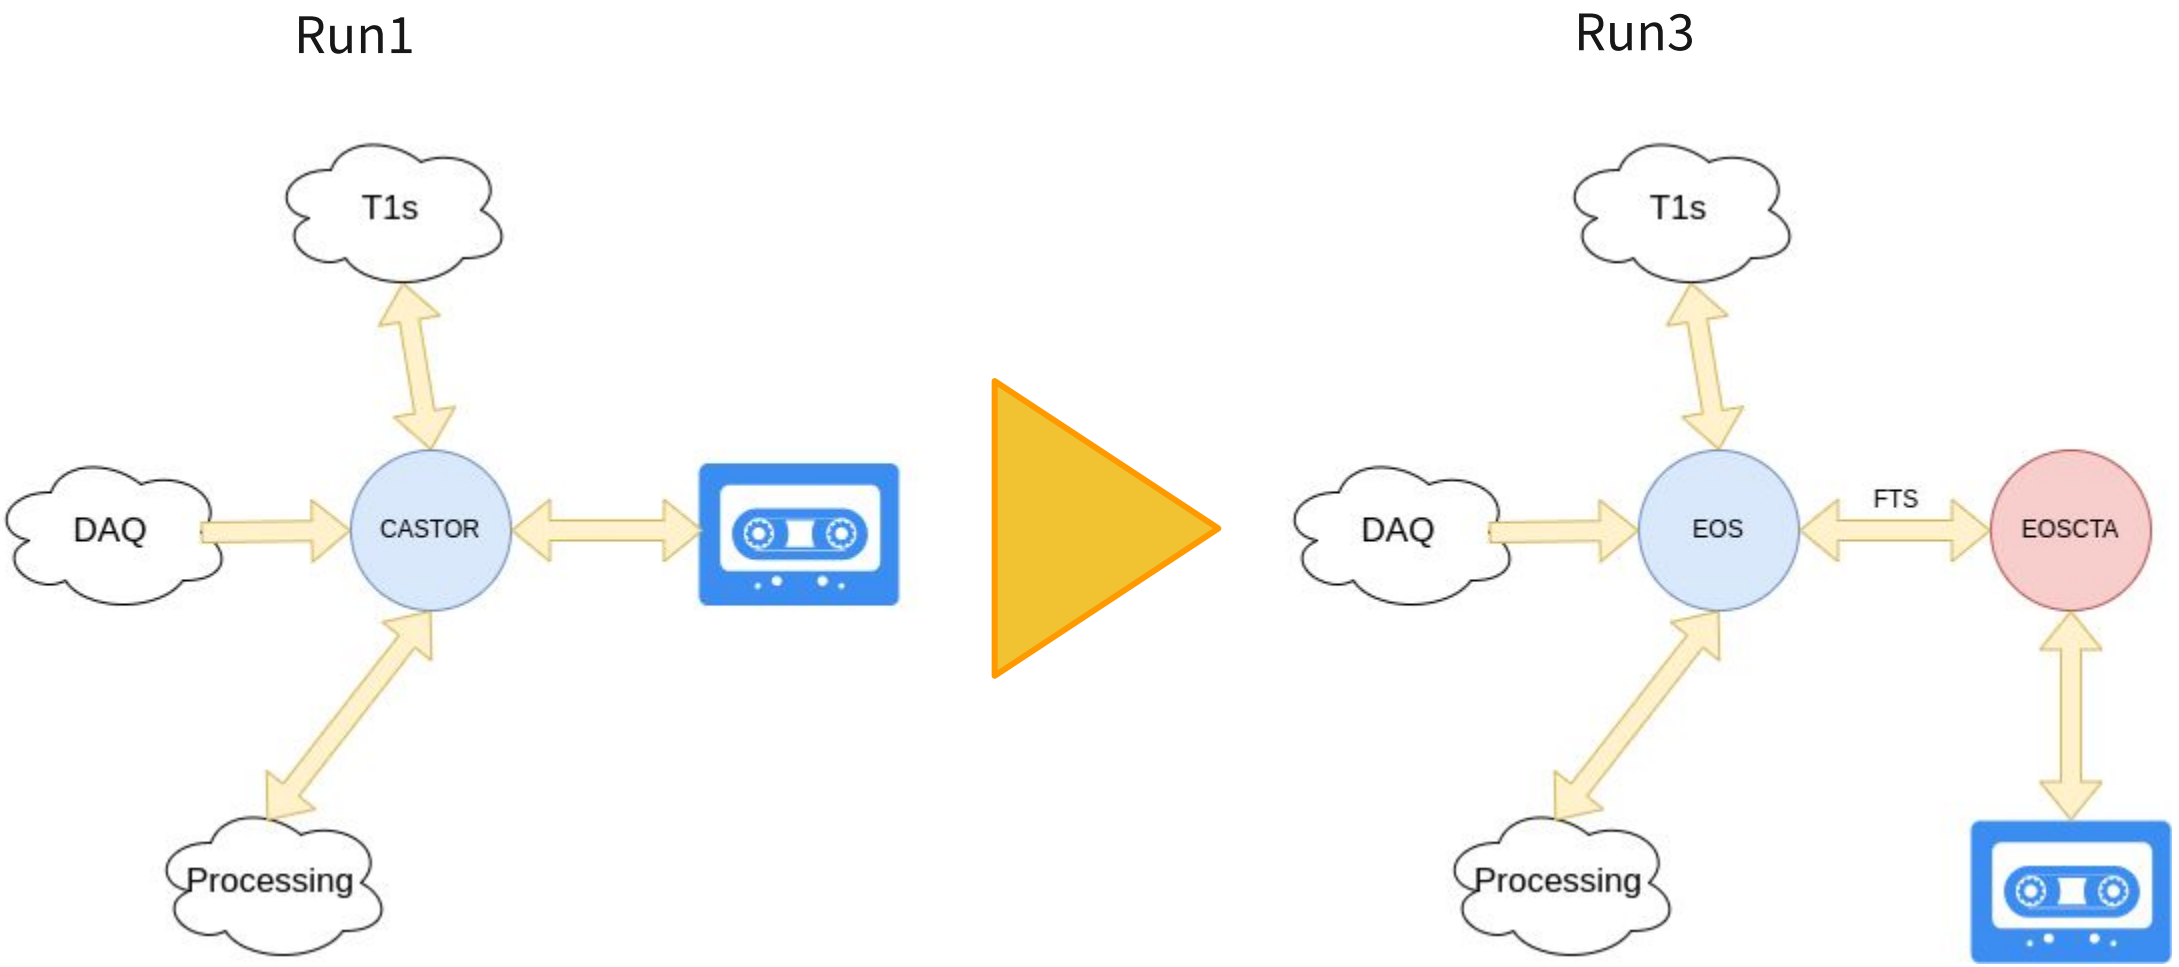

EOS+FTS+EOSCTA Architecture

- **EOSCTA is a pure tape system**
  - **Transfer to eoscta instance is a transfer to tape, strictly requires users to check that transferred files are on tape (FTS *archive_timeout*)**
- **Disk cache duty consolidated in main EOS instance.**
- **Operating tape drive at full speed full time efficiently requires a single replica SSD based buffer: EOSCTA**

EOSCTA service resource hierarchy

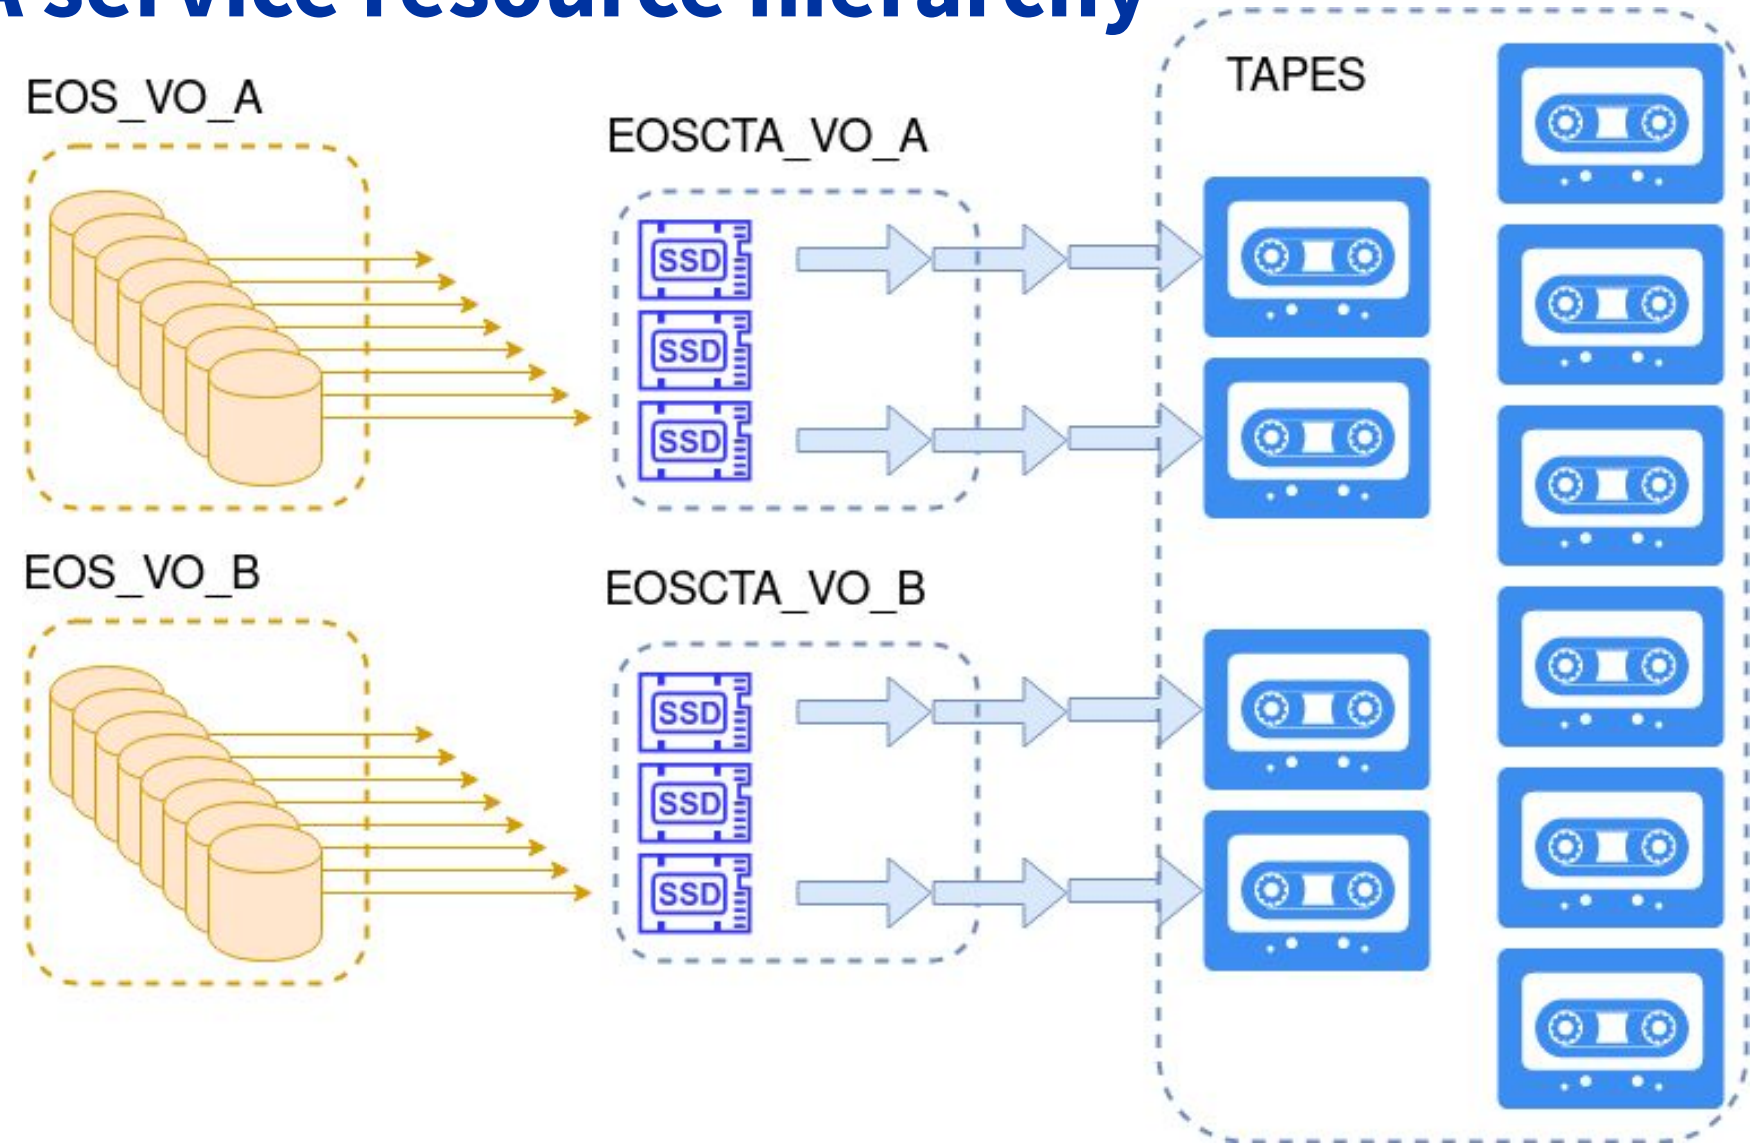

Match eoscta instance for LHC use cases

- **Consolidate use cases per eoscta instance**
  - 1 eoscta per large LHC experiment:
    - ALICE, ATLAS, CMS, LHCb - **strict SLA high performance**
    - **Data rates follow LHC schedule for DAQ with some variations per VO**
  - **Per VO SLA is controlled with writemaxdrives and readmaxdrives: SLA of 10GB/s means no more than 35-40 tape for VO distributed between read and write**
    - **Additional Best Effort throughput**

Match eoscta instance for non LHC use cases

- **Consolidate use cases per eoscta instance**
  - shared eoscta instances grouped by use case:
    - eosctapublic: Small and Medium Experiments (SMEs) using FTS - **documented workflows with single point of contact**
    - eosctapublicdisk: SMEs using legacy workflows, reading data directly on eoscta instance, preserved retired experiments, … - **low traffic, low volume, higher support needs**
    - eosctabackup: backup use cases running custom archive/retrieve scripts with dedicated BACKUP_*VO* (AFS data, EOS NS, EOSCTA NS, HDFS data, … - **high traffic, high data rotation, more technical support**
  - **Per VO SLA is controlled with writemaxdrives and readmaxdrives: SLA of 10GB/s means no more than 35-40 tape for VO distributed between read and write**

EOSCTA *default* instance flavor

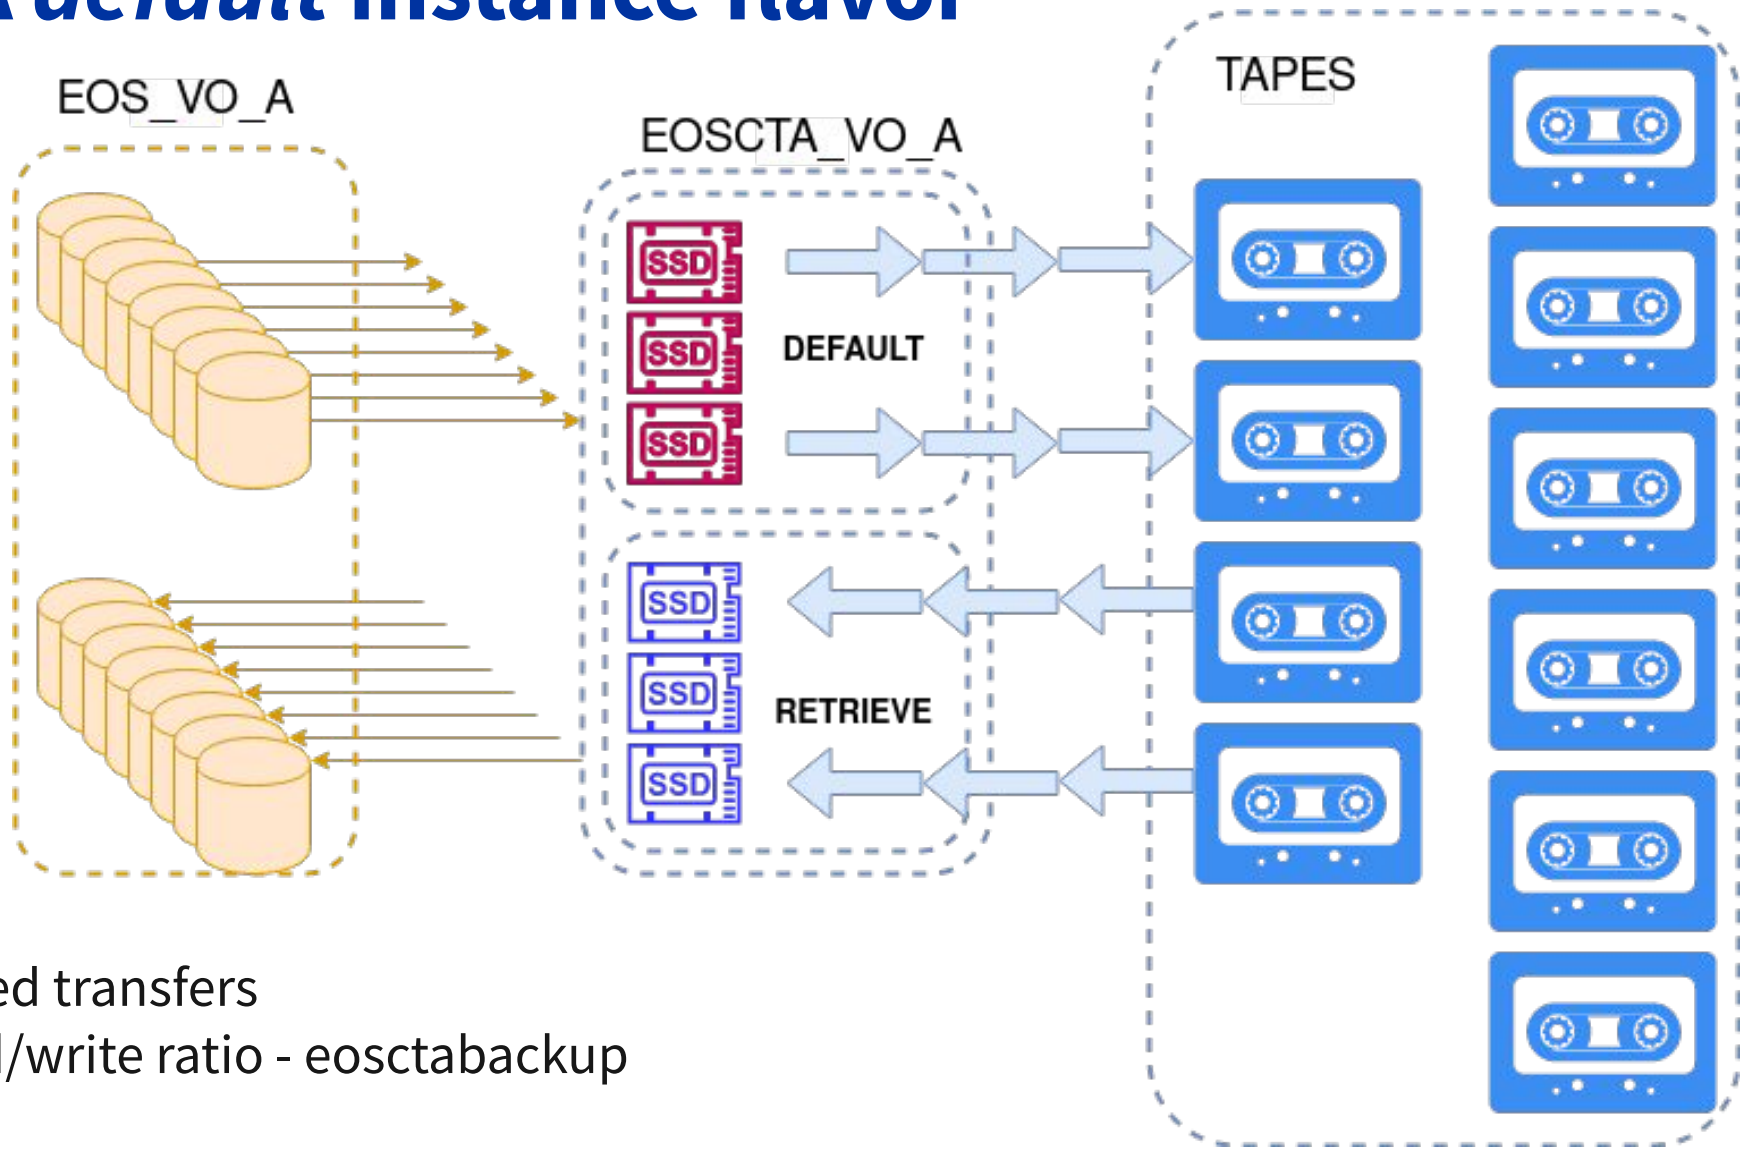

- FTS based transfers
- low read/write ratio - eosctabackup

EOSCTA *spinners* instances flavor

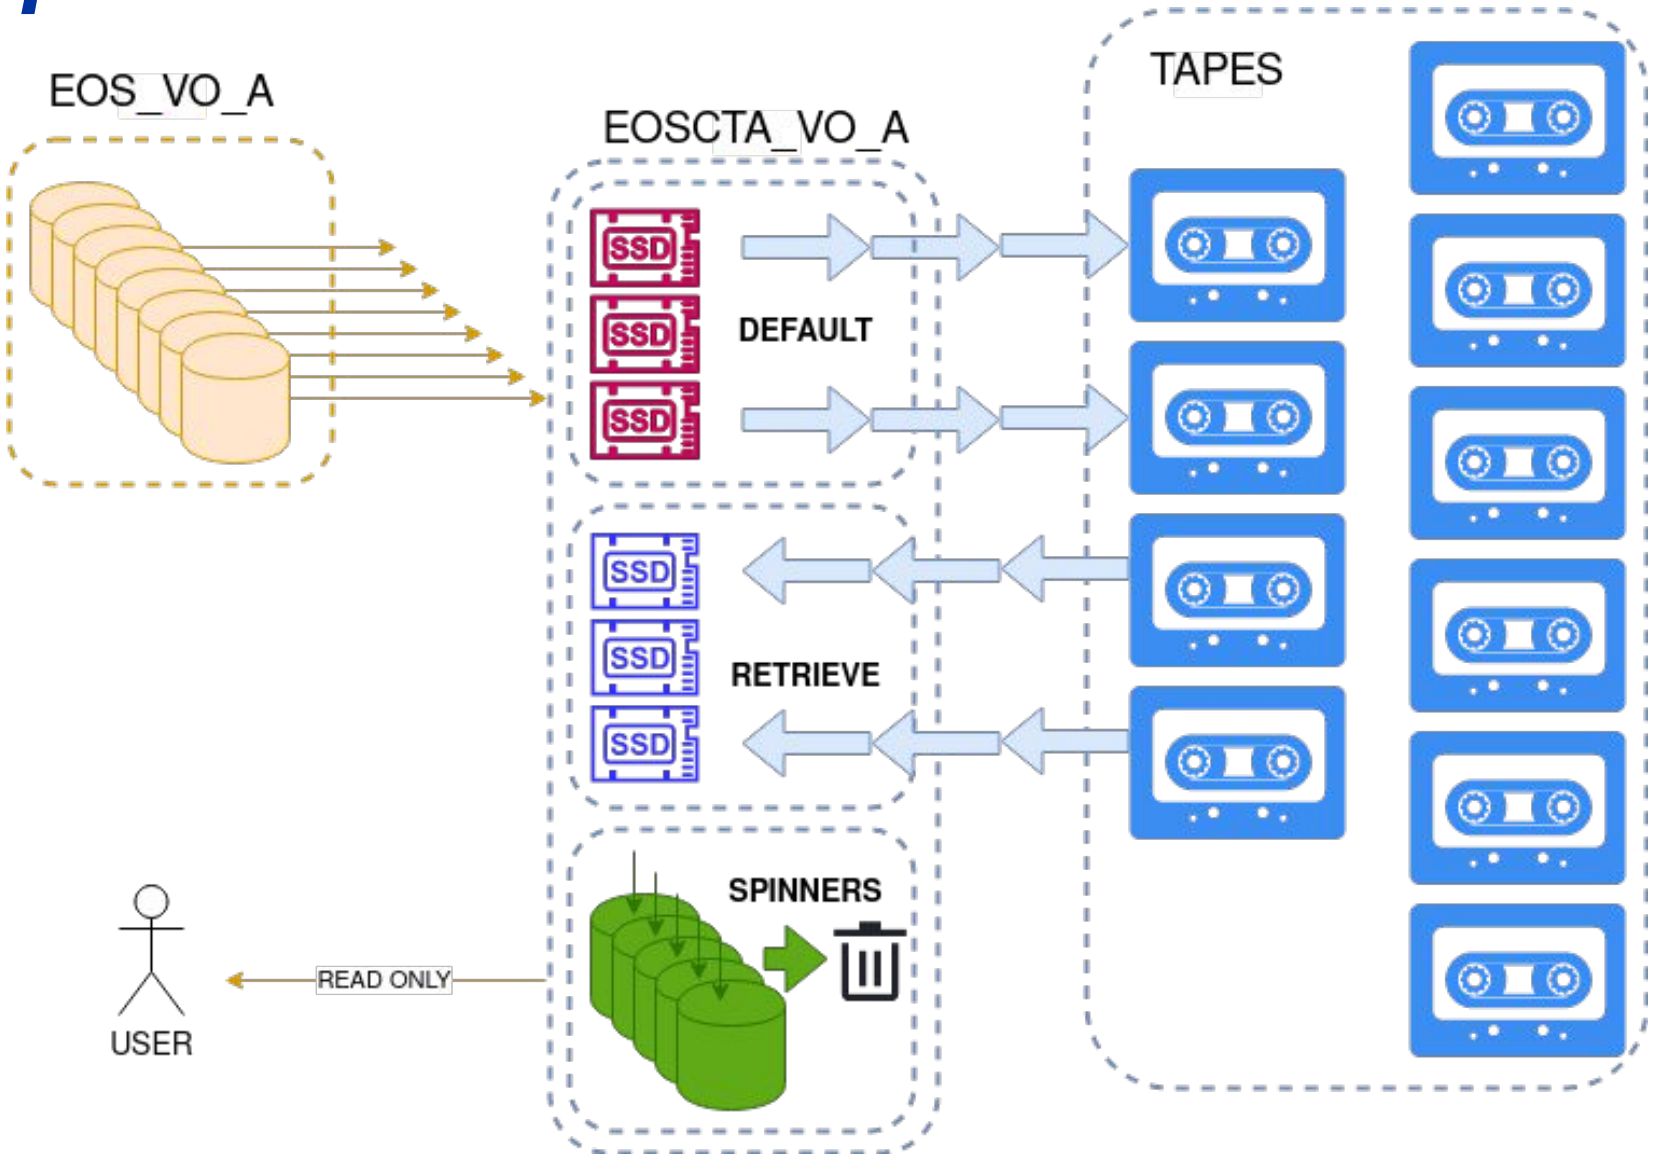

Efficient use of SSDs

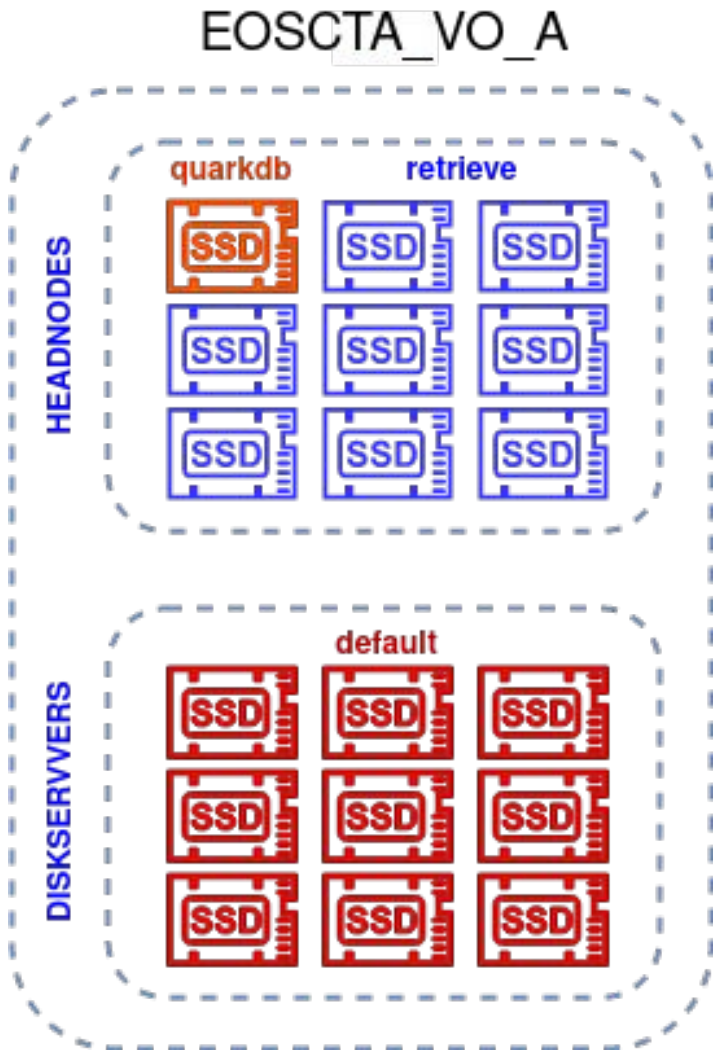

Minimise operation pain

- run *master* NS services on dedicated server or collocate with retrieve space

At T0 hardware model for headnodes and diskservers is identical

Balance bandwidth evenly between SSDs (strict RR stream allocation across FSes)

- **1 EOS filesystem per SSD**
- **1 scheduling group per EOS filesystem**

Currently using a *flat* geoscheduler tree

- **All production servers have the exact same hardware**
  - **all FSes have the same performance**
- **In the next EOS *flat scheduler* bandwidth will be simpler to balance between file systems**

Archive/Staging bandwidth allocation

Standard LHC eoscta instance: 10GB/s archival SLA for CTA T0

- **Archive boost needed during data taking, tape flushing, Heavy Ion run**
- **Staging boost during Year End Technical Stop (YETS) Heavy Ion data duplication to T1s/T2s**
- **Change eoscta bandwidth allocation accordingly**

Run3 LHC planning

Archived volume per month

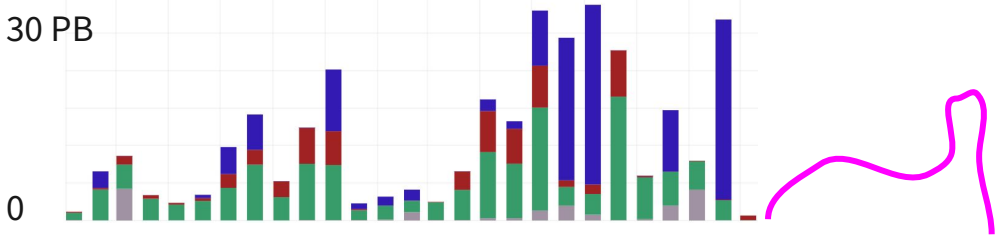

Staged volume per month

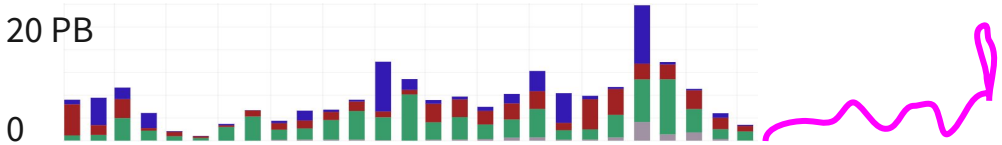

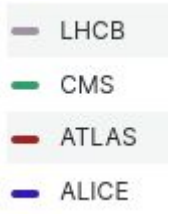

AI/quantum computing generated interpolation

Archive/Retrieve bandwidth allocation

Standard LHC eoscta instance: 10GB/s archival SLA for CTA T0

- **10 SSD servers**
- **Archive boost during data taking, tape flushing, Heavy Ion run**
- **Staging boost during Year End Technical Stop (YETS) HI data duplication to T1s/T2s**
- **Configure CTA ALICE VO writemaxdrives, readmaxdrives accordingly**
- **Measure bandwidth to/from tape buffer AND INDIVIDUAL TAPE DRIVE EFFICIENCY**

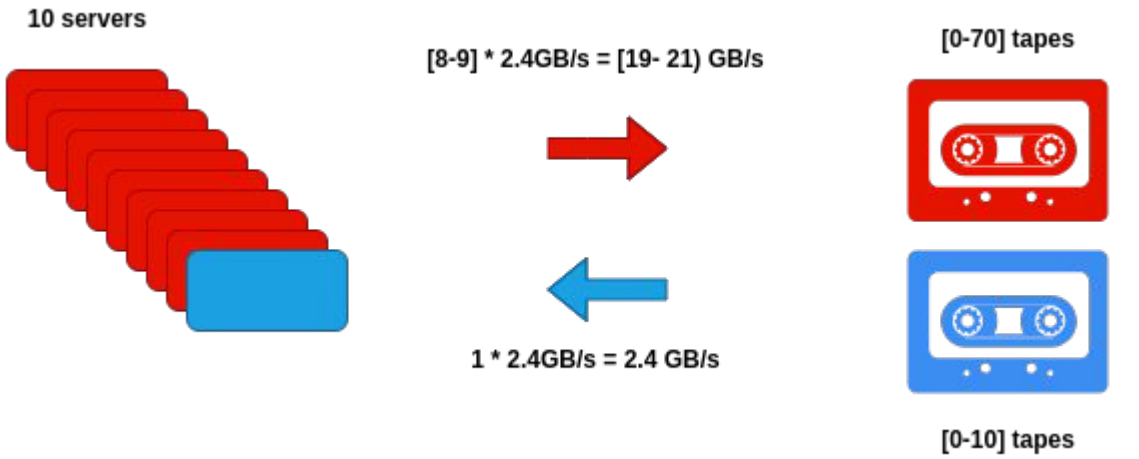

Data taking configuration

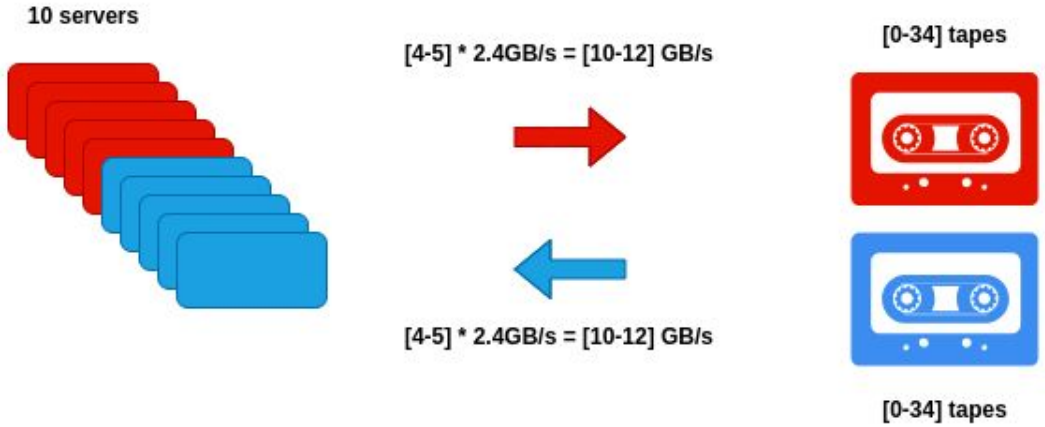

YETS configuration

Archive/Retrieve bandwidth allocation: *ALICE DC24*

22PB transferred in 21 days

- 1.5PB per day when Steady

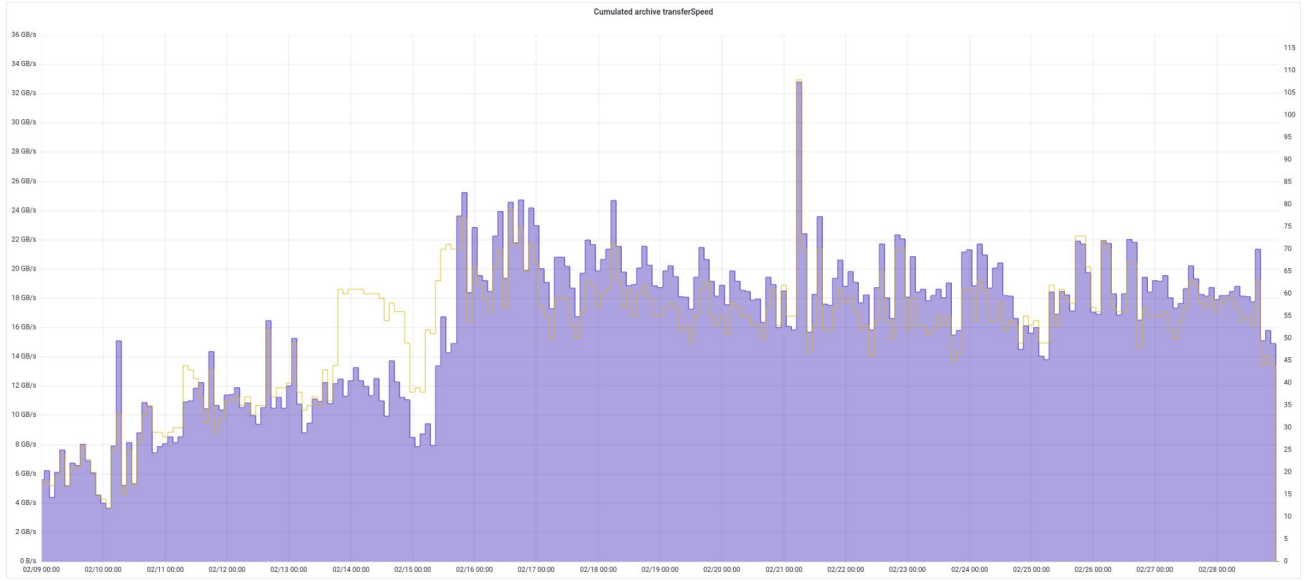

A Large Ion Collider Experiment

Time evolution T0

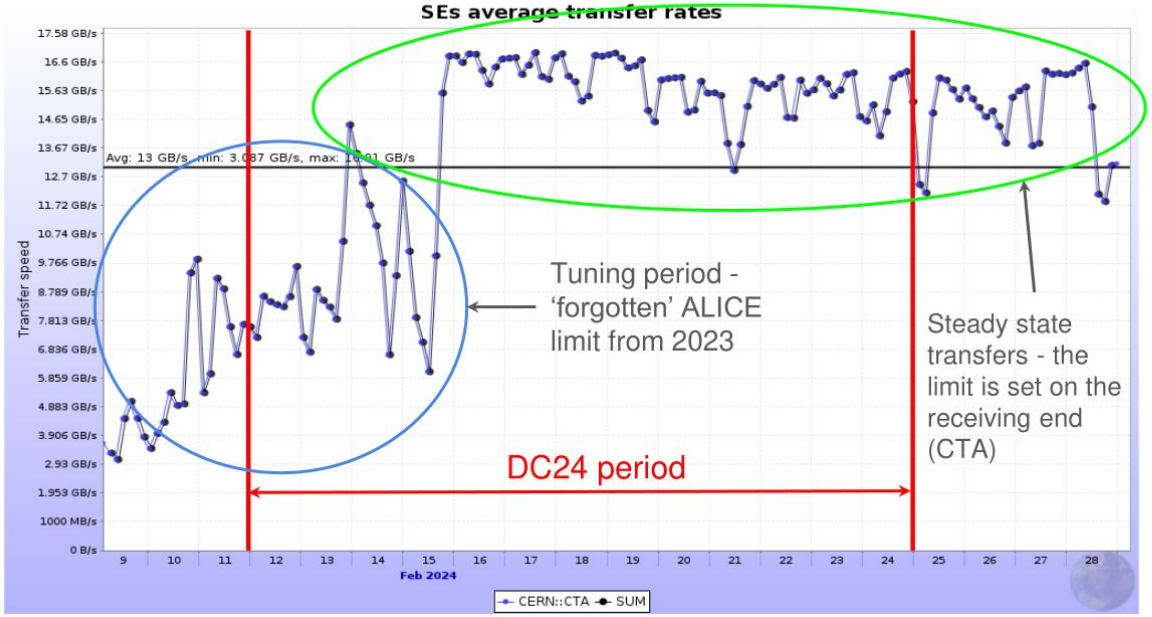

Target SE	Status	Progress	Files	Total size	Started	Ended
ALICE::CERN::CTA	Done		3341292	21.87 PB	08 Feb 2024 22:52	29 Feb 2024 11:13

Archive/Retrieve bandwidth allocation: *ALICE DC24*

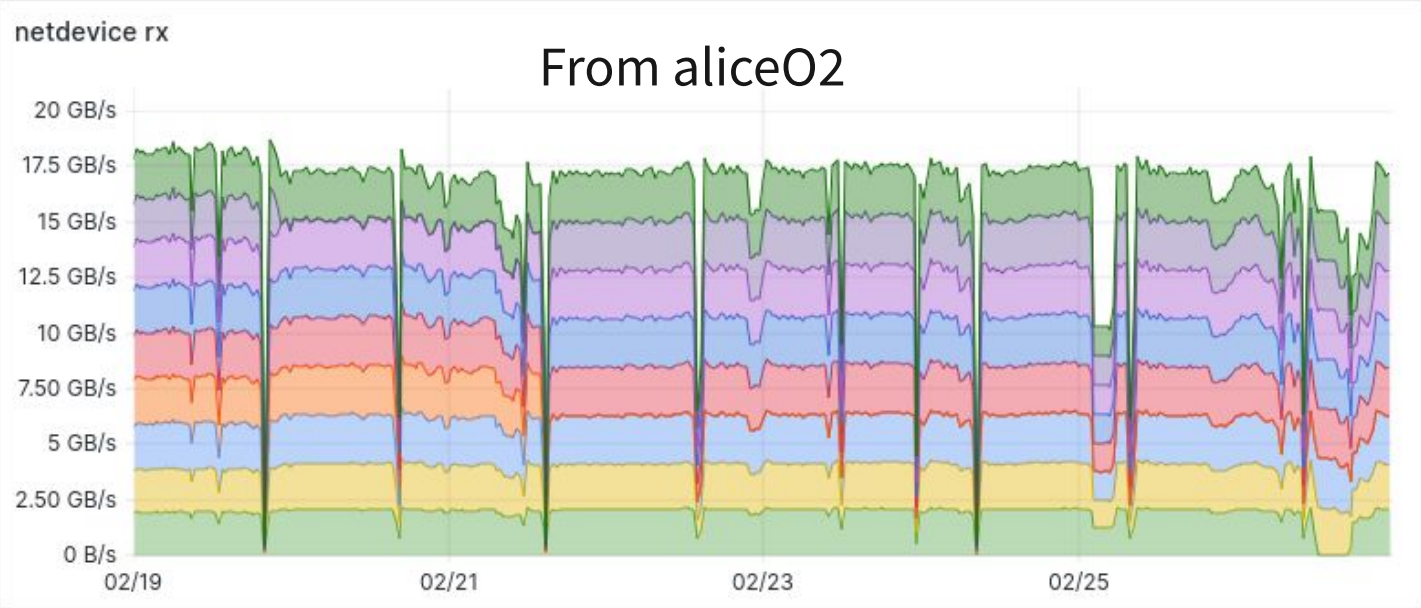

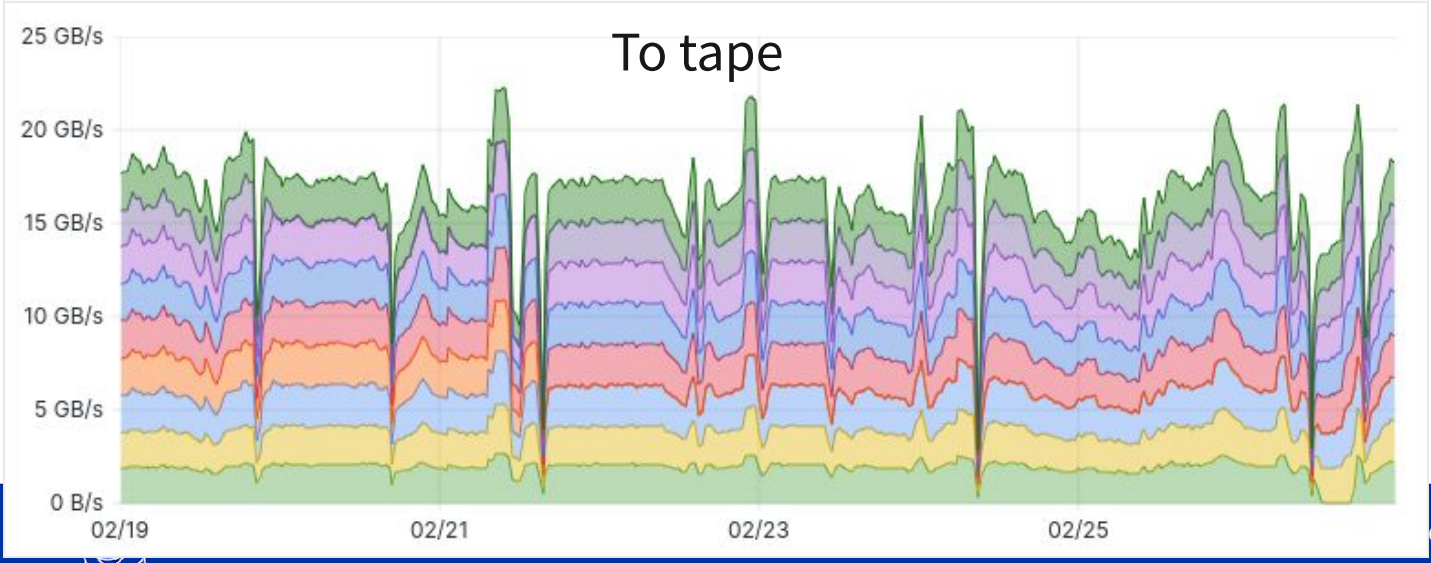

- Buffer servers >90% peak network throughput
- Tape drives > 85% peak write throughput

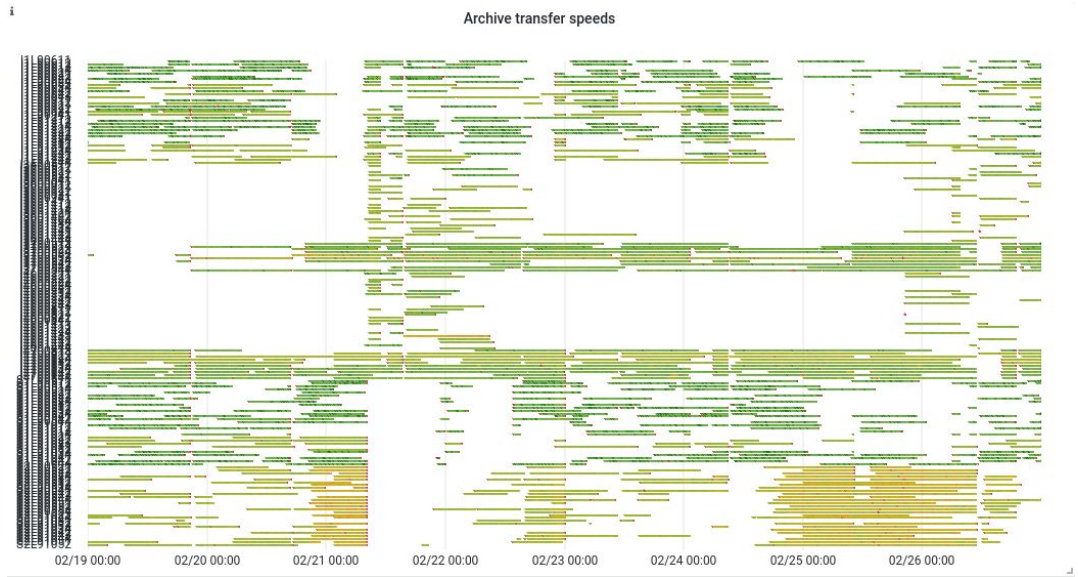

Know your weaknesses

When you find a problem:

- **Understand it**
  - Document and fix it!
- **Live with it**
  - Make sure you systematically spot it
    - **Create an alert triggered every time it happens**
  - Document procedures to go back to nominal performance

For example: Objectstore queuing performance during stress test

- **Always the same shape, we won't improve it**
  - Target performance: *IT SHOULD NOT GET WORSE!*
  - 9h to archive and stage 2.5M files on our well defined stress test machine
    - Longer is not acceptable

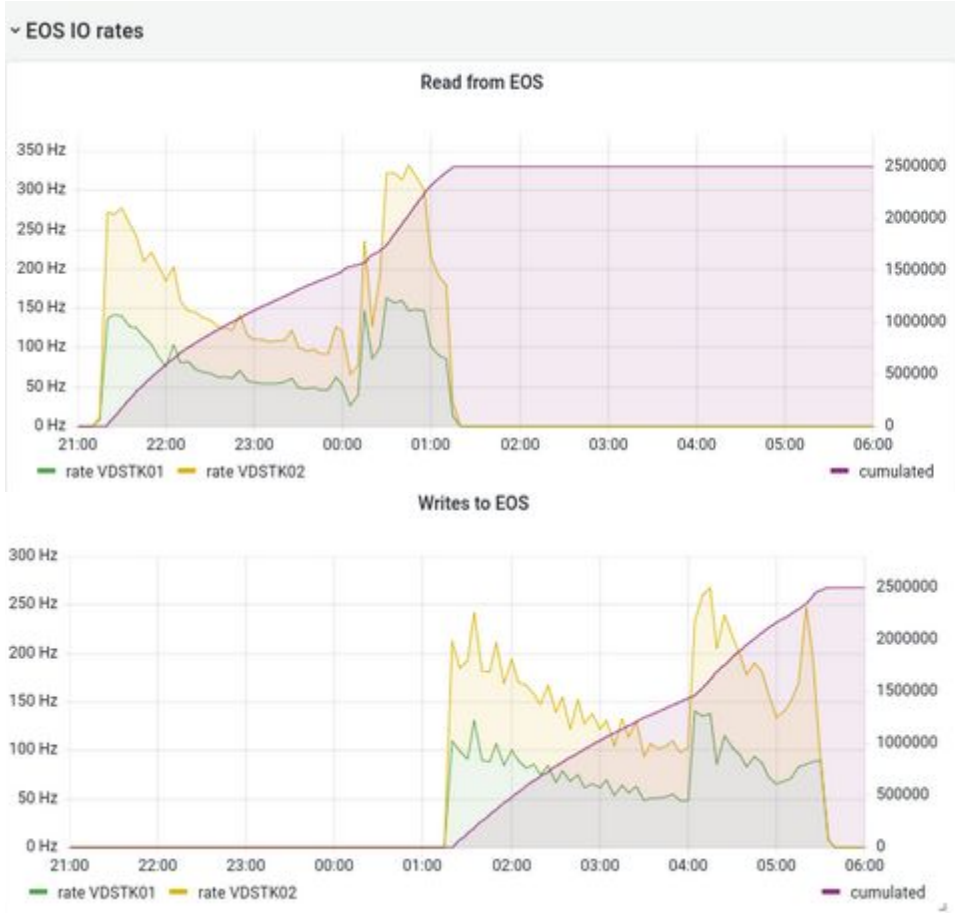

Know your weaknesses

Repack objectstore performance

- **Bad and just get worse over time**
  - Base measurement: state update time for UP drives
- **Contain with automation (ATRESYS)**
- **Separate backends user/repack**
  - Ongoing work
- **Document procedure to come back to nominal performance**
  - Dismount repack jobs
  - Delete ^Repack objects from objectstore
  - Recreate RepackIndex
    `cta-objectstore-create-missing-repack-index`
  - Clear repack requests and recreate some

RUN IT WHENEVER NEEDED BEFORE TRYING ANYTHING MORE EXPENSIVE ON THE OPERATIONS SIDE (+TRACK OCCURENCES?)

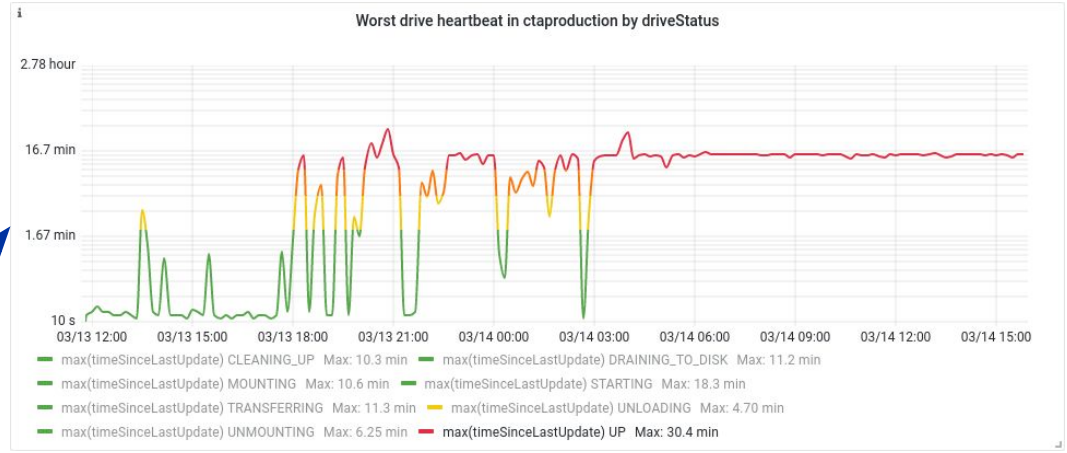

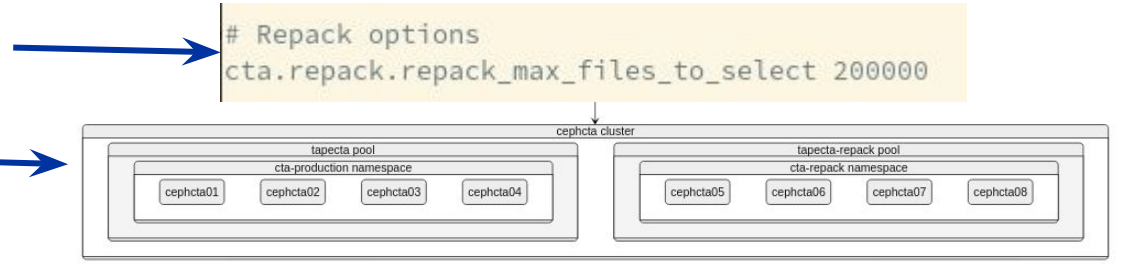

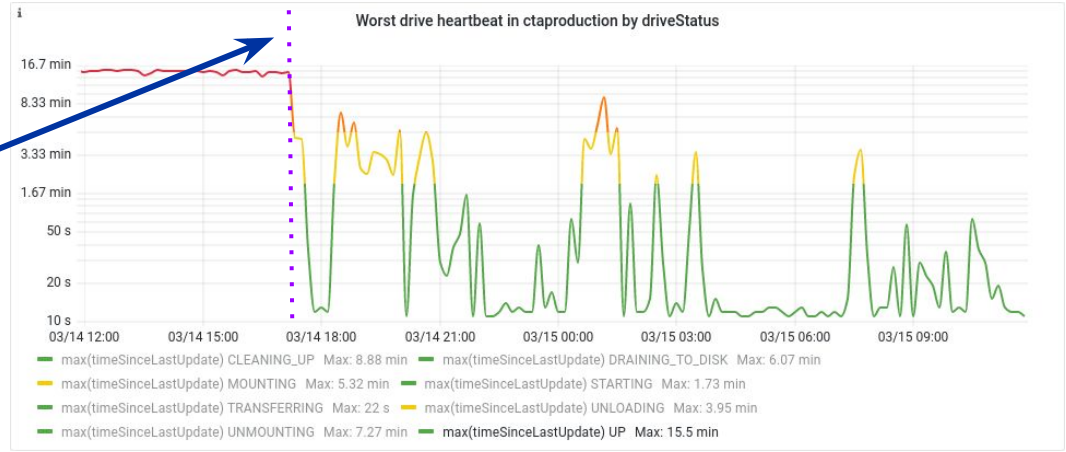

Meetings

Wednesday Operations Meeting tracks issues in the problem space (impact on operations, troubleshooting, monitoring, logs)

Dev issue needed

Prioritise each issue based on its
1. Likelihood of recurrence
2. Impact

priority high
workflow Assigned
Needs Discussion

Friday Developer Meeting tracks issues in the software development space (features needed, bugs to be fixed)

CTA Issues and software development

Conclusion

CTA meetings work really well for T0 needs, maybe not so great for community feedback…

- Interesting feedback from *Meet the team* session last week
- **What about a monthly *CTA community operations* meeting?**

Write more common recipes together

- What is *your* CTA base model?
  - How do you use your tape infrastructure
- How to work together on common operations tools?

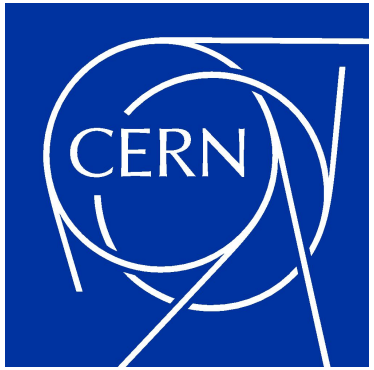

home.cern