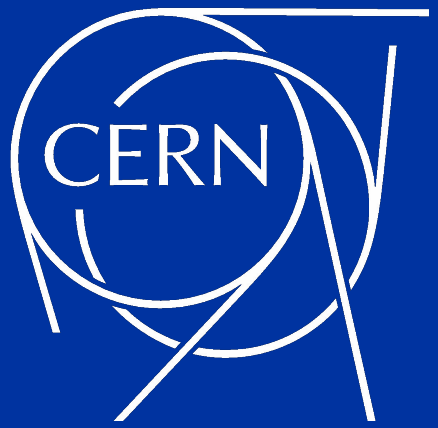

nTOF dataflow to tape review

Julien Leduc
on behalf of IT-SD-TAB

2024-08-22

EOS+CTA Architecture *spinners* Addon

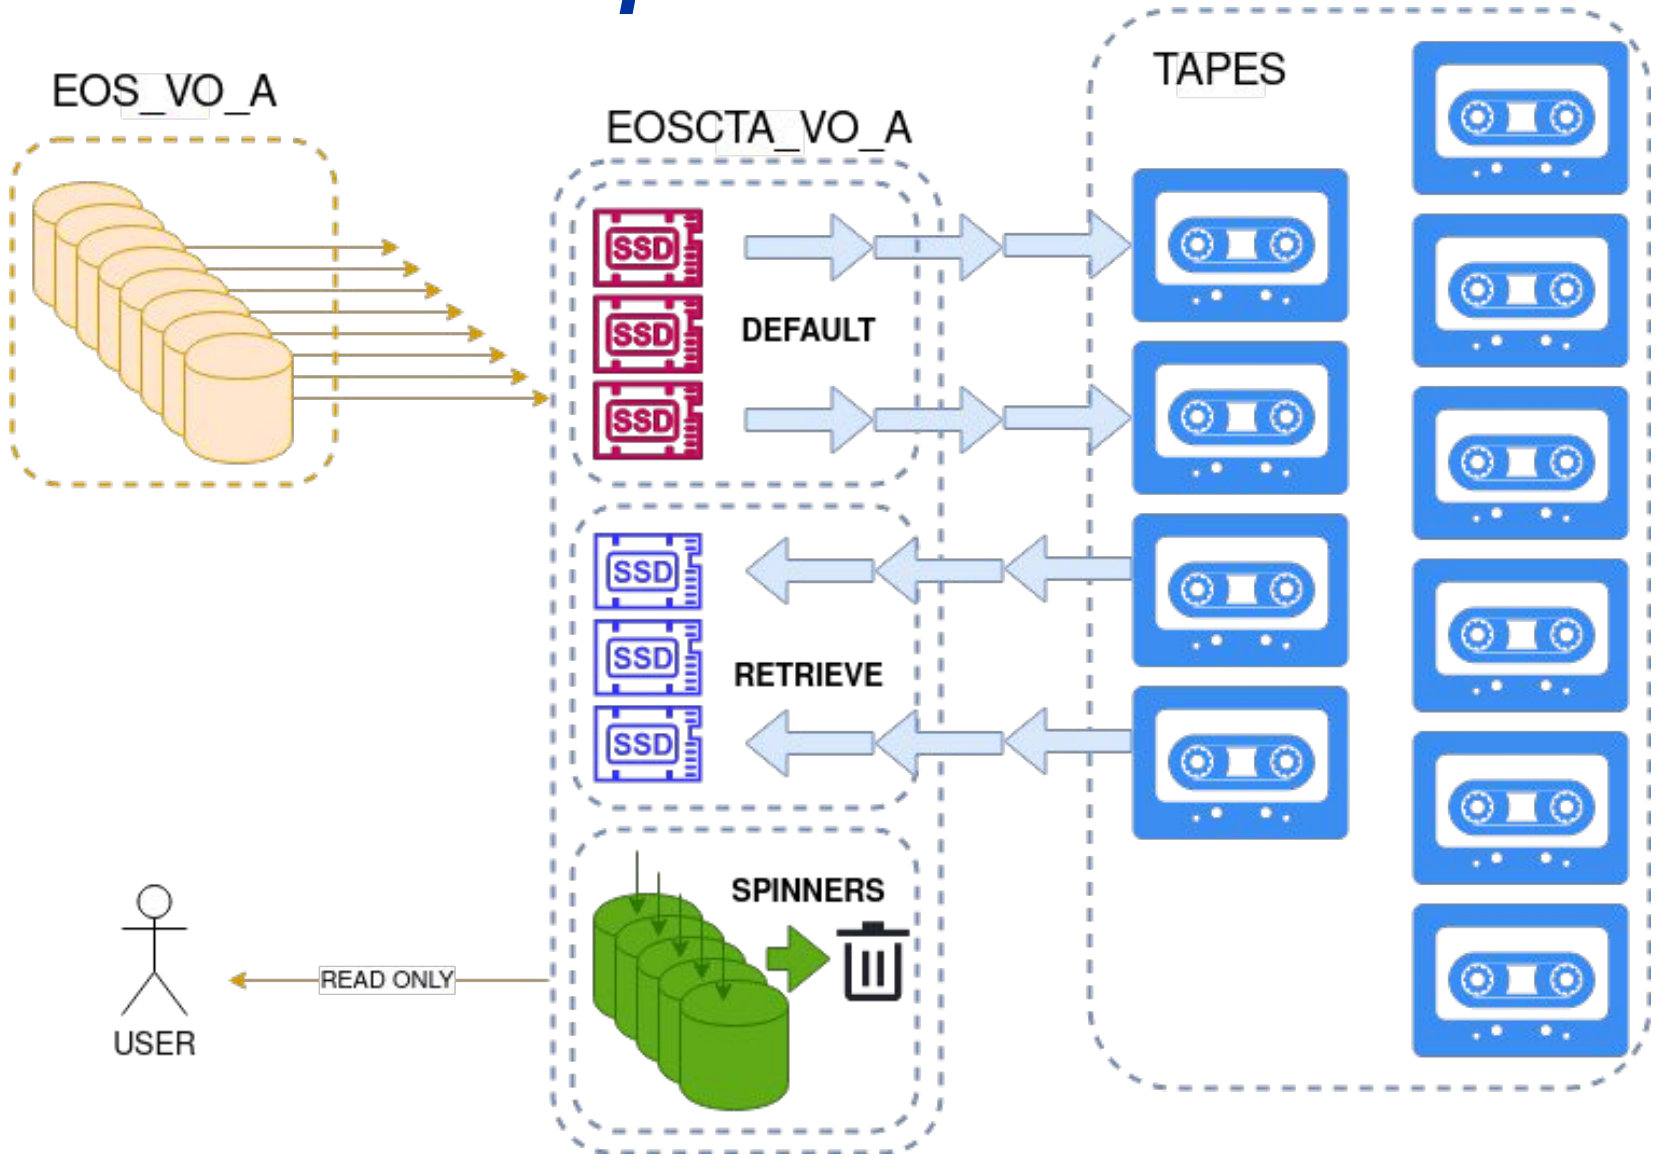

EOS+CTA Architecture *spinners* Addon

nTOF writes to eosctapublic *spinners* space

eosctapublicdisk:

- 2022: 200 disks
  - 2024: 450 disks
- 2022: 6 disk servers
  - 2024: 10 disk servers

eospublic:

- 9000 disks
- 142 disk servers

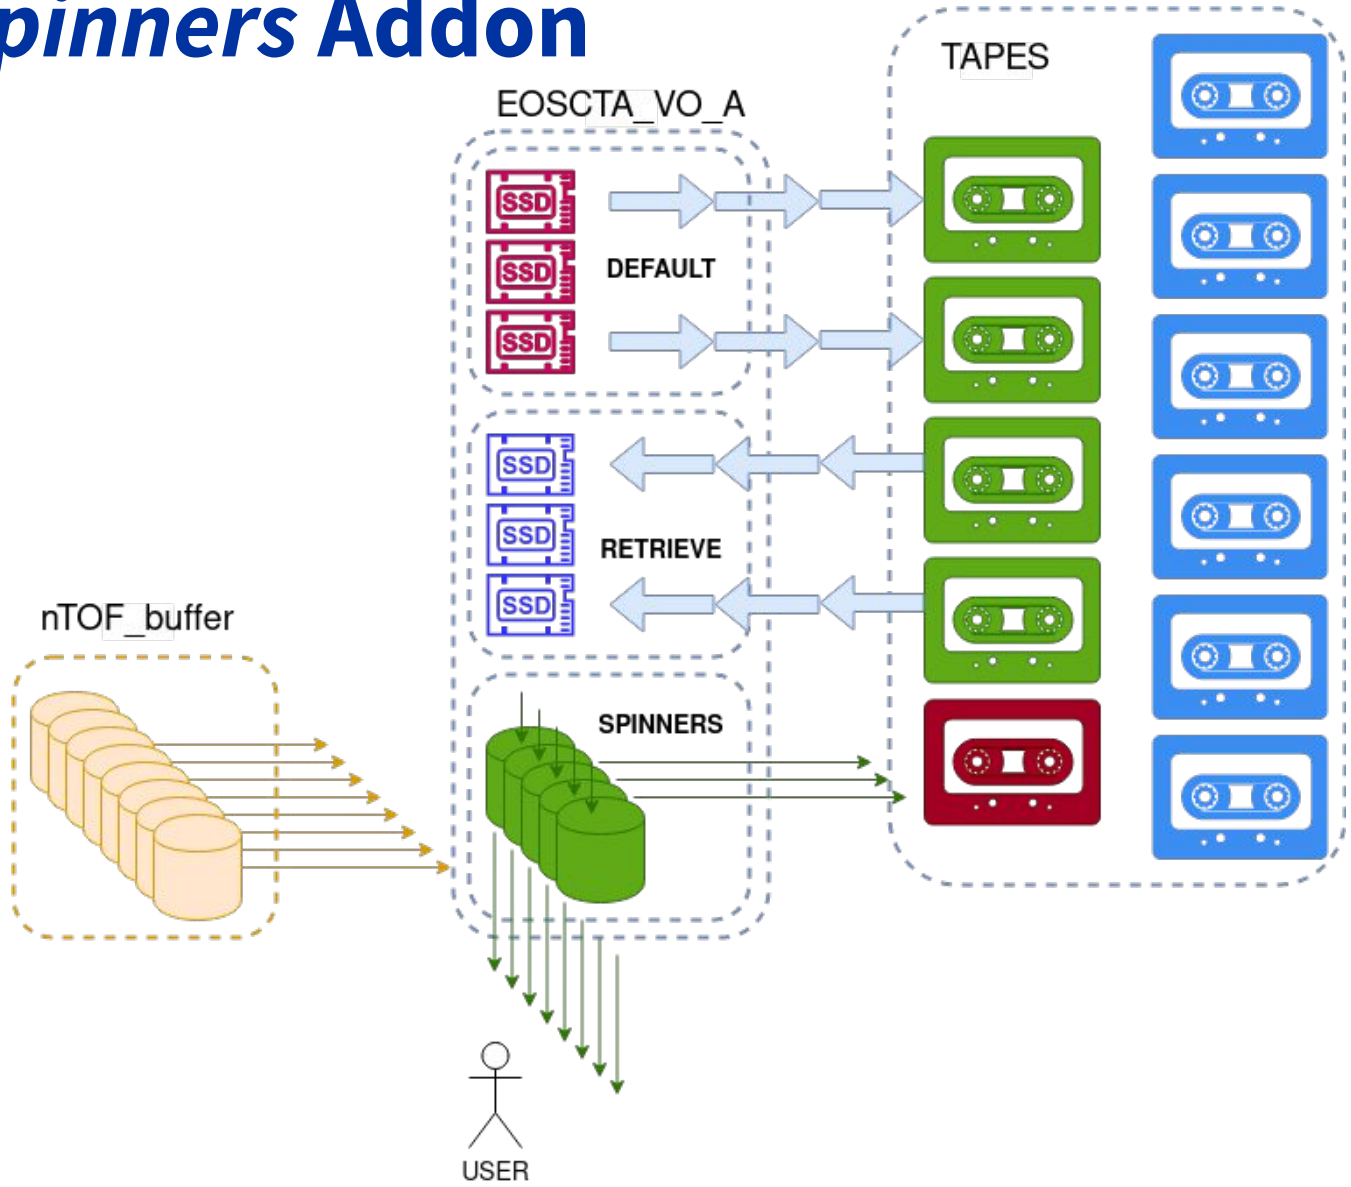

Issues with nTOF dataflow

- **nTOF is the only experiment writing to spinners in CTA**
  - using CTA exactly like CASTOR
    - write a file to spinners space explicitly
    - issue 2 staging request per file to pin it on disk so that it can be read in place
      - **CTA is not supposed to respect staging requests on archiving file**
      - **very likely to break as we are reworking tape eviction part**
    - Files that did not go to tape are copied to <filename.backup> and are stuck in spinners space forever
      - **spinners is supposed to only hold files read from tape with a tape copy**
      - these stuck files are complicating spinners operations for CTA team: disk replacements, OS upgrades…
  - spinners disks are almost full (with LRU garbage collection)
    - as files are only supposed to be written from tape we keep all disks almost full
      - nTOF concurrent writes to spinners conflicts with CTA retrieve backpressure logic complicating CTA operations
      - as disks are almost full nTOF writes are more likely to fail (INC3462454) complicating further CTA operations…
  - **Concurrency usage of spinners kills nTOF performance**
    - luckily everyone else is currently taking data and writing to CTA default space as designed…
      - but if some are finished nTOF will again face performance degradation

2022 nTOF archive throughput to tapes

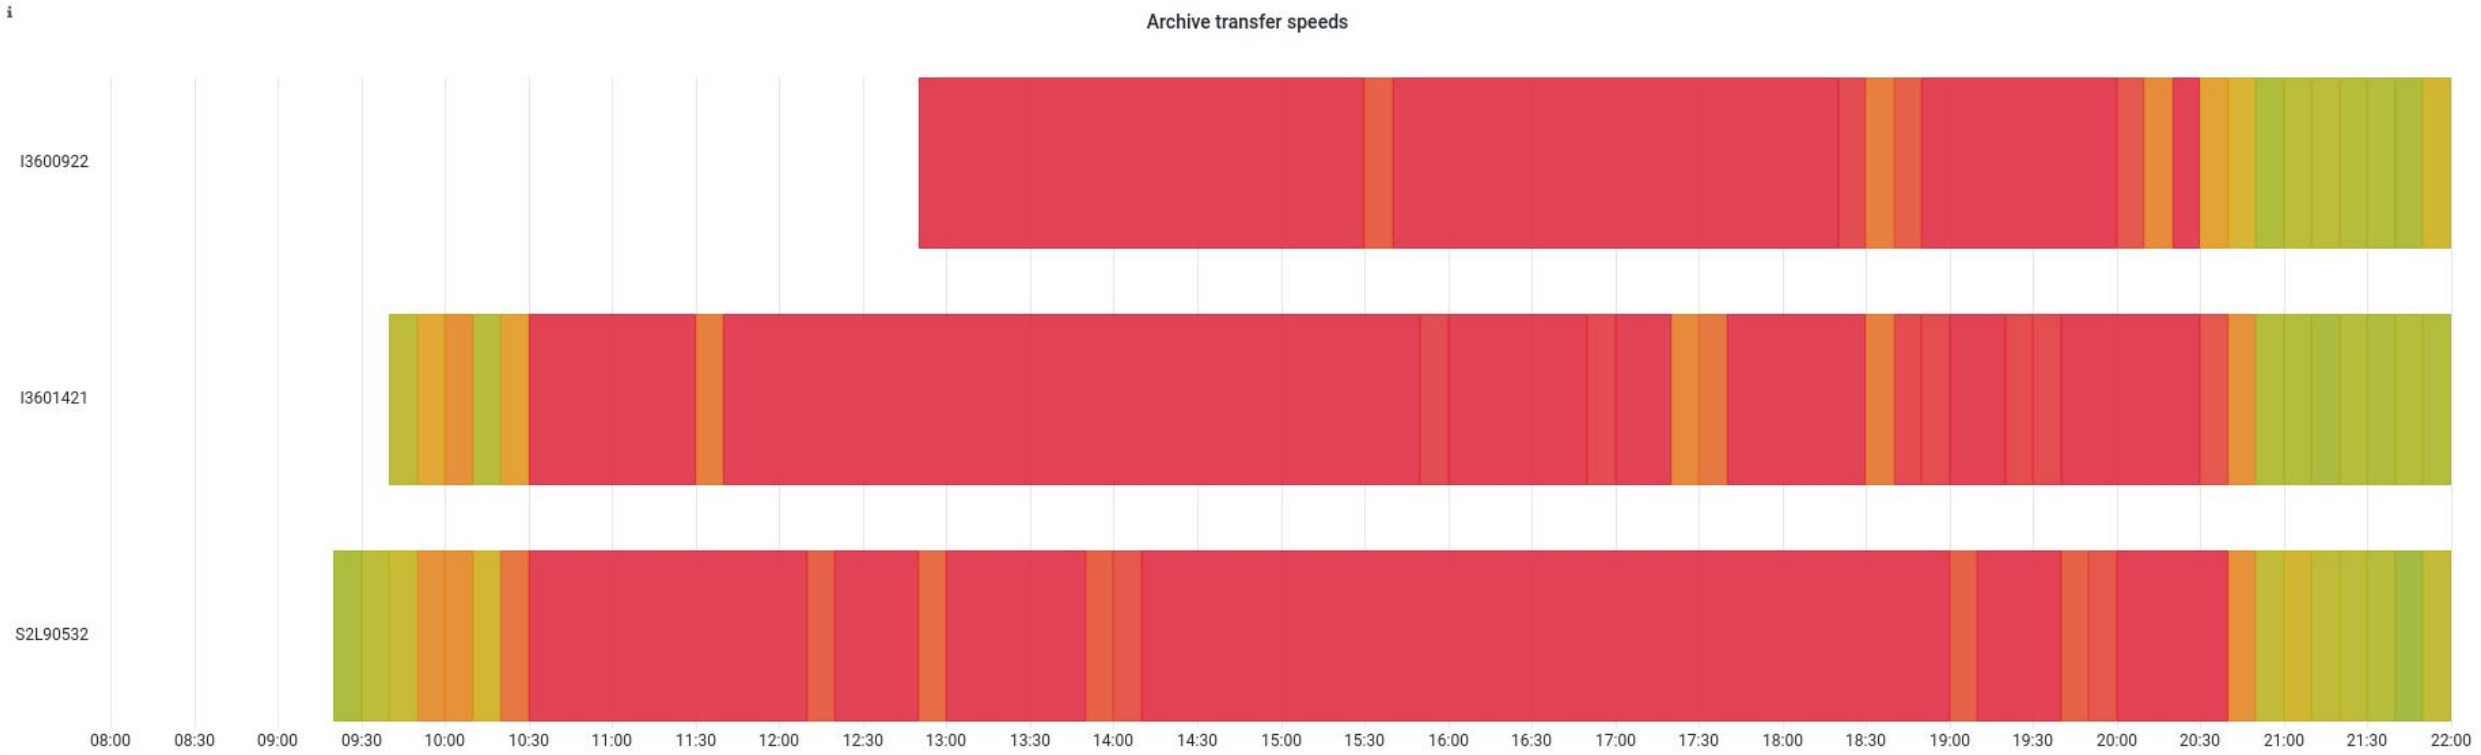

2024 nTOF archive throughput to tapes

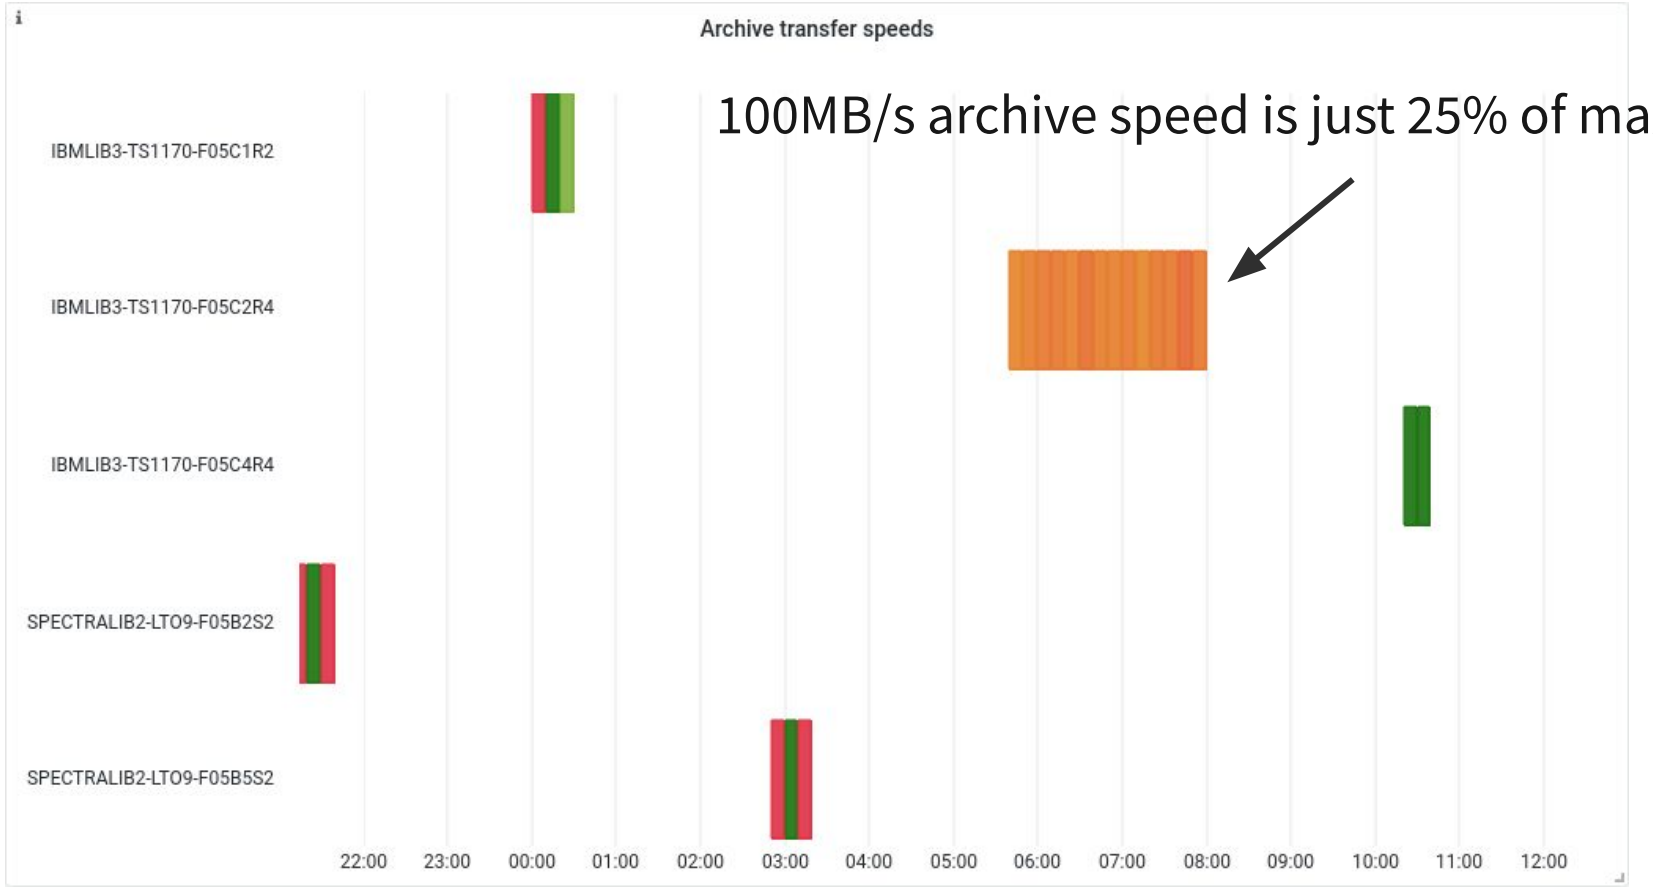

100MB/s archive speed is just 25% of max drive speed

For comparison LHC average archive speed per drive is over 350MB/s archiving to 45 drives at the same time

2022 nTOF Archive throughput VS user activity

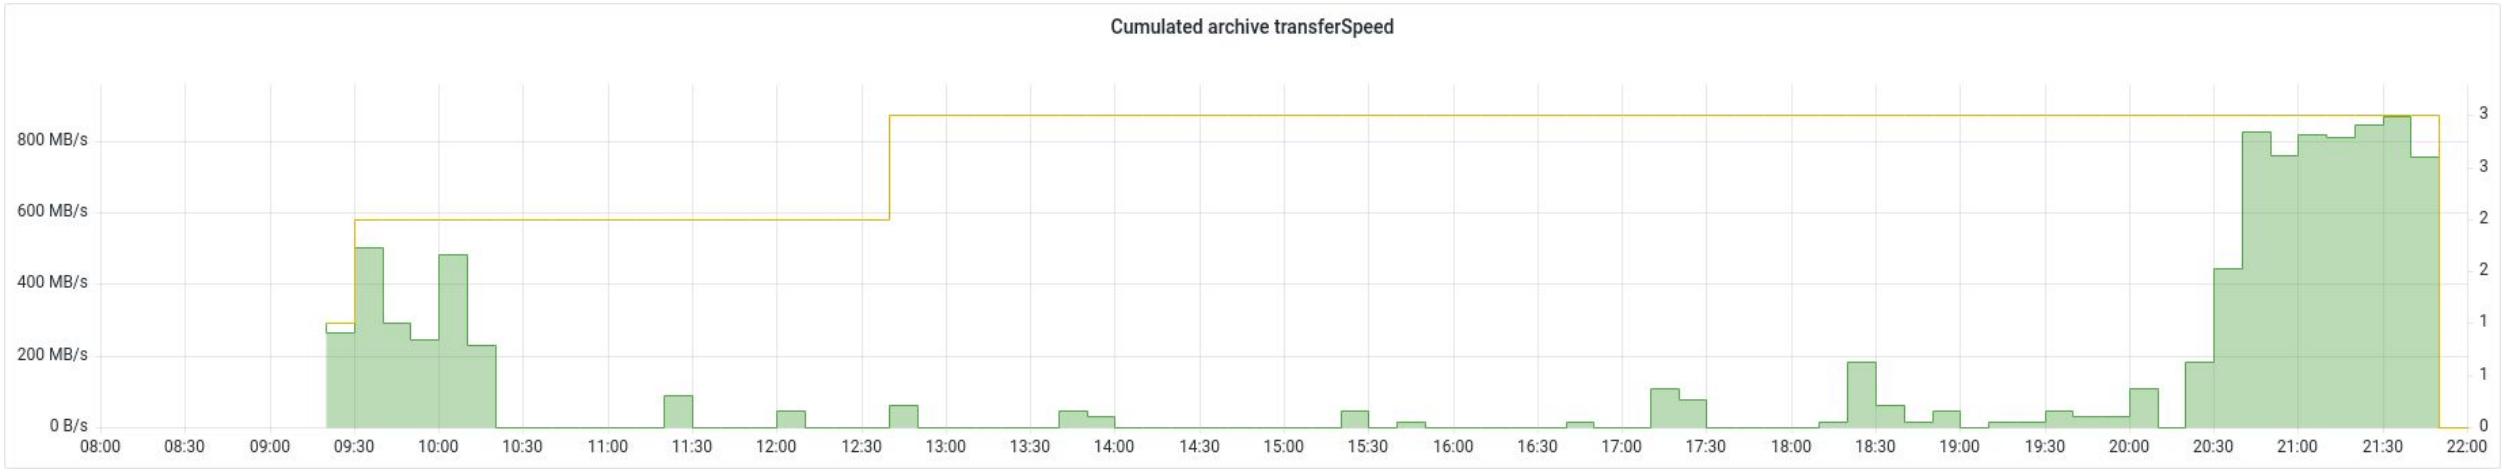

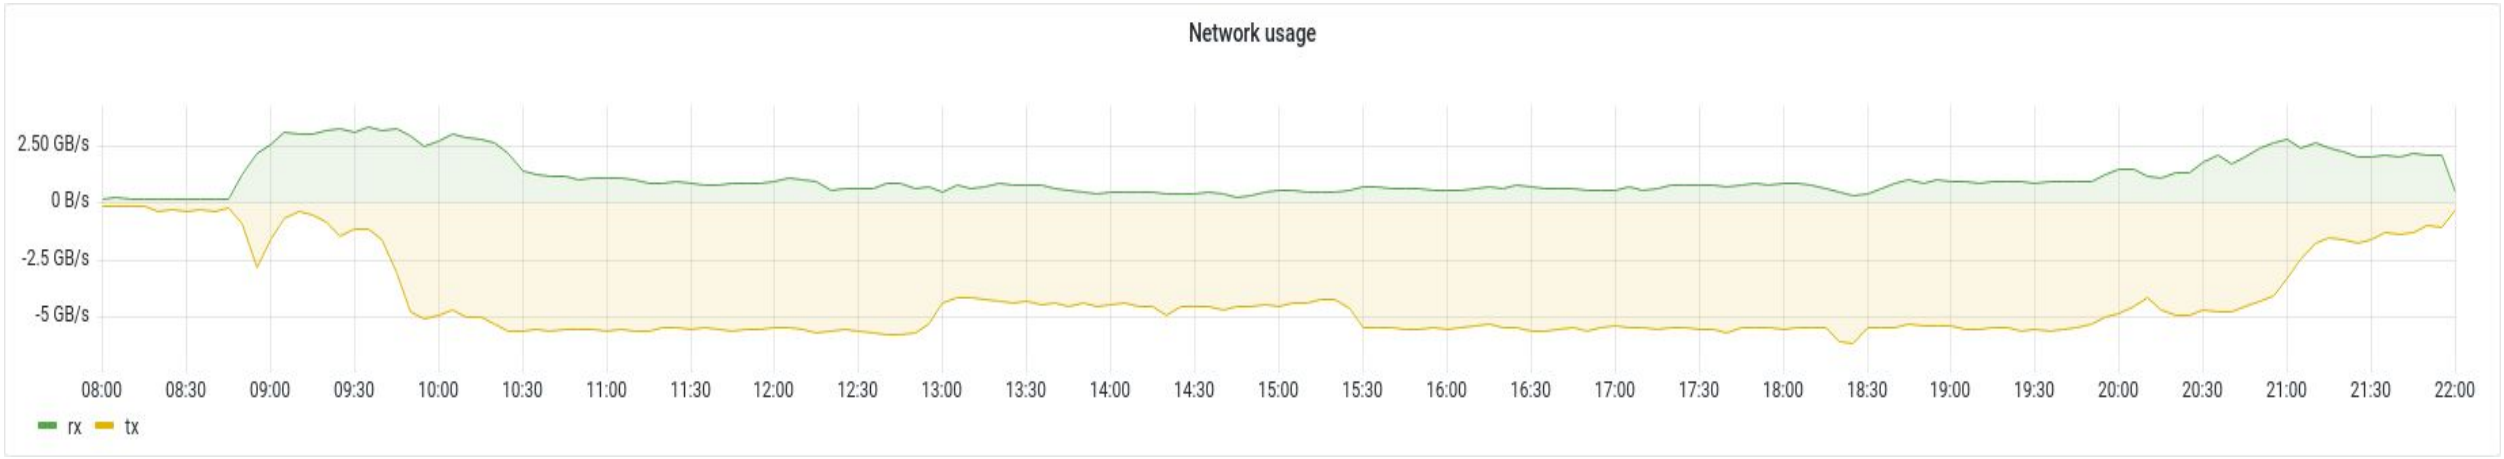

Conclusion

This must be the last nTOF data taking to spinners

- **When is the current data taking over?**
- **When can tests of the new workflow start?**
  - Contact?
- **Agreed hard deadline for deprecation of *the spinners dataflow*?**