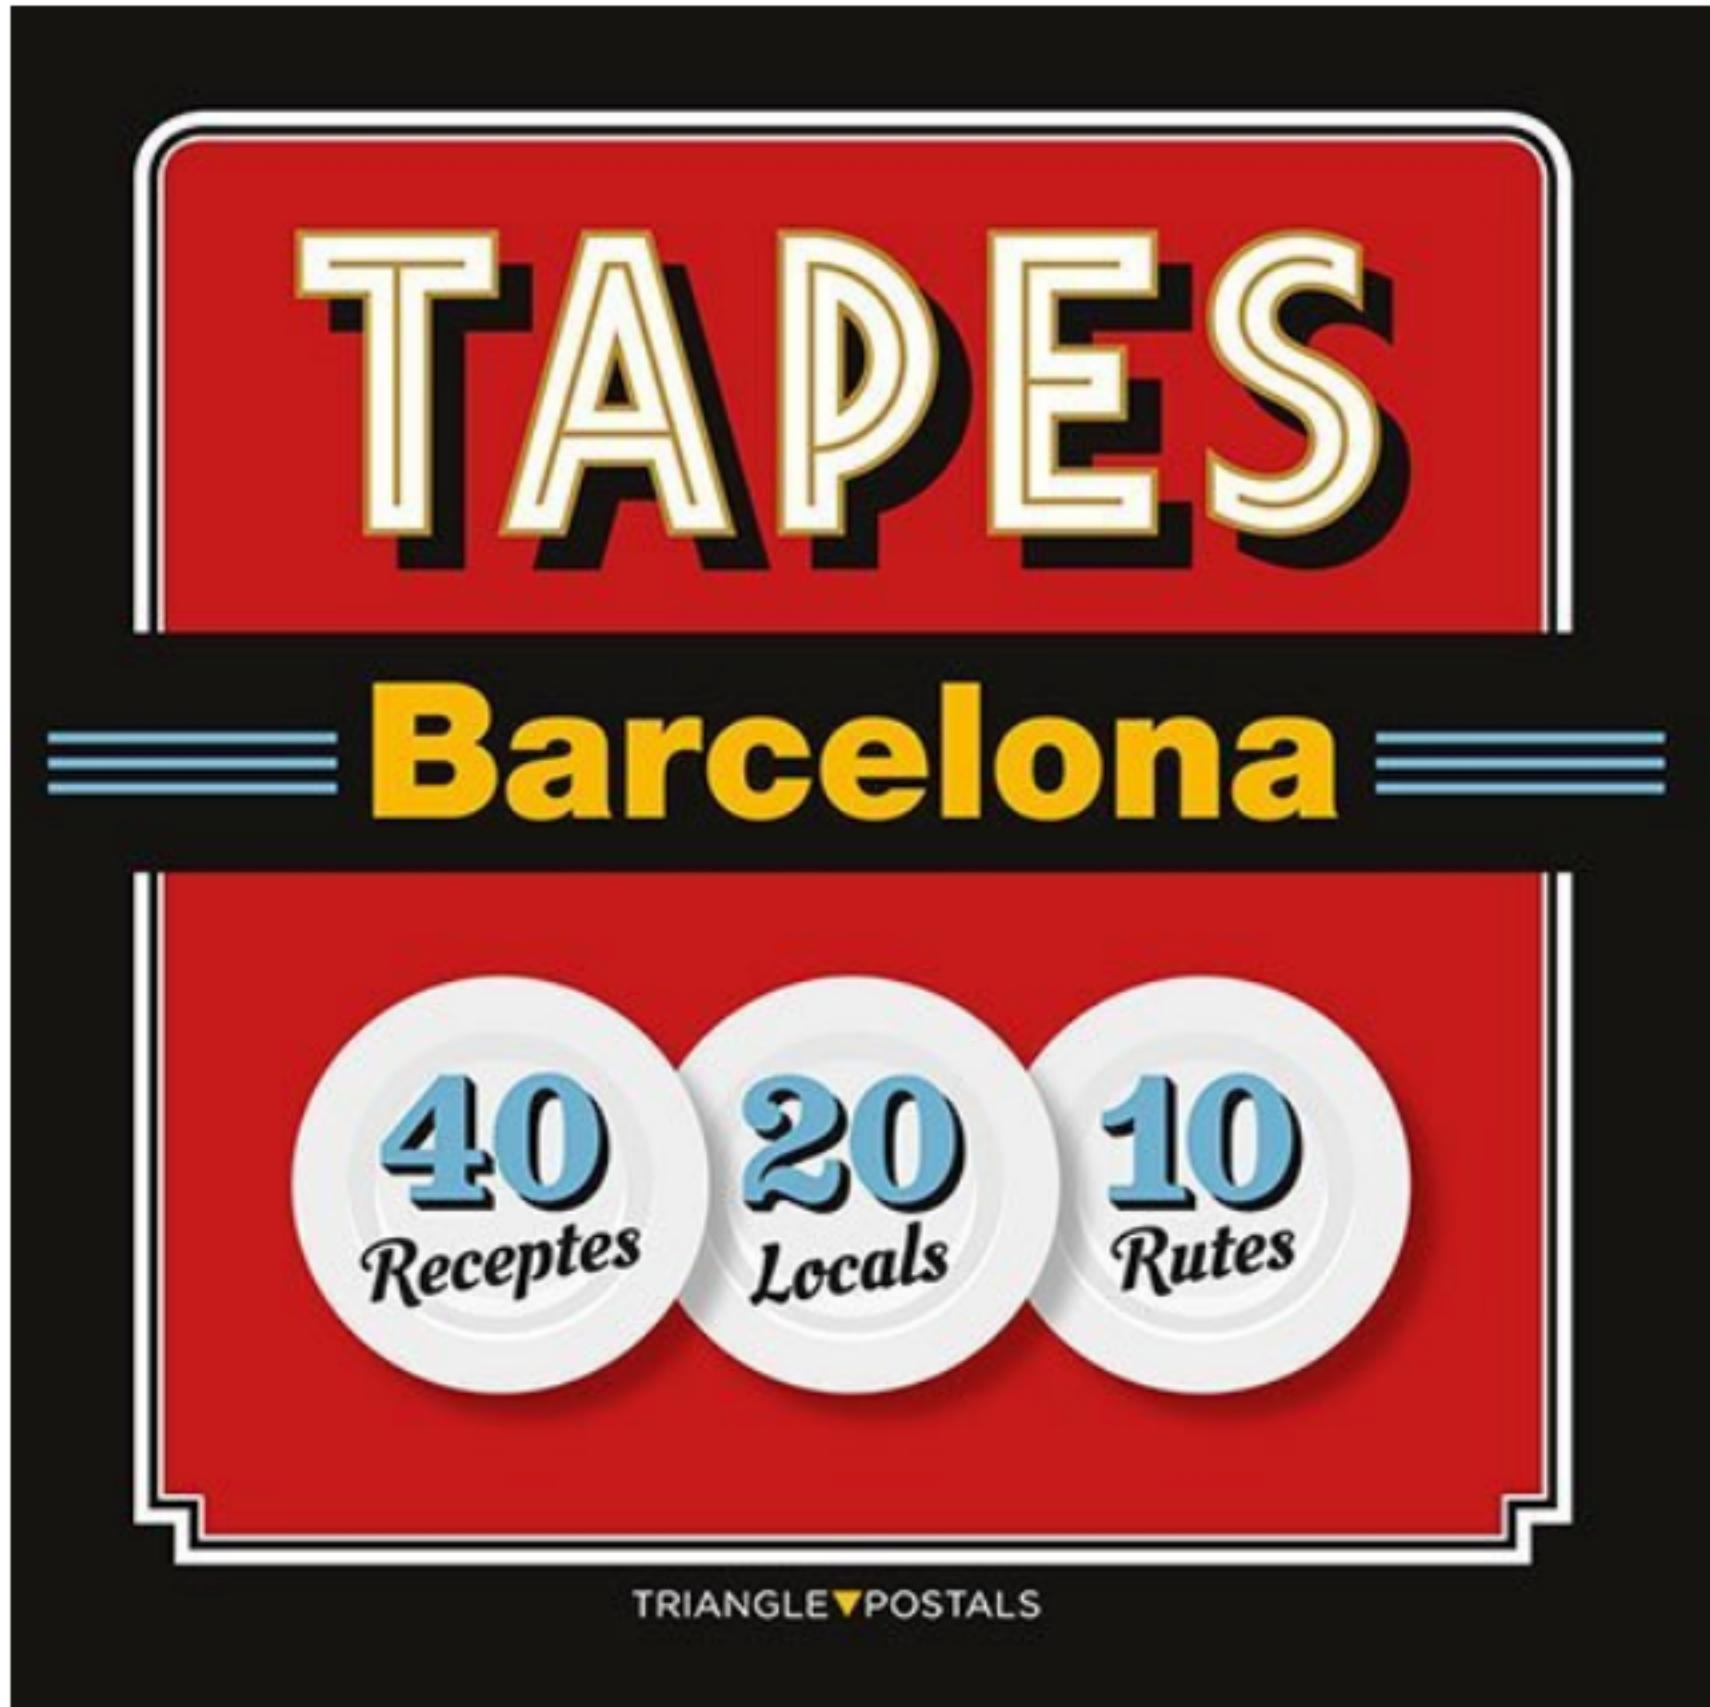




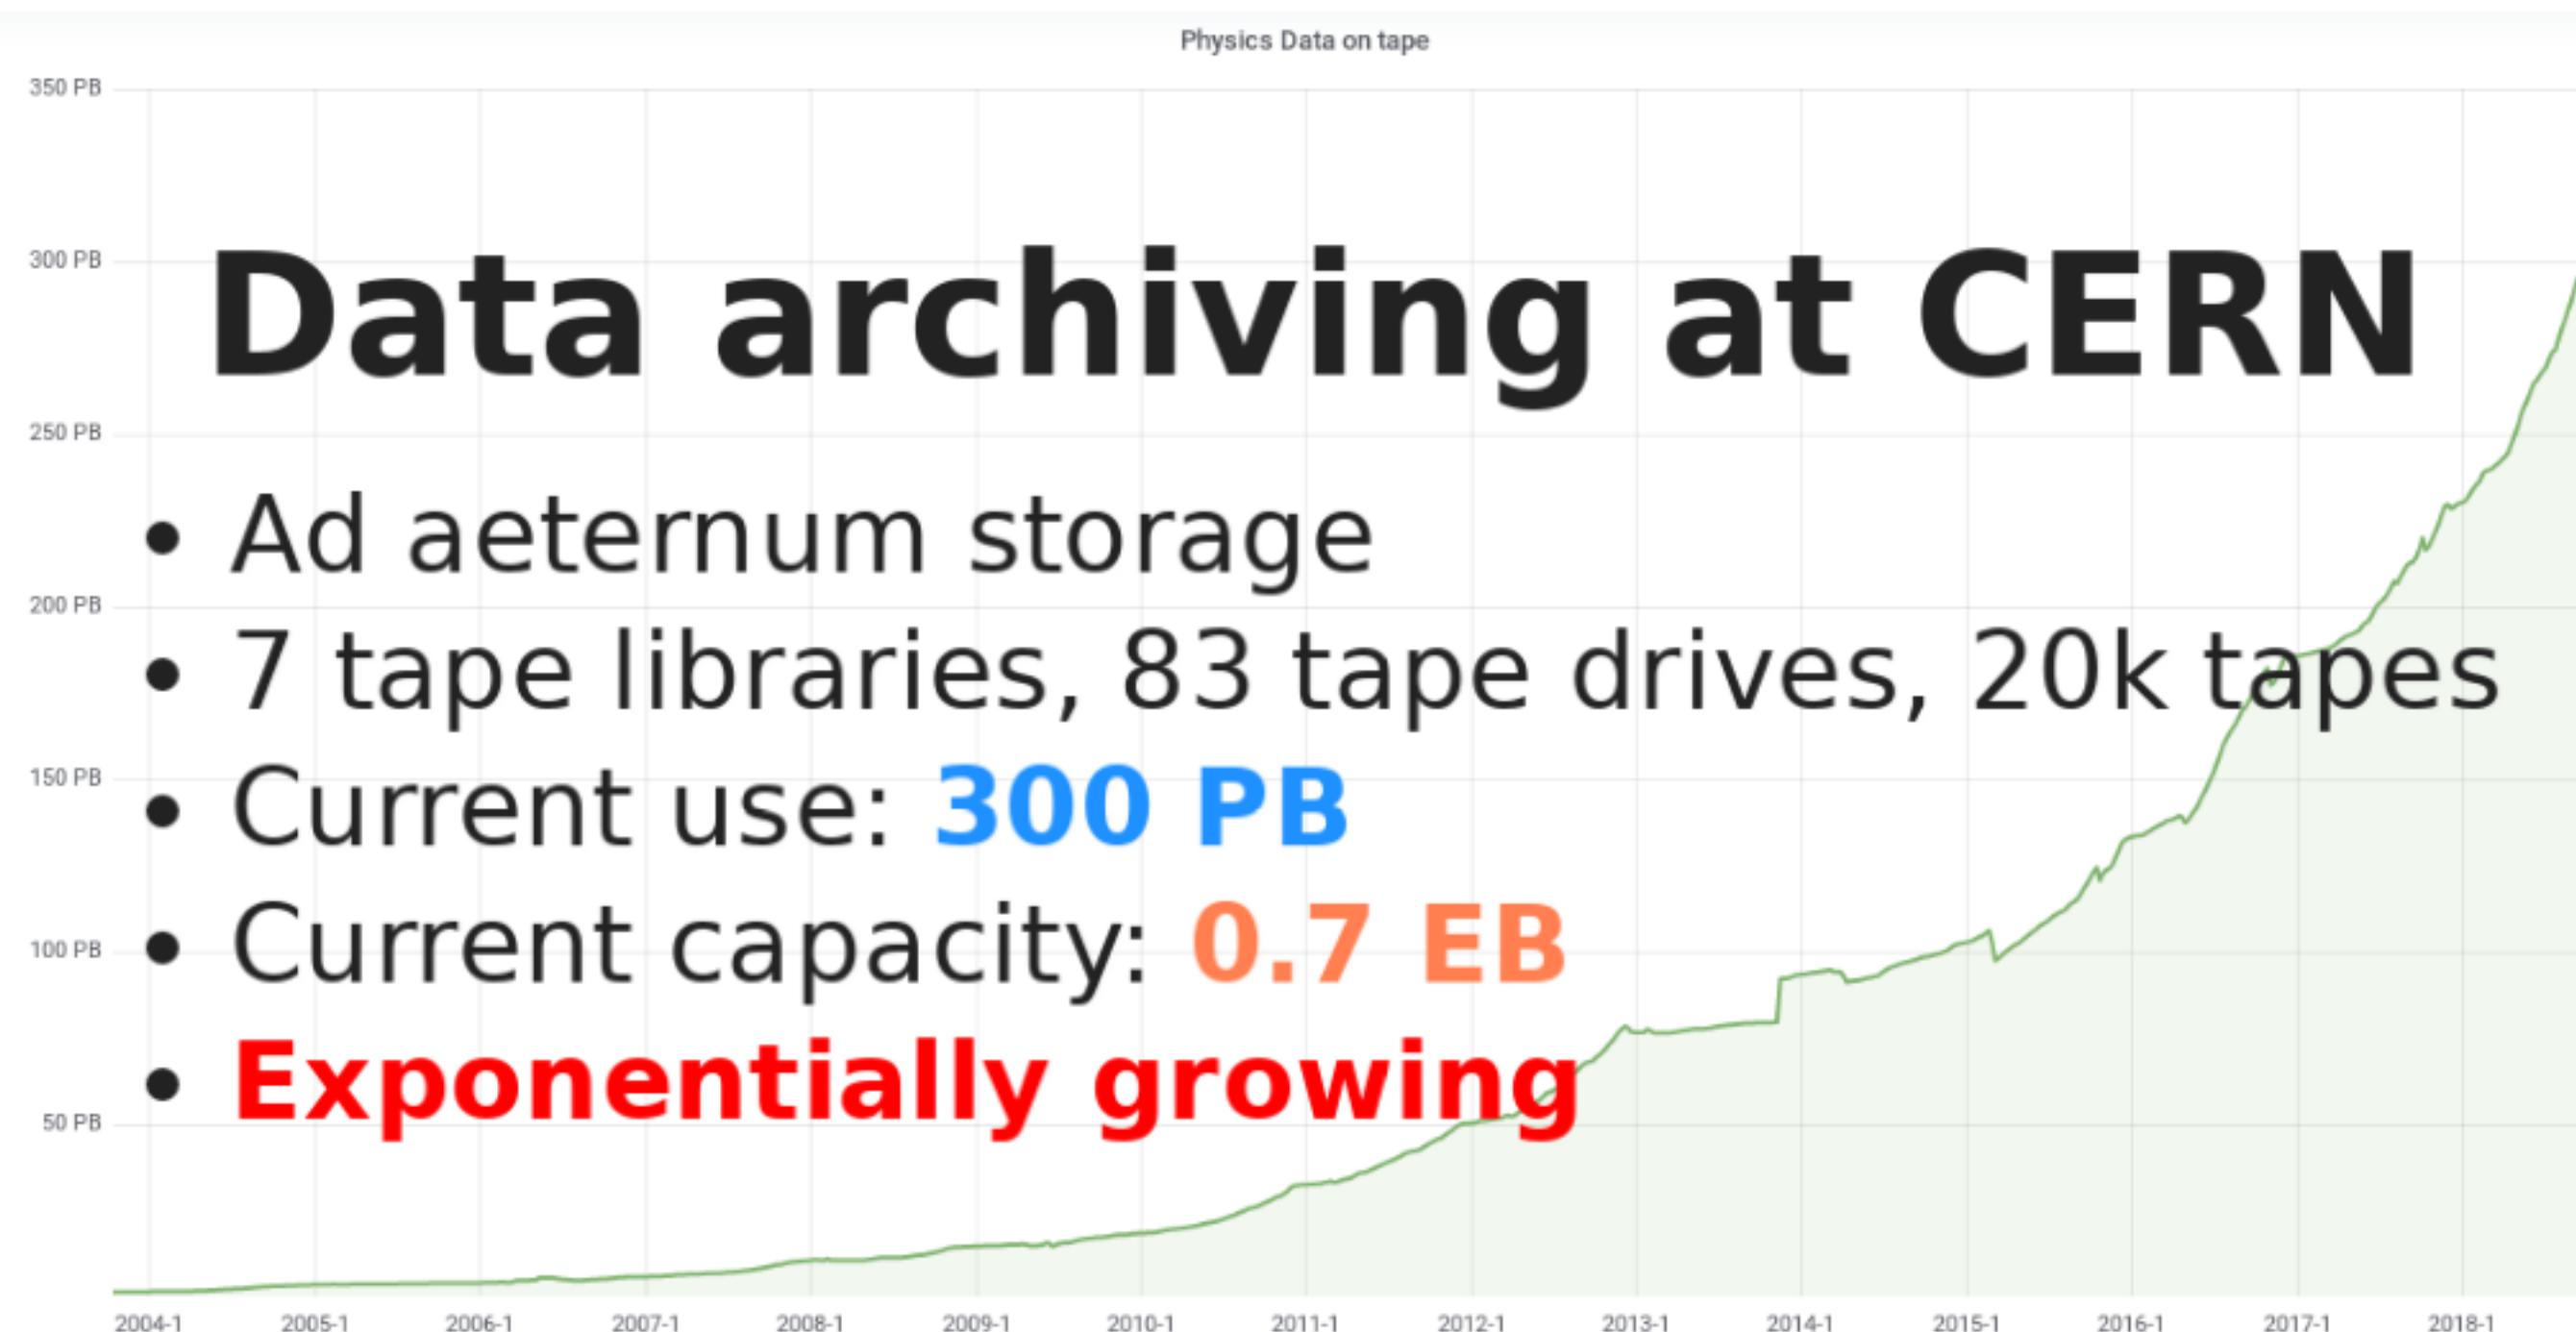
# initial deployment

Julien Leduc



# Data archiving at CERN

- Ad aeternum storage
- 7 tape libraries, 83 tape drives, 20k tapes
- Current use: **300 PB**
- Current capacity: **0.7 EB**
- **Exponentially growing**

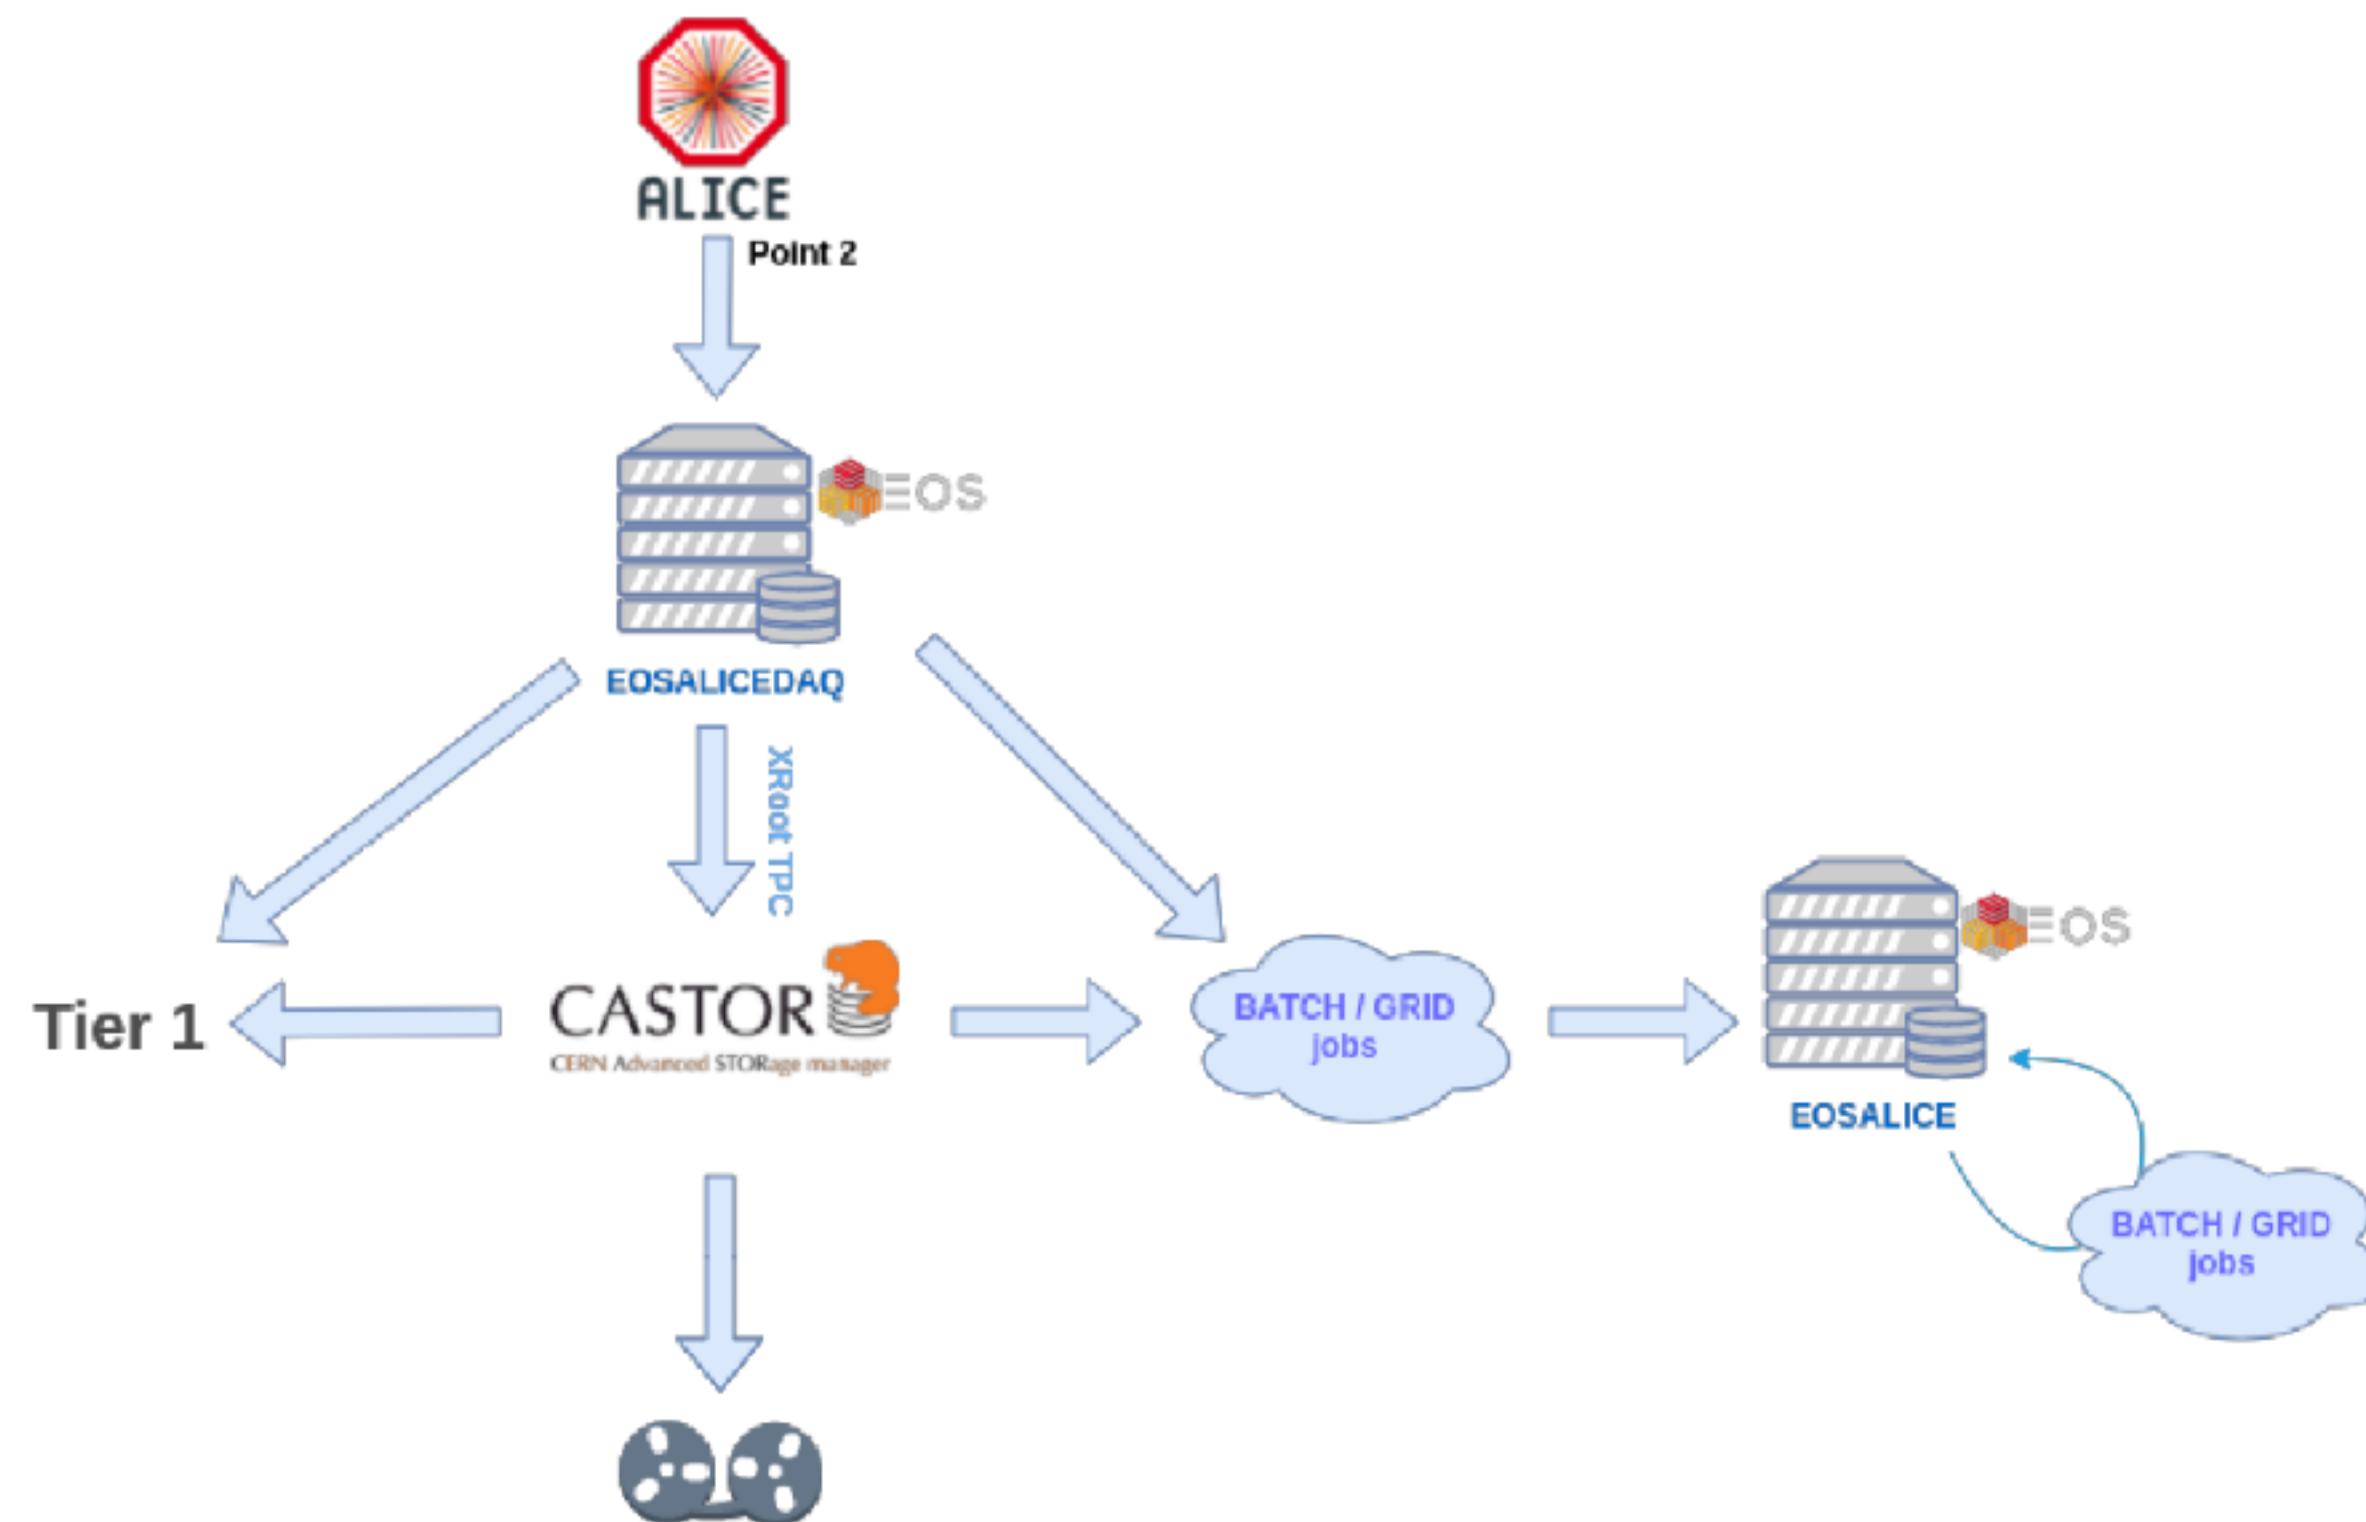


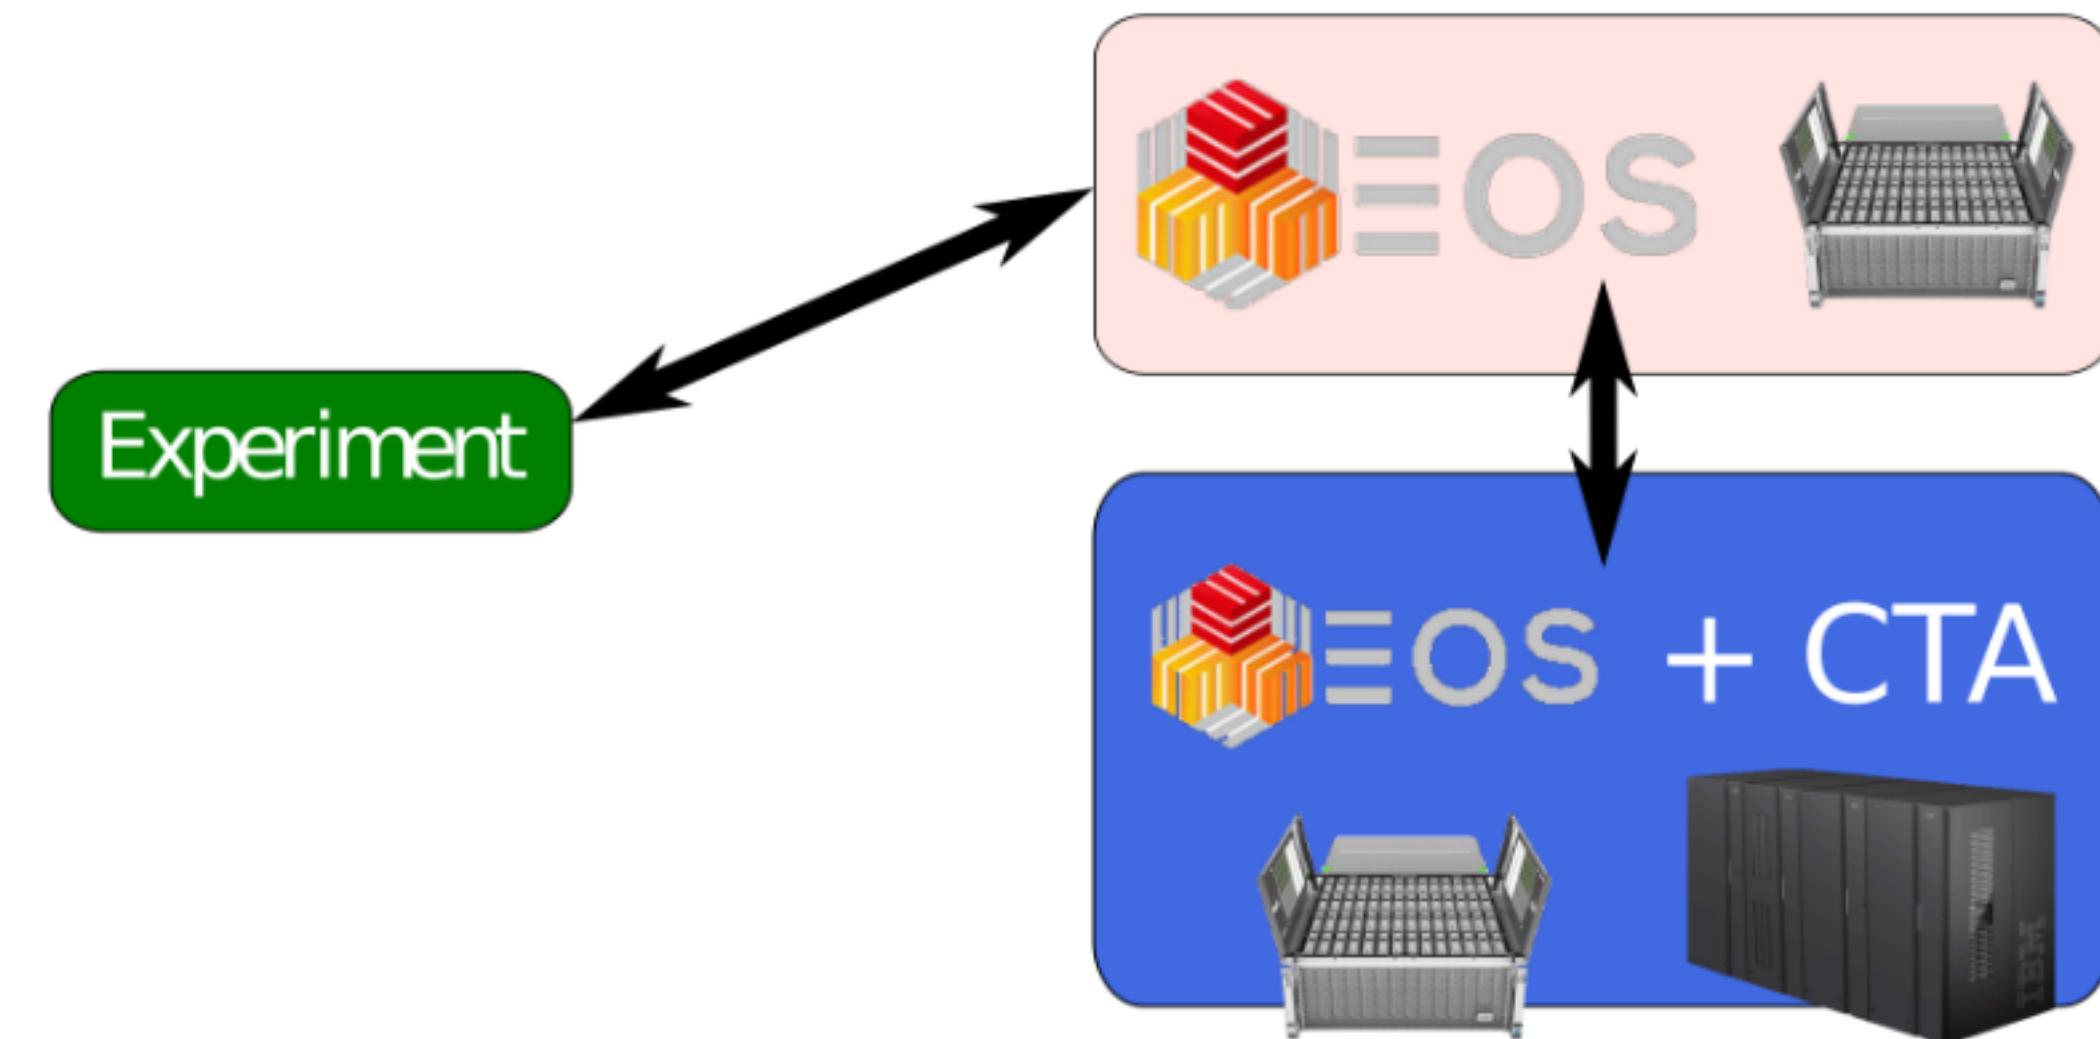
# Data Archiving at CERN

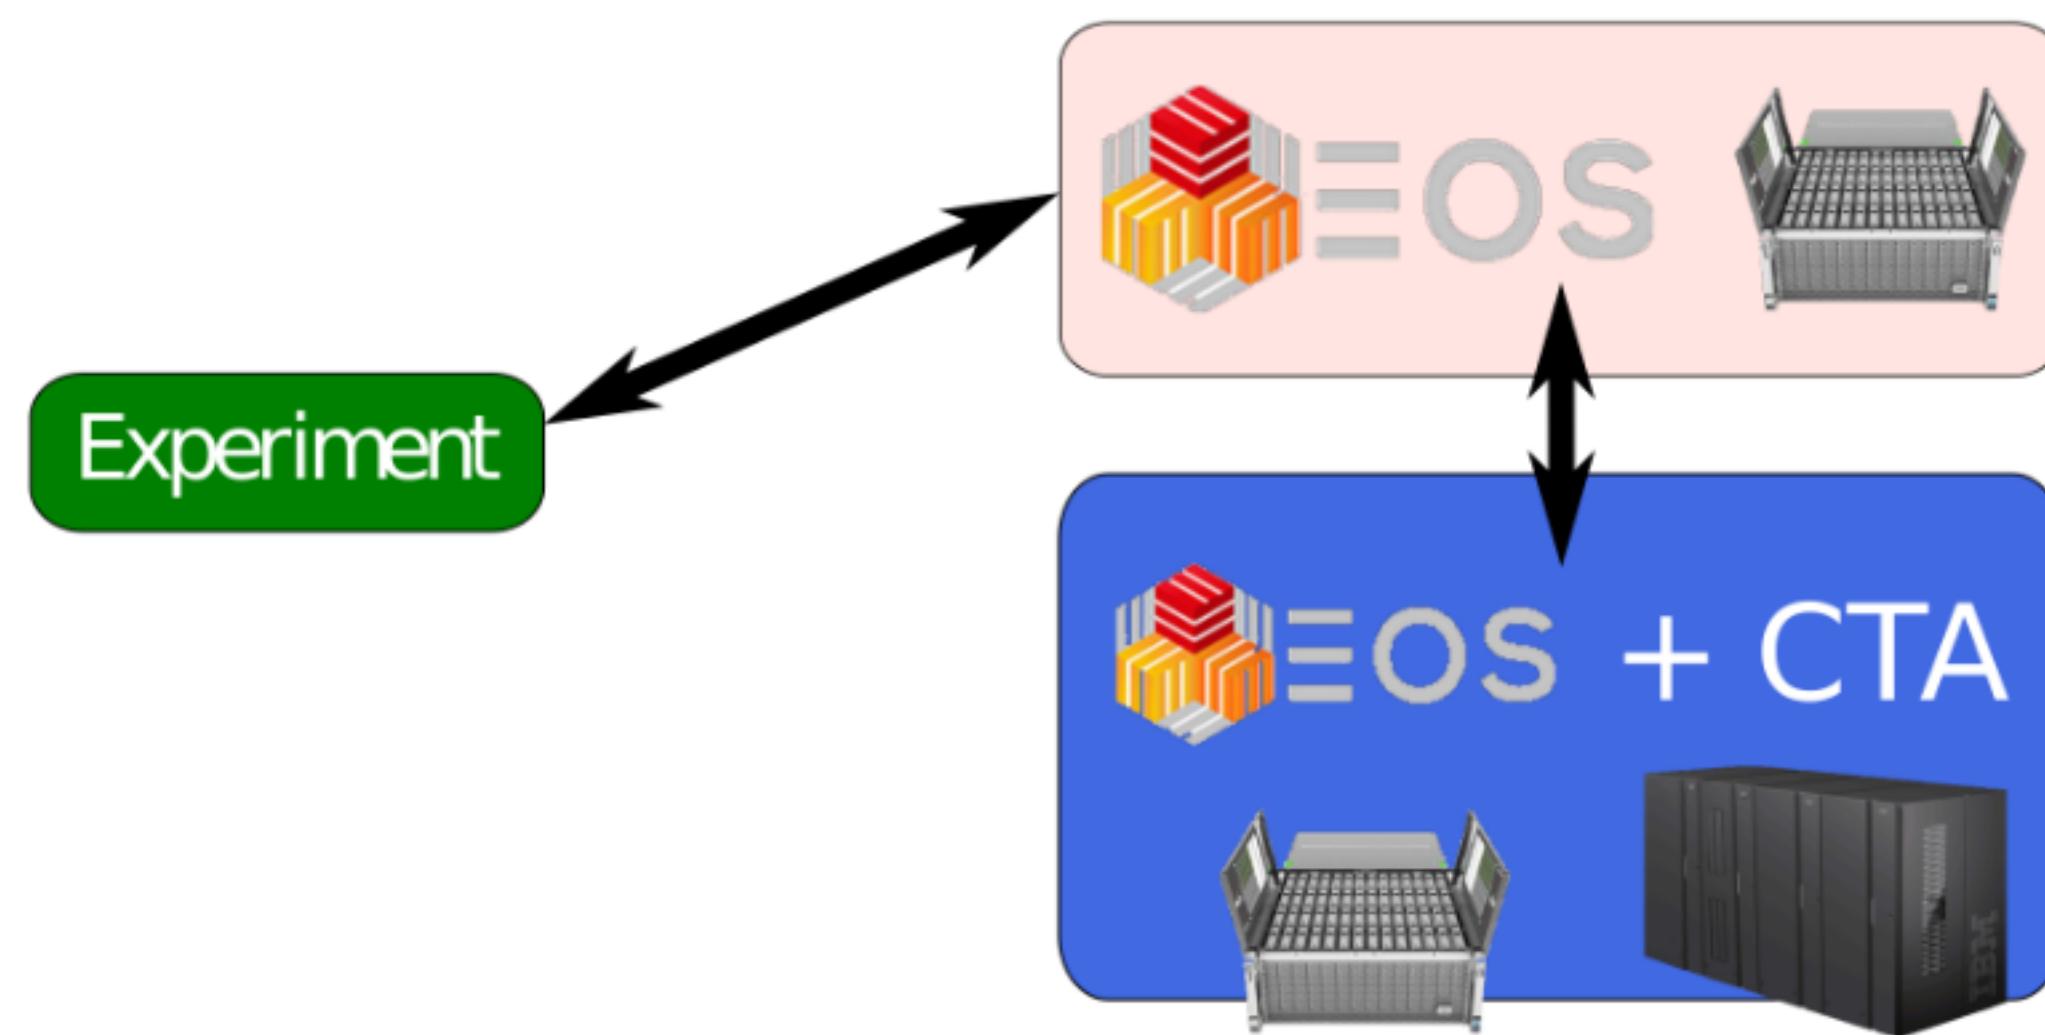
## *Evolution*

- EOS + tapes...
  - EOS is CERN strategic storage platform
  - tape is the strategic long term archive medium
- EOS + tapes = ❤
  - Meet CTA: CERN Tape Archive
  - Streamline data paths, software and infrastructure

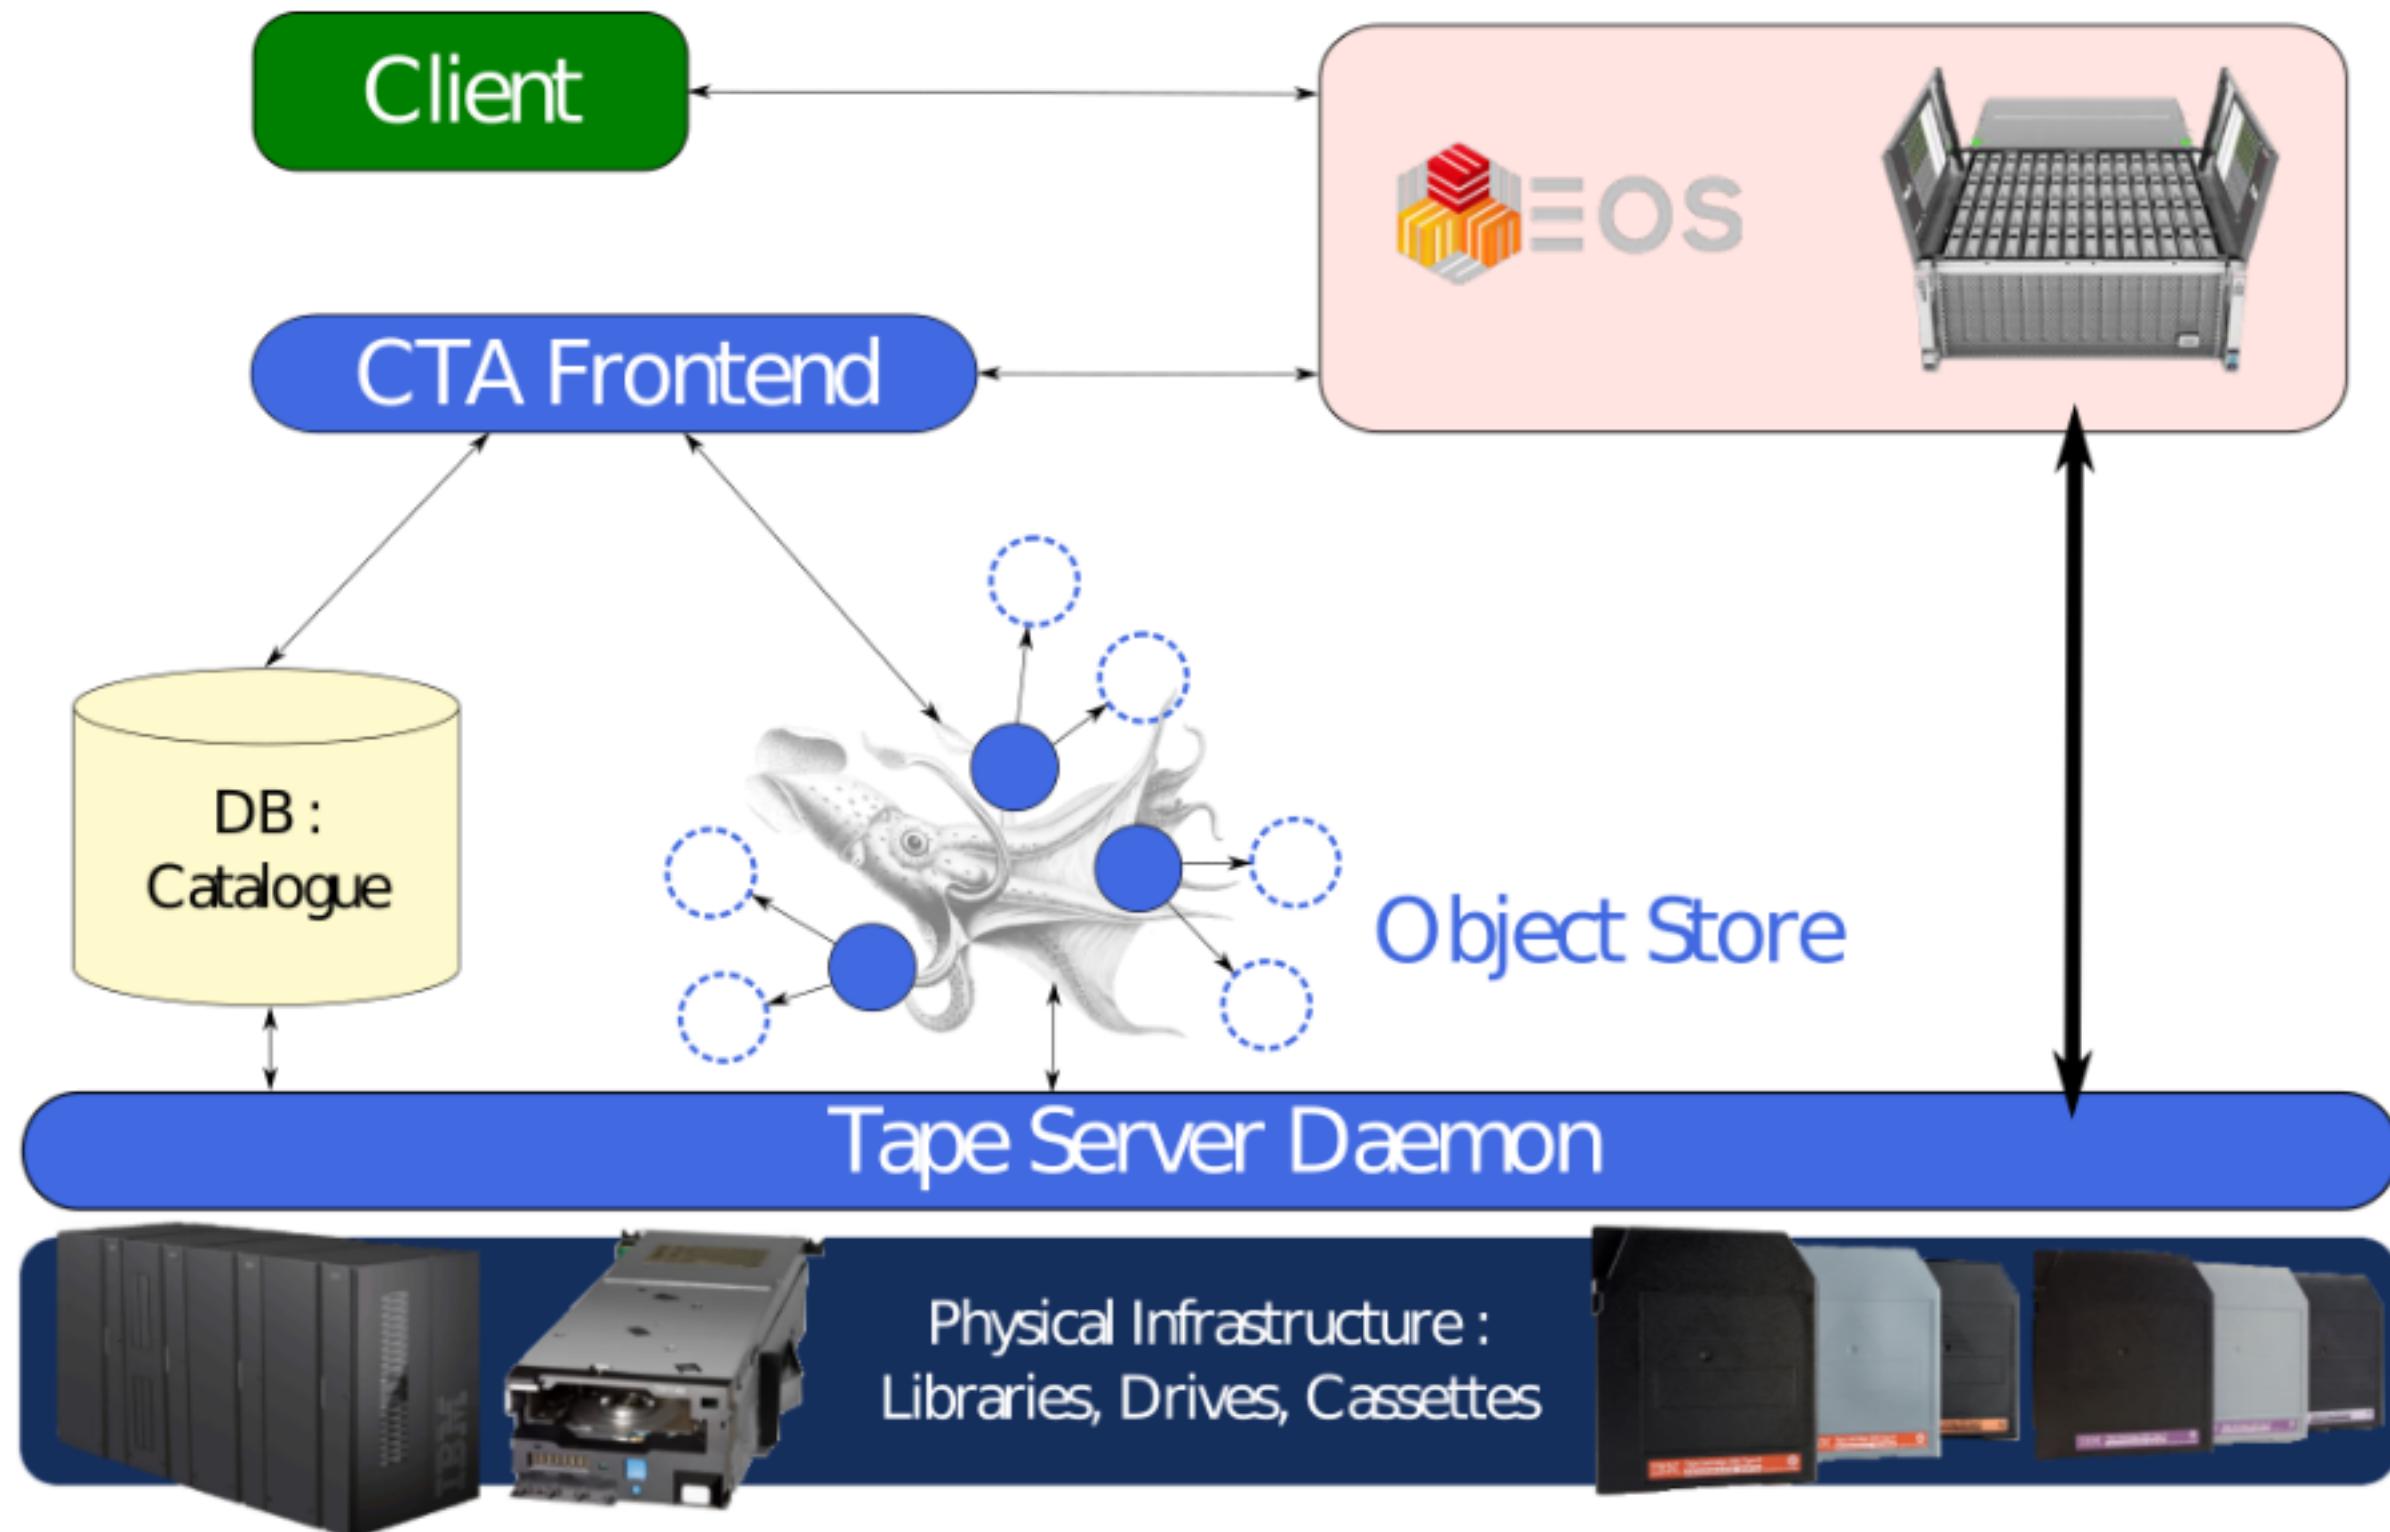
# **EOS+CTA** *Deployment*







# **EOS+CTA** *Architecture*



# **EOS+CTA** *Timeline*

2018  
Run 2

CTA FIELD  
TESTING

DATA TAKING AND  
RETRIEVAL : CASTOR

2019  
LS2

Migration during period  
of reduced data taking.

Migration is a metadata-only  
operation. No physical  
movement of data on tape.

2020  
LS2

DATA TAKING AND  
RETRIEVAL : CTA

2021  
Run 3

# **EOS+CTA** *Dev&oper*

Tightly coupled software ⇒ tightly coupled developments

Extensive and systematic testing is paramount to limit regressions

Extensive monitoring in place to ease debugging and target high performance from day 1

All 1,000+

Pending 0

Running 0

Finished 1,000+

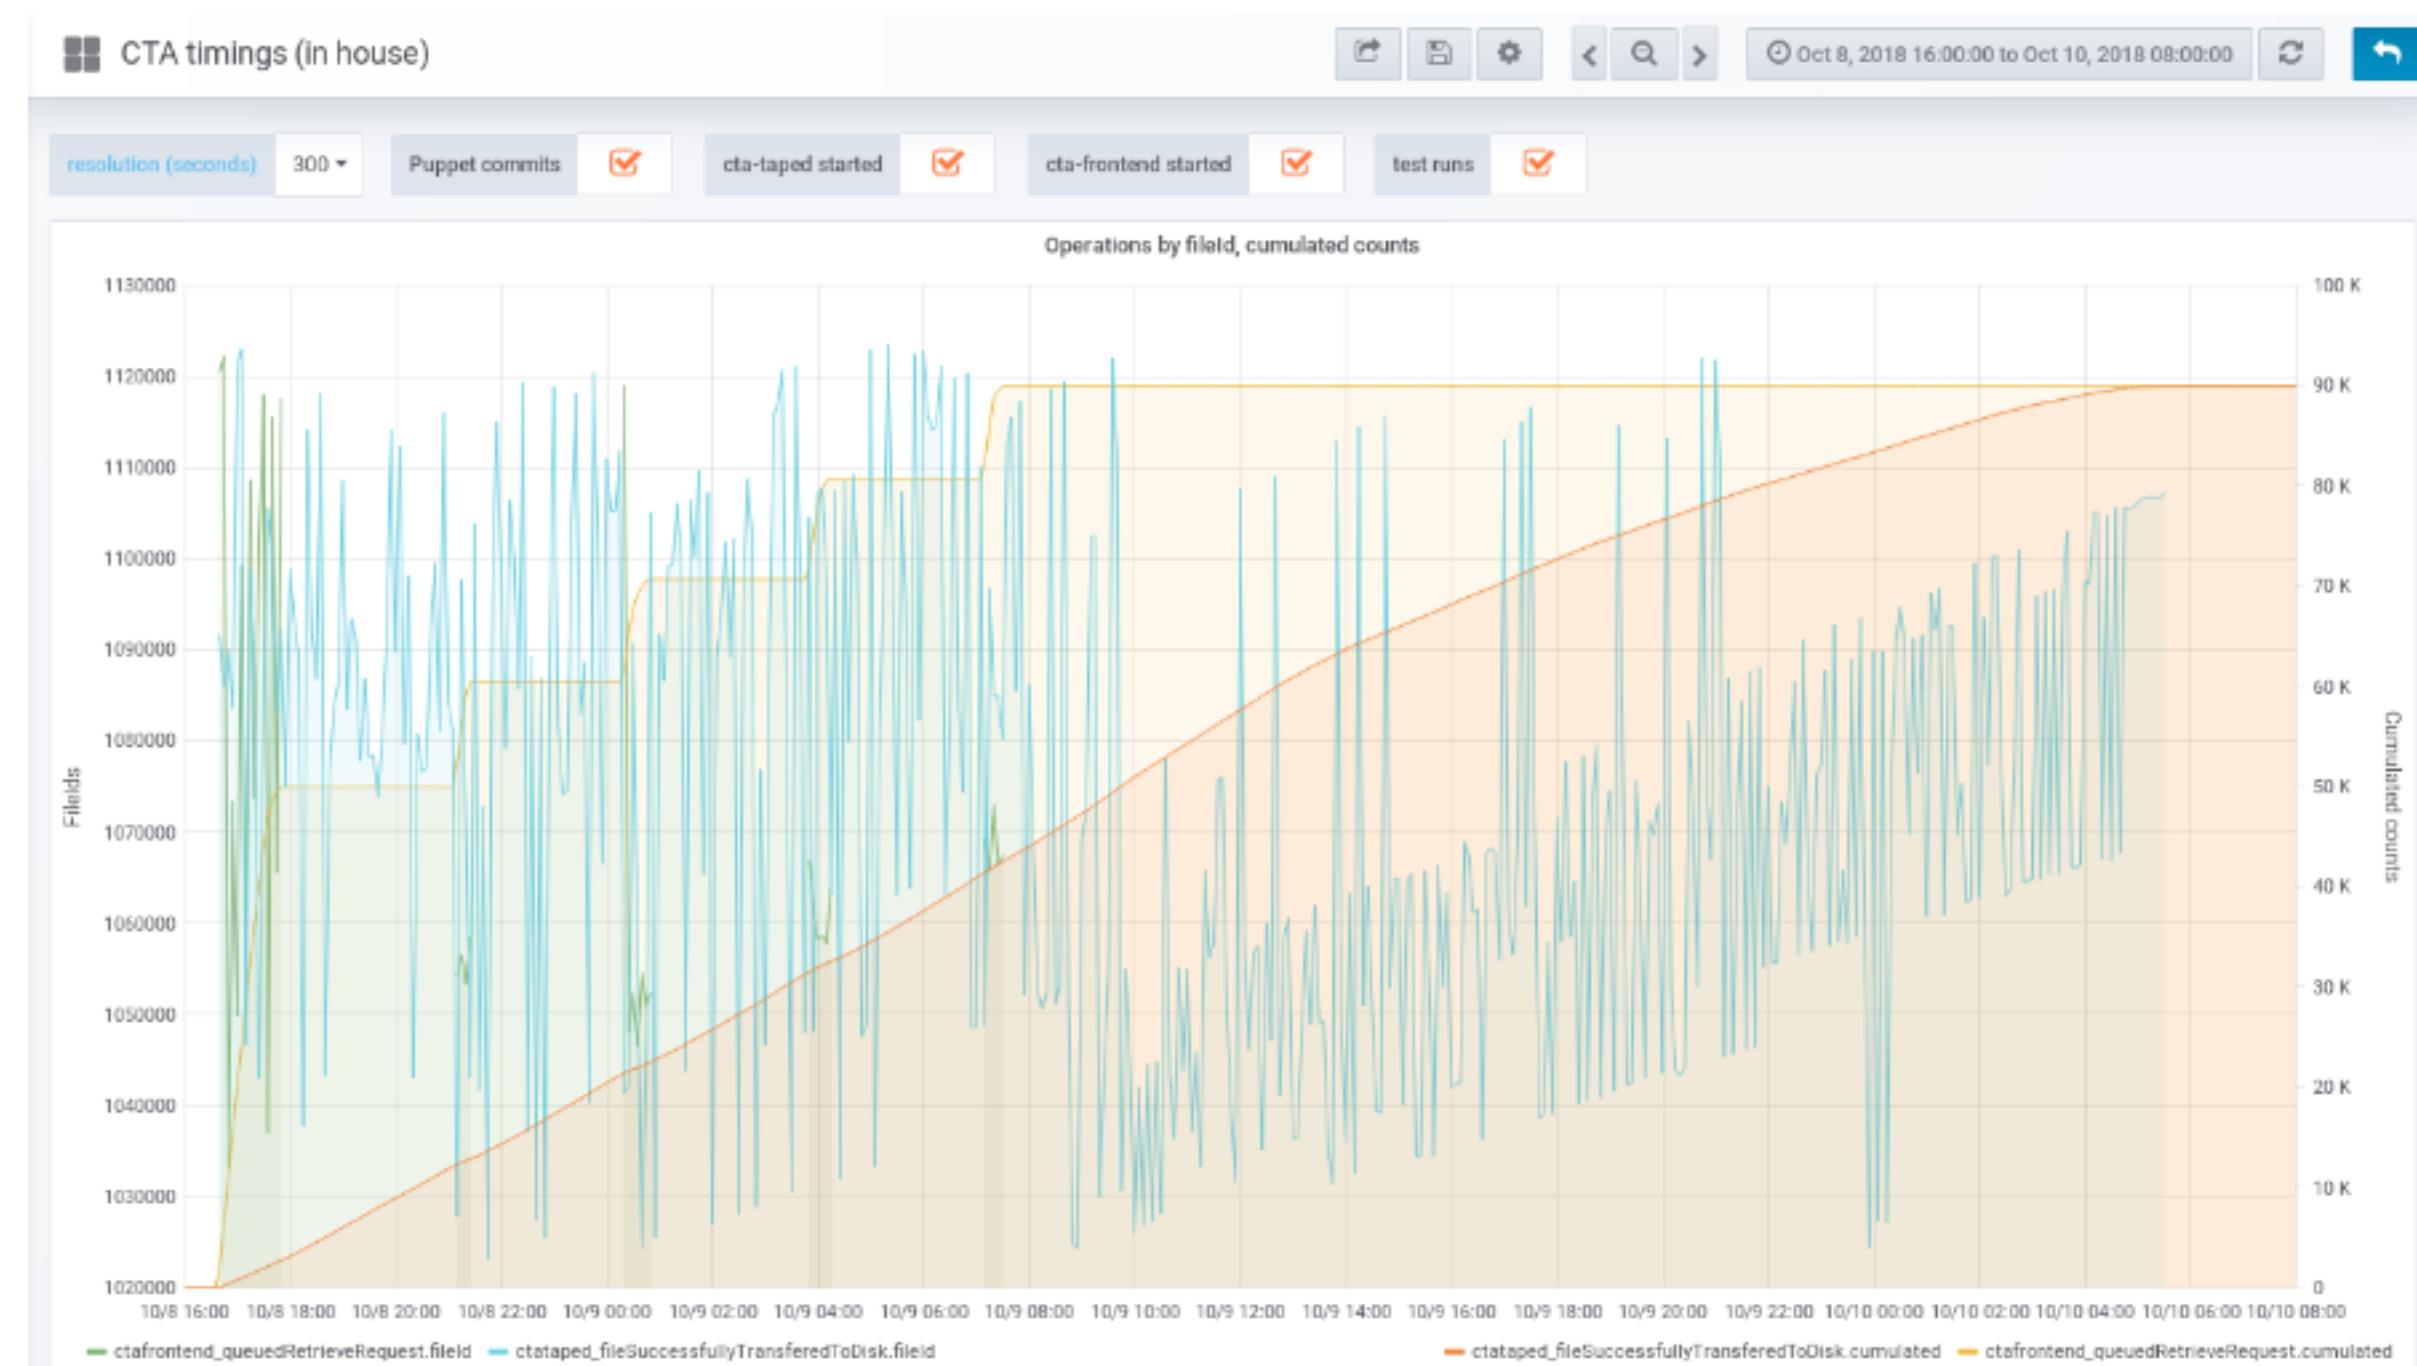
Branches

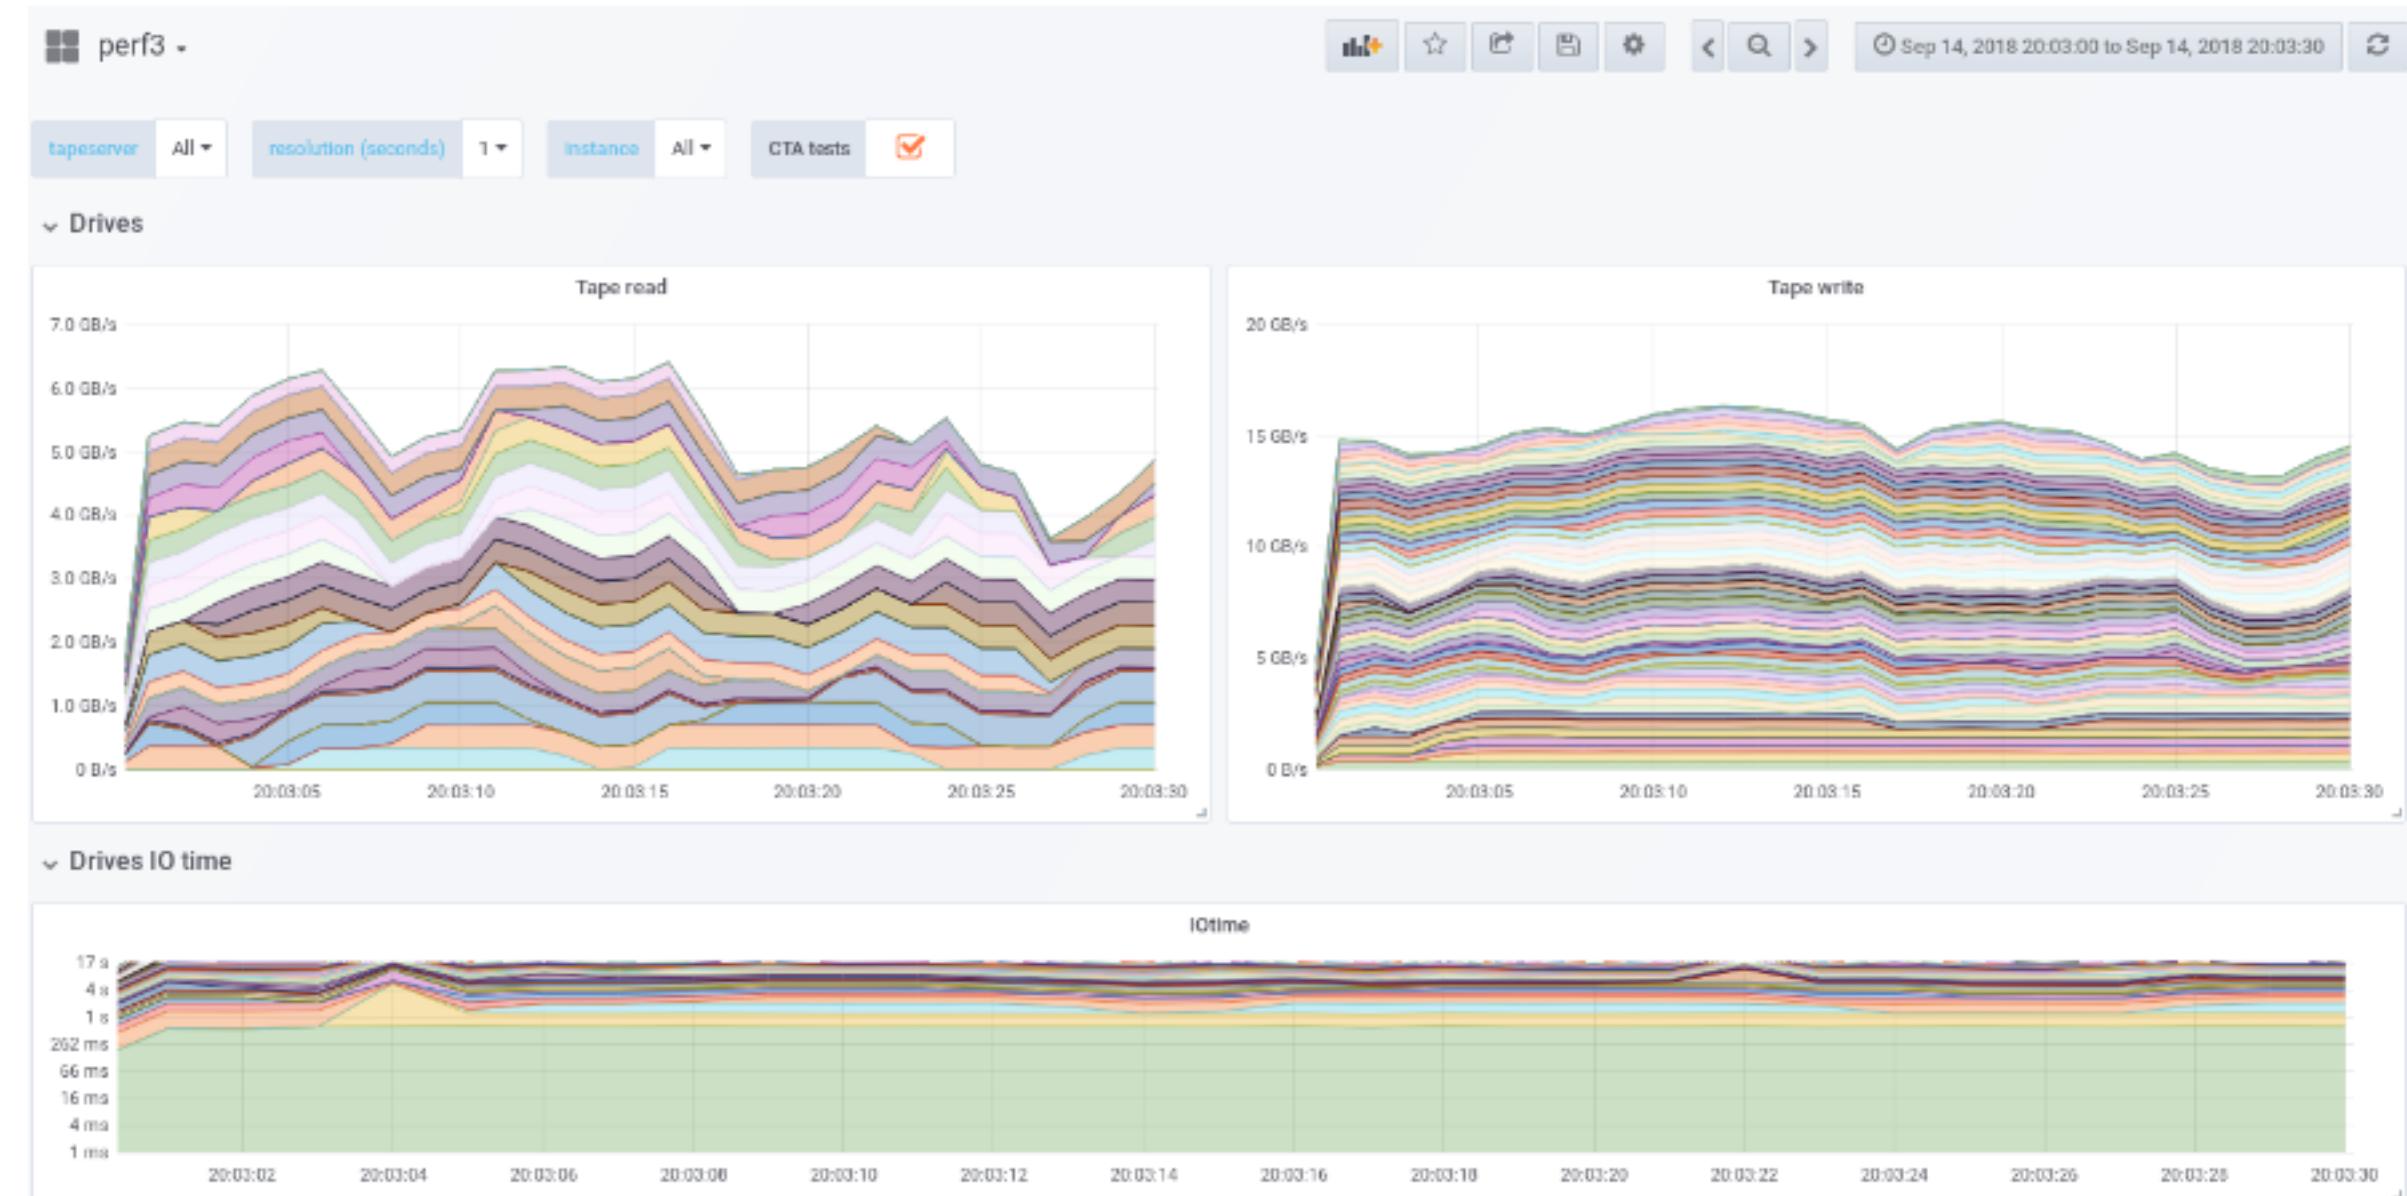
Tags

Run Pipeline

Clear Runner Caches

Status	Pipeline	Commit	Stages	
<span>passed</span>	#540338 by  latest	Y master -o 62f71584 Upgraded to xrootd-4.8.5...		01:21:09 4 hours ago
<span>passed</span>	#539010 by  latest	Y master -o 62f71584 Upgraded to xrootd-4.8.5...		00:44:23 20 hours ago
<span>passed</span>	#538956 by  latest	Y xrootd-4.8... -o 62f71584 Upgraded to xrootd-4.8.5...		00:43:44 21 hours ago
<span>passed</span>	#538771 by	Y master -o 25e80133 cta/CTA#310 tape pool wi...		01:08:01 1 day ago
<span>passed</span>	#538471 by  latest	Y v0.0-155 -o 25e80133 cta/CTA#310 tape pool wi...		00:08:51 1 day ago
<span>passed</span>	#538377 by	Y master -o 25e80133 cta/CTA#310 tape pool wi...		00:49:43 1 day ago





# **CTA metadata stress tests**

## **Small EOS instance VTL backed**

- Few retired 1Gb/s disk servers and a decent MGM for [EOS](#)
- A few VMs for the [CTA tape part](#)

**Each VTL tape server exposes 20 virtual tape drives and can sustain  
 $>1\text{kHz}$  1kB files**

# CTA metadata stress tests *Stage in*

Tests until CTA routinely archived 1M files

- **10M file test:**



# CTA metadata stress tests *Stage out*

Tests until CTA routinely retrieved **1M files**

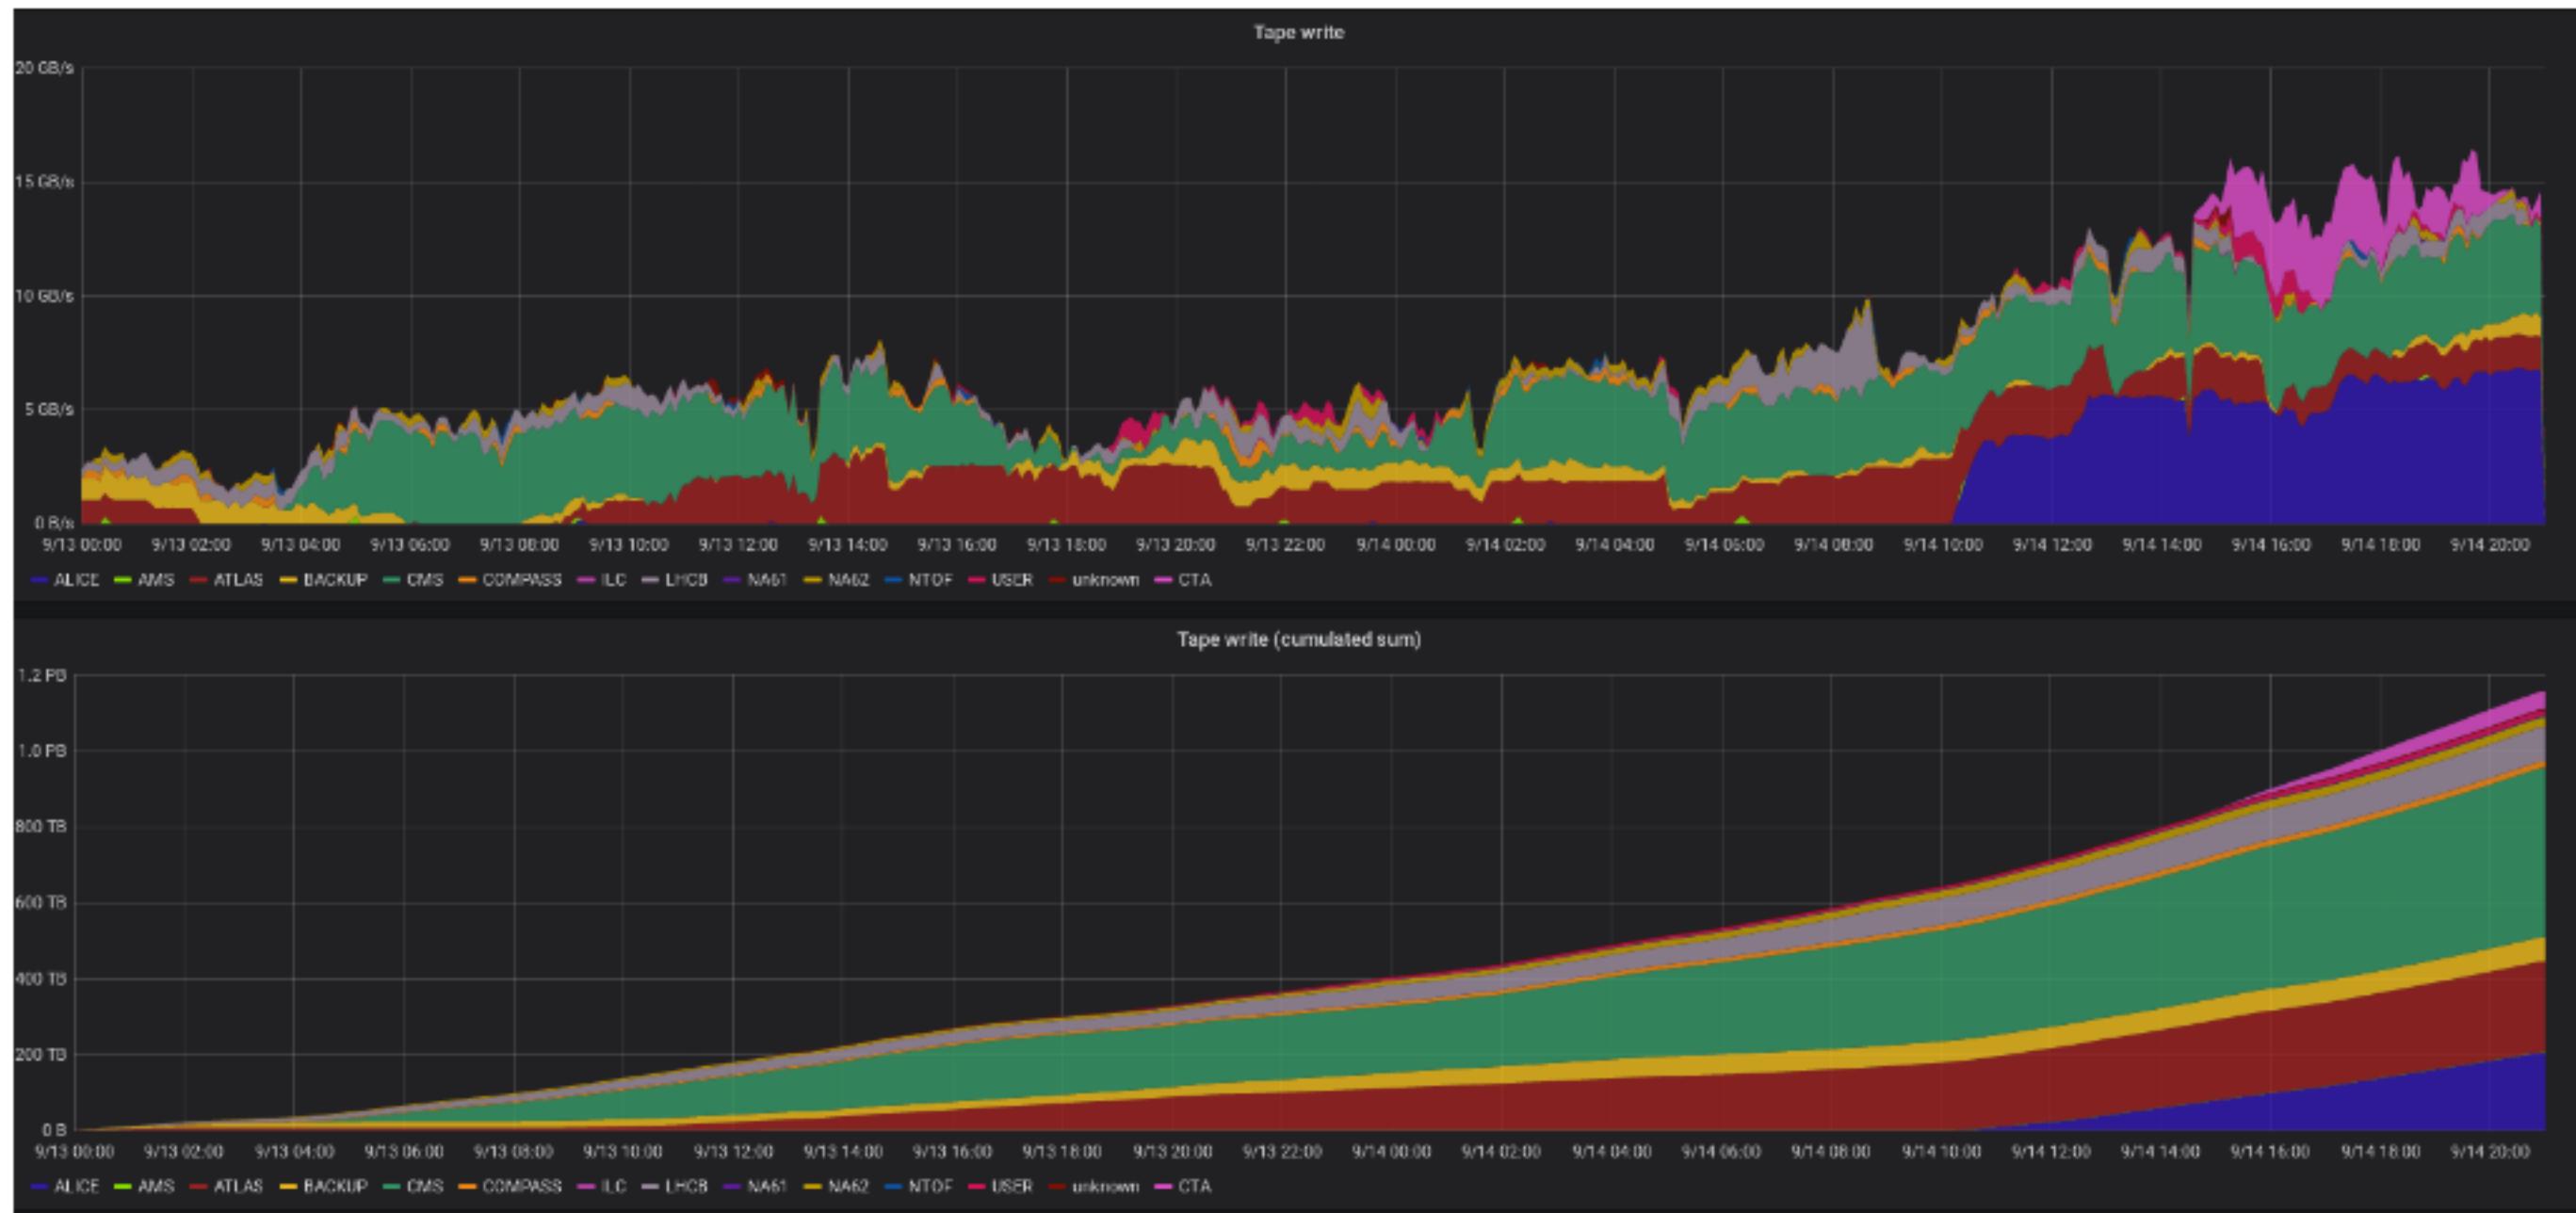
- **4M file test:**



# **CTA HI Data challenge**

## **Bigger EOS instance**

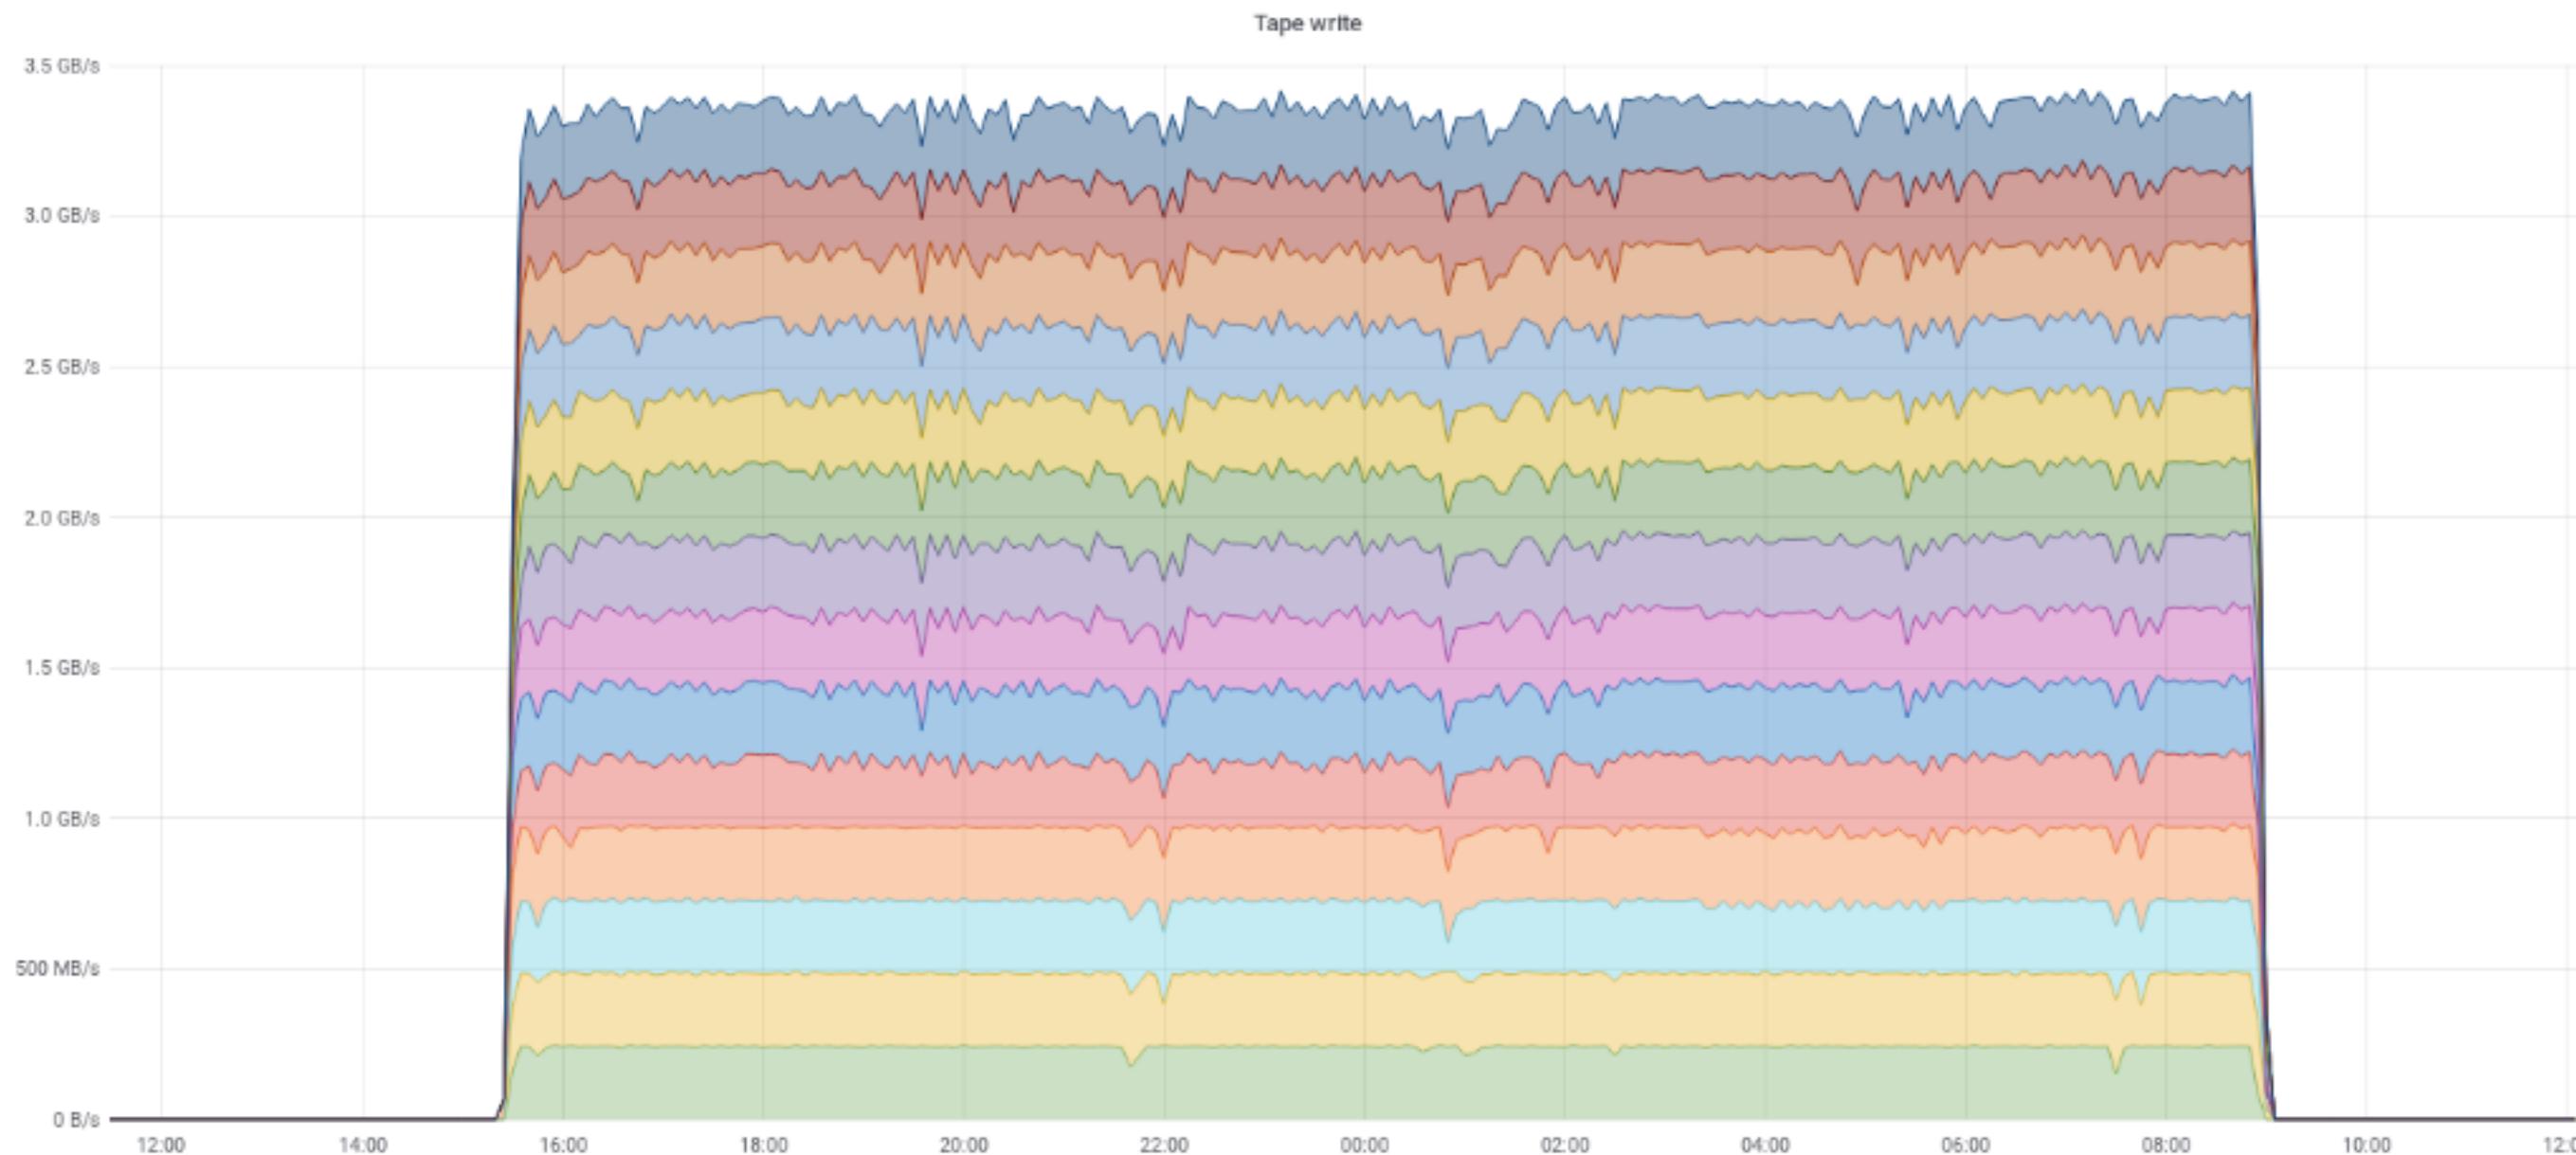
- 10x10Gb/s disk servers and a decent MGM  
for **EOS** (~2.4PB)
- 14xT10KD tape drives for the **CTA tape part** (3.5GB/s)



# **CTA HI Data challenge**

***Extended***

**Can we saturate the tape  
drive speed longer?**



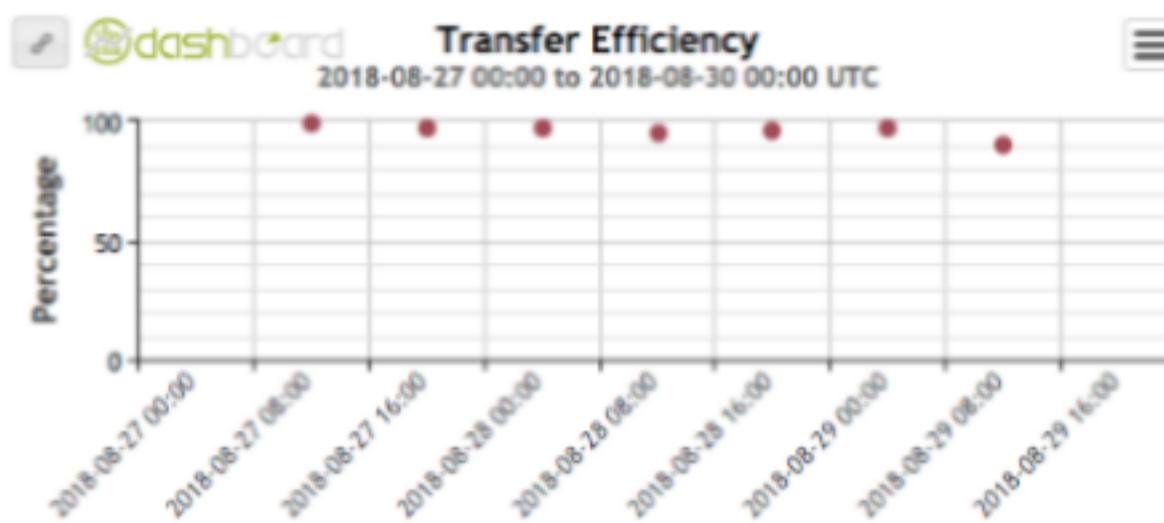
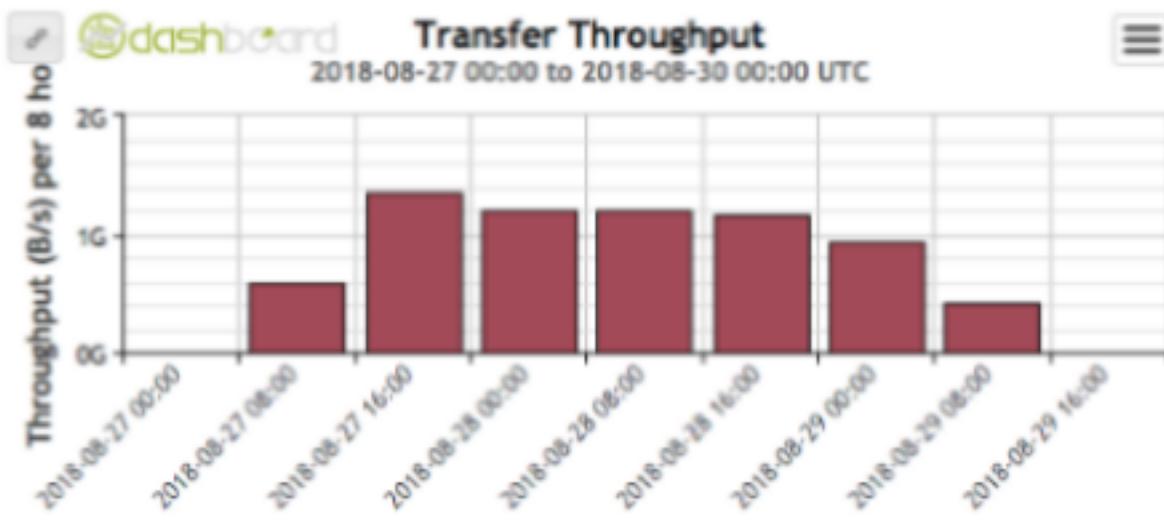
**CTA VS**  
**experiment**  
**data transfers**

# ATLAS stage in

Several tests conducted with Atlas DDM  
team using Rucio and FTS.

- 2 stage in tests of 200TB each
  - ~90k files of 2.6GB archived to tape
  - sub-optimal EOS instance (2 slow disk servers)

# ATLAS stage in

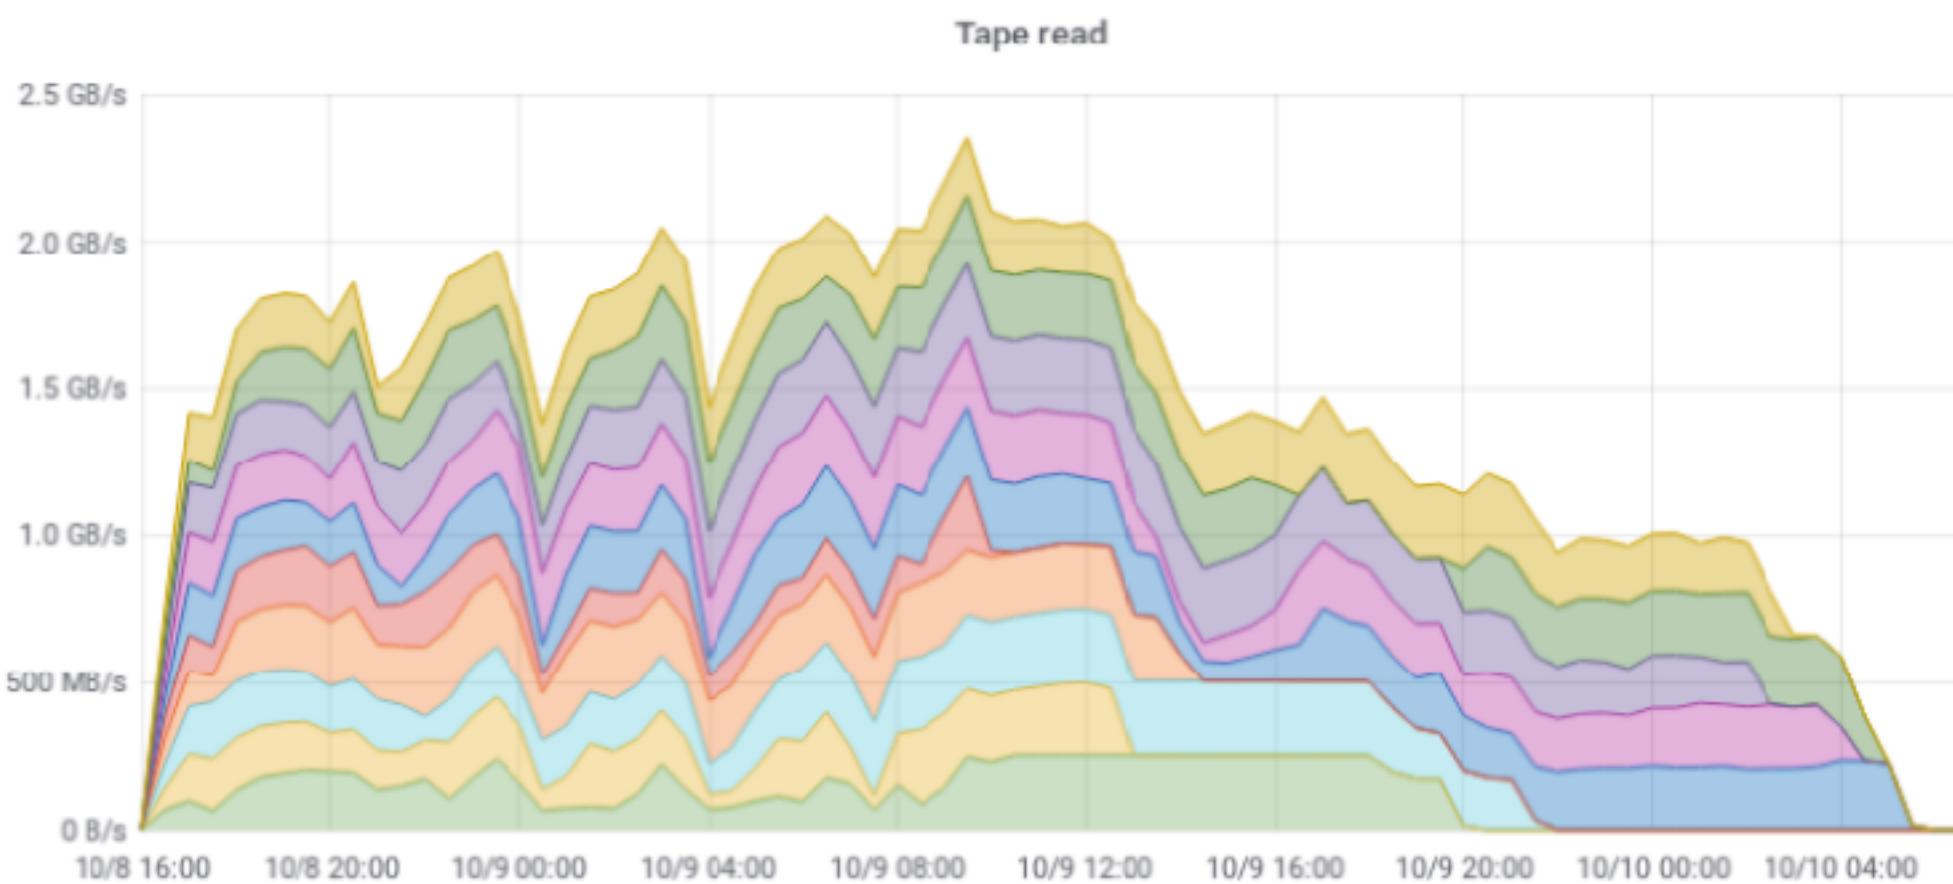


# ATLAS stage out

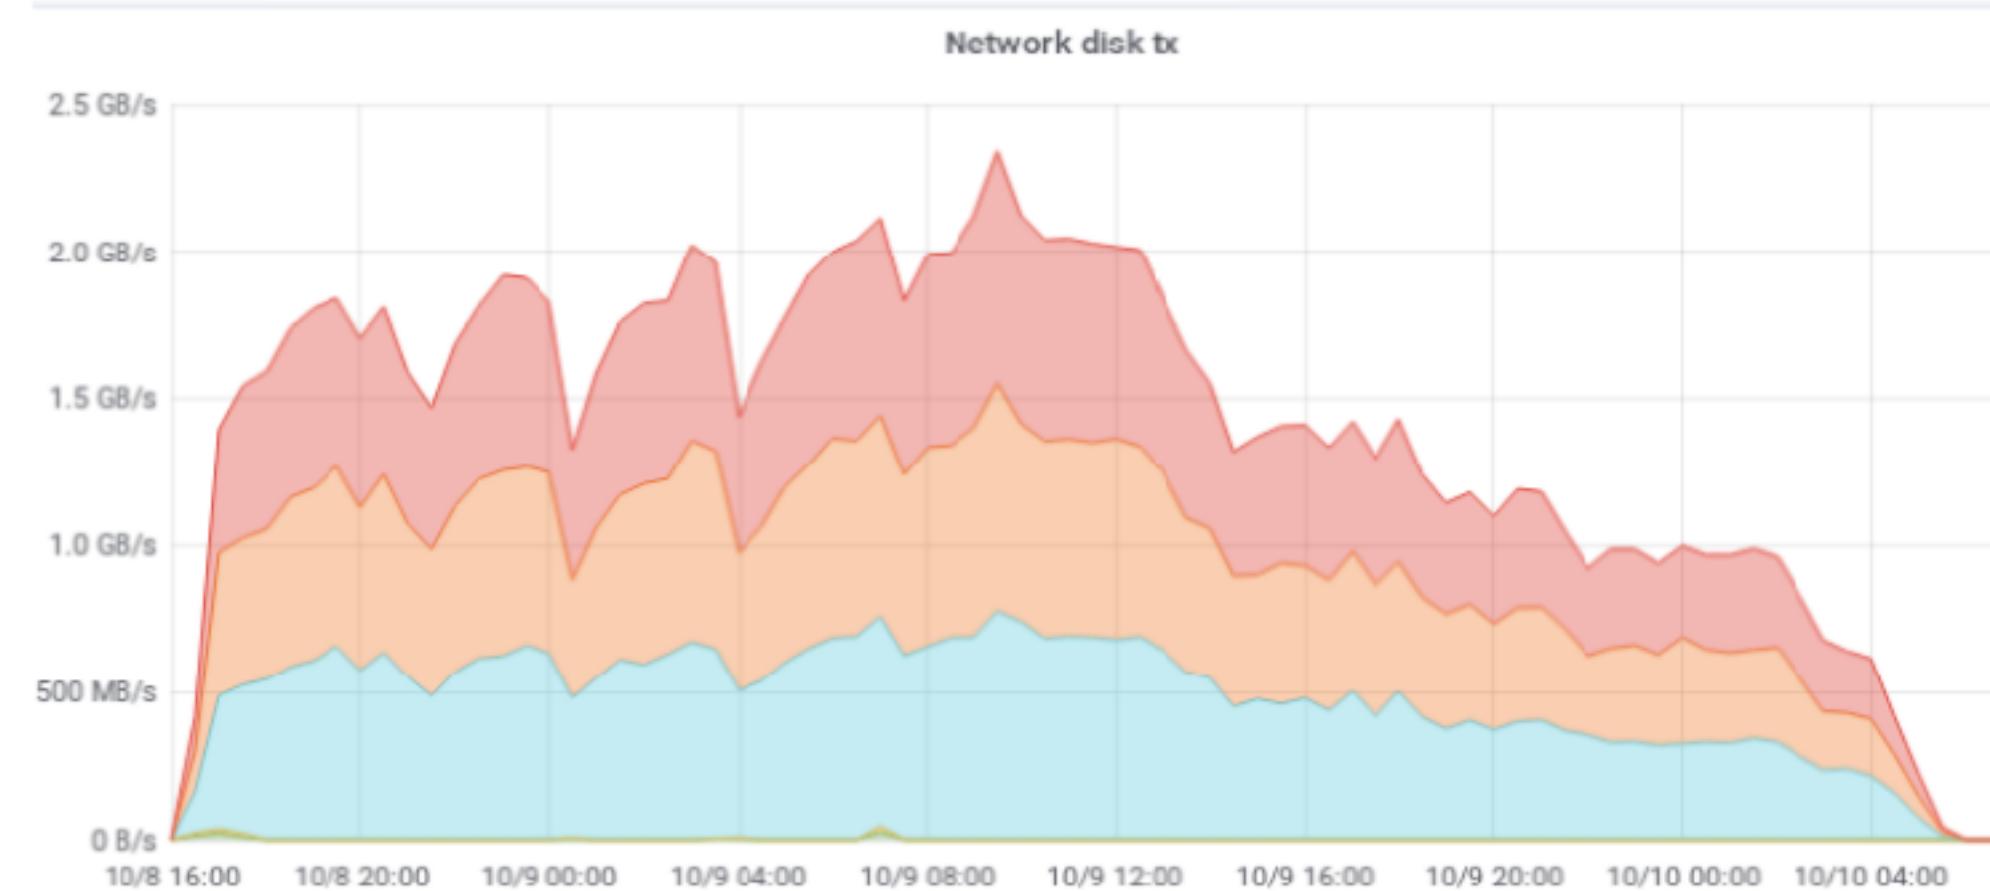
aka *Tape carousel* test took place during  
Hepix:

- 3 x EOS disk servers (~3x260TB of raw JBOD space)
- 6-10 x T10KD tape drives
- 90k files retrieved from EOSCTAATLASPPS (tape) to EOSATLAS by rucio through FTS

# ATLAS stage out



# ATLAS stage out



# ATLAS stage out DDM

