



Joint Institute for Nuclear Research

CTA Evaluation at JINR

Nikita Balashov

On behalf of the JINR storage team

3rd CTA Workshop

19 March 2024

Tape Storages at JINR

- 3 Tape Libraries
 - IBM TS3500 - 5961x2.5 TB tapes – LTO6, 12 drives (160 MB/s)
 - IBM TS4500 - 4500x20 TB 3592-JE tapes, 12 TS1160 drives (400 MB/s)
 - IBM TS3200 (Testbed) - 45x1.5 TB tapes, 4 drives (140 MB/s)
- EOS/CTA
 - TS3500 and TS4500
 - 432 TB all-flash buffer (6 servers)
- dCache/Enstore
 - TS4500
 - 2.65 PB HDD-based buffer (8 servers)
- Testbed
 - IBM TS3200
 - 4 drives (140 MB/s), 45x1.5 TB tapes

TS3500

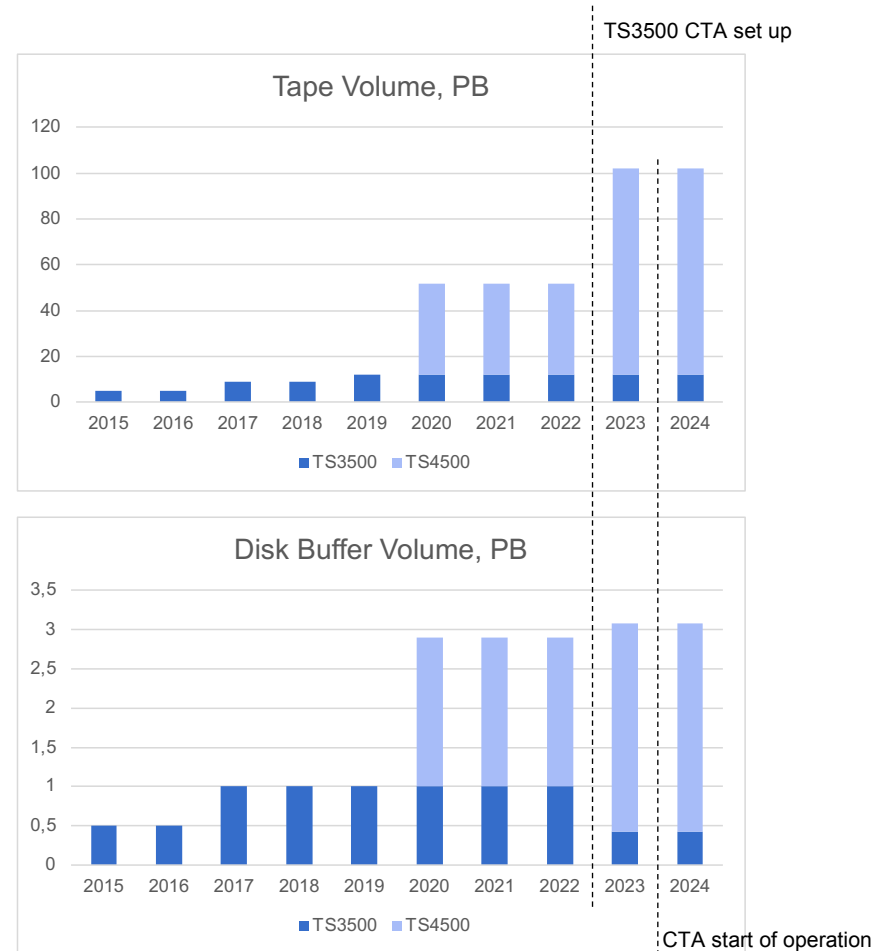


TS4500



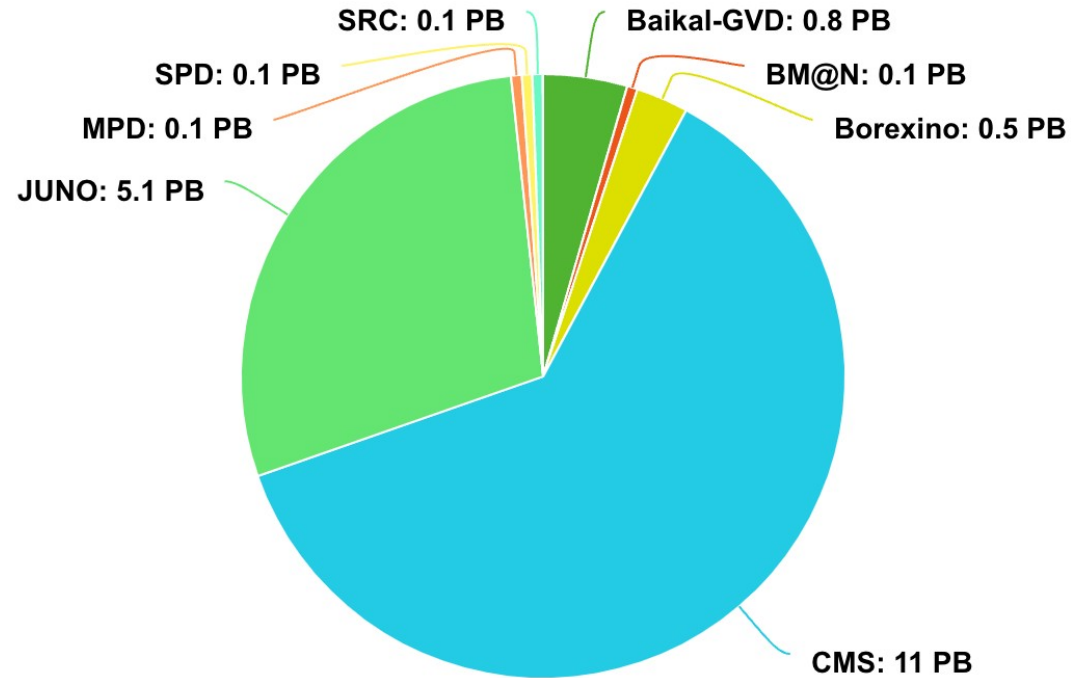
Currently Operating Libraries Timeline

- TS3500 in operation since 2015
 - Powered by Enstore till the end of 2023
 - CTA now
 - Changed HDD buffer to all-flash
- TS4500 in operation since 2021
 - Used by both dCache/Enstore and EOS/CTA
- TS3500 → TS4500 data migration
 - Took 365 days to finish (April 2020 – April 2021)
 - ~8 PB and ~3M files transferred
 - Only 3 files lost during migration



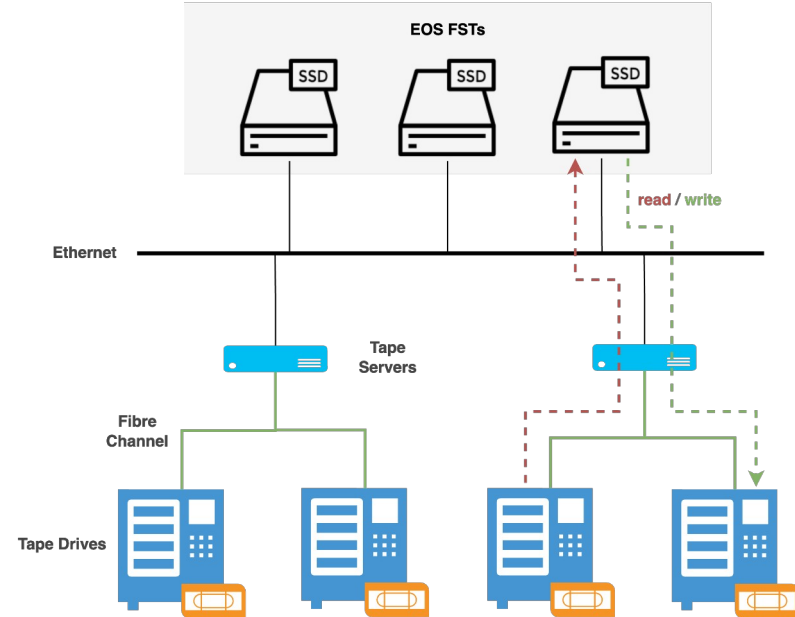
Tape Usage

- LHC: CMS
- Experiments on NICA:
 - BM@N
 - MPD
 - SPD
 - SRC
- Neutrino experiments:
 - Baikal-GVD
 - JUNO
 - Borexino
- Non-experimental data (backups, etc)



EOS-CTA Configuration Overview

- In operation since Autumn 2023
- Uses both TS3500 and TS4500
- All-flash disk cache based on EOS
- Standard EOS configuration
 - Active-passive HA for MGM/MQ
 - RAFT HA for QuarkDB
 - Replica layout
- Software
 - Scientific Linux 7 everywhere
 - EOS 5.1.9
 - CTA 5.10
 - Postgres 14.0



	CPU	Memory	Disk	Network
6 x Tape servers	2 x Xeon Gold 6126, 40 cores	128 GB	2 x 480 GB SSD (RAID1)	2 x 16 Gb/s FC, 40 Gb/s Ethernet
2 x Tape servers	1 x EPYC 7443P, 24 cores	128 GB	2 x 480 GB SSD (RAID1)	2 x 16 Gb/s FC, 40 Gb/s Ethernet
6 x EOS FST servers	2 x Xeon Gold 6348, 56 cores	256 GB	480 GB SSD (RAID1), 9 x 7.68 TB NVMe	100 Gb/s Ethernet

Humidity Issues

- Plot shows relative humidity (RH) in one of the cold aisle with near to constant temperature (~17 C)
- TS4500 triggers tape alert when RH goes out the allowable RH 20 – 50% range
 - TS4500 continues functioning
 - CTA blocks write operations
- We are working on improving environment conditions
- Meanwhile, it would be great if there was a way to make CTA ignore tape alerts




Table 1-20 Equipment environment specifications for the TS4500 tape library

Product operation					Product power off			
Dry-bulb temperature ^a		Relative Humidity (Non-condensing)		Maximum Wet-bulb temp.	Maximum elevation	Dry-bulb temp.	Relative humidity (% RH)	Maximum Wet-bulb temp.
Allowable	Recommended	Allowable (% RH)	Recommended (% RH)					
16 - 32°C	16 - 25°C	20 - 80%	20 - 50%	26°C	3050 m	5 - 45°C	5 - 80%	28°C

a. Derate maximum dry-bulb temperature 1°C/300 m (1.8°F/984 ft.) above 900 m (2 953 ft.).

Tape Library Monitoring

- Environmental conditions control
- Visual control to track manually triggered operations and fallen tapes
- Software
 - Grafana
 - Prometheus
 - MariaDB
- Hardware
 - Raspberry Pi 3
 - DHT22 and BME280 Sensors
 - POE IP Camera
 - USB camera connected via Pi3



Tape Future at JINR and Conclusion

- We are considering 3 scenarios in context of Enstore discontinuation
 - Move from dCache/Enstore to EOS/CTA completely
 - Continue operating both configurations
 - Evaluate other disk-tapes configurations: EOS/Enstore and dCache/CTA
- We have already started porting Enstore to Python 3 and consider taking over its support

Conclusions

- Currently, it is too early to make any conclusions about EOS/CTA
- The only issue so far is TapeAlert flag treatment by CTA, we are interested in implementing a CTA option to ignore it
- We haven't given up on Enstore, there's light at the end of the tunnel
- We'll continue operating both dCache/Enstore and EOS/CTA for some time
- Move to new OS due to SL 7 EOL (presumably AlmaLinux 9)

Thanks!

Nikita Balashov
balashov@jinr.ru