


From year 2022 to now


21.8 Bil files read for **5.15 EB**

2.82 Bil files written for **773 PB**


Avg read throughput 205GB/s, max 531GB/s (23/10/2022)



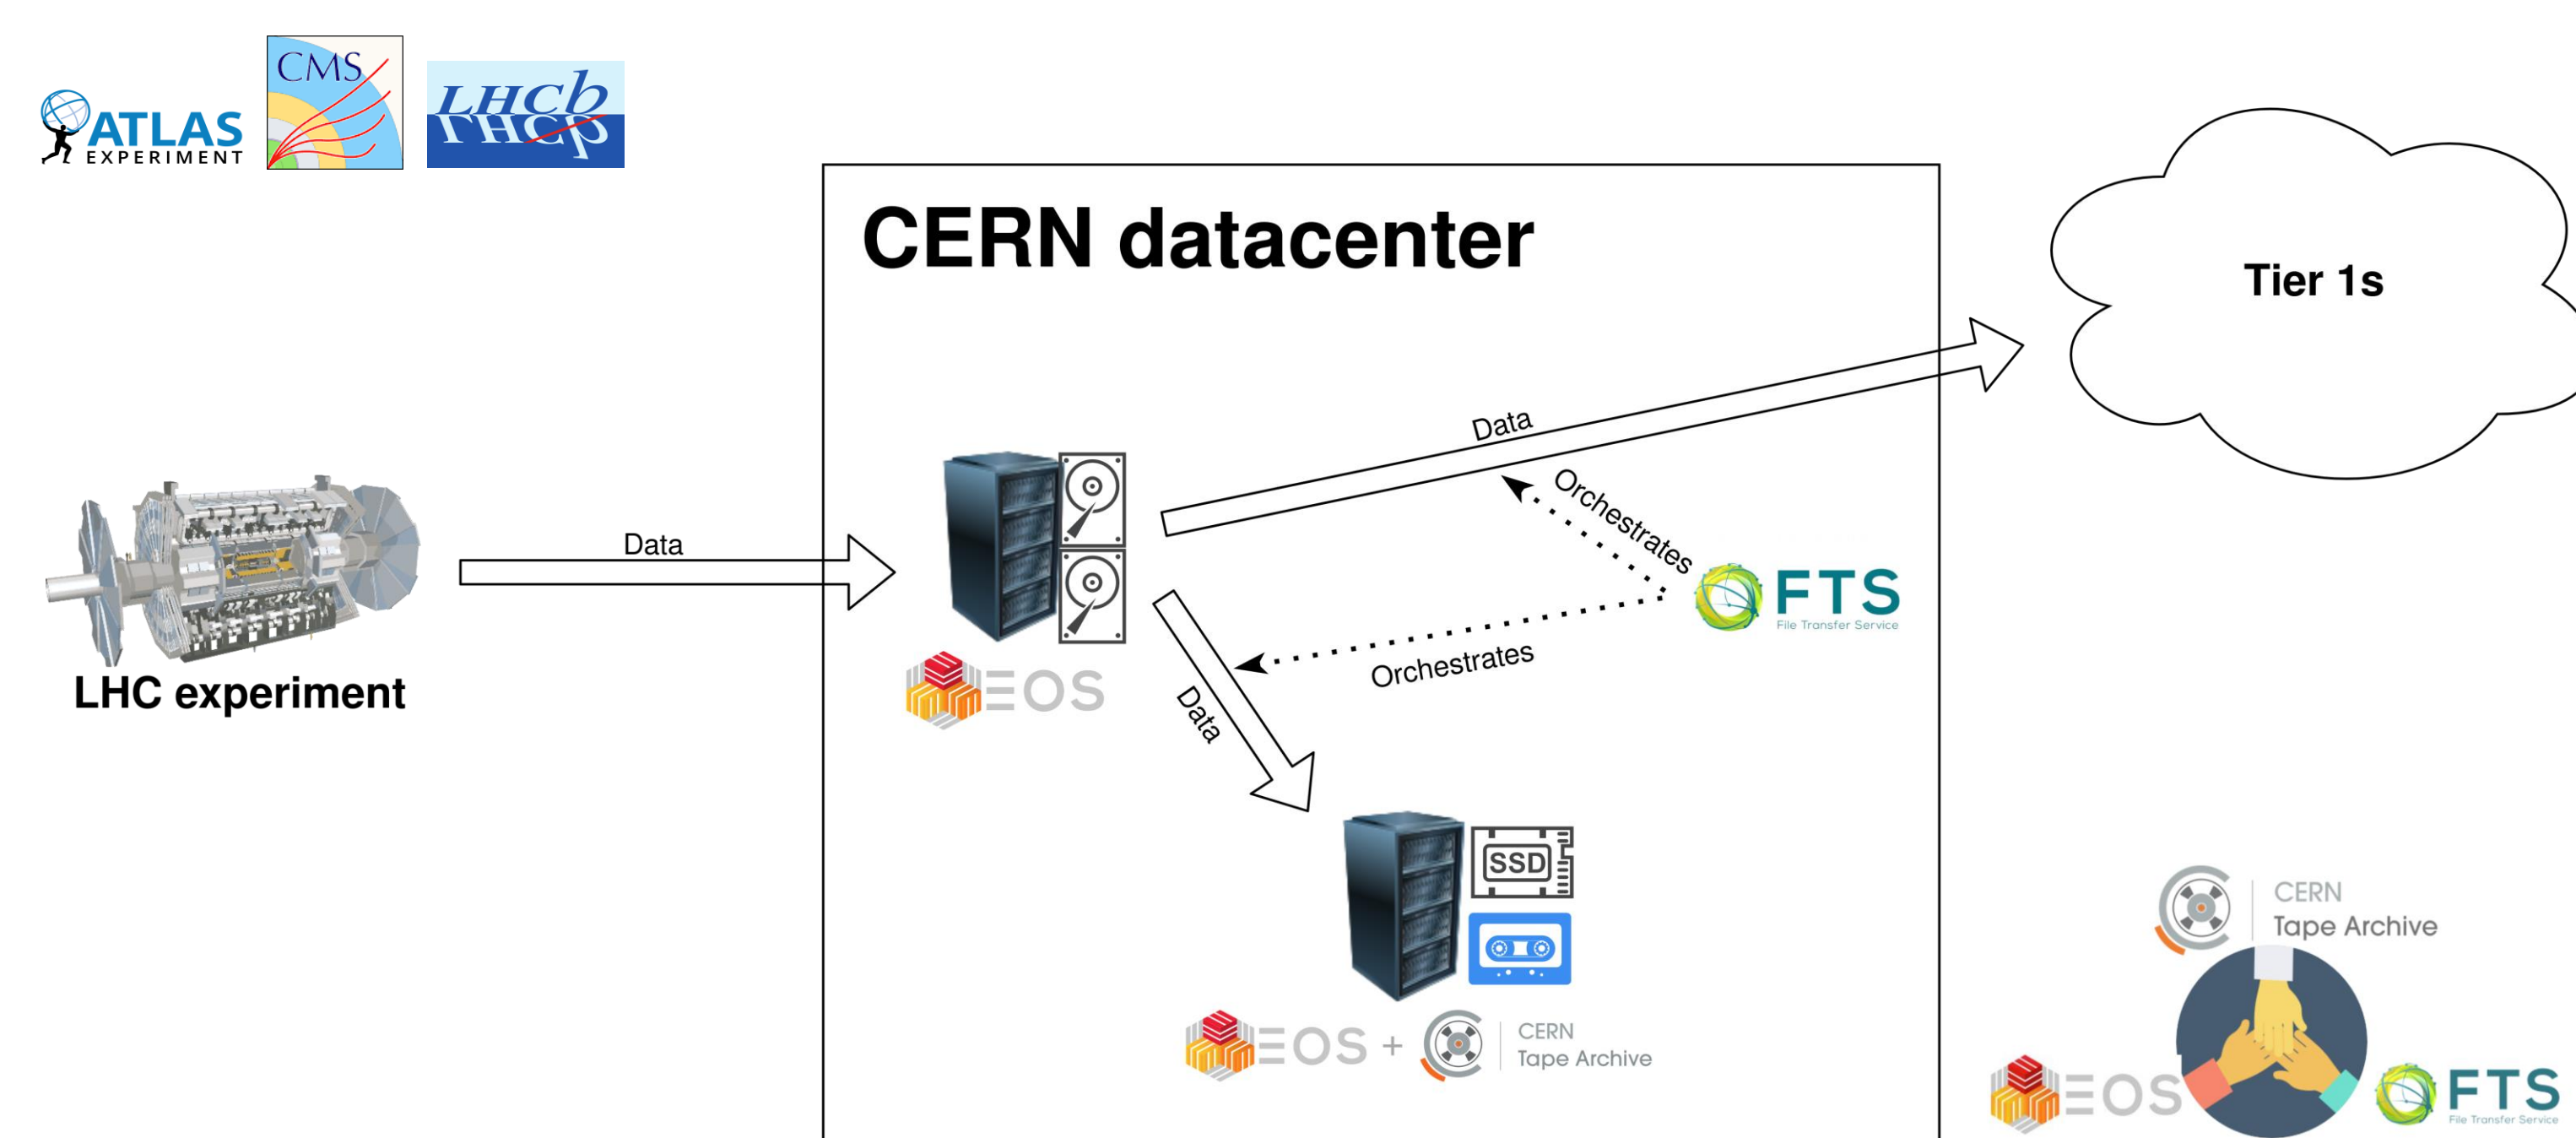
XRootD is the #1 interface used for data transfers
FUSE is the #2



Avg write throughput : 75 GB/s (including replication), max 289GB/s (19/11/2022)



General LHC data taking workflow



- The experiments' data is sent from their DAQ system to their dedicated main EOS instance
- FTS orchestrates data transfers between the main EOS instance, the tape archival storage (CTA) and Tier-1s

Hardware

NAMESPACE (MGM + QuarkDB)

- 386 GB of RAM
- 64 cores CPU Intel Xeon(R) Silver 4216 CPU @ 2.10GHz

Storage nodes


Storage nodes	Disks	Total raw space
~1300	~60000	780 PB

CentOS 7 } Default OS deployed **88%**

~150 machines under **AlmaLinux 8** } **12%**

- Alma 8 deployed on new hardware and results are promising
- A campaign to upgrade to Alma 9 will come in a close future

Deployment of erasure coding (EC)




63.79 PB optimized

- Instead of storing 2 replicas for each file, split it into chunks and add parity stripes
- Ensures files availability by saving space compared to mirror replication

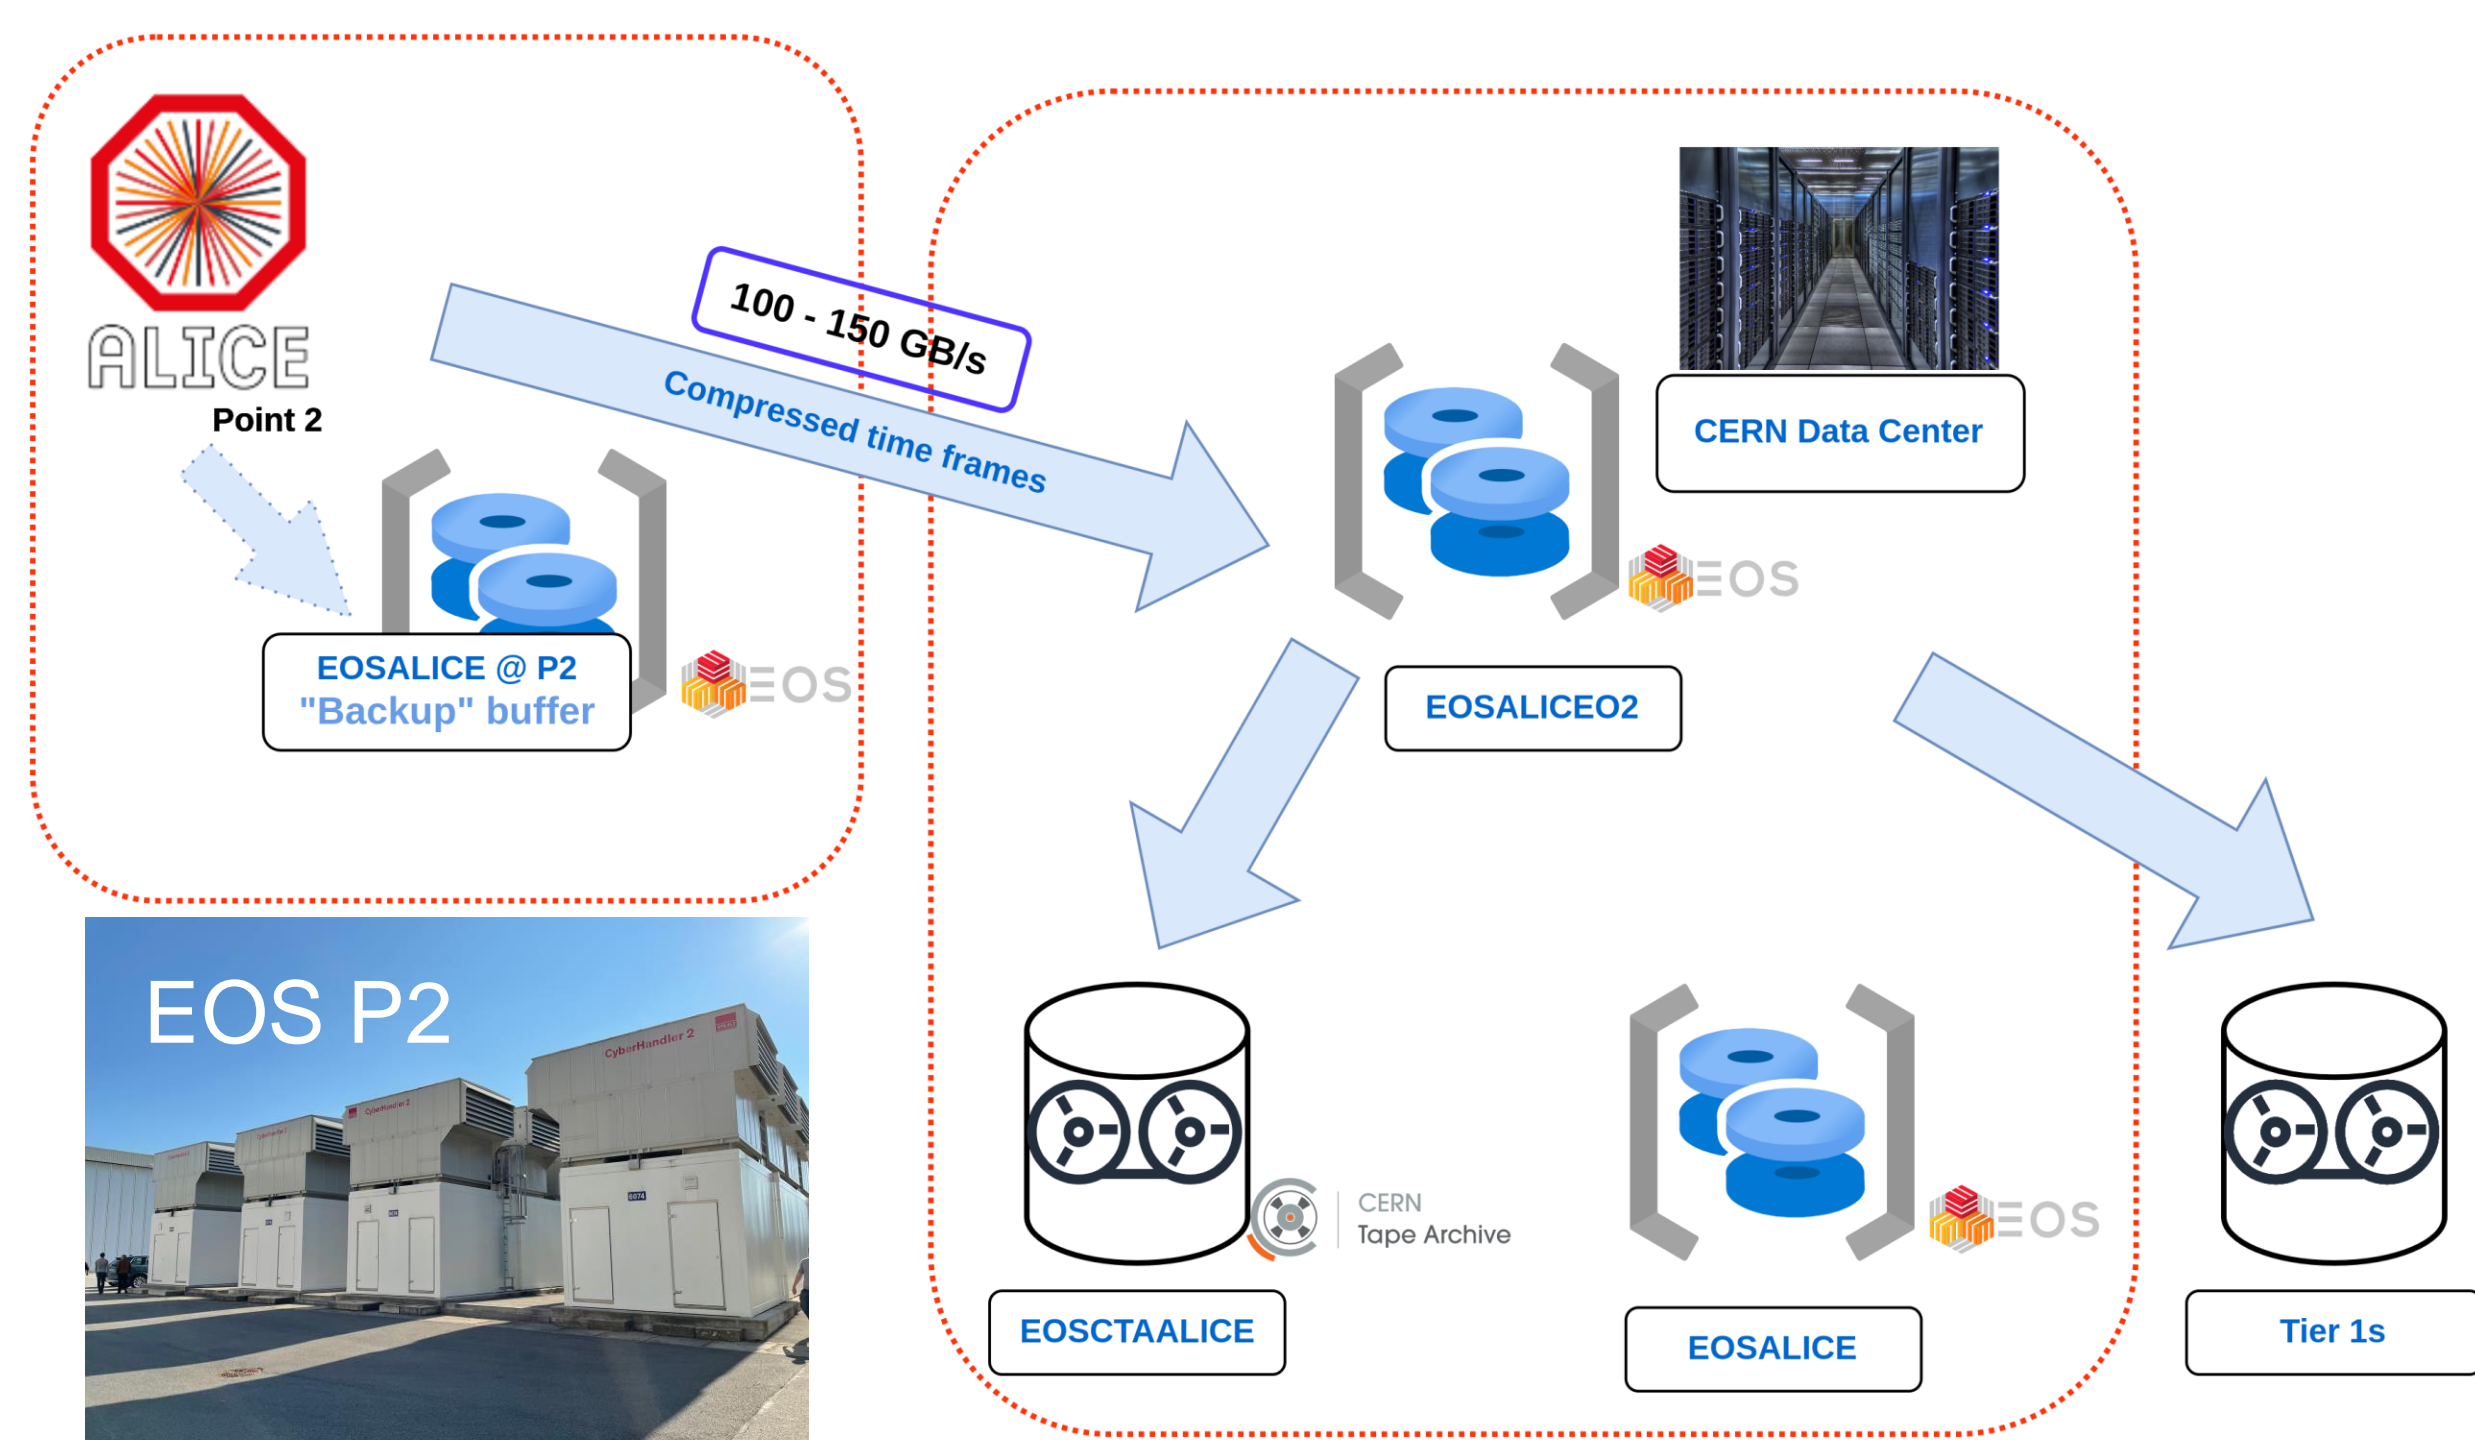
Latest hardware delivered

- 256 GB of RAM
- 96 x 18 TB disks
- 100 Gbps NIC

4 JBOD x 24 disks



The ALICE setup



- The EOS instance *EOSP2* has been configured as a backup instance in case there is a link failure between ALICE point 2 and CERN's data center.
 - ~1.5 day of buffer (13.5 PB) at 100 GB/s in single replication
- ALICE02 fully erasure coded, delivers up to 300GB/s of ingestion rate

