

The background of the slide is a close-up, low-angle photograph of a server rack. The perspective is from the bottom left, looking up and across the rows of server units. The units are dark, possibly black or dark grey, with a repeating pattern of rectangular cutouts. A bright blue light strip is visible on the left side, illuminating the scene and creating a strong contrast with the dark background. The overall mood is technical and modern.

Storage Hardware Procurement

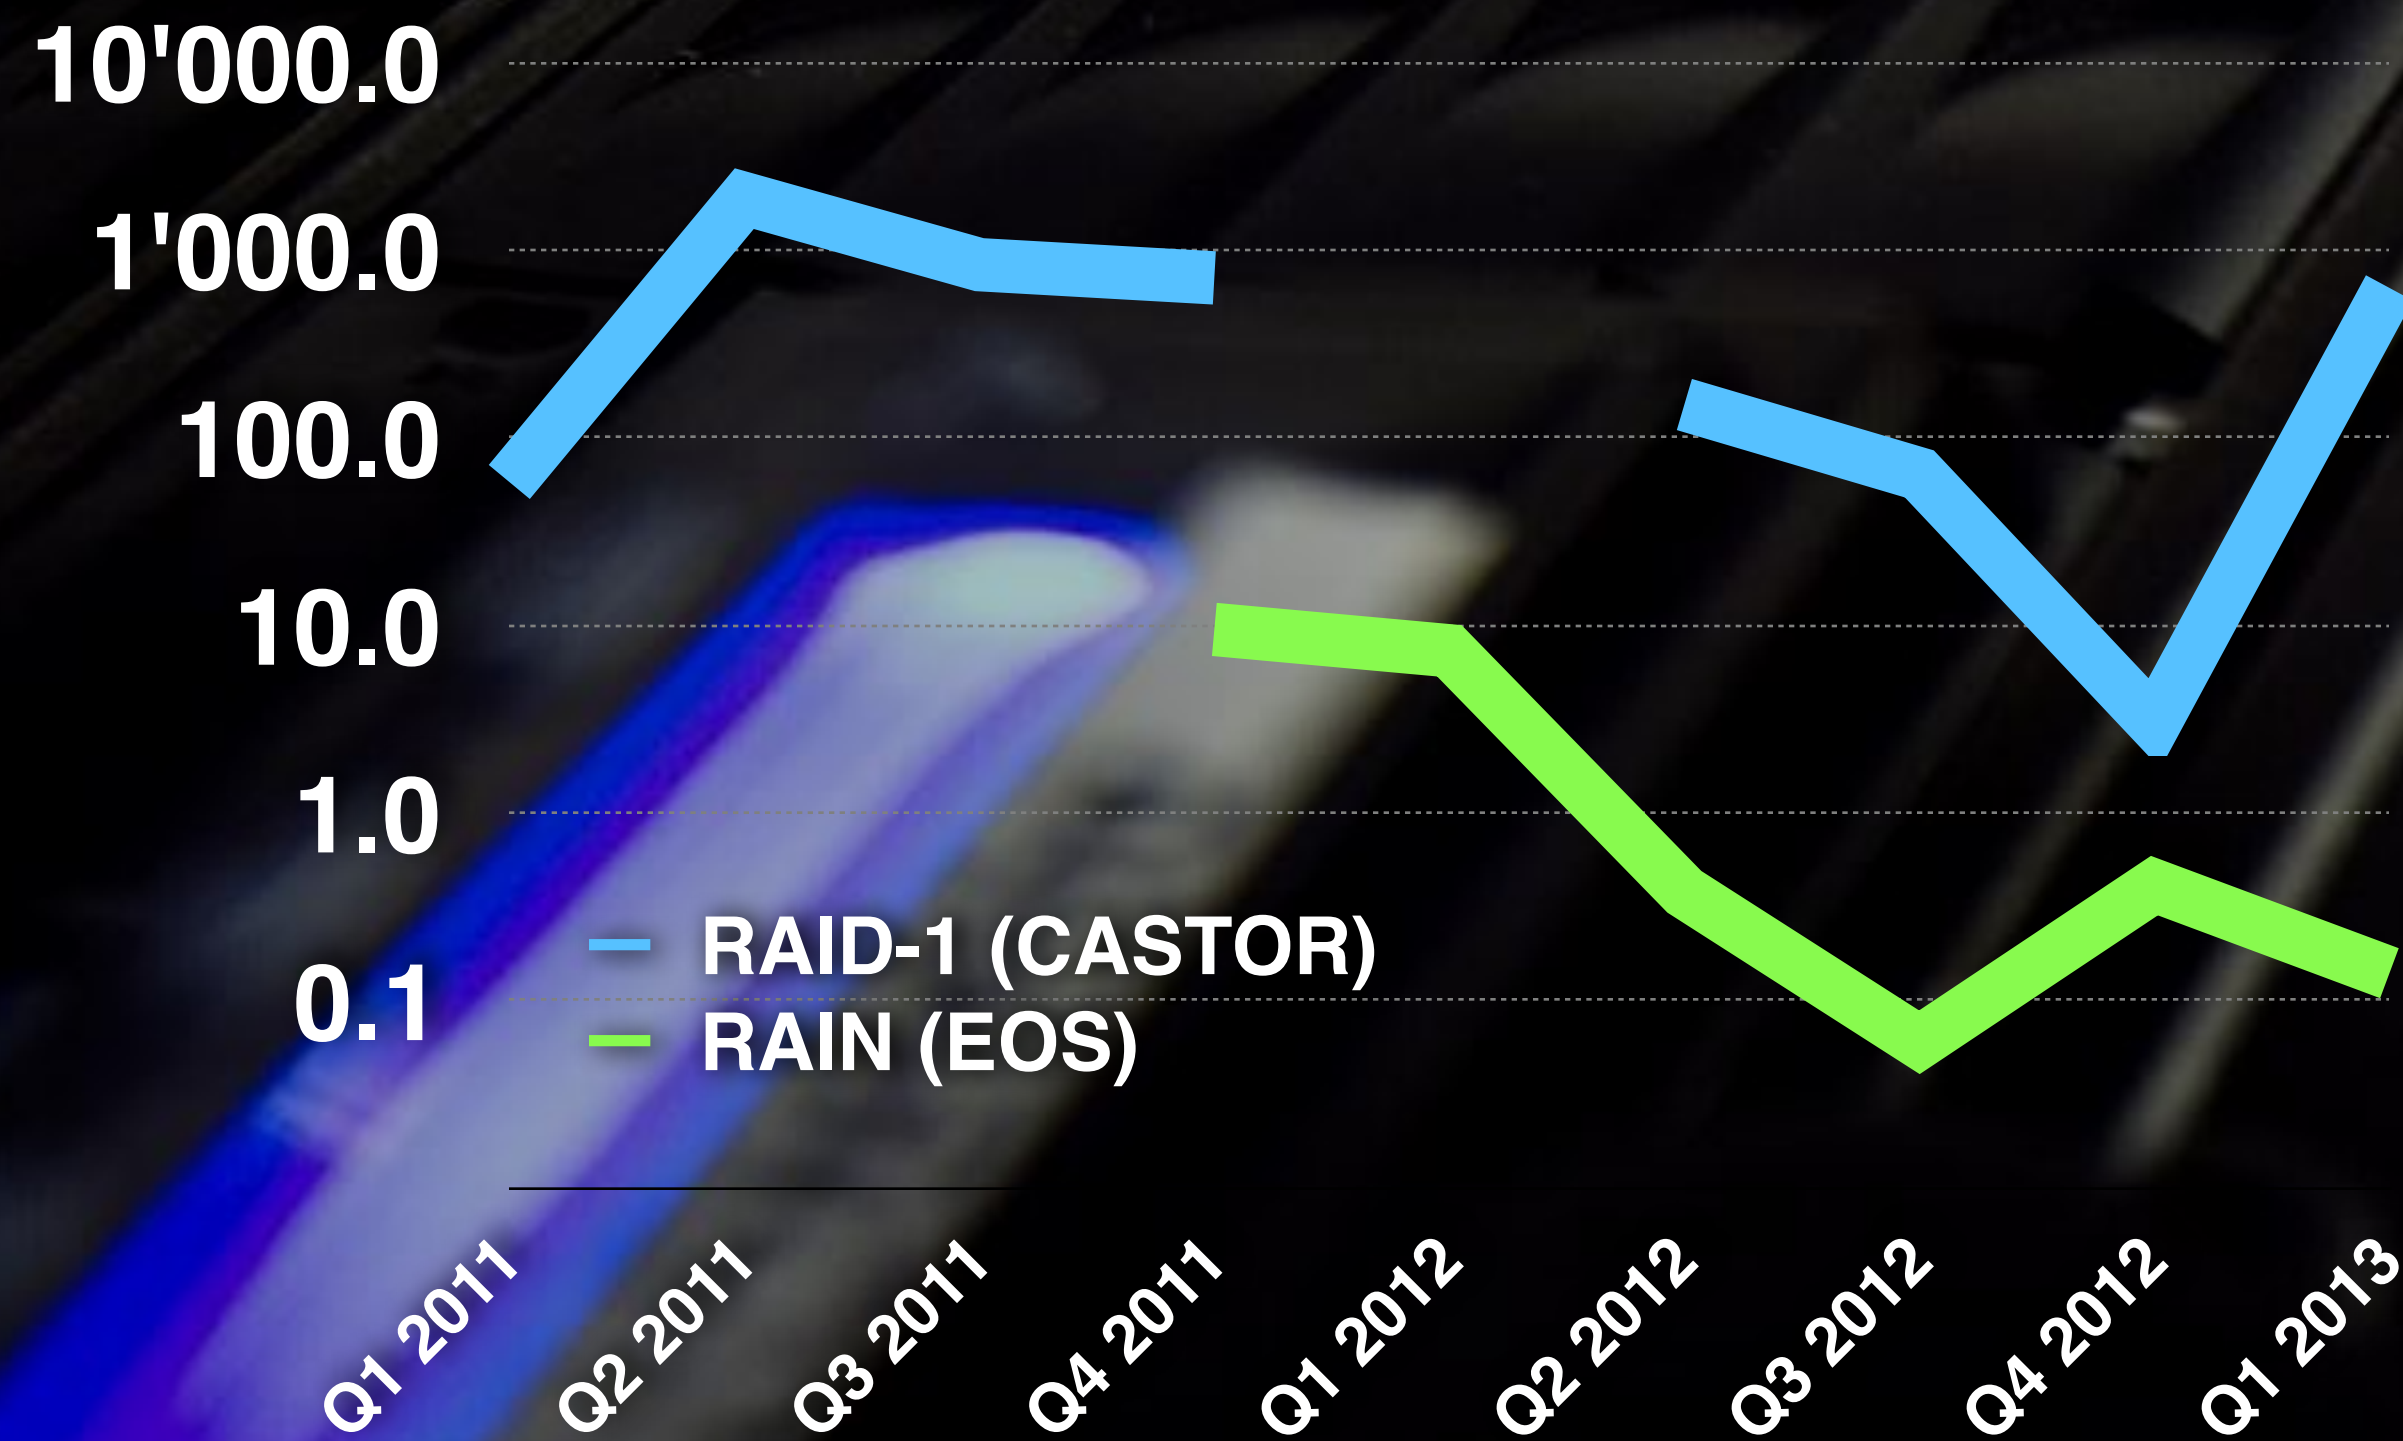
Luca Mascetti
Storage and Data Management Group

A bit of history...


Back in 2010-2013 for disk-only files the majority of issues was caused by faulty hardware controller, impacting mostly HW RAID Setup (RAID-1)

Thanks to EOS replication over multiple storage nodes (RAIN) we optimise our systems to only use JBODS and effectively changed our operational mode and our storage purchase strategy

Files lost per million




The new storage “Building Block”



QUAD + SAS Arrays

The new storage “Building Block”



QUAD + SAS Arrays

The new storage “Building Block”



Over the years we commissioned and operate multiple solutions:

- Server + 2x 24-bay SAS Arrays
- Server + 4x 24-bay SAS Arrays
- Server + 8x 24-bay SAS Arrays

Storage Server in 2014: 200 TB

Storage Server in 2023: 1700 TB

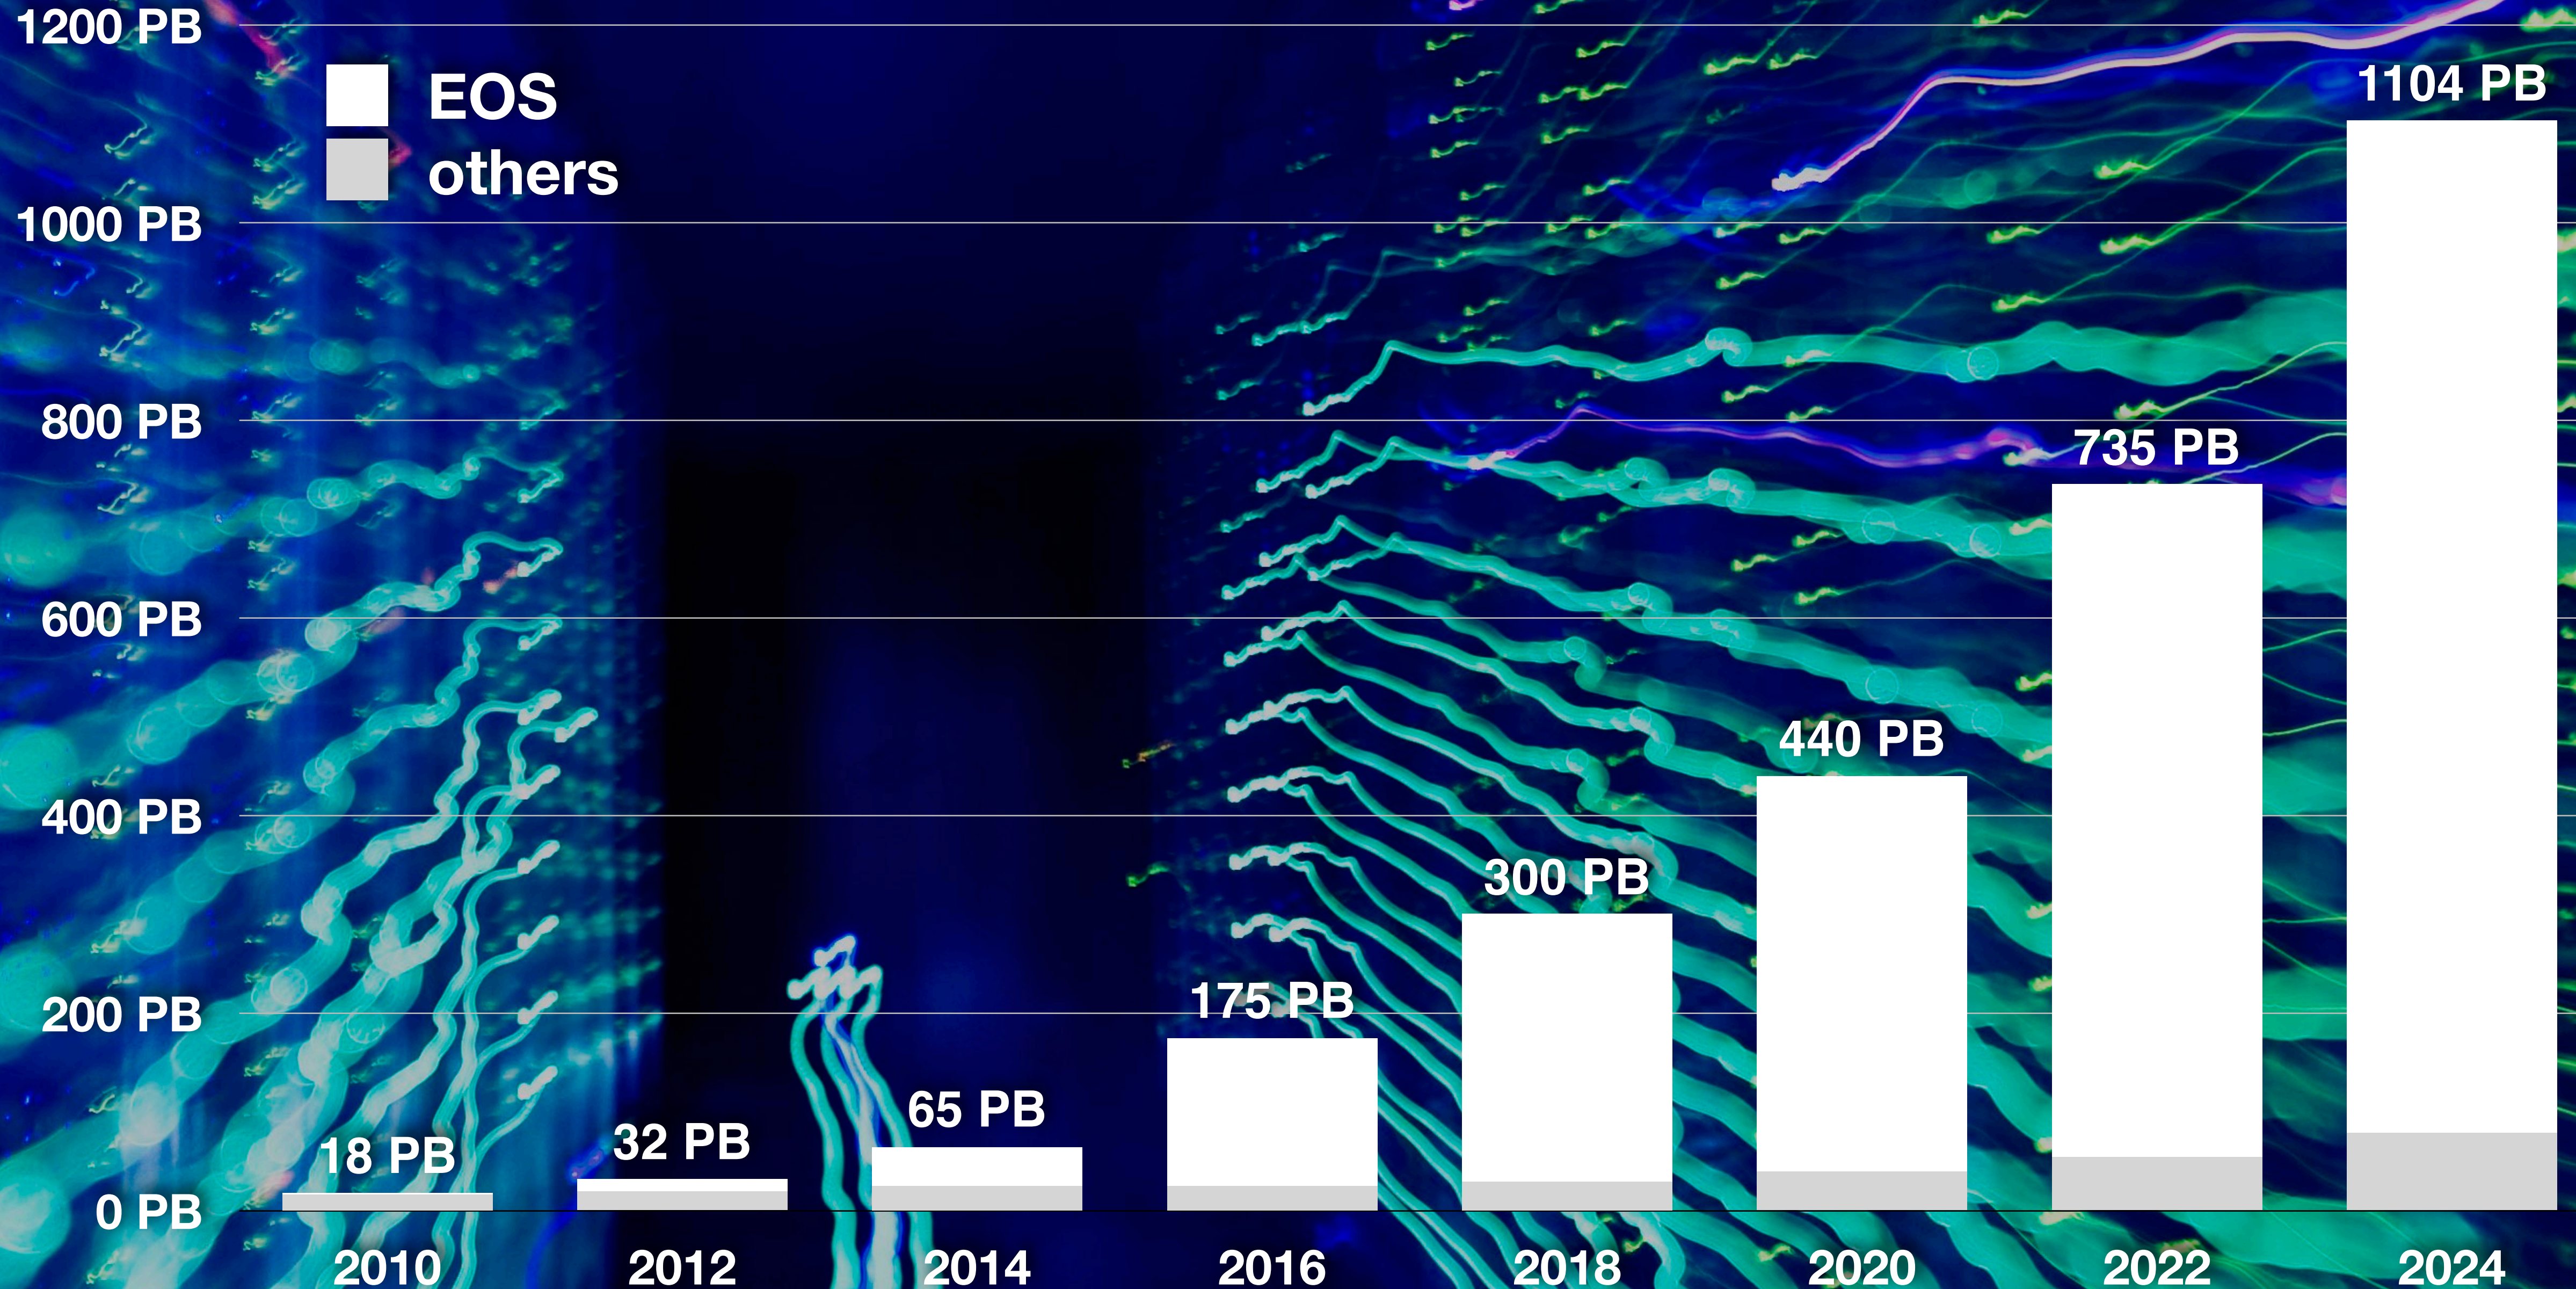
Storage Server in 2024: 2300 TB

Networking Evolution in the last 10Y

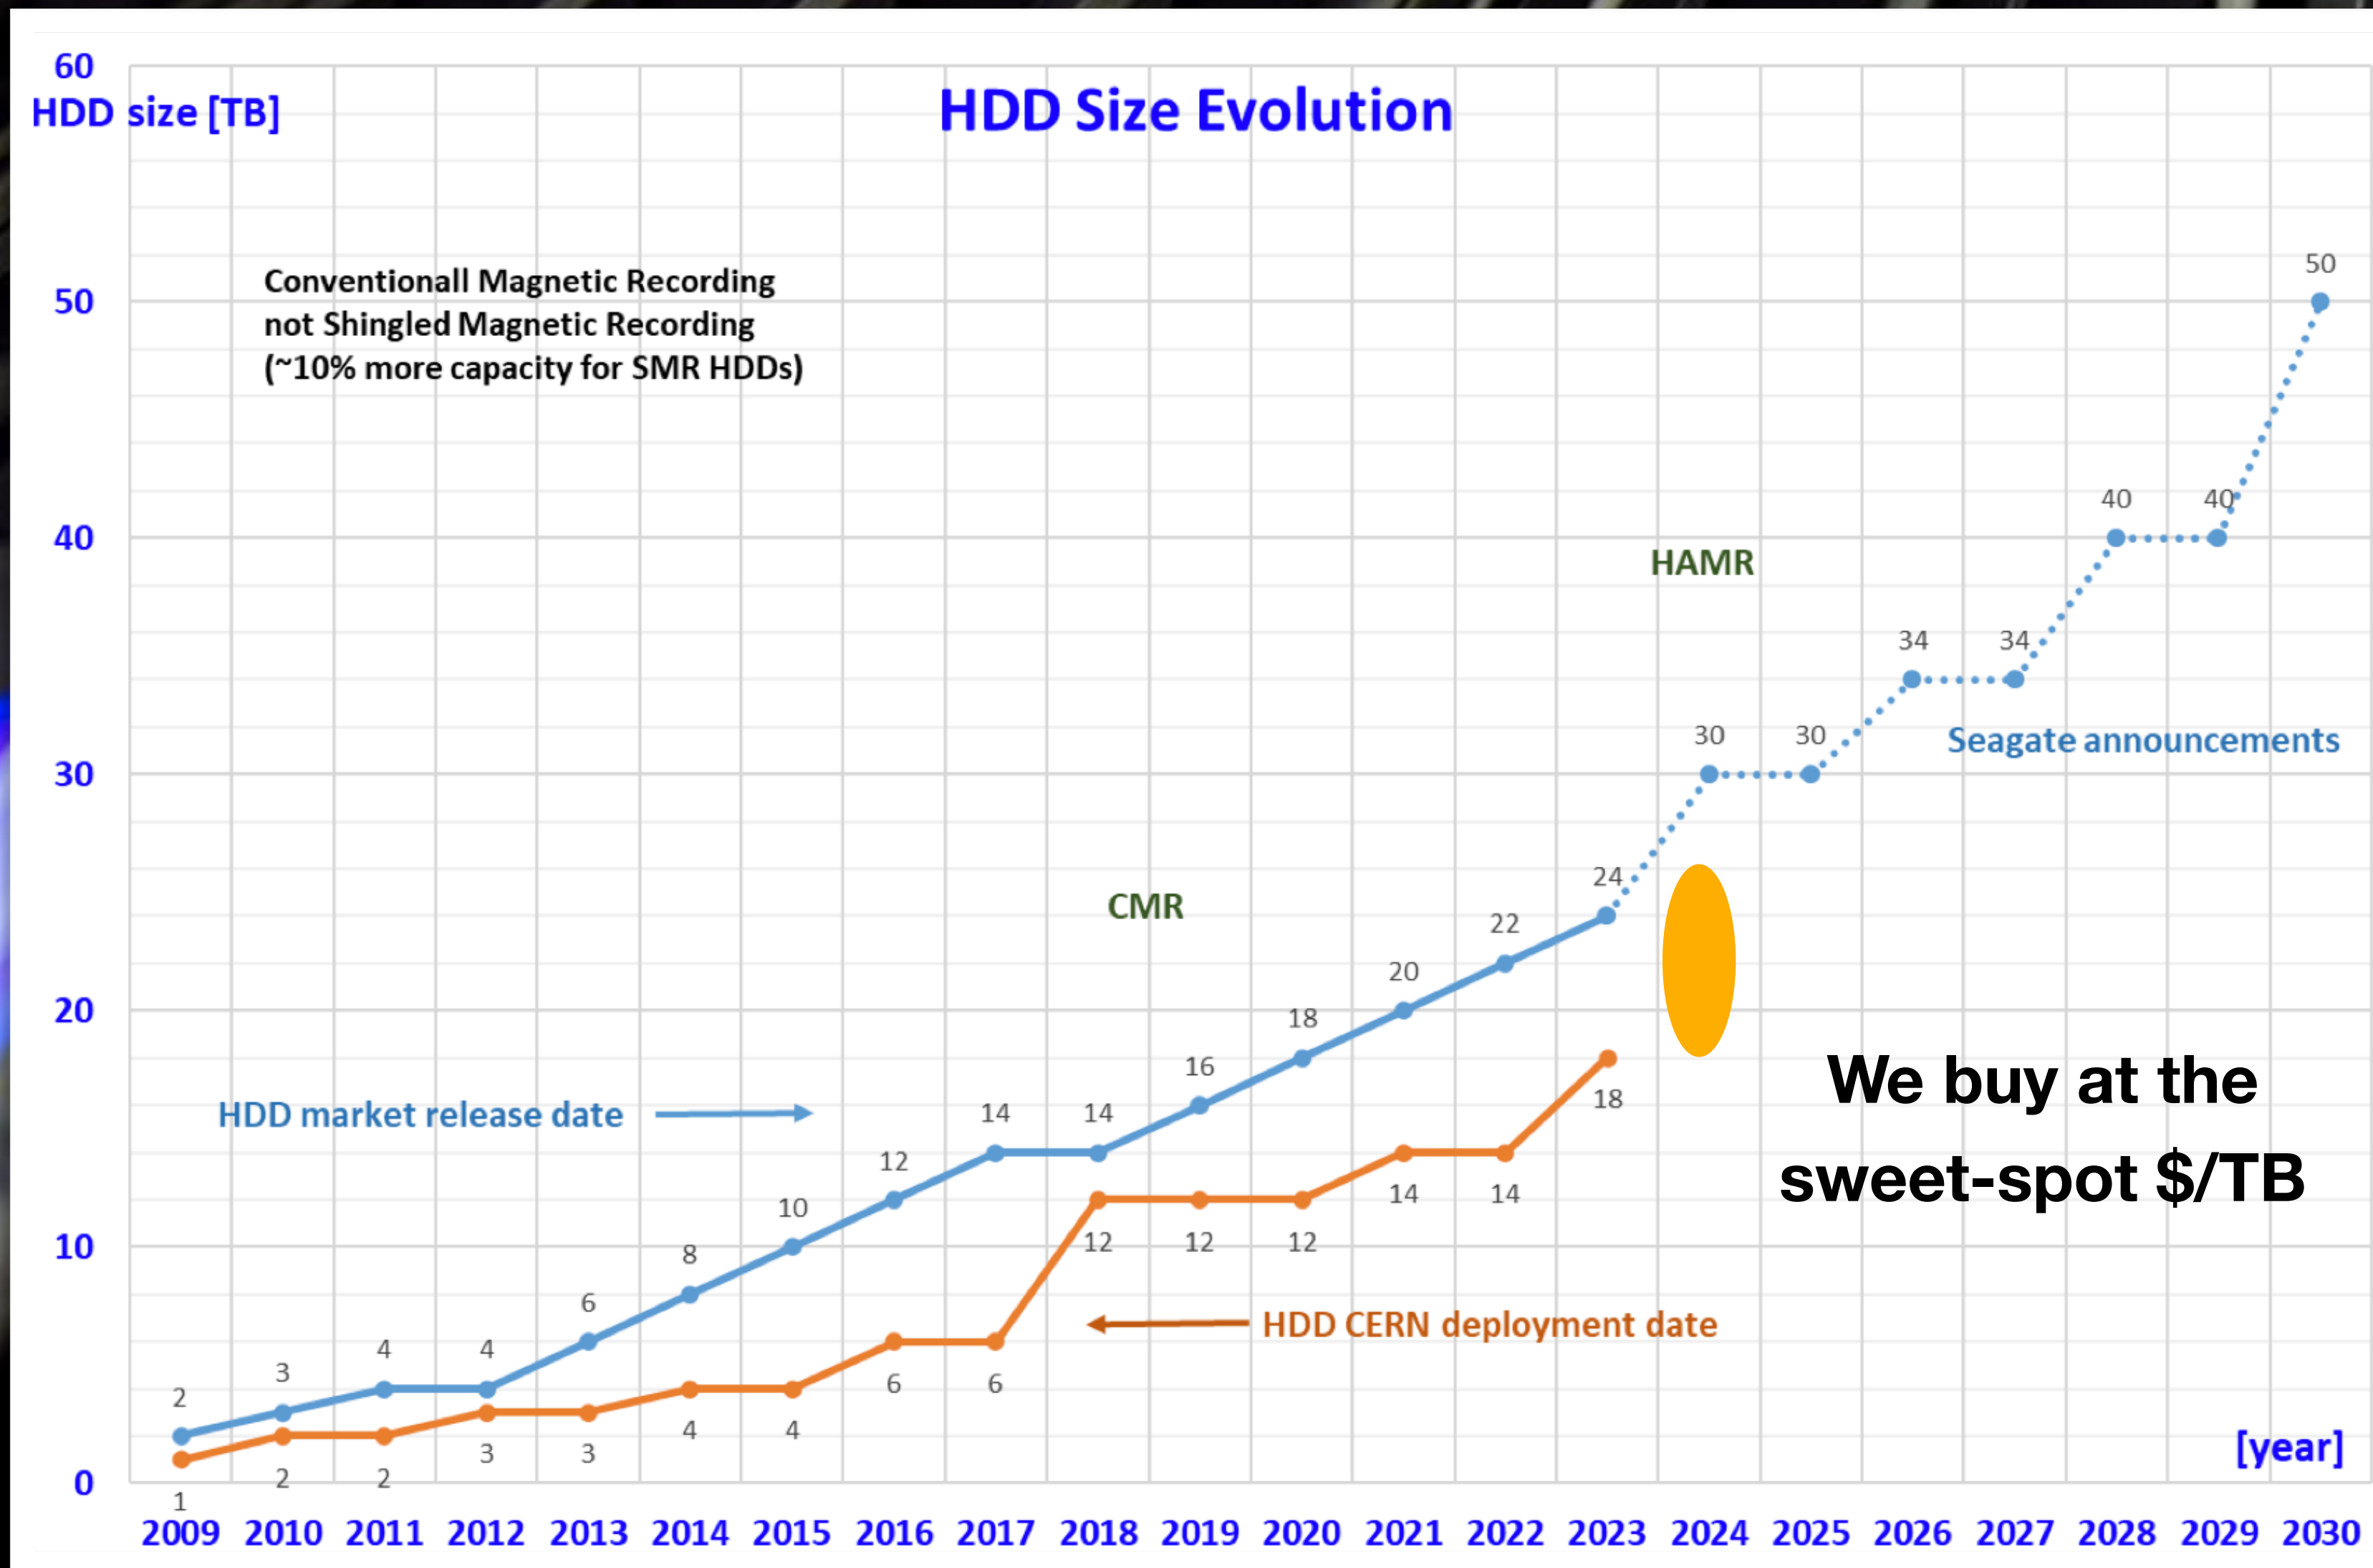
1Gb → 10Gb → 25 Gb → 40Gb → 100Gb

CERN IT - Operated Disk Storage Capacity

■ EOS
■ others



HDD Size Evolution



EOS Physics - Drives Distribution

25000


20000

15000

10000

5000

0



99.6 PB

258 PB

275 PB

169 PB

12.3 PB

3.8 PB

100 TB

180 TB

438 TB

1 TB

2 TB

3 TB

4 TB

6 TB

8 TB

12 TB

14 TB

18 TB

Next Generation Production Drives will be 24 TB


Storage Enclosures

Traditionally we operate 24x SAS/S-ATA drives

- Connection via SAS12
 - 12x SFF-8644 cables at 48Gbps

Lately we are looking into high-density enclosures


- 60 drives
- 108 drives



In the future with SAS24 and PCIe we expect up to 63 GB/s per enclosure



SAS12
cable



HBA 4-ports



Hard Disk Drive


OptiNAND + UltraSMR
HelioSeal technology

EOS @ Point 2



EOS @ Point 2

16 PB installation with
high-density enclosures



EOS @ Meyrin Data Centre




EOS @ Meyrin Data Centre


FRONT

24-bays
storage arrays

BACK



EOS @ Meyrin Data Centre



1/4 of the EOSALICE02 cluster

Future... Hard Disk Drives

Conventional
Magnetic
Recording (CMR)

```
graph LR; CMR[Conventional Magnetic Recording (CMR)] --> SMR[Shingled Magnetic Recording (SMR)]; CMR --> HAMR[Heated Assisted Magnetic Recording (HAMR)]; SMR --- HAMR_PLUS_SMR((HAMR+SMR));
```

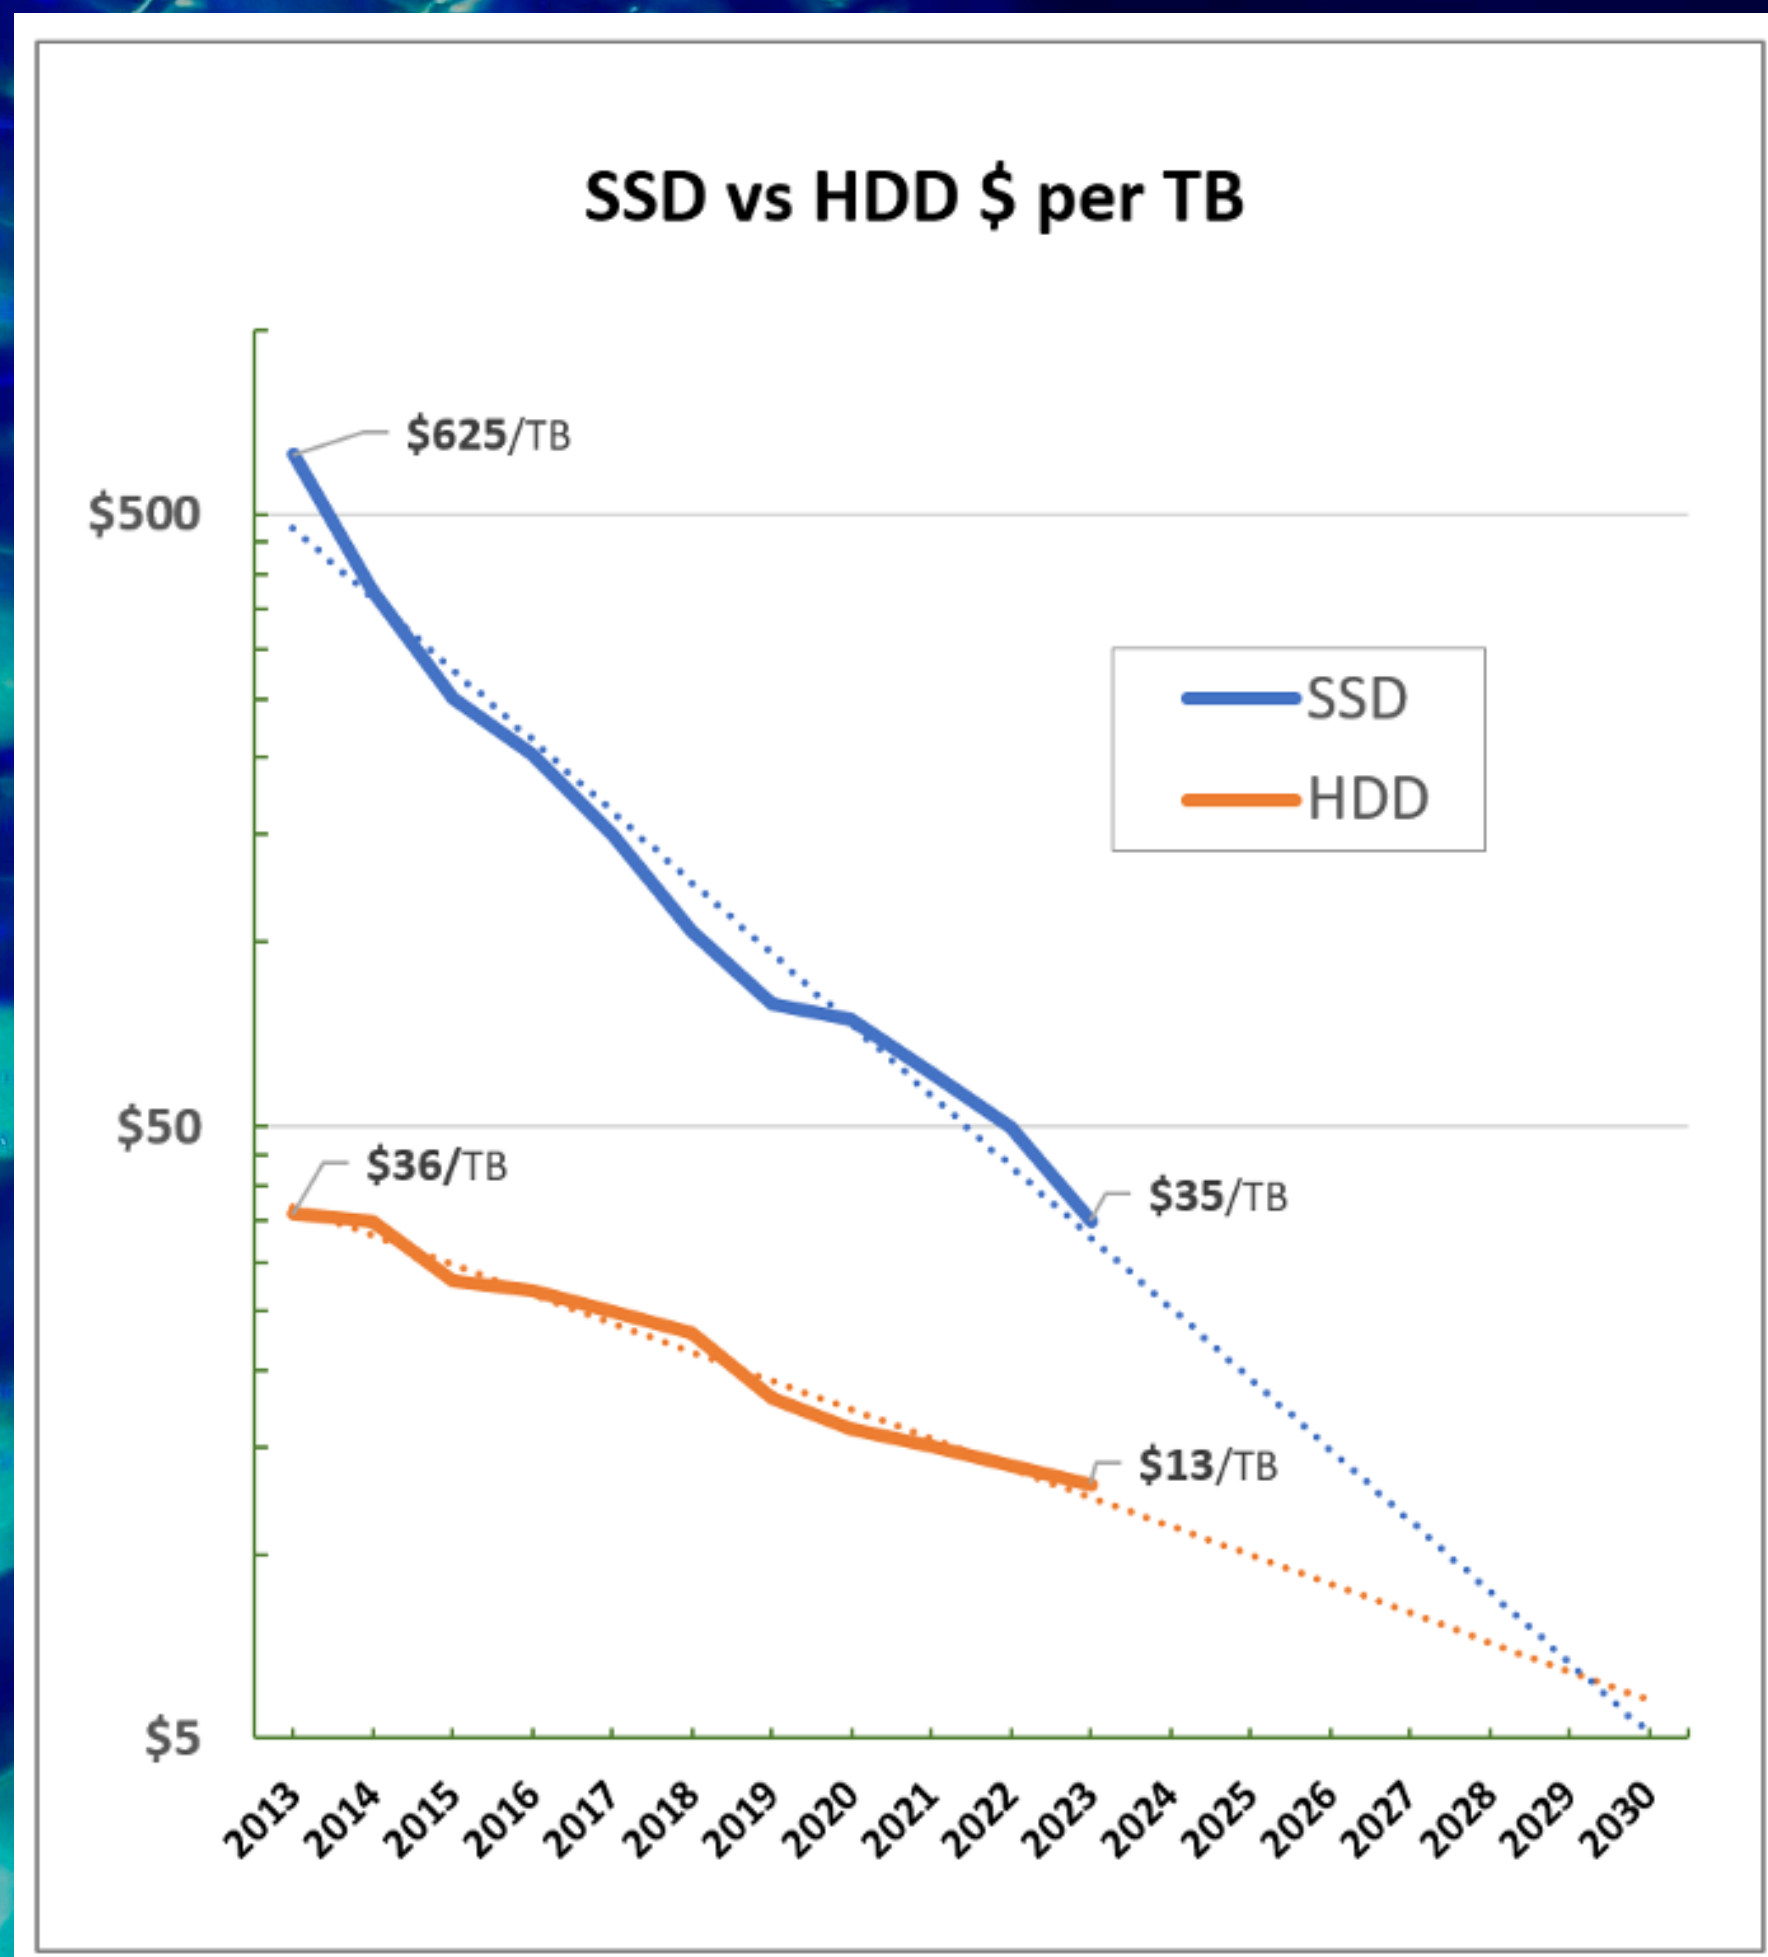
The diagram illustrates the future of hard disk drives, showing the evolution from conventional magnetic recording (CMR) to shingled magnetic recording (SMR), heated assisted magnetic recording (HAMR), and a hybrid HAMR+SMR technology. The background features a vibrant, abstract pattern of blue and green lines, resembling a microscopic view of magnetic tracks or data flow.

Shingled Magnetic
Recording (SMR)

HAMR+SMR

Heated Assisted
Magnetic Recording
(HAMR)

Future... SSD/NVMe roles...



"Toshiba projects that the cost-per-gigabyte of solid-state drives will get lower than that of low-speed (7200rpm), large-capacity nearline hard disk drives (3.5") in about 2025."

- Kitguru.net, quoting Yasuo Naruke, Executive Vice President Toshiba

Toshiba exec claims hard drives are 7X cheaper than SSDs and will continually evolve for large datacenters

News By Aaron Klotz published December 21, 2023

His statement directly counters predictions that hard drives will die in five years



Open Storage

Thanks for the attention!

**Luca Mascetti
CERN IT Department
Storage and Data Management**

