

TOSHIBA

Case Study

How three generations of Toshiba Hard Disk Drives have helped the world's largest physics lab keep track of the hundreds of terabytes of data that are generated by the LHC experiments every second.



supplier

Introduction

CERN's Large Hadron Collider (LHC) is at the forefront of physics research. The data output from its "Run 1" and "Run 2" phases have already been used to demonstrate the existence of a previously undetected subatomic particle and extend our understanding of the universe and how it formed. Notably in 2012 it confirmed the existence of the Higgs boson.

The scale of CERN is astounding. From the size of the large hadron collider – a circular particle accelerator with a radius of 4.3 km – to the rate of particle collisions - up to 1 billion particle collisions can take place every second inside the LHC experiment's detectors.

But it is the data that is most impressive, with the collisions generating 1 petabyte (PB) of data per second. Even after filtering only the interesting events, the facility requires approximately 10PB of new data to be stored for analysis each month.

This data is stored in the CERN Data Centre and is shared with a network of about 170 data centres for analysis, thanks to the Worldwide LHC Computing Grid (WLCG). The current storage setup at CERN consists of HDD buffers with 3,200 JBODs carrying 100,000 hard disk drives (HDDs) providing a total of 350PB.

LHC Runs are set to continue, and with each new "Run", data storage increases significantly. Following upgrades, CERN's Run 3 is scheduled for 2021.


Toshiba Electronics Europe GmbH's hard disk drives are used by CERN to manage huge volumes of data since 2014, with three generations of Toshiba hard drive technology giving it the capacity increases it requires. But, can this continue when, as CERN's Manager of the Facility Planning and Procurement Section at the IT Department, Eric Bonfillou, puts it: "The planned upgrades of the LHC machine will require scaling of compute and storage resources beyond what today's technology can offer."



A timeline of how CERN has scaled its storage with Toshiba

CERN Project Timeline

"So far our IT infrastructure, in terms of computing power and storage capacity, has scaled well with the scientific needs, making maximum use of Toshiba's high capacity and reliable hard disk drives." (Eric Bonfillou, CERN)



2014 - CERN purchases Toshiba hard drives

Since 2014, CERN uses Toshiba's hard drives. At the heart of CERN's needs are server-grade HDDs with high performance, high reliability and optimised for highest capacities. During 2013-14, CERN held a scheduled shutdown phase to prepare for Run 2. During this period, it upgraded its storage systems, adding arrays of 24-bay 4U JBODs populated with Toshiba's MG03SCA400 hard drives.

With 4TB of capacity per HDD, each spinning at 7200rpm and receiving data via a 6 GB/s Interface, CERN gained 96TB per JBOD expansion unit. The HDD's mean time to failure (MTTF) was 1.2 Mio hours, which translates to a predicted annualised failure rate (AFR) of 0.72%.

2015/2016 - CERN begins Run 2: Data storage needs increased significantly

With the start of the "Run 2" in 2015, the need for data storage increased immensely.

CERN adds new storage capacity using Toshiba's new 6TB SATA model, the MG04ACA600E. Again, installed in a 24 Bay 4U Front loader JBOD, the total raw capacity was increased to 144TB per unit.

The MTTF for this model improved to 1.4 Mio hours, giving a predicted AFR of 0.62%.

2018 - Toshiba launches MG07 helium-filled model to extend capacity to 14TB per HDD

After 2016, Toshiba accelerated the development and introduction of new high capacity enterprise HDD models to serve the world's need for cloud data storage - launching two new generations - the MG05 and MG06. The MG06 series became available with capacities up to 10TB per drive. And although the construction was still air-filled, Toshiba was able to increase the MTTF to 2.5 Mio hours, which translates into an AFR of 0.35%, the lowest currently on the market for air-filled HDDs.

For Conventional Magnetic Recording (CMR) technology in a 3.5" form-factor, air-filled HDD technology reaches a limit at 7 platters (and therefore 10TB) of capacity.

And while it is theoretically possible to increase this using thinner platters the air would be too heavy a gas, resulting in severe fluttering. Alternatively, Shingled Magnetic Recording (SMR) technology can be used, but this requires special handling to avoid serious performance issues, and is limited by the need for special file systems tailored to the technology's intricacies.

CERN's 2018 storage upgrade coincided with the launch of Toshiba's MG07 range, the world's first server HDDs with up to 14TB. With the purchase of the 12TB variant, CERN doubled its capacity per JBOD - to 288TB.

The MG07 increases the capacity without changing the form factor by using helium rather than air. This allows the use of thinner platters without the associated flatterring, with up to nine per drive.

By using CMR, the MG07 is suited to any workload without the loss of performance associated with SMR. Additionally, helium

creates less friction than air, significantly reducing the energy required to rotate the platter stack. With precise optimization of the spindle motor, the operating power has reduced by a third (from c.11W to <7W) for the helium-based model of the MG07 series.

The MG07's MTTF significantly extended to 2.5 Mio hours.

The increased data capacity demands from Run 3

In 2019, the LHC machine shut down again to install further upgrades before restarting for Run 3 in 2021. A significant increase in storage demands is expected, with the acceleration in data generated.

As Eric Bonfillou, states: "Toshiba's products and support have met CERN's stringent requirements. Our IT infrastructure, in terms of computing power and storage capacity, has scaled well with the scientific computing needs, making maximum use of Toshiba's high capacity and reliable hard disk drives."

Toshiba's planned launches for CMR (conventional magnetic recording) and SMR (shingled magnetic recording) based drives that use the same 3.5" form factor will give CERN access to 16TB and 18TB drives adding 432 TB of new capacity per JBOD.

"Toshiba Products are well suited for large scale datacenter storage and the successful deployment and operation of three generations of Toshiba Enterprise HDDs in CERN's challenging IT environment is a perfect reference case" says Larry Martinez-Palomo, General Manager of the HDD Business Unit at Toshiba Electronics Europe.

As for longer term R&D activity, Toshiba is developing a next-generation magnetic recording technology that will further extend capacities to exceed 20TB per HDD, while still maintaining the 3.5" form factor.

Martinez-Palomo adding: "We are confident that our HDD next generation technologies will contribute to solve CERNs future scaling challenges in terms of storage capacity, investment budget, power consumption and reliability."



Toshiba Electronics Europe GmbH

Hansaallee 181
40549 Düsseldorf
Germany

info@toshiba-storage.com
toshiba-storage.com

Copyright © 2019 Toshiba Electronics Europe GmbH. All rights reserved. Product specifications, configurations, prices and component / options availability are all subject to change without notice. Product design, specifications and colours are subject to change without notice and may vary from those shown. Errors and omissions excepted.