

Storage for the LHC: operations during Run 3

Thursday 3 November 2022 12:05 (25 minutes)

The CERN IT Storage group operates multiple distributed storage systems to support all CERN data storage requirements. Among these requirements, the storage and distribution of physics data generated by LHC and non-LHC experiments is one of the biggest challenges the group has to take on during Run-3.

EOS, the CERN distributed disk storage system is playing a key role in LHC data-taking. Since the beginning of the year 2022, more than 380PB have been written by the experiments and 2.4EB have been read so far. With the start of Run-3, the requirements in terms of data storage will be higher than what has been delivered so far.

The year 2022 was also marked by the decommissioning of the CASTOR service and its successful replacement by the CERN Tape Archive (CTA). This distributed tape storage system offers low-latency tape file archival and retrieval. CTA currently stores over 375 PB and this is expected to increase to more than 1 EB during Run-3.

The large-scale distribution of the data of the LHC stored on EOS and CTA across the WLCG is mainly ensured by the File Transfer Service (FTS). It offers a reliable, flexible and smart way of initiating data transfers and managing data archives between different storage endpoints all around the world.

In this presentation, we will show how all these different components interact with each other and the architecture and workflows in place to deliver the Run-3 resources in terms of data storage and provision.

Desired slot length

Speaker release

Yes

Presentation will be held...

in the conference venue

Primary author: CAFFY, Cedric (CERN)

Presenter: CAFFY, Cedric (CERN)

Session Classification: Storage and Filesystems

Track Classification: Storage & Filesystems