



Contribution ID: 18

Type: not specified

Enabling LHC Run 3 data storage workflows at CERN

Tuesday 16 April 2024 09:25 (25 minutes)

CERN Storage and Data Management group is responsible for ensuring that all data produced by physics experiments at CERN is safely stored and reliably accessible by the user community. 2023 Run-3 and especially the Heavy Ion Run have pushed further the previous records in terms of data volume and transfer rates delivered by the main LHC experiments. The targets anticipated by the data management coordinators were successfully accommodated by the main storage solutions provided by the Storage group, namely: EOS, CTA (tape storage) and FTS (data transfer orchestration). The EOS service is the main entry point for all data acquisition workflows and has demonstrated reliable operation and excellent peak performance throughout Run 3. In practice, this means that target rates of 10-20 GB/s were regularly surpassed with peaks of 25-30 GB/s for CMS, ATLAS and ALICE experiments.

Distribution of all this data as well as storing the custodial copy on tape required the orchestration capabilities of FTS that successfully met user's expectations and also ensured a good utilization of the tape infrastructure. The most demanding use-case was represented by ALICEO2 which achieved data rates of over 150GB/s using erasure encoded layouts and its own workflow for data distribution. In this presentation, we go over the archived rates and general performance for both the disk and the tape services. Looking towards the restart of Run-3, we detail foreseen challenges and lessons learned during this exceptional period of data taking.

Desired slot length

15

Speaker release

Yes

Author: SINDRILARU, Elvin Alin (CERN)

Presenter: SINDRILARU, Elvin Alin (CERN)

Session Classification: Storage and file systems

Track Classification: Storage & Filesystems