



Contribution ID: 148

Type: **Poster presentation**

Streamlining CASTOR to manage the LHC data torrent

Monday, 14 October 2013 15:00 (45 minutes)

This contribution describes the evolution of the main CERN storage system, CASTOR, as it manages the bulk data stream of the LHC and other CERN experiments, achieving nearly 100 PB of stored data by the end of LHC Run 1.

Over the course of 2012 the CASTOR service has addressed the Tier-0 data management requirements, focusing on a tape-backed archive solution, ensuring smooth operations of the required experiments' workflow (data taking, reconstruction, export to other sites), and guaranteeing data safety by enforcing strong authentication. This evolution was marked by the introduction of policies to optimize the tape sub-system throughput, going towards a cold storage system where data placement is managed by the experiments' production managers. More efficient tape migrations and recalls have been implemented and deployed where bulk metadata operations greatly reduce the overhead due to small files. A repack facility is now integrated in the system and it has been enhanced in order to automate the repacking of several tens of petabytes, required in 2014 in order to prepare for the next LHC run. Finally the scheduling system has been evolved to integrate the internal monitoring for a more efficient dynamic usage of the underlying disk resources.

To efficiently manage the service a solid monitoring infrastructure is required, able to analyze the logs produced by the different components (~1 kHz of log messages). A new system has been developed and deployed, which uses a transport messaging layer provided by the CERN-IT Agile Infrastructure and exploits technologies including Hadoop and HBase. This enables efficient data mining by making use of MapReduce techniques, and real-time data aggregation and visualization.

We will also present the outlook for the future. A further simplification of the CASTOR system is foreseen: xrootd will take the role of the main native protocol, and only few disk pools for each experiment would serve as staging areas for the Tier-0 workflow. The adoption of xrootd as main protocol opens the possibility to exploit its upcoming features, including for instance the support of xroot federations or the HTTP protocol. Directions and possible evolutions will be discussed in view of the restart of data taking activities.

Primary author: Dr LO PRESTI, Giuseppe (CERN)

Co-authors: IERI, Andrea; FIORINI, Benjamin (Centre National de la Recherche Scientifique (FR)); SINDRILARU, Elvin Alin (CERN); CANO, Eric (CERN); PONCE, Sebastien (CERN); MURRAY, Steven (CERN); Dr ESPINAL CURULL, Xavier (CERN)

Presenter: Dr ESPINAL CURULL, Xavier (CERN)

Session Classification: Poster presentations

Track Classification: Data Stores, Data Bases, and Storage Systems