



Contribution ID: 489

Type: **Parallel**

Overview of storage operations at CERN

Monday, 21 May 2012 14:45 (25 minutes)

Large-volume physics data storage at CERN is based on two services, CASTOR and EOS:

* CASTOR - in production for many years - now handles the Tier0 activities (including WAN data distribution), as well as all tape-backed data;

* EOS - in production since 2011 - supports the fast-growing need for high-performance low-latency (i.e. diskonly) data access for user analysis.

In 2011, a large part of the original CASTOR storage has been migrated into EOS, which grew from the original testbed installation (1 PB usable capacity for ATLAS) to over 6 PB for ATLAS and CMS. EOS has been validated for several months under production conditions and has already replaced several CASTOR service classes.

CASTOR also evolved during this time with major improvements in critical areas - notably the internal scheduling of requests, the simplifications of the database structure and a complete overhaul of the tape subsystem.

The talk will compare the two systems from an operation's perspective (setup, day-by-day user support, upgrades, resilience to common failures) while taking into account their different scope. In the case of CASTOR we will analyse the impact of the 2011 improvements on delivering Tier0 services, while for EOS we will focus on the steps to achieve to a production-quality service.

For both systems, the upcoming changes will be discussed in relation with the evolution of the LHC programme and computing models (data volumes, access patterns, relations among computing sites).

Primary authors: IVEN, Jan (CERN); LAMANNA, Massimo (CERN)

Co-authors: Mr PETERS, Andreas Joachim (CERN); Dr LO PRESTI, Giuseppe (CERN); Mr REGUERO, Ignacio (CERN); HEFFERMAN, John (CERN); MASCETTI, Luca (CERN); PONCE, Sebastien (CERN)

Presenters: IVEN, Jan (CERN); LAMANNA, Massimo (CERN)

Session Classification: Computer Facilities, Production Grids and Networking

Track Classification: Computer Facilities, Production Grids and Networking (track 4)