

Scalla As A Full-Fledged LHC Grid SE

Tuesday, March 24, 2009 2:00 PM (20 minutes)

Scalla (also known as xrootd) is quickly becoming a significant part of LHC data analysis as a stand-alone clustered data server (US Atlas T2 and CERN Analysis Farm), globally clustered data sharing framework (ALICE), and an integral part of PROOF-base analysis (multiple experiments). Until recently, xrootd did not fit well in the LHC Grid infrastructure as a Storage Element (SE) largely because it did not provide Storage Resource Manager (SRM) access; making it difficult to justify wider deployment. However, Scalla's extensible plug-in event-based architecture as well as its storage system view, were instrumental in integrating SRM access, using BestMan and Linux/FUSE, in a relatively short time. Today, Scalla can provide full SRM functionality in a way that is independent of the data access system itself. This makes it trivial to migrate to newer versions or run multiple versions of the SRM as they become available.

This paper discusses the architectural elements of the SRM-Scalla integration, where new code had to be written, how key SRM features (e.g., static space tokens and Grid quotas) leveraged the pre-existing Scalla infrastructure, and the overall flow of data as well as management information through system. We also discuss a side-effect of this effort, called xrootdFS, which offers a single file system view of an xrootd cluster; including its performance and what needs to be done to improve it.

Primary authors: SIMS, Alex (LBNL); Mrs HANUSHEVSKY, Andrew (SLAC National Accelerator Laboratory); FURANO, Fabrizio (Conseil Europeen Recherche Nucl. (CERN)); Mr YANG, Wei (SLAC National Accelerator Laboratory)

Presenter: Mrs HANUSHEVSKY, Andrew (SLAC National Accelerator Laboratory)

Session Classification: Distributed Processing and Analysis

Track Classification: Distributed Processing and Analysis