

XrootdFS, a posix file system for XrootD

Thursday, 13 October 2016 16:30 (15 minutes)

When we first introduced XRootD storage system to the LHC, we needed a filesystem interface so that XRootD system could function as a Grid Storage Element. The result was XRootDfs, a FUSE based mountable posix filesystem. It glues all the data servers in a XRootD storage system together and presents it as a single, posix compliant, multi-user networked filesystem. XRootD's unique redirection mechanism requires special handling of IO operations and metadata operations in the XRootDfs. This includes a throttling mechanism to gracefully handle extreme metadata operations; handling of returned results from all data servers in a consistent way; hiding delays of metadata operations, including storage media latency; enhancing the performance of concurrent IO by multiple applications; and using an advanced security plugin to ensure secure data access in a multi-user environment. Over the last several years XRootDfs have been adopted by many XRootD sites for data management as well as data access by applications that were not specifically designed to use the native XRootD interface. Many of the technical methods mentioned above can also be used to glue together other types (i.e. non-XRootD) data servers to provide seamless data access.

Primary Keyword (Mandatory)

Storage systems

Secondary Keyword (Optional)

Distributed data handling

Tertiary Keyword (Optional)

Primary author: YANG, Wei (SLAC National Accelerator Laboratory (US))

Co-author: HANUSHEVSKY, Andrew Bohdan (SLAC National Accelerator Laboratory (US))

Presenters: HANUSHEVSKY, Andrew Bohdan (SLAC National Accelerator Laboratory (US)); YANG, Wei (SLAC National Accelerator Laboratory (US))

Session Classification: Posters B / Break

Track Classification: Track 4: Data Handling