



Contribution ID: 173

Type: Poster

Evaluation of software based redundancy algorithms for the EOS storage system at CERN

Tuesday, May 22, 2012 1:30 PM (4h 45m)

EOS is a new disk based storage system used in production at CERN since autumn 2011. It is implemented using the plug-in architecture of the XRootD software framework and allows remote file access via XRootD protocol or POSIX-like file access via FUSE mounting. EOS was designed to fulfill specific requirements of disk storage scalability and IO scheduling performance for LHC analysis use cases. This is achieved by following a strategy of decoupling disk and tape storage as individual storage systems. A key point of the EOS design is to provide high availability and redundancy of files via a software implementation which uses disk-only storage systems without hardware RAID arrays. All this is aimed at reducing the overall cost of the system and also simplifying the operational procedures. This paper presents advantages and disadvantages of redundancy by hardware (most classical storage installations) in comparison to redundancy by software. The latter is implemented in the EOS system and achieves its goal by spawning data and parity stripes via remote file access over nodes. The gain in redundancy and reliability comes with a trade-off in the following areas:

- Increased complexity of the network connectivity
- CPU intensive parity computations during file creation and recovery
- Performance loss through remote disk coupling

An evaluation and performance figures of several redundancy algorithms are presented for simple file mirroring, dual parity RAID, Reed-Solomon and LDPC codecs. Moreover, the characteristics and applicability of these algorithms are discussed in the context of reliable data storage systems. Finally, a summary of the current state of implementation is given, sharing some experiences on migration and operation of a new multi-PB disk storage system at CERN.

Primary authors: Dr PETERS, Andreas (CERN); Mr SINDRILARU, Elvin Alin (CERN)

Presenter: Dr PETERS, Andreas (CERN)

Session Classification: Poster Session

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