



Contribution ID: 501

Type: poster presentation

Building a bridge between cloud storage and GPFS

The possibilities of cloud storage for use in HEP computing have been the topic of many studies and trials. The typical cloud storage values, easily accessible and expandable, relatively cheap and with a light weight interface have become available for local storage as well. Initially as part of larger environments like Open Nebula or OpenStack Swift, vendors now offer value storage with integrated object based interfaces. Although cloud storage itself can be accessed with little effort, the integration with existing storage is done at the application level. However, this approach ties the storage to the application and more importantly, is not usable for all applications because these require regular POSIX semantics. The German Tier 1 GridKa as well as the Large Scale Data Facility (LSDF) at KIT, uses GPFS as the underlying file system for most of its storage. In recent versions GPFS offers a built-in rule and policy engine that allows complex data management operations including migrating data to external storage systems. The policy engine can drive hierarchical storage management (HSM) in which a condition that matches a defined characteristic triggers an external event such as the copy of a file or the movement of files from and to the file system. Interfaced to GPFS, cloud storage has the potential to be used transparently for all applications.

In this paper we present the first results of experiments where data is migrated from and recalled to a GPFS system that serves as storage for an xrootd disk pool. The data moves between the GPFS and local cloud storage and as a result, the xrootd system operates transparently on disk-based virtual storage that is several times larger than the actual file system. The results of the presented solution are compared with the traditional disk to tape HSM in use at GridKa. The experiences with the automated migration will be very helpful for the design of large and cost effective storage systems.

Author: PETZOLD, Andreas (KIT - Karlsruhe Institute of Technology (DE))

Co-author: VAN WEZEL, Jos (KIT - Karlsruhe Institute of Technology (DE))

Presenter: PETZOLD, Andreas (KIT - Karlsruhe Institute of Technology (DE))

Track Classification: Track7: Clouds and virtualization