

Managing small files in Mass Storage systems using Virtual Volumes

Wednesday 15 February 2006 14:40 (20 minutes)

Efficient hierarchical storage management of small size files continues to be a challenge. Storing such files directly on tape-based tertiary storage leads to extremely low operational efficiencies. Commercial tape virtualization products are few, expensive and only proven in mainframe environments. Asking the users to deal with the problem by “bundling” their files leads to a plethora of solutions with high maintenance costs. Part of the problem is that data processing environments have evolved towards the illusion of an infinite file store with a subdirectory structure, eliminating the concept of a volume, be it physical or logical. Research has been undertaken to deal with these issues at the data center level, but the outcome is quite simple and can be used in general. Results are presented of prototype implementations of a paradigm termed “Virtual Volumes”, which combines standard operating system tools such as symbolic links and auto-mounters together with techniques to represent a volume as a file such as the ISO 9660 specification. Virtual Volumes allow large number of files to be handled as a single item in tertiary tape storage systems, whilst maintaining the infinite file store illusion towards the user by mounting these single items as branches within a file system. Whereas a totally general implementation of Virtual Volumes would require quite complex coding, the prototypes presented are optimized for the Write-Once-Read-Many (WORM) environment often found in scientific data applications. The choice to base the Virtual Volume implementation on standard operating system tools and techniques means that it can be easily combined with already existing or future tools used in HEP experiments and Grid Infrastructures. Examples are given of handling HEP data using Virtual Volumes integrated into the existing data frameworks of several experiments.

Primary author: Prof. DELFINO REZNICEK, Manuel (Port d’Informació Científica)

Co-authors: Dr PACHECO, Andreu (Institut de Física d’Altes Energies); Prof. HERNÁNDEZ, Emilio (Universidad Simón Bolívar); Ms ACCIÓN, Esther (Port d’Informació Científica)

Presenter: Prof. DELFINO REZNICEK, Manuel (Port d’Informació Científica)

Session Classification: Computing Facilities and Networking

Track Classification: Computing Facilities and Networking