Contribution ID: **400** Type: **poster**

ScratchFS: A File System To Manage Scratch Disk Space For Grid Jobs.

Wednesday, February 15, 2006 9:00 AM (20 minutes)

Managing the temporal disk space used by jobs in a farm can be an operational issue. Efforts have been put on controlling this space by the batch scheduler to make sure the job will use at most the requested amount of space, and that this space is cleaned up after the end of the job. ScratchFS is a virtual file system that addresses this problem for grid as well as conventional jobs at the file system level. Designed to be on top of a real file system, the virtualization layer is responsible for controlling the access to the underlying file system. For conventional (i.e. local) jobs, the access control to the disk space is based on standard Unix file permissions. In the case of grid jobs, ScratchFS uses only the grid credentials of the job to grant or deny access to the scratch disk space in a fully transparent way from the application point of view. Therefore, several grid jobs submitted by different individuals can be executed in the same worker node, even if their grid identities are all mapped to the same local user identifier, while preserving the confidentiality of each job scratch space. Designed to be extensible, it can be configured to collect information about the file system usage and to enforce quotas, among other things. In addition, a simplified interface is provided to ease the integration with a batch scheduler. In this paper we present how this virtualization layer is used by our local batch scheduler and we explain the implementation of the system in detail, along with some comments on its security and benchmark results compared to a real file system.

Primary author: Mr FRANCO, Leandro (IN2P3/CNRS Computing Centre)

Co-author: Mr HERNANDEZ, Fabio (IN2P3/CNRS Computing Centre)

Presenter: Mr FRANCO, Leandro (IN2P3/CNRS Computing Centre)

Session Classification: Poster

Track Classification: Grid middleware and e-Infrastructure operation