Contribution ID: **201** Type: **oral presentation**

Implementation of a reliable and expandable on-line storage for compute clusters

Monday, September 27, 2004 4:50 PM (20 minutes)

The HEP experiments that use the regional center GridKa will handle large amounts of data. Traditional access methods via local disks or large network storage servers show limitations in size, throughput or data management flexibility.

High speed interconnects like Fibre Channel, iSCSI or Infiniband as well as parallel file systems are becoming increasingly important in large cluster installations to offer the scalable size and throughput needed for PetaByte storage. At the same time the reliable and proven NFS protocol allows local area storage access via traditional Ethernet very cost effectively.

The cluster at GridKa uses the General Parallel File System (GPFS) on a 20 node file server farm that connects to over 1000 FC disks via a Storage Area Network. The 130 TB on-line storage is distributed to the 390 node cluster via NFS. A load balancing system ensures an even load distribution and additionally allows for on-line file server exchange.

Discussed are the components of the storage area network, specific Linux tools, and the construction and optimisation of the cluster file system along with the RAID groups. A high availability is obtained and measurements prove high throughput under different conditions. The use of the file system administration and management possibilities is presented as is the implementation and effectiveness of the load balancing system.

Primary author: VANWEZEL, J. (FORSCHUNGZENTRUM KARLSRUHE)

Presenter: VANWEZEL, J. (FORSCHUNGZENTRUM KARLSRUHE)

Session Classification: Computer Fabrics

Track Classification: Track 6 - Computer Fabrics