



Contribution ID: 57

Type: **not specified**

## Update on Chimera and NFS v4.1

*Tuesday, 6 November 2007 16:00 (30 minutes)*

With the start of the Large Hadron Collider, presumably in 2008, the largest share of the produced and generated physics data outside CERN will be managed by the dCache storage element. Beside the challenge to handle storage in the order of tens of Petabyte in a single installation, the number of file operations per second is certainly even more challenging. Two years ago we identified PNFS, the dCache file system component, as to be the part which may become a bottleneck at the point in time when the global LHC data chain will be operated at its full capacity. There are various reasons for this assumption. Due to the fact, that pnfs existed before dCache was designed, dCache and pnfs interact by means of the nfs2 protocol, which introduces an nonessential layer of complexity. Moreover, although pnfs is using the Postgres database, it doesn't make use of enhanced modern database functionalities. Another striking issue is certainly the fact that single read/write locks are used internally, protecting large collection of file which leads to unnecessary queuing of file operation requests. To circumvent those and many other issues we designed and implemented Chimera, a file system name space engine based on the Java programming language as well as on modern database technologies. Furthermore, Chimera is an inevitable prerequisite in dCache to support the NFS 4.1 protocol. NFS 4.1 is currently in its final state of specification. As to our current plans NFS 4.1 will be supported in dCache in the foreseeable future. This presentation will give more insight in the advantages of Chimera and the NFS 4.1 plans of the dCache team.

**Primary author:** Dr FUHRMANN, Patrick (DESY)**Presenter:** Mr VAN DER REEST, Peter (DESY - IT)**Session Classification:** Storage II**Track Classification:** Storage