

## **Investigation of storage options for scientific computing on Grid and Cloud facilities**

In recent years, several new storage technologies, such as Lustre, Hadoop, and BlueArc, have emerged. While several groups have run benchmarks to characterize them under a variety of configurations, more work is needed to evaluate these technologies for the use cases of scientific computing on Grid clusters and Cloud facilities. This paper discusses our evaluation of the technologies as deployed on a test bed at FermiCloud, one of the Fermilab infrastructure-as-a-service Cloud facilities. The test bed consists of 4 server-class nodes with 40 TB of disk space and up to 50 virtual machine clients, some running on the storage server nodes themselves. With this configuration, the evaluation compares the performance of some of these technologies when deployed on virtual machines and on "bare iron". In addition to running standard benchmarks such as IOZone to check the sanity of our installation, we have run I/O intensive tests using experiment specific applications. This paper presents how the storage solutions perform in a variety of realistic use cases of scientific computing.

Primary authors : Dr. GARZOGLIO, Gabriele (Fermilab) ; Dr. HESSELROTH, Ted (Fermilab) ; Dr. NORMAN, Andrew (Fermilab) ; Dr. PEREVALOV, Denis (Fermilab) ; Mr. STRAIN, Doug (Fermilab) ; Dr. TIMM, Steve (Fermilab)

Co-authors : Dr. CHADWICK, Keith (Fermilab) Presenter : Dr. CHADWICK, Keith (Fermilab)