Testing GlusterFS in a Grid Site for SuperB collaboration

Description of the work

In the proposed poster we show the results of the test obtained by a practical experience made on an high performance hardware. The setup which we used is mostly close to one possible production cluster. The aim of the tests is understand at first the global reliability of this kind of configuration, the performance and the scalability limits of GlusterFS. We proceeded in a systematically way by measuring the network occupation during large data transfer between the nodes, then we measure how the GlusterFS demon affect the CPU of each node. Finally we use standard I/O benchmark to measure the read/write limits.

After general evaluation we measure realistic behaviour to respect the SuperB usage, by sending burst of analysis job over simulated data coming from the previous MonteCarlo production created by the collaboration and store in the GlusterFS area.

Wider impact of this work

The interest for GlusterFS is due to the high performance, the possibility to setup hard disk in JBOD, the standard posix interface and the auto replication feature inter and intra cluster. Moreover, the Gluster community roadmap plans to make available soon new API to individuate the disks that contain a single file, that open the possibility to implement a file affinity scheduling in the future.

Printable Summary

The aim of this work is to present our experience and future plans to use GlusterFS at site level. The goal is to create a storage area that is shared between storage element and worker nodes in a Grid Cluster. GlusterFS is an open source cluster file system that allows creating a single namespace and managing the data redundancy and data replication automatically. In order to evaluate the system performance, we have setup a test infrastructure composed of 12 servers over a 10Gbit/s network. Each node shares its disk space through GlusterFS, and mounts a unique file system as the sum of every single 'brick'. The performed tests show interesting results in terms of performance, CPU usage, Network, I/O and Grid Integration.

The presented activities are performed within the SuperB R&D program to study the future computing model.

Primary author: Prof. PARDI, Silvio (INFN and UNINA)

Track Classification: Resource Infrastructure services (Peter Solagna: track leader)