

THE ALICE STORAGE SYSTEM: an Analysis of the Impact on the Performance of the Configuration Parameters and of the Load of Concurrent Streams

The ALICE data acquisition requires an aggregate event building bandwidth of up to 2,5 GBytes/s and a storage capability of up to 1,25 GBytes/s, giving a total of more than 1 PBytes of data every year. This makes the performance of the mass storage devices a dominant factor for the overall system behavior and throughput.

In this presentation, we present an analysis of the performance of the storage system used in the ALICE experiment by studying the impact of different configuration parameters on the system throughput.

The aim of this analysis is to determine the storage configuration which gives the best system performance. In particular, we show the influence of file and block size on the writing and reading rates and we present the relative performance of a clustered file system (Stornext) and a regular journaled file system (Xfs) based on disk arrays in a Fiber Channel storage area network (SAN) .

We will also present a comparative analysis between different parametrizations of Redundant Arrays of Inexpensive Disks (RAID) varying the parity level and the numbers of disks in each array.

At the end, we will conclude with the aggregate performance when concurrent writing and reading streams are sharing the system.