

Investigating Hard Disk Drive Failure Through Disk Torture

Author: *Daniel Perry*, Cornell University

Mentors: *David Bonnie* and *Garrett Ransom*

Hard disk drives (HDDs) are typically expected to fail after several years of use, with the probability of a disk failing increasing as the disk ages. High performance computing (HPC) storage systems at Los Alamos National Laboratory (LANL) experienced instances where HDDs began failing at rates much higher than those expected at these systems' points in their hardware's lifecycle. In order to determine the causes for these past failures, and to mitigate similar failures on future systems, we began to investigate how external variables such as temperature, vibration, and disk workload affects the health of HDDs. Our test bed for experimentation is a Dell R7425 server connected to four 84 bay JBOD enclosures containing 4 terabyte Seagate HDDs, totaling 336 hard disks. The initial stages of this ongoing project are composed of two main objectives: establishing a monitoring framework to collect HDD failure data and creating artificial workloads to simulate high-stress environments. Our implemented monitoring system collects disk temperature, health, and performance data in addition to temperature data from the enclosures. The collected data is stored in a Prometheus time series database, where it can be accessed through Grafana visualization tools. Development of high-stress disk operation workloads led to an in-depth analysis of the internal physical geometry of the test HDDs, to facilitate the creation of tests focusing a disproportionate workload on one disk head.

LA-UR-20-25704