

Analysis and Workload Characterization of the CERN EOS Storage System

Modern, large-scale scientific computing runs on complex exascale storage systems that support complex data workloads. Understanding the data access and movement patterns is vital for informing the design of future iterations of existing systems and next-generation systems. Yet we are lacking in publicly available traces and tools to help us understand even one system in depth, let alone correlate long-term cross-system trends.

In this work, we investigate the workload characteristics of the CERN EOS filesystem, analyzing over 2.49 billion events containing over 300 PB in reads and 150 PB in writes across 11 months. We contrast our findings with analyses from other scientific storage systems, allowing us to observe larger trends that appear over the years and revisit and question conventional wisdom such as "write once, read maybe" and the influence of user actions on system-wide data movement. By studying trace capture mechanisms across these systems, we motivate a standardized trace collection and analysis toolset, so that future researchers can more easily study existing systems to aid in system design.