Storage Research and Development at Large HEP Computing Sites

Robert Illingworth, Bo Jayatilaka*
Fermilab Scientific Computing Division

Storage for statistical data from the largest HEP experiments (e.g., LHC detectors) currently is achieved using a two-tiered storage system with a primary archive on magnetic tape and nearline access provided by magnetic hard drives. These storage resources can be distributed across multiple physical sites in some cases. Data storage scales per experiment currently reach O(100PB). Both the scale and complexity of the storage requirements for upcoming Particle Physics experiments will be significantly greater than in past generations. The High Luminosity LHC will produce unprecedented volumes of data, while the need to perform processing at HPC centers, machine learning, and the cloud represent a much more heterogeneous environment than previously. Certain workloads - for example, neutrino detector data - will likely benefit from using alternative storage paradigms, such as object stores, instead of the traditional HEP software stack. Data storage scales for these experiments will exceed O(1000 PB) annually in many cases.

The systems used for bulk storage needs to evolve in order to provide archival storage of the massive volumes of data and remaining affordable for both hardware and support costs, while at the same time providing sufficiently fast access for processing. The two-tiered paradigm currently in place was optimized for the technological landscape present at the outset of the LHC, over a decade ago; a landscape which has changed considerably since then. Fermilab currently provides custodial storage for 40% of the CMS data volume as well as most of the Fermilab-based neutrino and muon physics experiments. The storage architecture in place will need to evolve over the next decade to meet these Exabyte-scale demands.The authors of this LOI will be studying the current technological landscape and evaluate architectures and demands and report, in a Snowmass white paper, a roadmap of how to deploy and maintain an efficient storage architecture for Fermilab and similar-sized HEP computing sites into the late 2020s and beyond.

*Contact author (boj@fnal.gov)