

# Data Access Pattern Analysis for dCache Storage System

Amanda Pereira<sup>1</sup>, Alex Sim<sup>2</sup>, John Wu<sup>2</sup>, Shinjae Yoo<sup>3</sup>, Hiro Ito<sup>3</sup>  
Drake University<sup>1</sup>, Lawrence Berkeley National Laboratory<sup>2</sup>, Brookhaven National Laboratory<sup>3</sup>

## I. INTRODUCTION

High-Energy physicists routinely analyze files containing particle collision data for their research. The US CMS collaboration [1] [2] has set up a high performance storage system (HPSS) for their data files and uses the dCache system [3] as a disk cache in front of the HPSS. Users retrieve this data through a request to the dCache system. Upon request, dCache searches the local cache for the file. Files stored locally can be accessed without involving the HPSS. Files that are not on disk need to be retrieved from the HPSS prior to access. Predicting when a file is needed and moving it to disk ahead of time can reduce access time. To enable such predictions, we plan to study the data access patterns of the dCache system. Numerous efforts have been reported on such data access patterns [4], however, these earlier studies primarily focus on individual files. In this work, we plan to focus on file groups that could be more easily managed by the system administrators.

## II. ANALYSIS

Data transfer logs and metadata information were used to determine patterns in the popularity of datasets during previous time periods. Files must first be grouped into datasets according to information found in the file path. Access patterns can be studied at the dataset level for more efficient analysis. Initial groupings are determined by the energy measurement of particles being studied using a unit called teraelectronvolts (TeV). Most files in the metadata information provide access patterns statistics for files containing information about particles with 13TeV of energy. Many of these files contain "mc15" in front of the "13TeV" in their file path. The research focus is narrowed to files with "mc15" and "13TeV" in their path. Over 3/4 of the files in this group have a path that contains "EVNT". Popularity calculation and analysis was done for "EVNT" datasets as a result of the higher access counts. Dataset popularity is calculated by dividing the number of times a dataset was accessed in a given time period by the total number of files in the dataset. The popularity for Mc15 13TeV group datasets containing "EVNT" on April 3-9, 2021 is displayed in figure 1. The peaks in figure 1 represent the more popular "EVNT" datasets during the week.

## III. CONCLUSION

Certain events have a popularity almost 5x that of other events in figure 1. These results suggest that for maximum storage system efficiency, popular event datasets should be stored locally during the time period in which they are popular,




Fig. 1. Mc15 13TeV event dataset popularity April 3, 2021

and moved to the HPSS for storage when they are not popular. The development of tools for past data access pattern analysis is useful for forecasting dataset access patterns in future time periods. Predicting future dataset popularity is necessary for caching the appropriate data and improving the application performance of the dCache system.

## ACKNOWLEDGMENT

This work was supported in part by the Office of Advanced Scientific Computing Research, Office of Science, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231, and also used resources of the National Energy Research Scientific Computing Center (NERSC). The author would like to thank my mentors, Alex Sim and John Wu, at Scientific Data Management (SDM) Group, Lawrence Berkeley National Laboratory. This work was also supported in part by the U.S. Department of Energy, Office of Science, Office of Workforce Development for Teachers and Scientists (WDTS) under the Science Undergraduate Laboratory Internship (SULI) program.

## REFERENCES

- [1] CMS Collaboration, "The cms experiment at the cern lhc," *Journal of Instrumentation*, vol. 3, no. 8, p. S08004, 2008.
- [2] L. Bauerdick, K. Bloom, B. Bockelman, D. Bradley, S. Dasu, J. Dost, I. Sfiligoi, A. Tadel, M. Tadel, F. Wuerthwein *et al.*, "XRootd, disk-based, caching proxy for optimization of data access, data placement and data replication," in *Journal of Physics: Conference Series*, vol. 513, no. 4. IOP Publishing, 2014, p. 042044.
- [3] M. Ernst, P. Fuhrmann, M. Gasthuber, T. Mkrtychyan, and C. Waldman, "dcache, a distributed storage data caching system," 2001.
- [4] Y. Wang, K. Wu, A. Sim, S. Yoo, and S. Misawa, "Access patterns to disk cache for large scientific archive," in *Proceedings of the 2021 on Systems and Network Telemetry and Analytics*, 2020, pp. 37–40.