Predicting Resource Usage Trends with Southern California Petabyte Scale Cache

Caitlin Sim¹, Kesheng Wu², Alex Sim², Inder Monga³, Chin Guok³

Frank Würthwein⁴, Diego Davila⁴, Harvey Newman⁵, Justas Balcas⁵

University of California, Berkeley¹, Lawrence Berkeley National Laboratory², Energy Sciences Network³

University of California at San Diego⁴, California Institute of Technology⁵

caitlinsim@berkeley.edu¹, {kwu,asim}@lbl.gov², {imonga,chin}@es.net³

{fkw,didavila}@ucsd.edu⁴, {newman,jbalcas}@hep.caltech.edu⁵

There has been a significant increase in data volume from various large scientific projects, including the Large Hadron Collider (LHC) experiment. The High Energy Physics (HEP) community requires increased data volume on the network, as the community expects to produce almost thirty times annual data volume between 2018 and 2028 [1]. To mitigate the repetitive data access issue and network overloading, regional data caching mechanism [2], [3], or in-network cache has been deployed in Southern California for the US CMS, and its effectiveness has been studied [4], [5]. With the number of redundant data transfers over the wide-area network decreasing, the caching approach improves overall application performance as well as network traffic savings.

In this work, we examined the trends in data volume and data throughput performance from the Southern California Petabyte Scale Cache (SoCal Repo) [6], which includes 24 federated caching nodes with approximately 2.5PB of total storage. From the trends, we also determined how much a machine learning model can predict the network access patterns for the regional data cache. The fluctuation in the daily cache utilization, as shown in Figure 1, is high, and it is challenging to build a learning model to follow the trends.



Fig. 1. Daily proportion of cache hits volume and cache misses volume from July 2021 to June 2022, with 8.02 million data access records for 8.2PB of traffic volume for cache misses and 4.5PB of traffic volume for cache hits. 35.4% of the total traffic has been saved from the cache.

The daily and hourly study also modeled the cache utilization and data throughput performance, with 80% of the training data and 20% of the testing data. Figure 2 shows the samples of our hourly study results. The root-mean-square error (RMSE) is measured and compared to the standard deviation of the input data values to provide a reference to determine how large the errors of predictions are. The relative error, ratio of testing RMSE and standard deviation, is less than 0.5, indicating the predictions are pretty accurate.

The study results can be used to optimize the cache uti-



(a) Hourly volume of cache misses; training set RMSE=0.16, testing set RMSE=0.40, std.dev=1.42



(b) Hourly throughput of cache misses; training set RMSE=25.90, testing set RMSE=18.93, std.dev=121.36

Fig. 2. Hourly cache miss performance.

lization, network resources, and application workflow performance, and become the base for exploring characteristics of other data lakes as well as examining longer term network requirements for the data caches.

REFERENCES

- B. Brown, E. Dart, G. Rai, L. Rotman, and J. Zurawski, "Nuclear physics network requirements review report," Energy Sciences Network, University of California, Publication Management System Report LBNL-2001281, 2020. [Online]. Available: https://www.es.net/assets/Uploads/ 20200505-NP.pdf
- [2] X. Espinal, S. Jezequel, M. Schulz, A. Sciabà, I. Vukotic, and F. Wuerthwein, "The quest to solve the hl-lhc data access puzzle," *EPJ Web of Conferences*, vol. 245, p. 04027, 2020. [Online]. Available: https://doi.org/10.1051/epjconf/202024504027
- [3] E. Fajardo, D. Weitzel, M. Rynge, M. Zvada, J. Hicks, M. Selmeci, B. Lin, P. Paschos, B. Bockelman, A. Hanushevsky, F. Würthwein, and I. Sfiligoi, "Creating a content delivery network for general science on the internet backbone using XCaches," *EPJ Web of Conferences*, vol. 245, p. 04041, 2020. [Online]. Available: https://doi.org/10.1051/epjconf/202024504041
- [4] E. Copps, H. Zhang, A. Sim, K. Wu, I. Monga, C. Guok, F. Wurthwein, D. Davila, and E. Fajardo, "Analyzing scientific data sharing patterns with in-network data caching," in 4th ACM International Workshop on System and Network Telemetry and Analysis (SNTA 2021), ACM. ACM, 2021.
- [5] R. Han, A. Sim, K. Wu, I. Monga, C. Guok, F. Wurthwein, D. Davila, J. Balcas, and H. Newman, "Access trends of in-network cache for scientific data," in 5th ACM International Workshop on System and Network Telemetry and Analysis (SNTA 2022), ACM. ACM, 2022.
- [6] E. Fajardo, A. Tadel, M. Tadel, B. Steer, T. Martin, and F. Würthwein, "A federated xrootd cache," *Journal of Physics: Conference Series*, vol. 1085, p. 032025, 2018.