

Comparing Cache Utilization Trends for Regional Data Caches

Alex Sim^{1,*}, Erica Wang^{2,**}, Ronak Monga^{3,***}, Kesheng Wu^{1,****}, Justas Balcas^{1,†}, Brendan White^{1,‡}, Chin Guok^{1,§}, Inder Monga^{1,¶}, Diego Davila^{4,||}, Frank Würthwein^{4,**}, and Harvey Newman^{2,††}

¹Lawrence Berkeley National Laboratory, USA

²California Institute of Technology, USA

³Indiana University, USA

⁴University of California at San Diego, USA

Abstract. The rapid growth of data volumes from large scientific collaborations, such as the Large Hadron Collider (LHC), presents significant challenges for the High Energy Physics (HEP) community. With annual data volumes projected to increase by a factor of thirty by 2028, efficient data management has become a critical concern. The HEP community's reliance on wide-area networks for global data distribution often results in redundant long-distance transfers, leading to network congestion and degraded application performance. This study investigates the effectiveness of regional data caches in mitigating network congestion and enhancing application performance, using a large-scale dataset of millions of access records from regional caches in Southern California, Chicago, and Boston, which serve the LHC's CMS experiment. Our analysis reveals the substantial potential of in-network caching to transform large-scale scientific data dissemination, enabling faster and more efficient data access for researchers and scientists. Additionally, neural networks trained on data from multiple regional caches demonstrate enhanced predictive accuracy, particularly benefiting caches with limited historical data through transfer learning, thereby validating their robust generalization capability.

1 Introduction

The dramatic growth in data volumes from large-scale scientific collaborations, particularly the Large Hadron Collider (LHC), presents significant technical challenges for High Energy Physics (HEP) data management and distribution systems. With projected annual data

*e-mail: ASim@lbl.gov

**e-mail: ewang2@caltech.edu

***e-mail: ronmonga@iu.edu

****e-mail: KWu@lbl.gov

†e-mail: jbalcas@es.net

‡e-mail: bmwt@es.net

§e-mail: chin@es.net

¶e-mail: imonga@es.net

||e-mail: didavila@ucsd.edu

**e-mail: fkw@ucsd.edu

††e-mail: newman@hep.caltech.edu

volumes increasing by a factor of thirty by 2028 [1], the development of efficient data management strategies has become critical. Currently, the HEP community’s data distribution infrastructure relies heavily on wide-area networks, leading to substantial redundancy in long-distance data transfers and suboptimal resource utilization.

Empirical analysis reveals that a significant proportion of network traffic consists of repeatedly accessed datasets within institutional or regional boundaries. This pattern of redundant data transfer introduces multiple inefficiencies: increased network congestion, elevated bandwidth consumption, and heightened data access latency, collectively degrading the performance of data-dependent applications. To address these challenges, we investigate the effectiveness of regional data caches [2], analyzing millions of access records from distributed cache deployments in Southern California, Chicago, and Boston that serve the LHC’s Compact Muon Solenoid (CMS) experiment [3].

Our comprehensive analysis of cache utilization patterns reveals region-specific characteristics that necessitate adaptive cache management strategies. Through the application of machine learning techniques, we demonstrate the feasibility of accurate cache usage prediction and resource requirement forecasting. The study presents quantitative evidence that in-network data caching can significantly optimize data transfer efficiency and access patterns in scientific computing environments. This optimization reduced network congestion, minimized latency, and enhanced application performance metrics.

Our experiments demonstrate that neural network models trained on data from multiple regional caches exhibit enhanced predictive accuracy, despite the caches’ distinct operational characteristics (e.g., access patterns, cache capacities, and latency profiles). This cross-cache training approach proved particularly effective for the Boston Cache, where limited historical data constrained standalone model training. Contrary to initial assumptions that cache-specific training data would be essential, our results reveal that transfer learning – leveraging shared patterns across Southern California (SoCal), Chicago, and Boston Caches – achieves comparable or superior performance to cache-isolated models. This finding underscores the robust generalization capability of neural networks in distributed systems, offering a practical solution for mitigating data scarcity in emerging cache deployments. The empirical success of this domain-adaptation strategy highlights neural architectures’ ability to capture latent, cache-agnostic features, thereby validating their utility in optimizing resource allocation across heterogeneous storage environments.

2 Background

The 3 caches mentioned in this study serve data for the CMS experiment, a large and highly distributed scientific collaboration that expands to 100+ institutions around the world [4, 5]. The SoCal Cache is a regional storage cache with approximately 2PB of total capacity with 23 federated caching nodes geographically distributed across: Caltech, UCSD and ESnet Sunnyvale. The longest physical distance among these nodes exists between UCSD and the ESnet Sunnyvale PoP with about 450 miles and an RTT of about 10ms. This cache supports computing jobs at UCSD and Caltech. Its measurement data has been collected from June 2020 to March 2024, consisting of 30.94 million data access records, a total volume of 38.5 PB, and a hit ratio of 69.3%.

The Chicago Cache has 6 federated caching nodes, 5 of which are deployed at the University of Wisconsin in Madison, and 1 in the ESnet Chicago PoP. Its measurements have been collected from October 2022 to March 2024, consisting of 9.05 million data access records, a total volume of 25.8 PB, and a hit ratio of 48.4%.

Finally the Boston Cache is one of the latest deployments, and consists of a single caching node at the ESnet Boston PoP. Measurements for this cache have been collected from August

2023 to March 2024, consisting of 27.75 million data access records, a total volume of 93.5 PB and a hit ratio of 6.6%.

Table 1: Summary of data access from SoCal, Chicago and Boston Caches. About 69.3%, 48.4%, and 6.6% of the requested data sizes are satisfied by the SoCal, Chicago and Boston Caches respectively.

		# of accesses	cache miss (TB)	cache hits (TB)
6/2020-3/2024 SoCal	Total	30,939,524	11,837.38	26,686.73
	Daily Average	22,114	8.46	19.08
10/2022-3/2024 Chicago	Total	9,054,942	13,315.37	12,494.41
	Daily Average	16,553	24.34	22.84
8/2023-3/2024 Boston	Total	27,754,028	87,311.58	6,173.02
	Daily Average	114,214	359.31	25.40

In the following, we refer to a "cache hit" when talking of a data access request that is served directly from the cache, whereas a "cache miss" triggers a data transfer from a remote storage to the cache before the request can be served. Table 1 shows the summary of data accesses during the data collection period for three regional caches.

3 Cache utilization trends

In this section we analyze the hit and miss patterns observed in the 3 regional caches. For each cache we present 2 pairs of figures: on the left pair we plot the distribution of hits and misses in terms of the number of accesses whilst on the right we do it in terms of the amount of data requested. Bottom plots are the normalized version of the plot directly above them. Every bar in these plots represents the daily sum of access counts (left) and access size (right).

Starting with the SoCal Cache (Figure 1) we can note 2 distinct patterns in 2 different time periods. The first one between October 2021 and July 2022, which shows a dominant pattern of the cache misses. This has been previously studied and concluded that the reason for this pattern was the presence of an unusual analysis workflow that requested unusually large files [6]. Then a second period, from July 2022 to March 2024 (the last 21 months of this study), that shows a dominant pattern of cache hits with nearly 95% of the data requested being directly served from the cache. In this same period and looking just at the pattern of the cache misses (orange color) we can observe a significant difference between left and right plots, which indicates cache misses were triggered mainly by small requests.




Figure 1: Daily activities and rates for cache hits and cache misses from SoCal Cache

In Figure 2, we can see the distribution of cache hits and misses in the Chicago Cache. We can observe a very similar pattern between access counts and access sizes, and see how the left and right plots are almost identical. This indicates a very consistent daily average request size with a similar daily variation as shown in Figure 4 (b).

Similarly to the Chicago Cache, the Boston Cache usage, depicted in Figure 3, shows a consistent average request size with a small variation as shown in Figure 4 (c) but more

importantly, it shows a dominant pattern of the cache misses. The reason for this is the total amount of disk space available to the cache. Compared to the Chicago Cache, which caches the same superset of data, this cache has only 60% of the disk space. In order to increase the cache hit/miss ratio, the cached superset for this cache has been recently reduced, and we expect to see an improvement in the future.




Figure 2: Daily activities and rates for cache hits and cache misses from Chicago Cache




Figure 3: Daily activities and rates for cache hits and cache misses from Boston Cache




Figure 4: Average data access size per request from each federated storage nodes in three regional caches. (a) SoCal Cache showing average data access size between 1GB and 4GB, (b) Chicago Cache showing average request size occasionally goes above 4GB, (c) Boston Cache showing average size between 3GB and 4GB.

4 Modeling the cache utilization

Our exploration further employed Long Short-Term Memory (LSTM) [7, 8], a machine learning model for predicting cache utilization patterns within regional data caches.

Table 2: RMSE of daily LSTM model results for cache utilization. The relative prediction errors, ratios of testing RMSE and standard deviation, are less than 0.8.

	SoCal		Chicago		Boston	
	Testing RMSE	standard deviation	Testing RMSE	standard deviation	Testing RMSE	standard deviation
Daily volume of cache hits (TB)	21.1	26.3	15.7	30.7	24.1	31.1
Daily count of cache hits	6571.9	12827.7	4583.6	9355.9	6993.5	9310.9
Daily volume of cache misses (TB)	2.1	17.8	14.7	72.5	172.1	411.6
Daily count of cache misses	2144.3	8811.3	5604.2	21979.2	49560.0	123120.0

Figure 5 shows the prediction result of the daily cache utilization for regional caches. Our model includes 6 features such as access counts, access size, cache hit counts, cache miss counts, cache hit size and cache miss size. The parameters include 128 units with *tanh* activation function, 0.04 dropout rates, and the number of epochs of 50. The training data comes from the first 80% of the study period, and the testing data comes from the last 20%. We observe that the predictions in cache misses are closer to the actual values than the predictions for cache hits. The daily cache misses for both request counts and sizes from all three regions show that the relative root-mean-square error (RMSE) error in the testing dataset is less than 0.41, much less than cache hits (0.80), as shown in Table 2.

Table 2 shows the RMSE of daily request counts and sizes of cache hits and misses. The 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 0.8 for all three regional caches, indicating strong predictive performance. The relative errors from the Boston Cache show larger than the other two caches because the number of input data for Boston Cache is relatively smaller as the Boston Cache is a recent deployment covering a shorter study period, and it’s not enough to train the model relative to the other two caches. To accommodate smaller training datasets, we next explore the transfer learning model [9, 10] for Chicago and Boston Caches.

Figure 6 shows the relative length of the study period of the three regional caches and the overview process of the transfer learning. Based on the model from SoCal Cache, 20% of the records from Chicago Cache and 15% of the records from Boston Cache are used to refine the model for each cache, instead of the previous models with 80% of the data. We have selected 20% and 15% of the records for training data size for each cache because those produced similar results to the previous models without the transfer learning. Figure 7 shows the prediction results from the transfer learning for cache hit size. The relative error for Chicago Cache is 0.45, similar to 0.47 of the previous model without the transfer learning, but with less training data. The relative error for Boston Cache is 0.64, slightly dropped from 0.77 of the previous model without the transfer learning, but with less training data. Daily cache hit size predictions for Boston Cache with the transfer learning in Figure 7b actually follow the trends better than the one in Figure 5c, and it shows the transfer leaning not only makes up the insufficiency of the length of the training data, but also results in better prediction with less number of training data.

5 Discussion

While comparing the logs from three regional caches, we observe that the cache capacity has a significant impact on the cache utilization. If the cache capacity is doubled in Chicago Cache, the cache hit rates would be improved a lot from the currently observed 0.48 during the study period. Further simulation work on the resource allocation is planned for the future work. Also, the site configuration at the Boston Cache is planned to be adjusted for the cached superset of data to improve the low cache hit rates under the same cache capacity.




Figure 5: LSTM prediction for daily cache activities from three different regional caches. Total data requests, cache hits and cache misses, for both request counts and requested data size, showing actual measurements (blue bars), training predictions (orange bars), and testing predictions (red bars).

Figures 1, 2 and 3 show some number of spikes for cache utilization. While the predictions capture those patterns in the spikes in Figure 5, they may not precisely match the heights, and large fluctuations are generally hard to predict. For modeling the usage trends, it is not necessary to capture such spikes precisely. We could use the alternative approaches




Figure 6: Measurement data length comparison from three regional caches and overview of transfer learning model.




Figure 7: Transfer leaning based prediction for daily cache hits size from Chicago and Boston caches, showing actual measurements (blue bars), training predictions (orange bars), and testing predictions (red bars).

such as moving average method to obtain a performance model for the general trends. Understanding the impact of these usage spikes in the long term resource utilization would help us provisioning of the future in-network cache deployments.

Prediction results from the LSTM models in Figure 5 offer insights for resource usage estimation, including storage cache and network requirements, and can provide improvements in application workflow performance within regional data cache systems. Results from the transfer learning in Figure 7 follow the usage trends better with the less number of training data, and show an approach to accommodate the shortage of the training data. Also, it gives us ways to explore the base foundation model further for characterizing the cache utilization and resource provisioning.

6 Conclusion

In this study, we present a comparative analysis of cache utilization patterns across three regional caches serving the LHC CMS experiment: Southern California, Chicago, and Boston. Our analysis reveals distinct usage patterns unique to each region, highlighting the importance of developing region-specific approaches for future cache deployments.

To predict cache utilization trends, we evaluated two machine learning approaches: LSTM models and transfer learning. Both approaches demonstrated strong predictive performance, with RMSE values significantly lower than the standard deviation of the input data. Notably, our transfer learning models achieved comparable or superior prediction accuracy while requiring less training data, suggesting their potential for forecasting future resource utilization.

The effectiveness of the caching system is particularly evident in the Southern California deployment, where approximately 95% of the requested data volume was served from the

cache during the 21-month study period. This high cache hit rate demonstrates the system's success in reducing wide-area network traffic and minimizing data access latency. Furthermore, our analysis provides valuable insights into cache utilization characteristics, contributing to improved resource efficiency and application performance.

Future work will focus on simulating cache utilization using observed characteristics from current deployments. This simulation-based approach will deepen our understanding of the caching system and allow us to further explore the predictability of cache utilization and resource provisioning.

7 Acknowledgments

This work was supported by the Office of Advanced Scientific Computing Research, Office of Science, of the US Department of Energy under Contract No. DE-AC02-05CH11231, and also used resources of the National Energy Research Scientific Computing Center (NERSC). This work was also supported by the National Science Foundation through the grants OAC-1836650, PHY-2323298, PHY-2121686 and OAC-2112167.

References

- [1] J. Zurawski, et al., *High energy physics network requirements review: Two-year update*, Tech. Rep. OSTI ID:2405935, Energy Sciences Network (2023), <https://doi.org/10.2172/2405935>
- [2] E. Fajardo, D. Weitzel, M. Rynge, M. Zvada, J. Hicks, M. Selmecı, B. Lin, P. Paschos, B. Bockelman, A. Hanushevsky et al., *Creating a content delivery network for general science on the internet backbone using XCaches*, EPJ Web of Conferences **245**, 04041 (2020)
- [3] C. Sim, K. Wu, A. Sim, I. Monga, C. Guok, D. Hazen, F. Wurthwein, D. Davila, H. Newman, J. Balcas, *Predicting Resource Utilization Trends with Southern California Petabyte Scale Cache*, in *26th International Conference on Computing in High Energy & Nuclear Physics (CHEP2023)* (2023)
- [4] A. Dorigo, P. Elmer, F. Furano, A. Hanushevsky, *Xrootd - a highly scalable architecture for data access*, WSEAS Transactions on Computers **4**, 348 (2005)
- [5] L. Bauerdick, D. Benjamin, K. Bloom, B. Bockelman, D. Bradley, S. Dasu, M. Ernst, R. Gardner, A. Hanushevsky, H. Ito et al., *Using xrootd to federate regional storage*, Journal of Physics: Conference Series **396**, 042009 (2012)
- [6] C. Sim, K. Wu, A. Sim, I. Monga, C. Guok, F. Wurthwein, D. Davila, H. Newman, J. Balcas, *Effectiveness and predictability of in-network storage cache for Scientific Workflows*, in *IEEE International Conference on Computing, Networking and Communication (ICNC2023)* (2023)
- [7] A. Sherstinsky, *Fundamentals of recurrent neural network (rnn) and long short-term memory (lstm) network*, Physica D: Nonlinear Phenomena **404**, 132306 (2020)
- [8] K. Greff, R.K. Srivastava, J. Koutník, B.R. Steunebrink, J. Schmidhuber, *LSTM: A search space odyssey*, in *IEEE transactions on neural networks and learning systems* (2016), Vol. 28, pp. 2222–2232
- [9] J. Yosinski, J. Clune, Y. Bengio, H. Lipson, *How transferable are features in deep neural networks?*, in *Proceedings of the 28th International Conference on Neural Information Processing Systems (NIPS'14)* (2014)
- [10] F. Zhuang, Z. Qi, K. Duan, D. Xi, Y. Zhu, H. Zhu, H. Xiong, Q. He, *A comprehensive survey on transfer learning*, Proceedings of the IEEE **109**, 43 (2021)