



Analyzing Transatlantic Network Traffic over Scientific Data Caches

Ziyue Deng
University of California, Berkeley
Berkeley, CA, USA
ziyue_deng@berkeley.edu

Alex Sim
Lawrence Berkeley National
Laboratory
Berkeley, CA, USA
asim@lbl.gov

Kesheng Wu
Lawrence Berkeley National
Laboratory
Berkeley, CA, USA
kwu@lbl.gov

Chin Guok
Energy Sciences Network
Berkeley, CA, USA
chin@es.net

Damian Hazen
Energy Sciences Network
Berkeley, CA, USA
dhazen@es.net

Inder Monga
Energy Sciences Network
Berkeley, CA, USA
imonga@es.net

Fabio Andrijauskas
University of California, San Diego
La Jolla, CA, USA
fandrijauskas@sdsc.edu

Frank Würthwein
University of California, San Diego
La Jolla, CA, USA
fkww@ucsd.edu

Derek Weitzel
University of Nebraska, Lincoln
Lincoln, NE, USA
dweitzel@unl.edu

ABSTRACT

Large scientific collaborations often share huge volumes of data around the world. Consequently a significant amount of network bandwidth is needed for data replication and data access. Users in the same region may possibly share resources as well as data, especially when they are working on related topics with similar datasets. In this work, we study the network traffic patterns and resource utilization for scientific data caches connecting European networks to the US. We explore the efficiency of resource utilization, especially for network traffic which consists mostly of transatlantic data transfers, and the potential for having more caching node deployments. Our study shows that these data caches reduced network traffic volume by 97% during the study period. This demonstrates that such caching nodes are effective in reducing wide-area network traffic.

CCS CONCEPTS

• **Networks** → **Network performance analysis**; • **Computing methodologies** → **Distributed computing methodologies**.

KEYWORDS

network cache, osdf, resource utilization, data pattern, xcache

ACM Reference Format:

Ziyue Deng, Alex Sim, Kesheng Wu, Chin Guok, Damian Hazen, Inder Monga, Fabio Andrijauskas, Frank Würthwein, and Derek Weitzel. 2023. Analyzing Transatlantic Network Traffic over Scientific Data Caches. In *Proceedings of the 2023 Systems and Network Telemetry and Analytics (SNTA '23)*, June 20, 2023, Orlando, FL, USA. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3589012.3594897>



This work is licensed under a Creative Commons Attribution International 4.0 License.

SNTA '23, June 20, 2023, Orlando, FL, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0165-8/23/06.

<https://doi.org/10.1145/3589012.3594897>

1 INTRODUCTION

Scientific experiments and simulations generate large volumes of data over time. Such data is shared by geographically distributed users, which creates a large amount of network traffic for data replication and access during analysis [2]. Data storage caches have been deployed for regional users engaged in related study topics [1, 5, 6, 8, 10]. These storage caches can hold a significant portion of the datasets close to user accesses, which reduces the data access latency and improves data analysis throughput [3, 7, 9].

One set of such storage caches is deployed by the Open Science Data Federation (OSDF). This study explores the efficiency of resource utilization of these OSDF caches for transatlantic data transfers. More specifically, we focus on two caching nodes for data transfers from US to Europe. Understanding their resource utilization could lead to more efficient management of future cache deployments.

The contributions of this paper can be summarized as follows: (1) our study finds that the data caches reduce the network traffic volume by 97% during the study period; (2) this network traffic reduction is from transatlantic traffic from the US to Europe, which is what the data cache nodes are designed for; (3) based on the observed network traffic reduction rates, we can plan for additional deployments of the caching nodes to benefit science communities.

2 BACKGROUND

The Open Science Grid (OSG) provides data access resources for High Throughput Computing (HTC) facilities [11]. The OSG data infrastructure is named the Open Science Data Federation (OSDF). It holds data files from several large experiments and many smaller independent projects. It supports massive amounts of data from National Science Foundation (NSF) funded projects. At the core of the OSDF are the concepts of "data origin," "data caches," and "data access redirector," all implemented as services via XrootD [4]. Figure 1 shows simplified schematics of these components.




Figure 1: Open Science Data Federation.

Table 1: Summary statistics for data accesses from the Cardiff node

	# of Requests	# of Cache Miss Requests	Total Size of Requested Bytes (GB)	Total Size of Cache Misses (GB)	% of Cache Miss Size
noVO	10,807	131	0.00698	0.000134	1.92%
osg	388,484	96,083	38,454	1,173.37	3.05%
gwdata	30	5	7.78	0.109	1.40%
Total	399,321	96,219	38,461	1,173.48	3.05%

Origins host the data. Multiple origins are tied into a tree structure to form the data federation. Applications access the OSDF generally via the closest cache to the compute host they execute on. The meaning of closest is determined via GeoIP [12]. All services are deployed as containers via Kubernetes. When a job requests a file from OSDF, usually this file request is satisfied with a replica from the cache. However, if the file is unavailable in the cache, the cache asks the redirector to locate a copy. This file is copied from the origin to the cache, and the cache sends it to the client. On all OSDF levels, file accesses are monitored.

XrootD has several streams that provide monitoring information, such as cache monitoring, transfer information, and file access monitoring. Caches and Origins use “shoveler” software to send the monitoring streams to the queue message system. The collector processes these messages and sends them back to the queue system. In the end, all the monitoring data goes to a central logging system managed by RabbitMQ, as shown in Figure 1.

This study examines logs from one cache node in Cardiff, UK, and another in Amsterdam, Netherlands. These two cache nodes are mainly designed to serve those data from the US to Europe. On 10 Gps networks, Cardiff and Amsterdam nodes have 81 TB and 30 TB of storage capacities respectively.

3 DATASETS

Our work is based on the logs collected from the Amsterdam and Cardiff nodes between October 2022 and January 2023. The collected information includes the following attributes about every data access request: timestamp, access count, file path, project group (VO), file size, attach time, detach time, remote origin, bytes hit cache, bytes miss cache, and bytes bypass cache. A total of 399,321 data access requests from the Cardiff node and 31,508,228 data access requests from the Amsterdam node are included in this study.

Table 1 shows basic statistics about the data accesses for the Cardiff node for each Virtual Organization (VO). A total of 399K data requests were received, and about 24% of the total requests are cache misses. When the requested file is not in the cache, it is

Table 2: Monthly summary statistics for data accesses from the Amsterdam node

	# of Requests	# of Cache Miss Requests	Total Size of Requested Bytes (GB)	Total Size of Cache Misses (GB)	% of Cache Miss Size
Oct 2022	9,290,400	2,931	294,017	79	0.03%
Nov 2022	13,160,620	19,729	473,367	348	0.07%
Dec 2022	6,140,170	26,765	140,843	327	0.23%
Jan 2023	2,917,038	13,680	83,343	617	0.74%
Total	31,508,228	63,105	991,570	1,371	0.14%

Table 3: Summary statistics for data accesses per VO from the Amsterdam node

	# of Requests	# of Cache Miss Requests	Total Size of Requested Bytes (GB)	Total Size of Cache Misses (GB)	% of Cache Miss Size
fermilab	6,854,234	6,401	109,840	104	0.10%
osg	24,493,597	39,125	877,349	928	0.11%
noVO	78,502	10,386	0.0789	0.0106	13.48%
gwdata	72,865	5,474	1.825	169	9.25%
icecube	9,030	1,719	2.556	169	6.61%

Table 4: Summary statistics for data accesses for DUNE and LIGO projects from the Amsterdam node

	# of Requests	# of Cache Miss Requests	Total Size of Requested Bytes (GB)	Total Size of Cache Misses (GB)	% of Cache Miss Size
DUNE	5,006,143	102	99929	0.24	0.00024%
LIGO	1,981,630	38,593	164,451	927	0.56%

retrieved from the remote origins which are mostly from the US. The total requested data volume was 38.4 TB, among which 1.17 TB was transferred over the wide-area network. The percentage of cache misses is calculated by (Total Size of Cache Misses) / (Total Size of Requested Bytes). For the Cardiff node, the percentage of cache misses is very low; less than 4%. The VO osg has the highest percentage of cache misses at 3.05%, and the gwdata (LIGO) project has the lowest percentage of cache misses at 1.40%.

Table 2 and Table 3 show basic statistics about data accesses to the Amsterdam node during the study period, grouped by month and VO respectively. Table 2 shows the highest data access requests in November 2022 with about 13 million requests, but the percentage of cache miss sizes is low at 0.07%. Cache misses by volume ranging from 0.03% to 0.74% is observed to be low across all months; less than 1%, which means that most of the requested data are already in the cache.

Table 3 shows that two VOs have a high number of requests with low cache misses by volume. This indicates that users from these two VOs have similar research interests and are accessing the same dataset. On the other hand, the other two VOs have a relatively lower number of data access requests with higher cache misses.

Table 4 shows summary statistics about data accesses for the Deep Underground Neutrino Experiment (DUNE) and the Laser Interferometer Gravitational-Wave Observatory (LIGO) projects on Amsterdam node. DUNE and LIGO projects are identified from the data paths. The two projects have a high number of requests with very low number and volume of cache misses. This indicates that users from these two projects tend access the same set of files.

4 RESOURCE UTILIZATION

Next, we study the effectiveness of the storage caches of OSDF by exploring two types of statistics: daily file requests and network traffic avoided by the use of these caches.




Figure 2: Daily number of data requests (in blue) and cache misses (in orange) in log scale from Oct. 2022 to Jan. 2023. Note that where the number of cache misses are typically orders of magnitude fewer than the number of files requested there are a few days where most of the file requests are cache misses on the cache at Cardiff.

4.1 Cache Utilization

Figure 2a and Figure 2b show the daily number of data requests (in blue) and number of cache misses (in orange) in log scale during the study period for the Cardiff node and Amsterdam node respectively. The daily number of data accesses for the Cardiff node has large variance throughout the study period, and the overall cache miss rate is 24.1%. For the Amsterdam node, the daily number of data accesses has some variance throughout the study period with low cache miss rates. Overall cache miss rate for data access requests on the Amsterdam node is 0.2%.

Figure 3a and Figure 3b show the daily number of data requests (in blue) and number of cache misses (in orange) in log scale on the Amsterdam node for DUNE and LIGO projects respectively. The daily number of data accesses is fairly consistent throughout the study period for the DUNE project, and it shows a very low number of cache misses during the study period. For the LIGO project, the number of data accesses has a large variance throughout the study period, and it has a higher number of cache misses than DUNE project.

4.2 Network Traffic Reduction

The file requests served from the cache nodes are reducing network traffic on the wide-area network. Next, we examine the cache hits and misses based on the volume of data involved.

Figure 4a and Figure 4b show the daily volume of requested data (in blue) and volume of cache misses (in orange) in log scale for the Cardiff node and Amsterdam node respectively. Daily requested data volumes on the Amsterdam node in Figure 4b follows a roughly similar pattern to the daily number of data accesses in Figure 2b. Daily volumes of cache misses on the Amsterdam node in Figure 4b




Figure 3: Daily number of data requests (in blue) and cache misses (in orange) in log scale for Amsterdam node from Oct. 2022 to Jan. 2023. Note there are a couple of days where most of the accesses to LIGO files are cache misses.




Figure 4: Daily volume of requested data (in blue) and cache misses (in orange) in log scale from Oct. 2022 to Jan. 2023.

also follow a similar pattern proportionally to the daily number of cache misses in Figure 2b, as the figures are in log scale. Overall cache miss rate for the data volume on the Amsterdam node is 0.14%, and the cache miss rate for data access requests is 0.2%. It indicates that the average data size for each cache miss is slightly smaller than the average data size for each data request during the study period. From the Table 2, the average file size for each cache miss is 21.7 MB, and the average data size for each data request is 31.4 MB.




Figure 5: Daily volume of requested data (in blue) and volume of cache misses (in orange) in log scale for Amsterdam node from Oct. 2022 to Jan. 2023

Figure 5a and Figure 5b show the daily volume of requested data (in blue) and volume of cache misses (in orange) in log scale for the DUNE and LIGO projects respectively. The daily volume of requested data and cache misses for the DUNE project in Figure 5b shows similar patterns as the daily numbers of data requests and cache misses in Figure 3a. For the DUNE project, the overall cache miss rate based on data volumes is 0.00024% and the cache miss rate based on file requests is 0.00203%. From the Table 4, the average file size for each cache miss is 2.4 MB, and the average file size for each file request is 20 MB. For the LIGO project, the overall cache miss rate in data volume is 0.56%, and the cache miss rate for data access requests is 1.95%. The average file size for each cache miss is 24 MB, and the average file size overall is 83 MB.

5 CONCLUSIONS

In this study, we explored the efficiency of cache utilization and the network traffic savings at the Cardiff and Amsterdam nodes from OSDF. Also, we studied data access and cache miss patterns for the DUNE and LIGO projects on the Amsterdam node. Network traffic savings would be mostly from transatlantic data transfers. Our study shows the potential for big improvements of backbone network performance from the impacts of the caching nodes. From a total of 31,907,549 data accesses logged on both Cardiff and Amsterdam nodes from Oct. 2022 to Jan. 2023, we observed a total of 1.03 PB of client data accesses, with 2.5 TB of data transfers for cache misses from the remote data sources to the local cache, and 1.027 PB of network traffic volume savings from the repeated shared data accesses. The Cardiff node reduced 96.95% of the network traffic volume, and the Amsterdam node saved 99.86% of the network traffic volume during the study period.

With additional deployments of the caching nodes, we plan to study the access trends for a longer period of time and better

understand the resource utilization which could lead to the efficient management of the data caches and long-term resource allocation planning.

ACKNOWLEDGMENTS

This work was supported 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). This work was also supported by the National Science Foundation through the grants OAC-2030508, OAC-1836650, MPS-1148698, PHY-1120138, and OAC-1541349. This work is also supported by the US CMS M&O 2121686, PNRP: NSF 2112167, and PRP: NSF OAC-1541349

REFERENCES

- [1] L. Bauerick, K. Bloom, B. Bockelman, D. Bradley, S. Dasu, J. Dost, I. Sfiligoi, A. Tadel, M. Tadel, F. Wuerthwein, A. Yafil, and the CMS collaboration. 2014. XRootd, disk-based, caching proxy for optimization of data access, data placement and data replication. *Journal of Physics: Conference Series* 513, 4 (2014).
- [2] Ben Brown, Eli Dart, Gulshan Rai, Lauren Rotman, and Jason Zurawski. 2020. *Nuclear Physics Network Requirements Review Report*. University of California, Publication Management System Report LBNL-2001281. Energy Sciences Network. <https://www.es.net/assets/Uploads/20200505-NP.pdf>
- [3] E. Coppins, H. Zhang, A. Sim, K. Wu, I. Monga, C. Guok, F. Wuerthwein, D. Davila, and E. Fajardo. 2021. 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.
- [4] A. Dorigo, P. Elmer, F. Furano, and A. Hanushevsky. 2005. XROOTD - A highly scalable architecture for data access. *WSEAS Transactions on Computers* 4, 4 (2005), 348–353.
- [5] X. Espinal, S. Jezequel, M. Schulz, A. Sciabà, I. Vukotic, and F. Wuerthwein. 2020. The Quest to solve the HL-LHC data access puzzle. *EPJ Web of Conferences* 245 (2020), 04027. <https://doi.org/10.1051/epjconf/202024504027>
- [6] E. Fajardo, A. Tadel, M. Tadel, B. Steer, T. Martin, and F. Würthwein. 2018. A federated Xrootd cache. *Journal of Physics: Conference Series* 1085 (2018), 032025.
- [7] R. Han, A. Sim, K. Wu, I. Monga, C. Guok, F. Wuerthwein, D. Davila, J. Balcas, and H. Newman. 2022. 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.
- [8] Ruth Pordes, Don Petravick, Bill Kramer, Doug Olson, Miron Livny, Alain Roy, Paul Avery, Kent Blackburn, Torre Wenaus, Frank Würthwein, Ian Foster, Rob Gardner, Mike Wilde, Alan Blatecky, John McGee, and Rob Quick. 2007. The open science grid. *Journal of Physics: Conference Series* 78, 1 (2007), 012057.
- [9] C. Sim, K. Wu, A. Sim, I. Monga, C. Guok, F. Wuerthwein, D. Davila, H. Newman, and J. Balcas. 2023. Effectiveness and predictability of in-network storage cache for Scientific Workflows. In *International Conference on Computing, Networking and Communication (ICNC 2023)*. IEEE, IEEE.
- [10] Derek Weitzel, Marian Zvada, Ilija Vukotic, Rob Gardner, Brian Bockelman, Mats Rynge, Edgar Hernandez, Brian Lin, and Mátyás Selmeci. 2019. StashCache: A Distributed Caching Federation for the Open Science Grid. *PEARC '19: Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines (learning)*, 1–7. <https://doi.org/10.1145/3332186.3332212>
- [11] Derek Weitzel, Marian Zvada, Ilija Vukotic, Rob Gardner, Brian Bockelman, Mats Rynge, Edgar Fajardo Hernandez, Brian Lin, and Mátyás Selmeci. 2019. StashCache: A Distributed Caching Federation for the Open Science Grid. In *Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines (Learning)* (Chicago, IL, USA) (PEARC '19). Association for Computing Machinery, New York, NY, USA, Article 58, 7 pages. <https://doi.org/10.1145/3332186.3332212>
- [12] Derek Weitzel, Marian Zvada, Ilija Vukotic, Rob Gardner, Brian Bockelman, Mats Rynge, Edgar Fajardo Hernandez, Brian Lin, and Mátyás Selmeci. 2019. StashCache: a distributed caching federation for the open science grid. In *Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines (learning)*, 1–7.