



Open Science Data Federation - operation and monitoring

Fabio Andrijauskas
fandrijauskas@ucsd.edu
University of California - San Diego
San Diego, CA, US

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

Frank K Würthwein
fkw@ucsd.edu
University of California - San Diego
San Diego, CA, US

ABSTRACT

Extensive data processing is becoming commonplace in many fields of science. Distributing data to processing sites and providing methods to share the data with collaborators efficiently has become essential. The Open Science Data Federation (OSDF) builds upon the successful StashCache project to create a global data access network. The OSDF expands the StashCache project to add new data origins and caches, access methods, monitoring, and accounting mechanisms. Additionally, the OSDF has become an integral part of the U.S. national cyberinfrastructure landscape due to the sharing requirements of recent NSF solicitations, which the OSDF is uniquely positioned to enable. The OSDF continues to be utilized by many research collaborations and individual users, which pull the data to many research infrastructures and projects.

CCS CONCEPTS

• **Computer systems organization** → *Grid computing*; • **Information systems** → **Distributed storage**; • **Human-centered computing** → *Information visualization*; • **Networks** → *Storage area networks*.

KEYWORDS

OSDF, data transfer, OSG, scientific data

ACM Reference Format:

Fabio Andrijauskas, Derek Weitzel, and Frank K Würthwein. 2024. Open Science Data Federation - operation and monitoring. In *Practice and Experience in Advanced Research Computing (PEARC '24)*, July 21–25, 2024, Providence, RI, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3626203.3670557>

1 INTRODUCTION

As the scale of data and science complexity increases, research must look for additional resources to manage storage and compute requirements in memory or processor speed [6, 10]. Regardless of whether projects need large computational throughput, e.g. Open Science Consortium (OSG) [12] or computational performance, e.g. ACCESS [1], or the National Research Platform [2] an unquestionable challenge is pervasive across all forms of computational access: what is the most efficient way to deliver data to a compute host when there are always limits on local infrastructure. For many collaborative research projects, the answer is to place input data

wherever there is storage capacity, maybe at the site of an experiment and then find ways to deliver that data to a compute site as needed by a computational pipeline. For a domain researcher, it is a challenge to bridge the geographical distance separating their data from a CPU that will process them [11]. The Open Science Data Federation (OSDF) is a proven data access framework that provides the infrastructure and the tools to access data globally, implementing the notion of "Any Data, Anytime, Anywhere". When we read our favorite newspaper online, we don't ask for it to first be delivered to our laptops, or phones. OSDF does the same for data access for computation. It replaces a focus on data transfer with a focus on data access by the application. Figure 1 shows the current global deployments of OSDF caches and origins.

OSDF is the evolution of an OSG project, StashCache [9, 12]. Together, the OSG and OSDF create a unique research environment, providing computational throughput capacity and efficient access to their data for users [10]. To meet demand and growth, the OSDF ecosystem needs to improve in tools and capacity continuously. It also needs to ensure that data are reliably delivered at compute sites and staff are empowered to troubleshoot potential shortfalls as early as possible. In this work, we show improvements in accounting, available storage space for caches and origins, and enhancements in monitoring the quality of the service delivery.

2 BACKGROUND

In order to serve data to computational workflows running on the distributed computing infrastructure, the OSG utilizes the Open Science Data Federation [12]. At the crux of this data delivery framework are the concepts of "origin," "caches," and "redirectors," all implemented as services via the XrootD [5, 8] software framework, which allows for low latency and scalable data access. Figure 2 highlights the integration of these components. Origins effectively refer to the backend storage hosting project data. In the context of OSDF, an origin refers to the XRootD configuration that allows access to the storage via a data transfer node that mounts a project directory. Multiple origins are tied into a tree structure that connects to a redirector which communicates with the cache network. Applications generally access the OSDF via the closest cache to a computing site. The determination of the closest cache is via GeoIP [12]. The most efficient way to deploy and manage OSDF software and services (at origins or caches) is with containers on a federated Kubernetes infrastructure. The federated model allows operators to monitor data access and debug issues effectively.

Figure 2 shows the primary file access. The black line represents a file request by a job, the purple dashed lines represent file access, and the orange shows the monitoring process. When a job requests a file, it queries the file from the nearby cache. If the file is on the cache, the job receives the file. However, if the file is unavailable on the cache, the cache queries the redirector of its location at



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

PEARC '24, July 21–25, 2024, Providence, RI, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0419-2/24/07
<https://doi.org/10.1145/3626203.3670557>



Figure 1: OSDF caches and origins worldwide.

an origin. Once the file location is determined, the cache requests the file from the origin and makes it available to the job. XRootD provides operators with the ability to get informational streams across all levels of interaction. We will discuss this briefly in Section 4. Data access on OSDF is enabled via multiple ways, with the majority being requested by jobs running in pools managed by the Open Science Consortium (OSG). The network of Caches and Origins can serve data for multiple projects and science collaborations to points of execution and facilitate the acceleration of scientific discovery. This capability is being promoted by the U.S. National Science Foundation (NSF) via the Campus Cyberinfrastructure (CC*) program that awards funds for storage infrastructure that will host OSDF origins and caches. <https://opensciencegrid.org/campus-cyberinfrastructure.html> and <https://beta.nsf.gov/funding/opportunities/campus-cyberinfrastructure-cc>. Most OSDF hosts are deployed using NRP and Kubernetes using a GitHub DevOps architecture.

3 NEW CACHES AND ORIGINS

Funding accessibility, along with the success of past deployments, help pave the way for deploying more federated caches and origins to geographically span all of the U.S. We provide below a list of recent deployments:

- (1) 1 cache (50 TB) and one origin (1.6TB): San Diego Supercomputer Center.
- (2) 1 cache (42 TB) and one origin (1.2 PB): University of Nebraska-Lincoln.
- (3) 1 cache (29 TB) and one origin (1.2PB): Massachusetts Green High-Performance Computing Center.
- (4) New caches on Internet2 hubs: Boise - Idaho (42 TB), Houston - Texas (in progress), Jacksonville - Florida (42 TB), Denver - Colorado (42 TB), and Northeastern University Boston. Existing Internet2 caches in Chicago, New York, and Kansas City.
- (5) 1 cache (7 TB) and one origin (7 TB): The University of Tokyo.
- (6) 1 cache (14 TB): The National Center for Atmospheric Research (NCAR), Boulder.
- (7) 1 origin (337 TB): University of Nebraska-Lincoln.

The location of OSDF infrastructure is primarily determined by proximity to a University campus, an institution, or an Internet2 point of presence (PoP). Furthermore, the institution must have an affiliation to a project or a collaboration. As shown in Figure 2, the monitoring streams from OSDF are collected in an Elasticsearch database, which yields an abundance of information that we can mine for performance and utilization metrics of the service [7].




Figure 2: The Open Science Data Federation uses XrootD, and several other tools to serve data to execution points via client requests.




Figure 3: The monthly evolution of cache storage per project namespace shows the steady growth of the OSDF utilization by research projects.

Metric	Value
Caches	37 across 14 Inst.
Origins	15 across 6 Inst.
Transfers	5,861,721,390
Total size of transfers	294.7PB
Number of checks/monitoring	352
Number of Kubernetes Pods	82

Table 1: OSDF top-level statistics, 2022-2023.

We show in Figure 3 the aggregated storage utilization values of the caches per month of top project namespaces; there are 375 project namespaces in OSDF at present. Table 1 showcases some of the top level statistics for the OSDF infrastructure during the past year, Table 2 shows OSDF top-level statistics for 2022-2023.

- LIGO - Laser Interferometer Gravitational-Wave Observatory <https://www.ligo.caltech.edu>
- NOvA - NuMI Off-axis νe Appearance - <https://novaexperiment.fnal.gov>
- MINERvA - (Main Injector Neutrino Experiment to study ν-A interactions) - <https://minerva.fnal.gov>
- DUNE - Deep Underground Neutrino Experiment - <https://lbnf-dune.fnal.gov>
- MicroBooNE - large 170-ton liquid-argon time projection chamber - <https://microboone.fnal.gov>

4 END-TO-END MONITORING OF OSDF

As it is critical that we ensure the operational stability and availability of the OSDF infrastructure, we employ extensive monitoring of all the components and perform regular tests that check the status

Project	Files request
LIGO	6,893,702,577
LIGO user-specific	565,677,512
fnal.gov - Nova	492,252,998
fnal.gov - Minerva	435,978,085
fnal.gov - Dune	412,684,786
OSG Collaborations	226,827,338
fnal.gov - uboone	156,274,029
LIGO gwdata - O3a - uboone	145,505,762
LIGO gwdata - O2 - uboone	103,851,076
OSG user	98,685,798
Other	766,798,555

Table 2: Files requested by projects between 2018 and 2023.

of the end-to-end service delivery. Figure 2 provides a graphical description of the service flow. As mentioned previously, when a client requests a file from a nearby cache and it is not there, then the cache queries the file location via the redirector and requests it from the origin. This flow of queries and decision-making is captured by an XRootD stream, one of several supported in XRootD. For example, the f-stream captures information about the file access. The g-stream has information about the cache hit and misses. All of these collections of messages from the cache or the origin are sent to a Shovel service. The Shovel [3] is a tool that collects the UDP XrootD streams and sends them to the RabbitMQ by TCP. After RabbitMQ receives the message, a monitoring collector logs the message. The queue of messages in RabbitMQ is then indexed in the Elasticsearch database, from where statistics and useful information can be mined.

In addition to the XRootD monitoring streams, we employ tests in Checkmk (<https://checkmk.com>) designed to check the end-to-end OSDF file access. These aim to detect issues with the infrastructure before they become problematic for users. For example, we would be interested for a view into the state of the network bandwidth between origins and caches. A daily cadence test checking transfer rates between origins to caches with sample data can reveal problems, allowing time for teams to respond before a production pipeline running at execution points starts suffering from performance degradation in the data delivery. We list below a comprehensive list of tests we perform on the OSDF infrastructure that ensures that staff get advanced warnings on the state of the service:

- (1) Monitor authenticated access using certificates and tokens (caches).
- (2) Check the protected files on OSDF. OSDF has a structure of protected files using SciTokens [13] or certificates. Checking if accessing the file without the correct credentials is possible (caches).
- (3) Monitor the number of messages the shoveler processed (caches and origins).
- (4) Check the size of the shoveler queue (caches).
- (5) Check if copying files from public and private origins (caches and origins) is possible.
- (6) Verify the transfer rate considering the host location (caches and origins).
- (7) Verify the SSL certificate (caches and origins).

- (8) Check the load of each cache and origin in NRP (caches and origins).
- (9) Check the CVFMS access (caches).
- (10) Verify if the redirector is running and operational.

5 CONCLUSION

We have presented a list of operational and monitoring upgrades of the Open Science Data Federation. Storage space tracking, monitoring checks, monitoring stream collections, and adding new caches has increased the robustness and capacity of this data delivery infrastructure. New monitoring processes provide a way to detect problems before the users. New caches expand nearby proximity access to compute sites. New origins increase storage capacity. File access streams allow for more efficient debugging and support. These improvements create a way of the data delivery experience in the computing ecosystem fostered by the Open Science Consortium.

New software and caching techniques are being added frequently to the Open Science Data Federation software. In the future, we will incorporate software from the Pelican [4] project, which will increase visibility for cache and origin owners, as well as provide better traceability of usage by clients.

ACKNOWLEDGMENTS

This work was supported in part by National Science Foundation (NSF) awards #1836650, CNS-1730158, ACI-1540112, ACI-1541349, OAC-1826967, OAC-2030508, OAC-2112167, CNS-2100237, CNS-2120019, PHY-2323298, the University of California Office of the President, and the University of California San Diego's California Institute for Telecommunications and Information Technology/Qualcomm Institute. Thanks to CENIC for the 100Gbps networks. Also, Fabio would like to acknowledge Erick Baltes and Paschalis Paschos.

REFERENCES

- [1] 2024. Home - Access — access-ci.org. <https://access-ci.org/>. [Accessed 20-Apr-2024].
- [2] 2024. National Research Platform. <https://nationalresearchplatform.org/>. [Accessed 20-Apr-2024].
- [3] 2024. OSDF Shoveler. <https://osg-htc.org/docs/data/xrootd/install-shoveler/>. [Accessed 20-Apr-2024].
- [4] 2024. Pelican Platform. <https://pelicanplatform.org/>. [Accessed 4-Jun-2024].
- [5] 2024. XRootD. <https://xrootd.slac.stanford.edu> [Accessed 20-Apr-2024].
- [6] Vitor R. Coluci, Fabio Andrijauskas, and Sócrates O. Dantas. 2023. 8 - Modeling thermal conductivity with Green's function molecular dynamics simulations. In *Modeling, Characterization, and Production of Nanomaterials (Second Edition)* (second edition ed.), Vinod K. Tewary and Yong Zhang (Eds.). Woodhead Publishing, 171–187. <https://doi.org/10.1016/B978-0-12-819905-3.00008-7>
- [7] 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 on Systems and Network Telemetry and Analytics (HPDC '23)*. ACM. <https://doi.org/10.1145/3589012.3594897>
- [8] Alvise Dorigo, Peter Elmer, Fabrizio Furano, and Andrew Hanushevsky. 2005. XROOTD-A Highly scalable architecture for data access. , 348–353 pages.
- [9] E Fajardo, A Tadel, M Tadel, B Steer, T Martin, and F Würthwein. 2018. A federated Xrootd cache. *Journal of Physics: Conference Series* 1085, 3 (sep 2018), 032025. <https://doi.org/10.1088/1742-6596/1085/3/032025>
- [10] David Schultz, Igor Sfiligoi, Benedikt Riedel, Fabio Andrijauskas, Derek Weitzel, and Frank Würthwein. 2023. IceCube experience using XRootD-based Origins with GPU workflows in PNRP. arXiv:2308.07999 [physics.comp-ph]
- [11] Derek Weitzel, Brian Bockelman, Duncan A. Brown, Peter Couvares, Frank Würthwein, and Edgar Fajardo Hernandez. 2017. Data Access for LIGO on the OSG. In *Proceedings of the Practice and Experience in Advanced Research Computing 2017 on Sustainability, Success and Impact* (New Orleans, LA, USA) (PEARC'17).

- Association for Computing Machinery, New York, NY, USA, Article 24, 6 pages. <https://doi.org/10.1145/3093338.3093363>
- [12] Derek Weitzel, Marian Zvada, Ilija Vukotic, Rob Gardner, Brian Bockelman, Mats Rynge, Edgar Fajardo Hernandez, Brian Lin, and Mátyás Selmei. 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>
- [13] Alex Withers, Brian Bockelman, Derek Weitzel, Duncan Brown, Jeff Gaynor, Jim Basney, Todd Tannenbaum, and Zach Miller. 2018. SciTokens: Capability-Based Secure Access to Remote Scientific Data. In *Proceedings of the Practice and Experience in Advanced Research Computing* (Pittsburgh, PA, USA) (PEARC '18). Association for Computing Machinery, New York, NY, USA, Article 24, 8 pages. <https://doi.org/10.1145/3219104.3219135>