



The Pelican in Flight: Delivering Throughput Data





Pelican and Throughput Computing

The Pelican Project builds a software platform that delivers data to throughput computing.

The OSDF (Open Science Data Federation) is a global deployment of the platform.



What is the OSDF ?

... By Analogy...



+



“Netflix for Science”

Allowing computational workloads to stream data



What is the **OSDF** ?



... By Analogy...



“Cloudflare for Science”

Infrastructure for scaling access to your storage



OSDF enabling science



- The LIGO collaboration stores their interferometer data at Caltech.
 - Scientists around the world want to access and use it from the computing *they* have access to.
 - LIGO wants to enable science beyond the (finite) compute capacity at Caltech.
- LIGO has connected their proprietary and public data to the OSDF.
 - Any browser, python notebook, or compute cluster access LIGO data via the OSDF.



In 2024, LIGO moved 27PB of objects; 200TB of unique data.



Not just “the big guys”

- At U. Hawaii, an individual PI - Curt Dodds – uses their CC*-funded storage to share data from the SPIN4D project
 - [doi:10.5281/zenodo.13879854](https://doi.org/10.5281/zenodo.13879854)
- If you use the SPIN4D notebook to explore their dataset, you use the OSDF!
 - <https://github.com/ifauh/spin4d-data/blob/v0.5/spin4d-data-exploration.ipynb>

Check out [Curt Dodd's talk on Friday](#).



Research Highlight – U. Hawaii



Decoding the Sun's Secrets With the Help of Pelican

Astronomers and computer scientists at the University of Hawai'i, Moana are unlocking the Sun's magnetic mysteries with the help of cutting-edge data caching and transfer applications provided by the Pelican Project and the Open Science Data Federation (OSDF).

<https://pelicanplatform.org/news/2024/12/20/sun-secrets>



Operated by PATH
(NSF #2030508)

OSDF - A national object delivery service



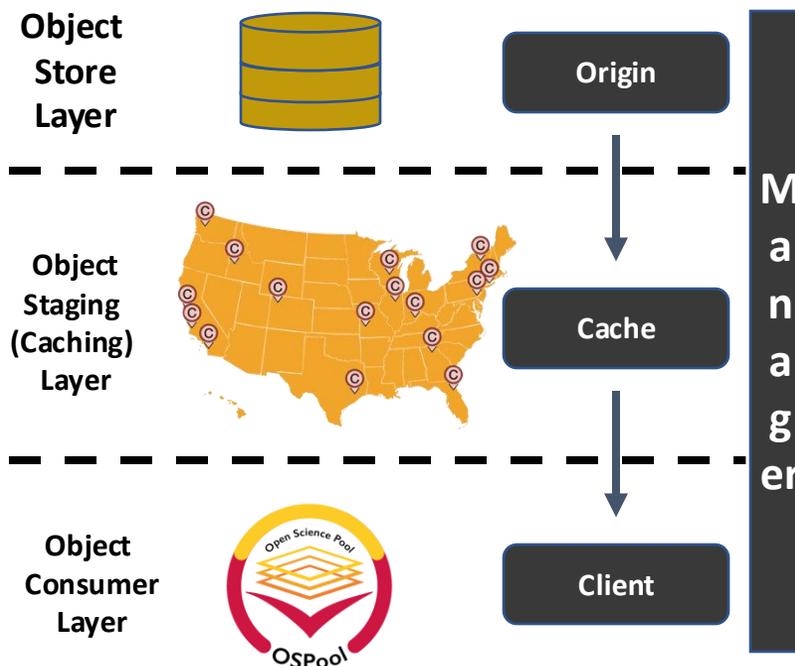
OSDF is powered by the Pelican software
(NSF #2331480).

Value Proposition

- **Shares** your data (Objects) with authorized consumers.
- **Manages** the load of Object reuse on your Object Store.
- **Fuses** Objects from diverse Object Stores through a common namespace and API.
- **Stages** Objects in network-embedded storage.

Who can benefit?

- **Open** to any US researcher, data (object) provider, or computing facility.



Who is using the OSDF?

- Workloads that have significant **reuse and redelivery** of objects at remote autonomous Object Stores.
- Object Stores with limited capacity or egress costs.
- **NOAA**: Via AWS OpenData
- **NSF**: NCAR, NRAO, LIGO, EHT
- **DOE**: FNAL (LHC, DUNE), JLab, Fusion

Key concepts

- **Object**: both data & non-data objects (containers, models, libraries).
- **Object Store**: a S3-compatible endpoint, filesystem, or HTTP webserver.
- **Origin**: connects an Object Store with national infrastructure. Objects are moved from (GET) and to (PUT) the store through Origin.

>140
PB moved

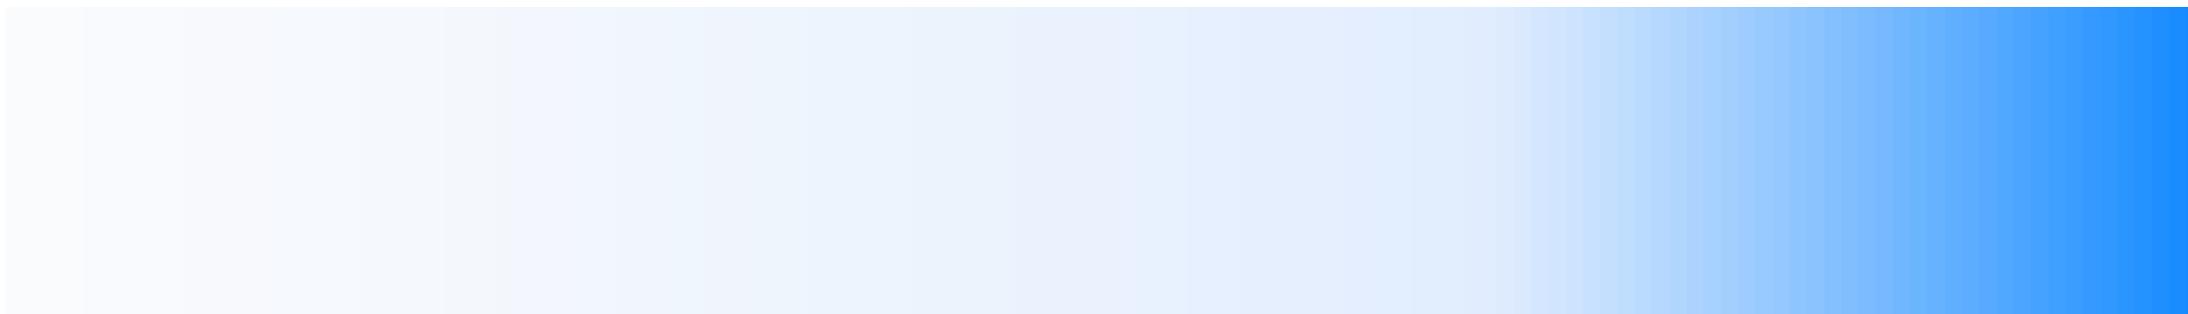
>3.5B
Objects delivered

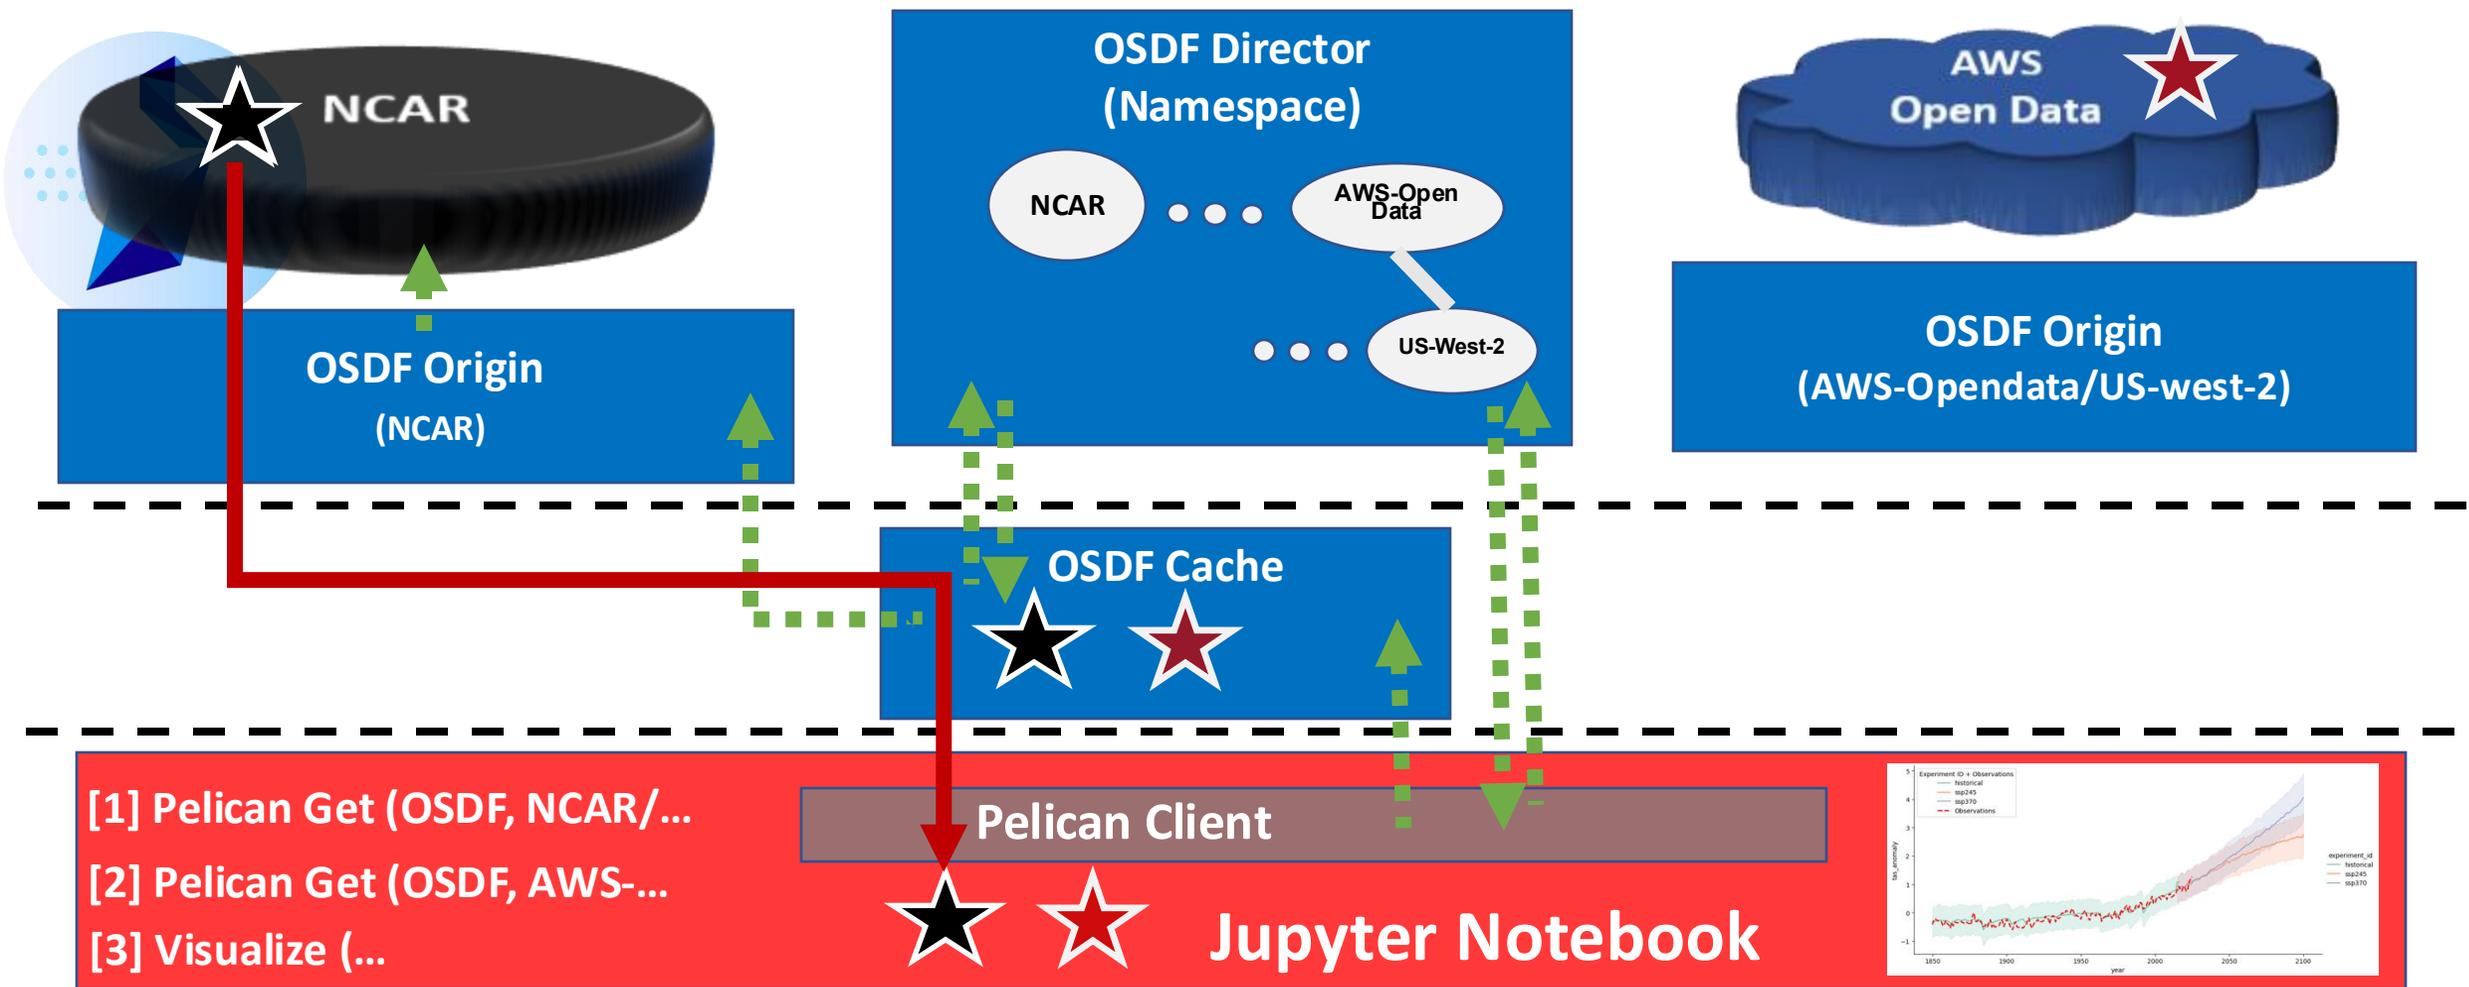
In the last 12 months.

34 caches throughout the world, at points of presence within major compute centers and the global R&E networks (ESNet, Internet2).



OSDF – Core Concepts





Researcher uses a Jupyter Notebook to create a visualization that requires two objects:

```

★ NCAR/rda/harshah/osdf_data/HadCRUT.5.0.2.0.analysis.summary_series.global.monthly.zarr
★ AWS-OpenData/US-West-2/cmip6-pds/CMIP6/CFMIP/NCAR/CESM2/aqua-4xCO2/r1i1p1f1/Amon/co2mass/gn/v20190816
  
```



What happened in the last year?





Clients

- **Command line:** End-to-end checksumming, token auto-discovery.
- **HTCondor Plugin:** Extensive monitoring reporting, end-to-end checksums.
- **Python library / FSSpec:** First official release! Integrates with the “data science” ecosystem.
- **Web Browser:** Simple object download. Browser-based OAuth2 support in-progress.

Explore the OSDF

at <https://osg-htc.org/services/osdf/data>

NCEI Water Column Sonar Data

NOAA collects and uses active acoustic (or sonar) data for a variety of mapping requirements. Water column sonar data focus on the area from near the surface of the ocean to the seafloor. Primary uses of these specific sonar data include 3-D mapping of fish schools and other mid-water marine organisms; assessing biological abundance; species identification; and habitat characterization. Other uses include mapping underwater gas seeps and remotely monitoring undersea oil spills. NCEI archives water column sonar data collected by NOAA line offices, academia, industry, and international institutions.

[View Datasets](#)

Organization

National Oceanic and Atmospheric Administration

Field of Research

Fishing and Fisheries Sciences and Management

Download a Public Object

With Pelican Client on the Command Line

```
pelican object get osdf:///noaa/fisheries-1/noaa-wcsd-pds/data/raw/Henry_B._Bigelow/HB2403/README_HB2403_EK80.md
```

From Your Browser

[Click to Download Public Object](#)

1

Dataset(s)

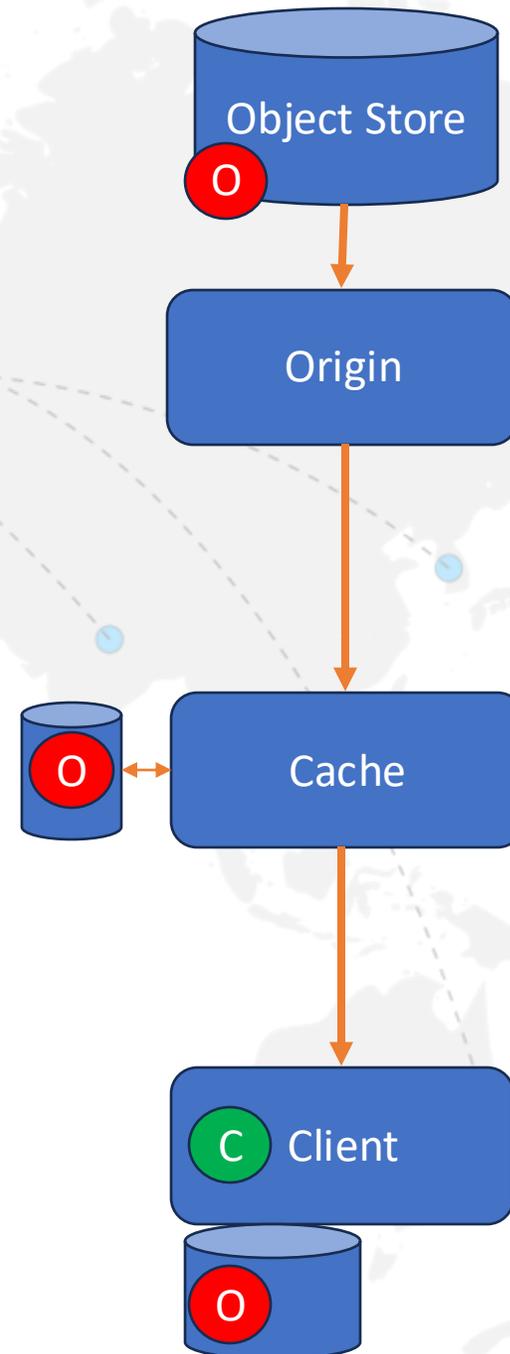
332

Terabytes



End-to-end checksumming

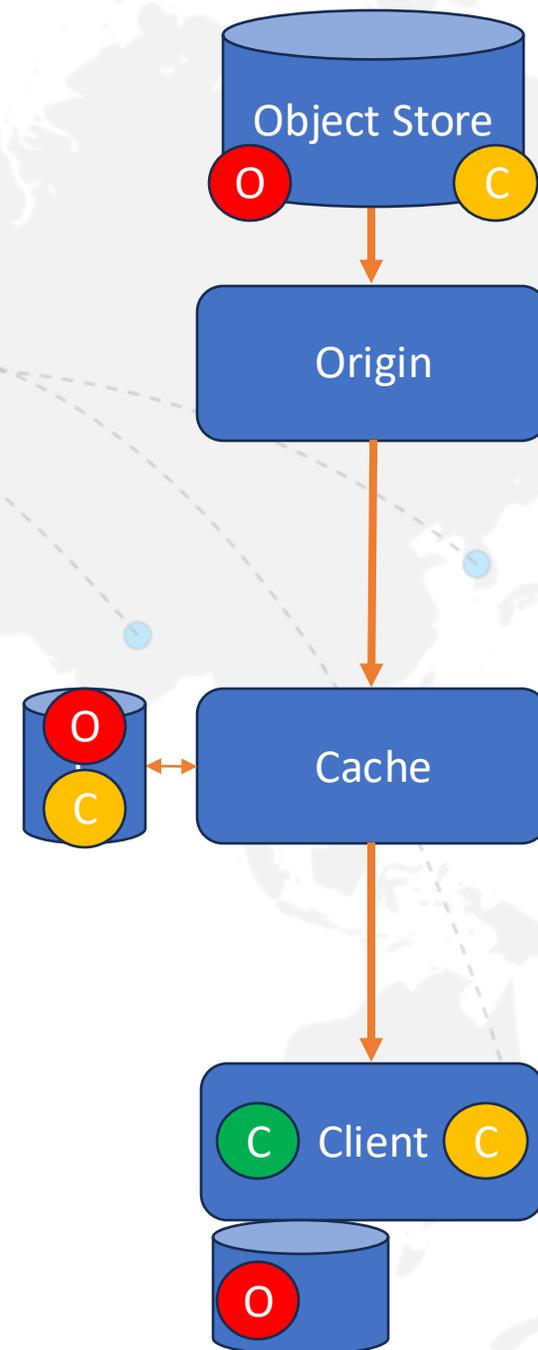
- Objects (**O**) travel from the object store through the data federation and to the client.
 - A copy may be made at the cache for repeated access.
- Ultimately, a copy of the object is made on disk by the client
- **New (v7.16)**: the client will also calculate the object's checksum (**C**) as it is downloaded.





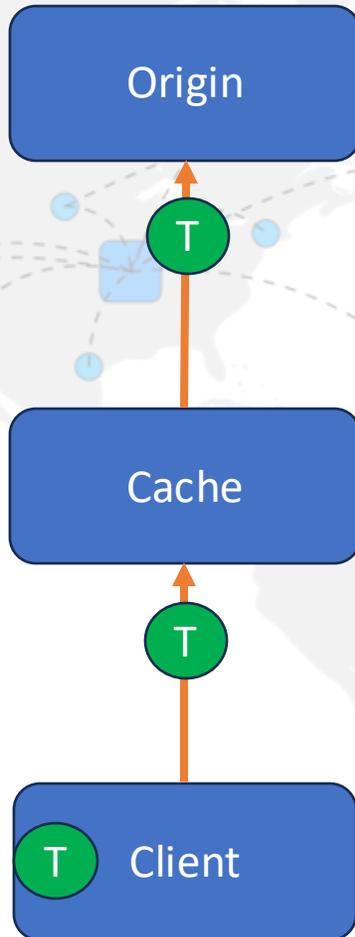
End-to-end checksumming

- Objects **(O)** travel from the object store through the data federation and to the client.
 - A copy may be made at the cache for repeated access.
- Ultimately, a copy of the object is made on disk by the client
- **New (v7.16)**: the client will also calculate the object's checksum **(C)** as it is downloaded.
- **New (v7.16)**: the client will request a copy of the server-side checksum **(C)** in the object store and compare the two, failing the transfer on mismatch.





Innovating in Authorization



Who is the origin talking to?

- Originally, the client acquired a token (T) which was used to authorize the request to the cache.
- If the cache does not have the object, it will forward the request to the origin.
- Conveniently, the same token can also be forwarded to the origin to authorize the cache's request.
- **Problem:** How does the origin distinguish a cache from a "badly behaved client"?
 - Cannot prevent cache bypass!



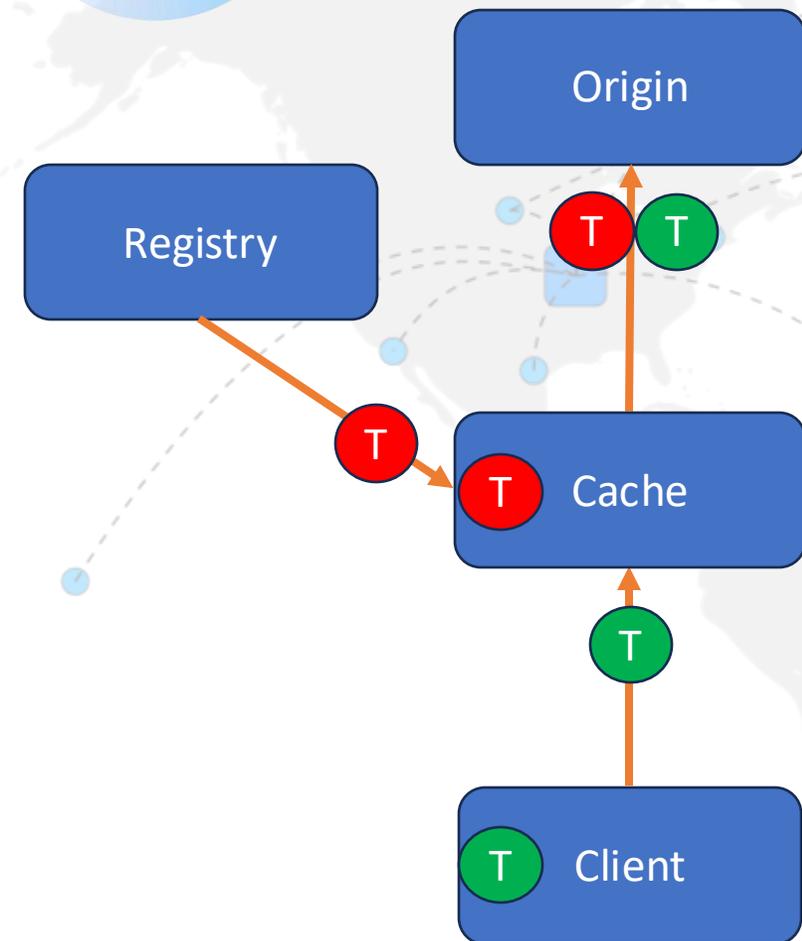
Innovating in Authorization

For more information about authorization, see [Justin's talk later today!](#)



New!

- Caches receive a token **(T)** from the registry service.
 - The token identifies the bearer as a service part of the federation.
- Origins can require:
 - A request be performed by a known service **(T)**,
 - On behalf of an authorized client **(T)**.
- Clients can no longer talk directly to origins, pretending to be a cache!





Powering Storage Backends

For more about our summer plans for the Globus backend, [see the presentation today](#) by our summer fellow, William.

Each origin is configured with a plugin based on the object store type:

- **S3**: Allows integration with AWS and compatible services. Support for write, object listings, and an overhaul of error handling in the last year.
- **HTTP**: Connects generic web servers to the data federation. Read-only.
- **Globus**: Connects a Globus collection to the federation. Read-only.
- **xroot**: Full-featured! The “native language” of the XRootD server.
- **POSIX**: Still the most popular!

What happens when an “overly enthusiastic user” launches a **large workload with little reuse**?

- The “throttling” component purposely slows the user requests to protect the object store.
- **Object store protection** is a core value proposition provided by the OSDF, *even without reuse*.
- New: The throttle is now multi-user aware. Only the “overly enthusiastic user” is slowed, not every user.

Much to do: When HTCSS manages the workload, there’s plenty of opportunities for the AP to be aware



Keeping a Close Eye on You

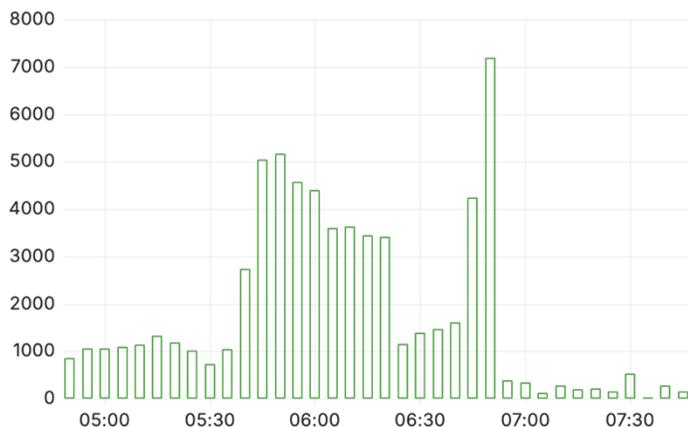
For a deeper dive in monitoring, [check out Patrick's talk](#) on Wednesday.



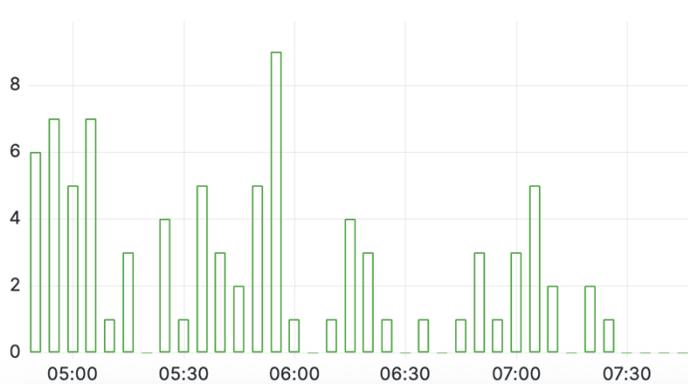
We've historically drowned in monitoring data – most of it not useful for operations.

- We leverage reports from the HTCondor plugin to understand the “typical user experience” for the OSPool.
- The OSPool APs aggregate these per-file documents and we can turn these into reports or charts.
- **Assumption:** The OSPool is distributed enough that it provides a “good proxy” for a random user.
 - Has been essential for spotting common failure cases which have identified ~half dozen serious bugs.
- **TODO:** Working to better separate “infrastructure problems” (random EOF) from “user problems” (404 file not found).

OSDF Successes (5 min)



OSDF Failures (5 min)



Per OSDF endpoint download (i.e. input transfer) statistics

Endpoint Name	Endpoint Institution	Total Attempts	Total Jobs	Successful Attempts	Successful Jobs	Failed Attempts	Pct Attempts Failed	Failed Attempts per Job	Num Jobs w/ Failed Attempts	Num Jobs Interrupted	Pct Jobs Interrupted
TOTALS		731,926	704,066	717,015	704,030	14,911	2.0%	0.02	1,487	1,14	0.2%
MGHPCC_NRP_OSDF_CACHE	Massachusetts Green High Performance Computing Center	234,469	234,319	234,411	234,307	58	0.0%	0.00	56	2	0.0%
CHTC_PELICAN_CACHE	University of Wisconsin-Madison	117,061	114,209	113,088	113,119	3,973	3.4%	0.03	1,015	294	0.3%
BOISE_INTERNET2_OSDF_CACHE	Internet2	106,785	106,192	106,772	106,183	13	0.0%	0.00	13	1	0.0%
JACKSONVILLE_INTERNET2_OSDF_CACHE	Internet2	97,903	95,601	97,846	95,568	57	0.1%	0.00	55	4	0.0%
Stashcache-Kansas	Internet2	87,728	83,374	82,934	82,271	4,794	5.5%	0.06	1,189	942	1.1%
Stashcache-Chicago	Internet2	30,689	26,969	26,661	25,935	4,028	13.1%	0.15	1,055	642	2.4%
Sunnyvale-I2-PRP	Unmapped endpoint Sunnyvale-I2-PRP	17,580	17,573	17,568	17,561	12	0.1%	0.00	12	0	0.0%



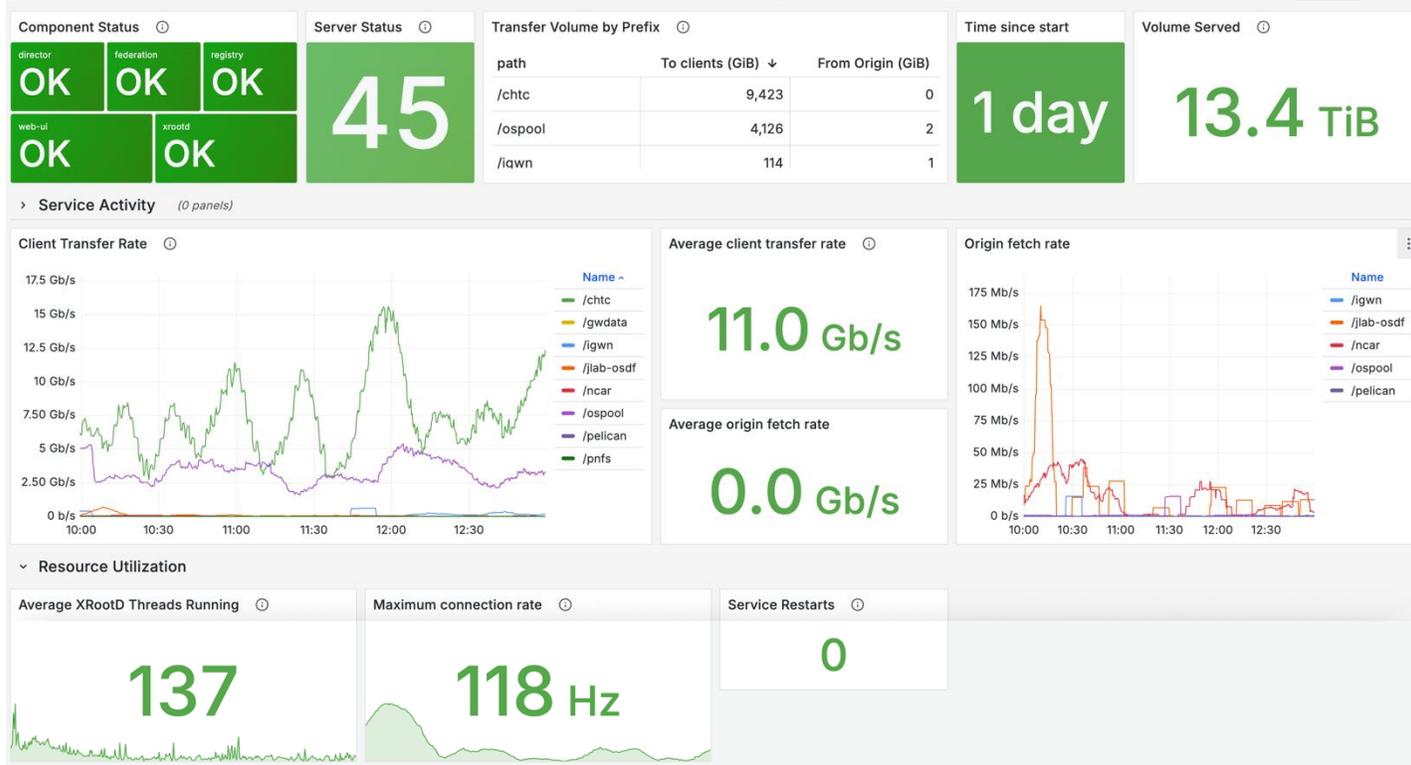
Keeping a Close Eye on You

For a deeper dive in monitoring, [check out Patrick's talk](#) on Wednesday.



We still maintain various internal dashboards, typically Grafana-based.

Pushing more toward identifying errors and problems – less “big pretty numbers”.





NDC Pathfinders

- Pelican was funded as part of NSF’s “National Data Cloud” for Climate initiative in 2023.
 - The NDC brings together a broad set of domain science and cyberinfrastructure projects, allowing one to assemble end-to-end approaches spanning the vertical stack of expertise.
- The NDC Pathfinders collaboration consists of Pelican and 4 earth science-driven projects that aim to show potential approaches (“paths”) through the NSF-funded cyberinfrastructure.

NDC Pathfinders Collaboration at Madison in 2024.

<https://ndc-pathfinders.org/>



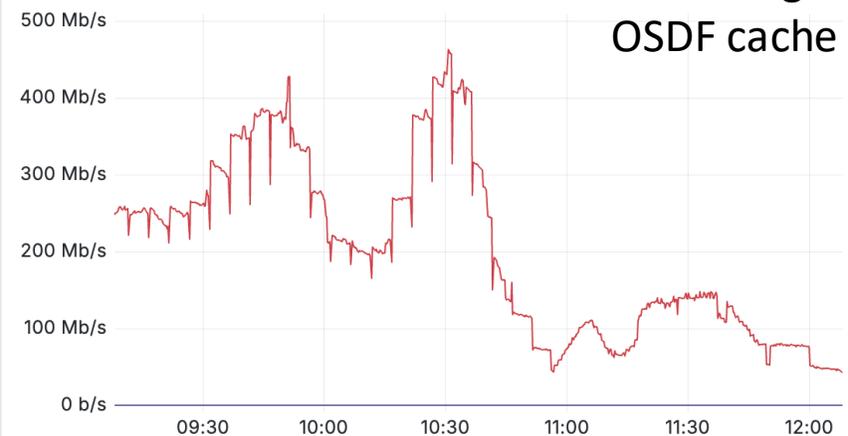


NCAR Research Data Archive

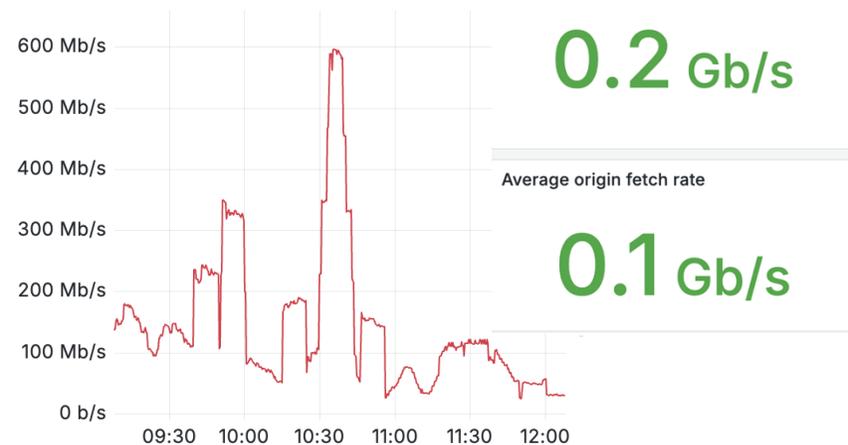
To learn more about using the NCAR data via OSDF, [see Harsha's talk tomorrow.](#)



Client Transfer Rate ⓘ View from a single OSDF cache

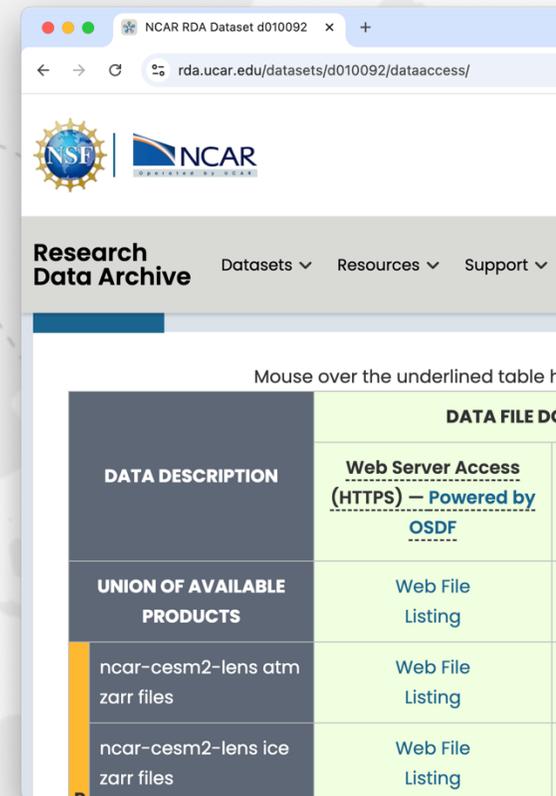


Origin fetch rate Average client transfer rate ⓘ



NCAR's part in the NDC Pathfinders is to make its data available through the OSDF and to inspire its community on how to leverage it.

- In May 2025, NCAR's Research Data Archive (RDA) switched all its HTTP-based download links to use the OSDF.
- More importantly, Harsha Hampapura has been developing material for the community demonstrating [how the OSDF can be leveraged](#) via environments including Jupyter, OSPool, and Dask.





NAIRR

For an NDP / OSDF / Pelican example, check out [Curt Dodd's talk on Friday](#).



Pelican is one of the data-centric projects of the National AI Research Resource (NAIRR) Pilot. We collaborate with the [National Data Platform](#), which builds out data analysis services on top of the [NRP](#).

Our plans:

- Provide technology (Python FSSpec library) to allow common AI tools (PyTorch) to stream from the OSDF.
- Ensure NAIRR datasets are available via the OSDF.
- With PATH, show the the vision of Throughput machine learning to train ensembles of models and high throughput inference.

[Watch this space!](#)



Retiring the old OSDF

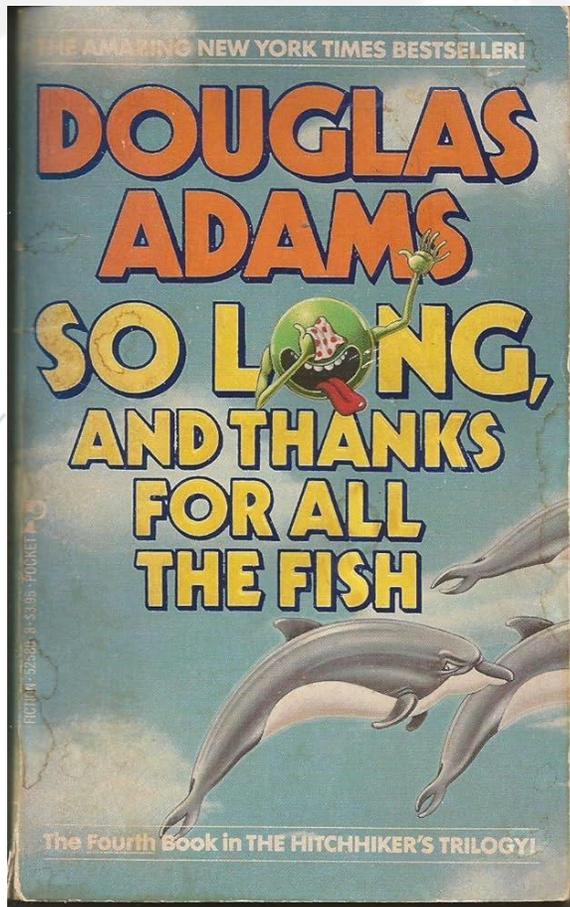


Image from <https://www.amazon.com/So-Long-Thanks-all-Fish/dp/0671525808>

Before the Pelican Project was funded, the OSDF was built “one patch at a time”:

- 2011: The AAA project (NSF #1104664) developed a file cache based on the XRootD software.
- 2015: Started by deploying XRootD-based caches for the [Stash filesystem @ UChicago](#).
- 2020: [Grew into a generic system](#) for multiple origins. Leveraged WLCG services for GeolP and OSG Topology for cache discovery.
- 2023: Pelican Project started, re-engineering the federation components while keeping OSDF operational.
- **2025: Retired the pre-Pelican components in OSDF!**



The View from Above





Pelican is now >1 yr old



The Condor Project is 40; the OSG Consortium is 20; the Pelican Project is 1.

- The “translational CS” loop is just starting!
 - The laboratory is churning out new ideas – still iterating on fundamental concepts.
 - The local – the OSDF – is now 100% Pelican.
 - We are rapidly building out the community.

From the Laboratory to the Community, Advancing Throughput Computing Through **Translational Computing**



Releases 145

Contributors 28

v7.16.5 Latest last week



+ 144 releases

+ 14 contributors

Commits over time

Weekly from Aug 19, 2023 to May 24, 2025





Who is the Community?

The “locale” is the data federation. Who is the community?

- **OSPool users** leveraging the OSDF for powering computational workflows across the nation.
- **UW-Madison researchers** using the ‘UWDF’ for on-campus data.
- **Collaborations**, like LIGO, distributing data to large-scale workloads.
- **Individual PIs** publishing & sharing their own datasets.
- Scientists reading out **community datasets** (NCAR data).

You are an essential part of this project!
Broadening the community makes Pelican better!



Alexander Tuna
Computational Data Scientist and Researcher
University of California-San Diego



Andrew Owen
Research Computing Facilitator
University of Wisconsin-Madison



Austin Schneider
Student Web Developer
Morgridge Institute For Resea



Howard Zhong
Research Software Engineer
Morgridge Institute for Research



Alja Mrak Tadel
Analytic Programmer
University of California San Diego



Matevz Tadel
Project Scientist
University of California San Diego



Brian Aydemir
Systems Integration Developer
University of Wisconsin-Madison



Patrick Brophy
Student Research Software Engineer
Morgridge Institute for Research



Brian Bockelman
Principal Investigator
Morgridge Institute for Research



Rich Wellner
SDx Director
San Diego Supercomputer Center



Brian Lin
OSG Software Area Coordinator
University of Wisconsin-Madison



Christina Koch
Lead Research Computing Facilitator
University of Wisconsin - Madison



Miron Livny
Co-Principal Investigator
University of Wisconsin-Madison



Frank Wuerthwein
Co-Principal Investigator
University of California San Diego



Tae Kidd
Project Manager
Morgridge Institute For Research



Mátyás Selmeci
Software Integration Developer
University of Wisconsin-Madison



William Swanson
Research Cyberinfrastructure Specialist
University of Wisconsin-Madison



Justin Hiemstra
Research Software Engineer
Morgridge Institute For Research



Cannon Lock
Web Developer
Morgridge Institute for Research



Emma Turetsky
Research Software Engineer
Morgridge Institute for Research

We Love Feedback!

Pelican sits at the “infrastructure layer” and is not always visible to the user.

We value your participation in events like HTC25 – please strike up a conversation with a Pelican!



Questions?

This project is supported by the National Science Foundation under Cooperative Agreements OAC-2331480. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.



Backup – Things Coming Up in 2025!

As you meet with us in the coffee breaks this week, ask about any of these upcoming development topics for the next year:

- Improved cache storage management as part of the Kingfisher project.
- Deployment of high-availability central components.
- Origins behind firewalls – no incoming connectivity.
- User-shareable “collections”.
- Write-through caches
- Cache integration with S3 storage
- Finish Globus integration