

# Open Science data federation and National Research platform

**Fabio Andrijauskas, PhD**

[fandrijauskas@ucsd.edu](mailto:fandrijauskas@ucsd.edu)

UCSD - SDSC



# WHY DO PIRATES LOVE TRIGONOMETRY AND DATA TRANSFER?

Because they need TRIGONOMETRY to calculate their bearings and be...

...and they need DATA TRANSFER to track their loot and analyze their success...

**$\pi$ -RATES!**


**RATES!**

$$A = 60^\circ \quad \Delta \pm 36^\circ =$$

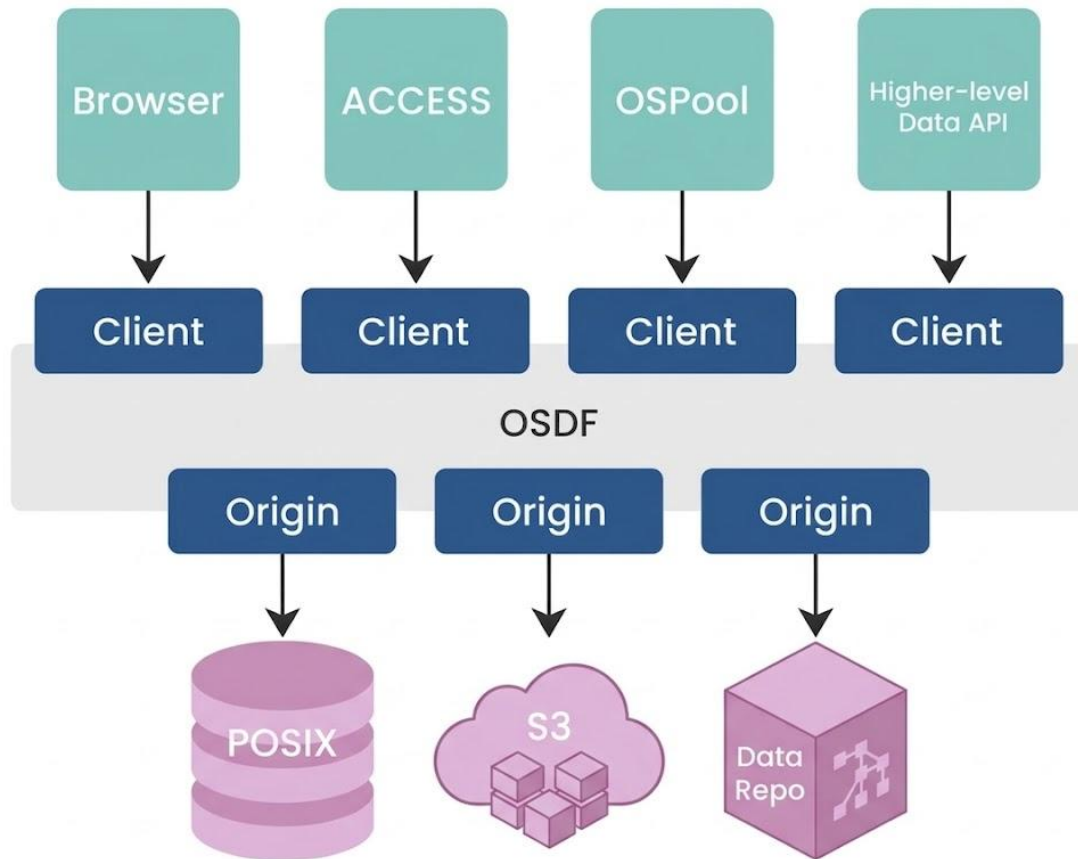
$$\Delta \pm 60^\circ = \sin \alpha \alpha$$



$$E = \sin \theta = \sqrt{\frac{\cos \sin \alpha}{\sin \alpha}}$$



# Pelican/OSDF



# OSDF map



# Pelican Origin Interface



SDSC\_NRP\_OSDF\_ORIGIN

## Status

CMSD

### Config-updates

Authfile and scitoken.cfg file updated successfully

### Director

Director object transfer test succeeded at: 2026-06-09 22:15:06

### Federation

### Oa4mp

### Prometheus

Server is ready to receive web requests.

### Registry

### Storage

All monitored filesystems have adequate storage

### Web UI

### XRootD

Self-test monitoring cycle succeeded at 2026-06-09T22:14:59Z

Last Updated: Jun 9, 2026, 5:15 PM

## Data Exports

### Origin

Registration Completed

### Namespaces

Federation Prefix

/ndp/burnpro3d

Storage Prefix

/mnt/origin/ndp/burnpro3d

PublicReads ✓

Reads ✓

Writes ✓

Listings ✓

DirectReads ✗

Federation Prefix

/caida/protected

Storage Prefix

/mnt/origin/caida/protected

PublicReads ✗

Reads ✓

Writes ✓

Listings ✓

DirectReads ✗

< 1 2 3 4 5 6 >

# OSDF: Object Transfer Log

```
fabioandrijauskas@macf % pelican -d object get osdf://nrp/cachetest/1g 1g
```

```
DEBUG[10:25:49] Trying object servers: [unl-cache, dtn-pas, osdf1.chic]
```

```
DEBUG[10:25:49] Queuing transfer...
```

```
DEBUG[10:25:49] Downloading object from osdf:///nrp/cachetest/1g to 1g
```

```
DEBUG[10:25:50] First available cache (unl-cache.nationalresearchplatform.org:8443) responded
```

```
DEBUG[10:25:50] Attempting download: job=019eb223-c71b-78d9-a223-d7cf03e596a4
```

```
138.94 MiB / 1.00 GiB
```

```
[=====>-----] 41s | 21.16 MiB/s
```

# New Caches or Origins

---



## RPM Installation

Running an RPM-based installation on various infrastructure environments:

- Bare Metal Servers
- Virtualization Platforms



## Local Container

Efficient deployment utilizing containerization for local environments:

- Using a local Docker setup



## NRP Integration

Leveraging the National Research Platform for distributed systems:

- Direct NRP integration

# NRP

## Distributed System

A distributed system spanning multiple locations

### Components:

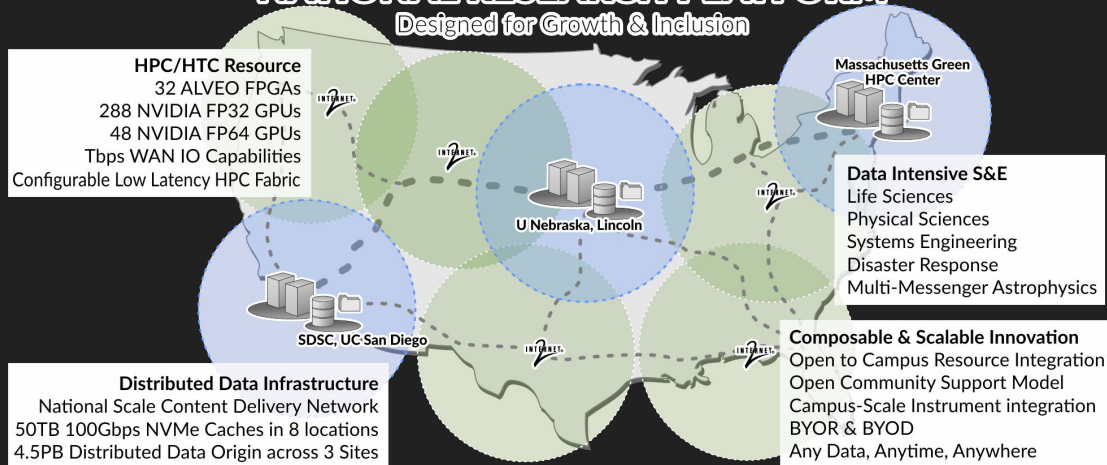
- GPU-based HTC
- GPU/FPGA HPC
- Storage

### Management:

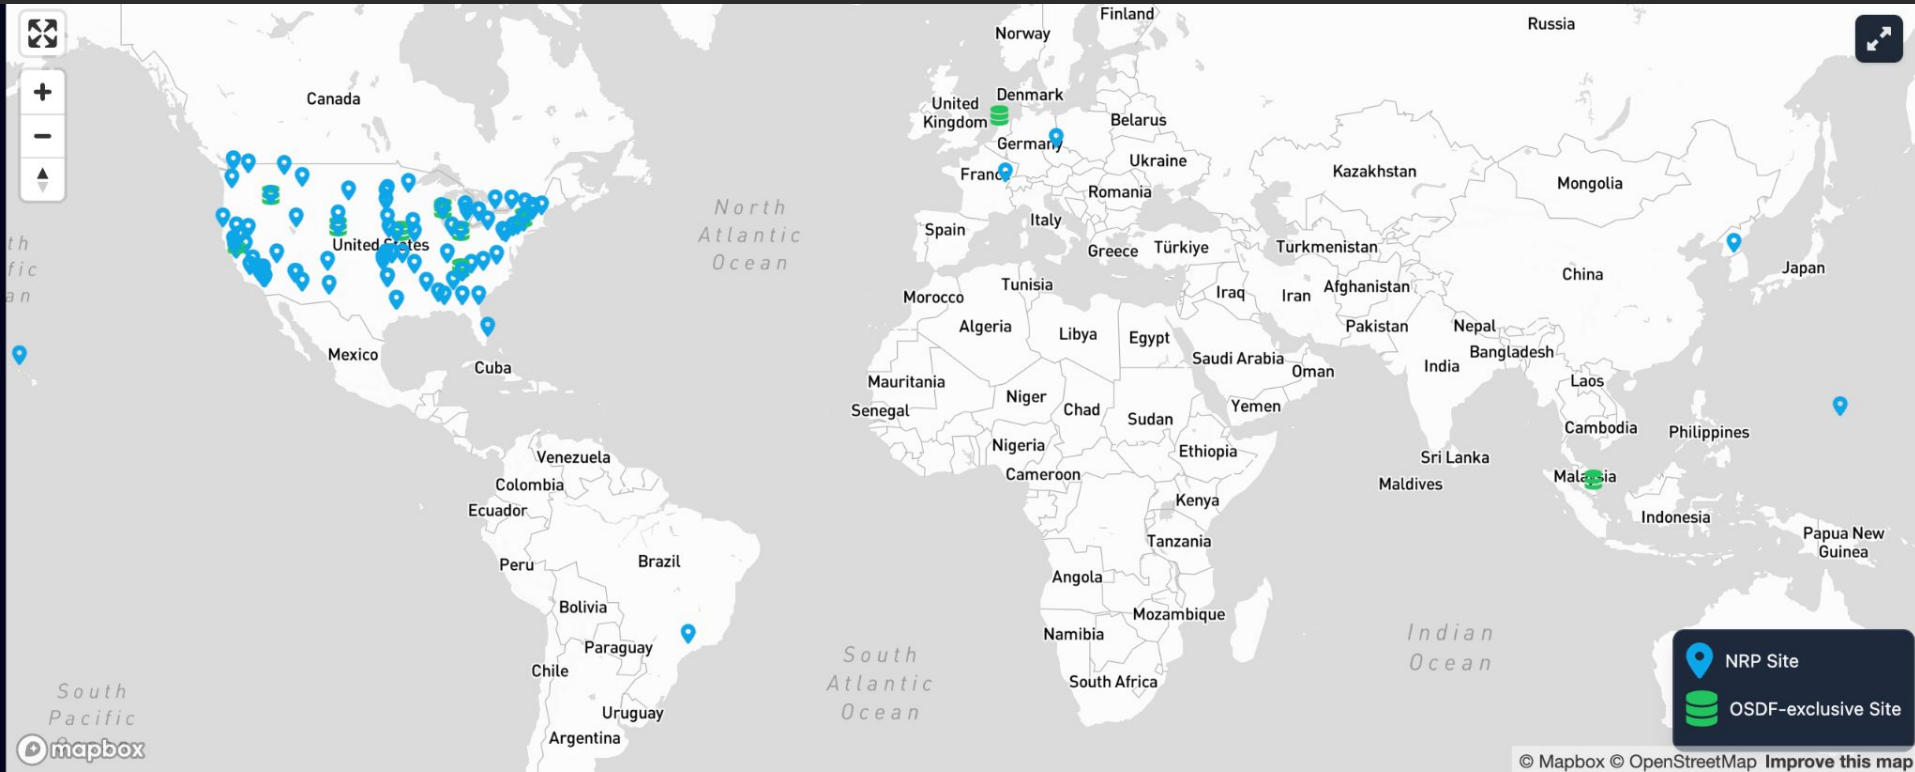
Managed via unified Kubernetes setup

## NATIONAL RESEARCH PLATFORM

Designed for Growth & Inclusion



# NRP map



# Using NRP

---

## Node Integration into NRP

- Public IP
- IPMI Access
- Root access
- Firewall controlled by NRP
- **Full Control Maintained**

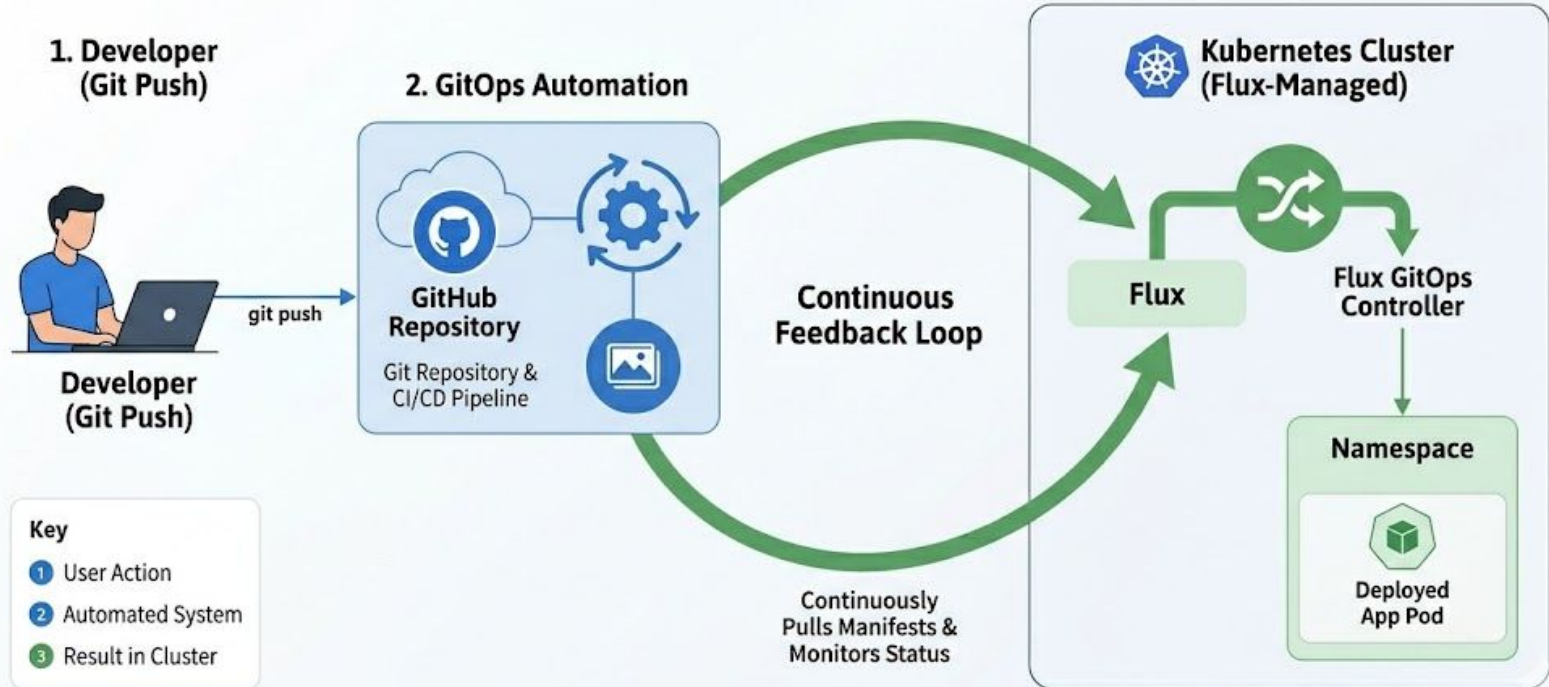


**Documentation:**


<https://nrp.ai/documentation/admindocs/participating/new-contributor-guide/>

# Deployment workflow







## Flux GitOps Workflow: GitHub to Kubernetes



# Deployment workflow

 **biozit** adding - name: "OSDF\_CACHE\_DATALOCATION" value: "/cache/area"

6088ce2 · 2 weeks ago  **History**

Name	Last commit message	Last commit date
 ..		
 amsterdam-i2-cache	Update pelican-cache image version to v7.24.3	2 weeks ago
 amsterdam-pelican-cache	Upgrading caches to 7.25.0	2 weeks ago
 boise-i2-pelican-cache	Upgrading caches to 7.25.0	2 weeks ago
 chicago-i2-pelican-cache	Upgrading caches to 7.25.0	2 weeks ago
 chicago-i2-squid	Swaps toleration on chicago squid	last year
 cincinnati-i2-pelican-cache	Upgrading caches to 7.25.0	2 weeks ago
 denver-i2-pelican-cache	Upgrading caches to 7.25.0	2 weeks ago

# NRP Pod view

Context: nautilus [RW]  
 Cluster: nautilus  
 User: oicd  
 K9s Rev: v0.50.9 ⚡v0.51.0  
 K8s Rev: v1.33.8  
 CPU: 12%  
 MEM: 9%

👁 all ⏪ Attach ⌘+K Kill 👁 Show Node  
👁 osg-services ⏪ Delete ⌘+L Logs 👁 Show PortForward  
👁 cms-admin ⏪ Describe 👁 Logs Previous 👁 Transfer  
👁 default ⏪ Edit ⌘+F Port-Forward 👁 YAML  
👁 ⏪ Help 👁 Sanitize  
⌘+J Jump Owner 👁 Shell



pods(osg-services)[68]													
NAME	PF	READY	STATUS	RESTARTS	CPU	%CPU/R	%CPU/L	MEM	%MEM/R	%MEM/L	IP	NODE	AGE
mo-ndp-yaml-private-osdf-pelican-origin-7ddb74cf8c-9f2hz	●	2/2	Running	0	25	0	0	5187	30	28	128.206.235.1	origin-g7r4704.rnet.missouri.edu	21d
mo-ndp-yaml-public-osdf-pelican-origin-6d79cc6669-tlbwq	●	2/2	Running	0	31	0	0	7788	46	42	128.206.235.1	origin-g7r4704.rnet.missouri.edu	21d
mo-osdf-pelican-origin-77d67cb4d4-td77t	●	2/2	Running	0	27	0	0	3618	10	10	128.206.235.1	origin-g7r4704.rnet.missouri.edu	13d
ncar-gdex-redirectorncar-redirect-gdex-66f65978f4-z84wt	●	1/1	Running	0	0	0	0	16	1	1	10.244.140.92	ncar-osdf00.hpc.ucar.edu	43d
ncar-osdf-pelican-cache-5fc46fddcd-lldfh	●	2/2	Running	0	1770	0	0	248989	27	26	128.117.251.130	ncar-osdf00.hpc.ucar.edu	13d
ncar-osdf-pelican-origin-5c65ff44d-fqrm5	●	2/2	Running	0	4917	15	15	174437	94	94	10.244.218.225	ncar-osdf01.hpc.ucar.edu	5h1m
ncar-redirect-9d894687d-6rhSz	●	1/1	Running	0	0	0	0	15	1	1	10.244.218.212	ncar-osdf01.hpc.ucar.edu	43d
ncar-s3-osdf-pelican-origin-56888cc4cd-cnpl	●	2/2	Running	0	1115	12	12	3481	18	18	10.244.218.235	ncar-osdf01.hpc.ucar.edu	13d
ncraska-nrp-osdf-pelican-origin2-64964fbf94-zbqs5	●	2/2	Running	0	3	0	0	3306	6	6	129.93.153.129	hcc-nrp-shor-c6005.unl.edu	13d
newyork-i2-osdf-pelican-cache-75d46fbc-ggq4	●	2/2	Running	0	5239	2	2	17764	2	2	163.253.29.27	osdf1.newy32aooa.net.internet2.edu	38m
nginx-deployment-8b7fcb74-sgx5s	●	1/1	Running	0	0	0	0	2	1	0	10.244.235.169	k8s-u200-00.calit2.optiputer.net	15d
nginx-deployment-8b7fcb74-x4st8	●	1/1	Running	0	0	0	0	1	1	0	10.244.200.98	nautilus2.sox.net	57m
nrdstor-yaml-osdf-pelican-origin-575999cbd5-4t5cg	●	2/2	Running	0	570	9	8	4079	7	7	129.93.153.195	hcc-nrdstor-osdf.unl.edu	13d
nsdf-climate1-osdf-pelican-origin-57db4c49c7-jh4fg	●	2/2	Running	0	0	0	0	3848	22	21	129.93.153.129	hcc-nrp-shor-c6005.unl.edu	13d
nsdf-climate2-osdf-pelican-origin-755f6cc9f-svc6b	●	2/2	Running	0	2	0	0	160	0	0	132.249.252.214	node-1-5.sdsc.optiputer.net	70m
nsdf-climate3-data-import-6877ff655-mw9w	●	1/1	Running	0	0	0	0	0	0	0	192.69.103.138	storage-01.nrp.mghpcc.org	13d
nsdf-climate3-osdf-pelican-origin-7579fb95f-taghq	●	2/2	Running	0	3	0	0	14678	87	80	192.69.103.138	storage-01.nrp.mghpcc.org	13d
osdf-pelican-director-644c75d7c4-xhg16	●	2/2	Running	0	0	0	0	6	0	0	10.244.1.236	hcc-nrp-shor-c5930.unl.edu	26d
osg-repo-0	●	1/1	Running	0	16	1	1	1267	44	44	10.244.94.226	hcc-chase-shor-c4719.unl.edu	14d
ospool-cbc-5b7f4fd887-n26gn	●	1/1	Running	0	94	1	1	600	3	2	10.244.218.4	k8s-haosu-15.sdsc.optiputer.net	219m
ospool-cm-8584b4dd74-9wx4b	●	2/2	Running	0	1814	20	10	8217	72	38	10.244.149.118	k8s-haosu-24.sdsc.optiputer.net	2d4h
ospool-images-osdf-pelican-origin-698d594cd-mp8dj	●	2/2	Running	158	35	1	0	6851	74	74	10.244.97.153	storage-01.nrp.mghpcc.org	40d
ospool-images-puller-7b6fd7f448-7blxw	●	1/1	Running	1	0	0	0	44	0	0	10.244.149.250	k8s-haosu-11.sdsc.optiputer.net	9d
ospool-images-web-75956f8d56-8nfa4	●	1/1	Running	1	2	0	0	110	1	1	10.244.52.94	node-2-7.sdsc.optiputer.net	26d
ospool-images-web-75956f8d56-42mck	●	1/1	Running	1	0	0	0	29	0	0	10.244.114.191	k8s-a100-01.suncorridor.org	19d
osstore-ehf-osdf-pelican-origin-66b86c77e8-6m9x	●	2/2	Running	0	3	0	0	3704	7	7	129.93.153.195	hcc-nrdstor-osdf.unl.edu	13d
osstore-epic-osdf-pelican-origin-7dde8dd49c-2s7lc	●	2/2	Running	0	17	0	0	3769	40	40	128.206.235.1	origin-g7r4704.rnet.missouri.edu	13d
osstore-hepsim-osdf-pelican-origin-6fb564c7b9-7d7cq	●	2/2	Running	0	3	0	0	3317	6	6	129.93.153.195	hcc-nrdstor-osdf.unl.edu	13d
osstore-kennesaw-redtop-osdf-pelican-origin-86bdb449c4-agstq	●	2/2	Running	0	155	1	1	36102	54	54	130.218.2.122	osgprdk8s01.kennesaw.edu	35m
pnfs-redirector-osdf-pelican-origin-7d89897f59-xd4jm	●	2/2	Running	0	2	0	0	212	1	1	163.253.29.25	osdf1.chic.net.internet2.edu	42m
sdsc-osdf-pelican-cache-b6f5567c-4k57d	●	2/2	Running	0	171	1	1	36832	23	23	132.249.252.215	node-2-12.sdsc.optiputer.net	67m
sdsc-s3-osdf-pelican-origin-6b9c7cf954-h45vd	●	2/2	Running	0	1	0	0	119	0	0	132.249.252.214	node-1-5.sdsc.optiputer.net	70m
sdsc-yaml-cernopendataosdf-pelican-origin-7dcff664dd-dx7cj	●	2/2	Running	0	1	0	0	136	5	3	132.249.252.214	node-1-5.sdsc.optiputer.net	70m
sdsc-yaml-osdf-pelican-origin-78f57645d-p24sq	●	2/2	Running	0	4	0	0	620	1	1	132.249.252.214	node-1-5.sdsc.optiputer.net	70m
singapore-osdf-pelican-cache-8e86d657ff-tpkq7	●	2/2	Running	0	701	0	0	136584	14	14	103.5.243.15	singapore.nationalresearchplatform.org	13d
sunnyvale-i2-osdf-pelican-cache-7bbb7fbb4-7g7bm	●	2/2	Running	0	585	1	1	185556	60	60	163.253.73.2	osg-sunnyvale-stashcache.nrp.internet2.edu	13d
taac-globus-archive-osdf-pelican-origin-85c94d5476-t195m	●	2/2	Running	0	4	0	0	177	0	0	10.244.254.183	osg-houston-stashcache.nrp.internet2.edu	37m
ucsd-osdf-pelican-cache-8488d6989b-678br	●	2/2	Running	0	443	2	2	68562	44	44	198.17.101.66	stashcache.t2.ucsd.edu	13d
unl-osdf-origin-xenon-5cddf6f678-7zsg7	●	2/2	Running	0	0	0	0	114	0	0	129.93.153.129	hcc-nrp-shor-c6005.unl.edu	26d
unl-osdf-pelican-cache-694f8c4c7-hgbtp	●	2/2	Running	0	533	3	3	87426	42	42	129.93.153.150	hcc-nrp-shor-c5832.unl.edu	13d

# Monitoring - CheckMK

UP	amst-fiona.nationalresearchplatform.org	☰	7	0	0	0	0
63.253.29.13	amst-osdf-xcache01.es.net	☰	7	0	0	0	0
UP	ccpelicanli01.in2p3.fr	☰	7	0	0	0	0
UP	daejeon-kreonet-net2.nationalresearchplatform.org	☰	7	0	0	0	0
UP	daejeon-kreonet-net.nationalresearchplatform.org	☰	8	0	0	0	0
UP	<u><a href="#">dtn-pas.bois.nrp.internet2.edu</a></u>	☰	7	0	0	0	0

# OSDF Publications

## **Open Science Data Federation - operation and monitoring**

Fabio Andrijauskas, Derek Weitzel, and Frank Wuerthwein. 2024. In PEARC '24. ACM, Article 63, 1–5.

<https://doi.org/10.1145/3626203.3670557>

## **Benchmarking the OSDF services to develop XRootD best practices**

Fabio Andrijauskas, Igor Sfiligoi, Frank Würthwein. EPJ Web Conf. 337 01256 (2025)

DOI: 10.1051/epjconf/202533701256

## **Using the OSDF for data distribution: Big Bear Solar Observatory use case**

Sydney Montiel, Alexandra Guadarrama, Fabio Andrijauskas. 2025. In PEARC '25. ACM, Article 94, 1–3.

<https://doi.org/10.1145/3708035.3736032>

## **IceCube experience using XRootD-based Origins with GPU workflows in PNRP**

David Schultz, Igor Sfiligoi, et al. 2024. In EPJ Web of Conferences, vol. 295, p. 11011.

## **Global Profiling of OSDF Services using the Pelican Platform**

Fabio Andrijauskas and Frank Würthwein. CHEP26

# OSDF Publications

## Analyzing Transatlantic Network Traffic over Scientific Data Caches

Ziyue Deng, Alex Sim, Kesheng Wu, Chin Guok, Damian Hazen, Inder Monga, Fabio Andrijauskas, Frank Würthwein, and Derek Weitzel. 2023.

In Proceedings of the 2023 on Systems and Network Telemetry and Analytics (SNTA '23). ACM, NY, USA, 19–22.

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

---

## Fabio Andrijauskas, PhD

fandrijauskas@ucsd.edu

UCSD - SDSC