

OSDF: A Federation for all of Open Science

Brian Bockelman, Fabio Andrijauskas

HTC23, Madison, WI

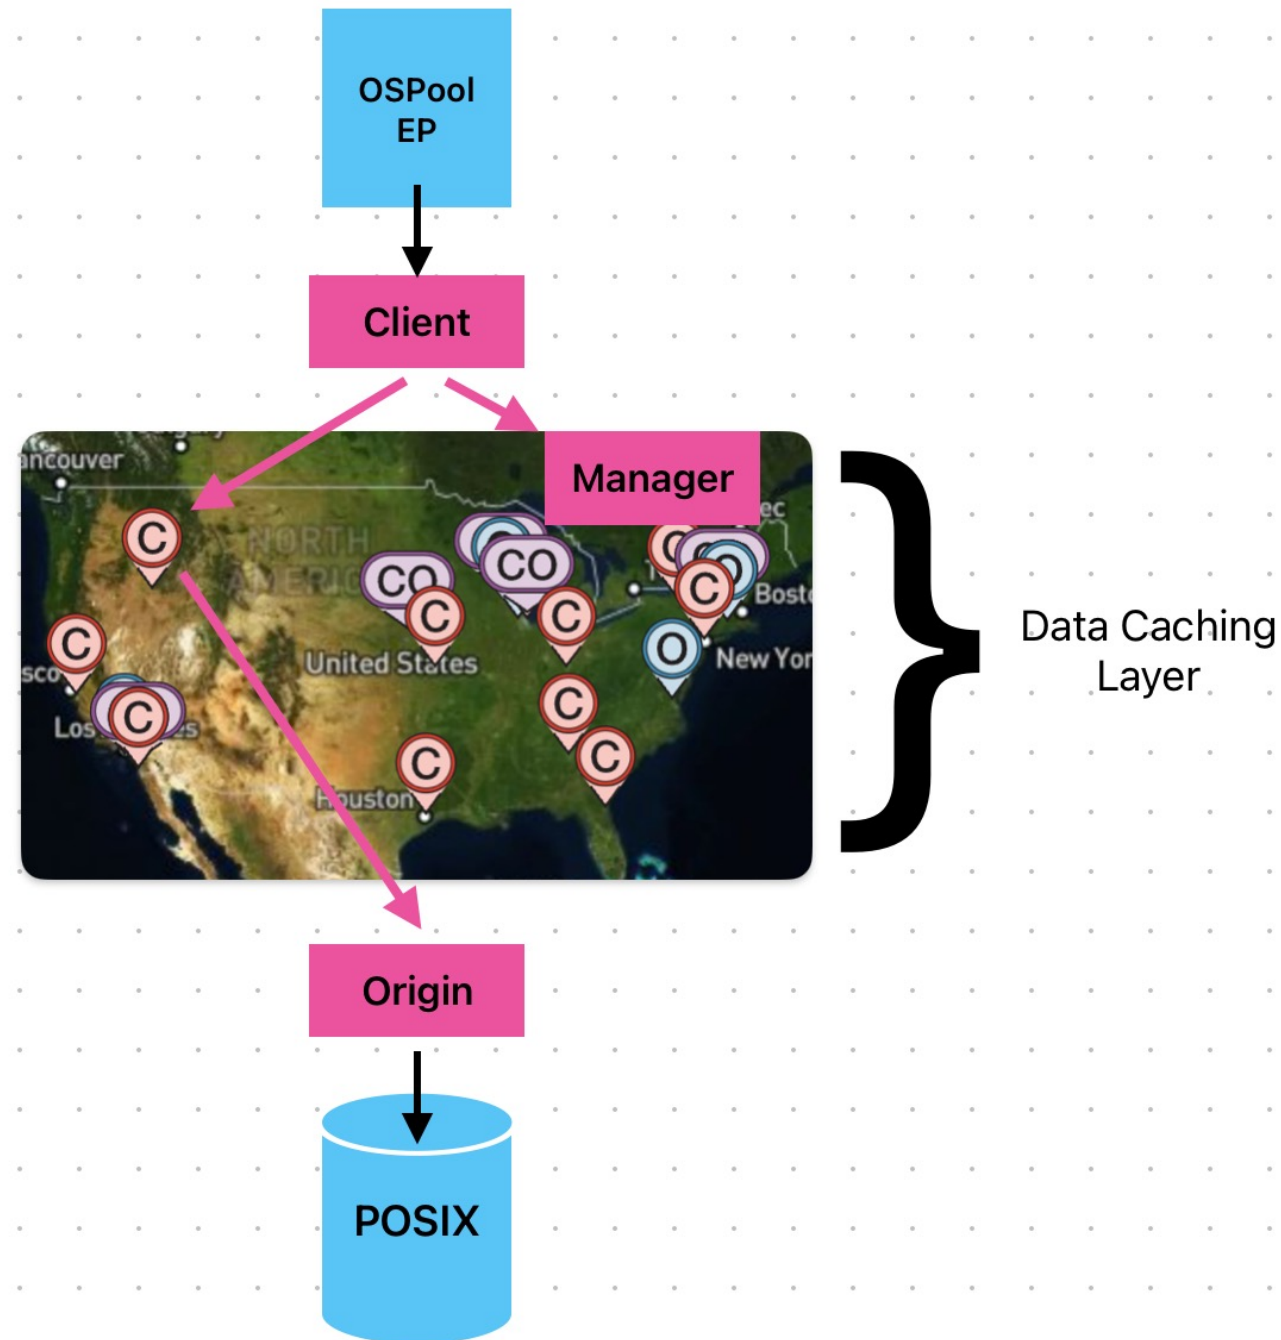
OSDF: In the OSG Fabric of Services

The Open Science Data Federation (OSDF) service federates datasets in disparate repositories into a coherent namespace and to deliver their objects to computational capacity through a network of caches.

- The OSDF connects dataset providers to the federation via the origin service.
- The value to the provider is a uniform set of client tools, authorization models, and help scaling data access.

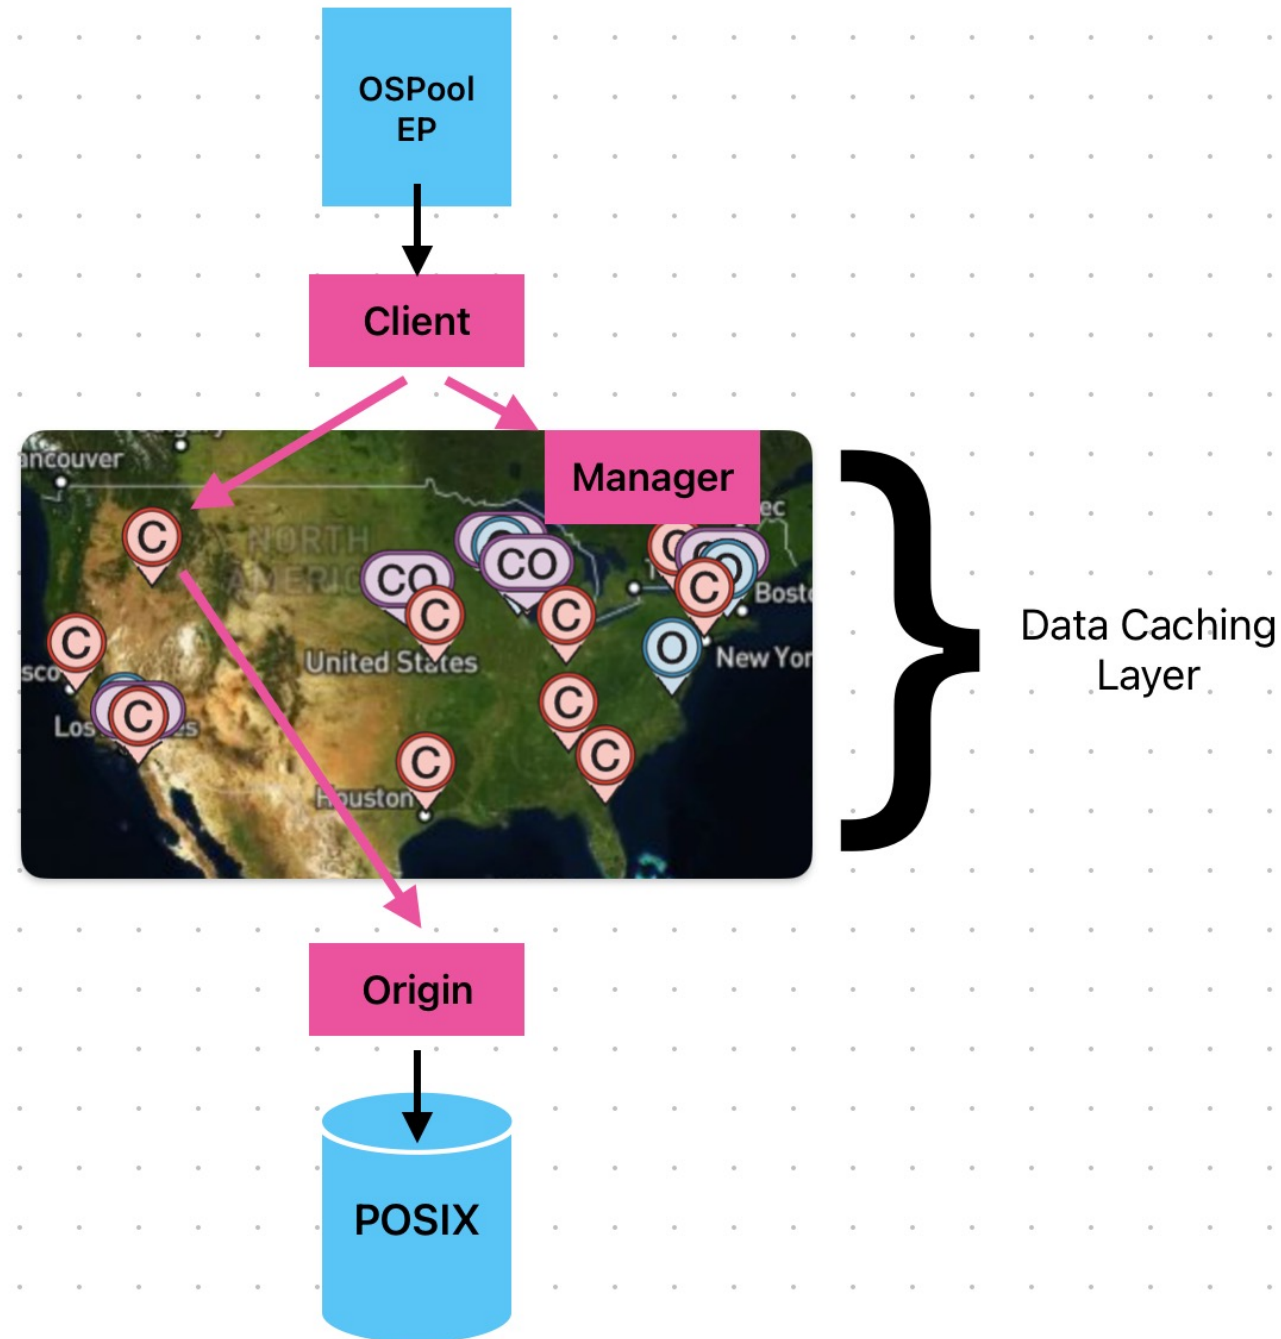
OSDF Architecture

- An origin service integrates the object store into the OSDF in the same way a CE integrates a batch system into the OSPool. Interfaces to move data and map authorizations.
- The cache service stores and forwards objects, providing scalability to the data access.
- The manager selects a source/sink of an object for clients and maintains the namespace.

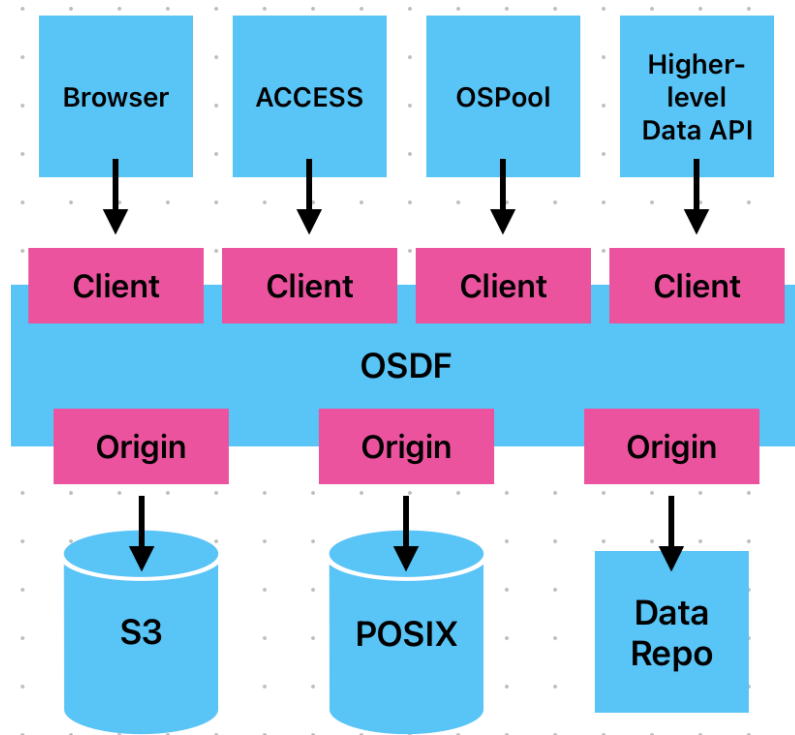


OSDF Architecture

- The OSPool is tightly integrated with the OSDF, giving jobs reliable, scalable means to move data in and out.
 - Same technique used by the CHTC, IGWN, PATh pools.
- OSDF is just starting to integrate more with campuses via the CC* Storage awardees and address non-compute cases.



OSDF Architecture - Vision



- Long-term vision: OSDF will become a platform, providing a ‘transport bus’, connecting a broad range of dataset providers to consumers.
 - Today, we just integrate POSIX backends which leaves out a wide range of storage types (S3, data repositories like DataVerse)
- It has the potential to become a platform beyond the CLI-centric, object delivery service it is today.

Connecting to the OSDF




Fabio Andrijauskas

OSDF Operator

University of California San Diego

- The best way to connect a dataset to the OSDF is through the 'hosted' origin service.
- Place a host in your ScienceDMZ and have it added to the NRP's Nautilus distributed Kubernetes cluster.
 - This gives the OSDF operations team the ability to run containers in your ScienceDMZ.
- Expose the storage area you want to the Kubernetes host; coordinate with the operations team to determine exports and access control.




OSPool and HTCondor

- The OSPool is PATH's flagship HTCondor system.
- How confusing would it be if we called the software and the service the “OSPool”?

That's exactly what we
do with the OSDF!




OSPool



HTCondor
Software Suite

OSDF & Pelican

- We split out the technology powering the OSDF and christened it “Pelican”.
- The software is the same – integrating technologies like XRootD, SciTokens, OA4MP, and the ‘stashcp’ client utilities.



Pelican Software Suite

- The Pelican Software Suite is barreling toward its first release.
 - It will be released as a single, statically-linked binary (no library dependencies) that has both client and server components. Server requires a local install of XRootD, of course.

```
$ pelican object get /test/query1 ./
```

```
pelican origin serve -f https://director.example.com -v /tmp/stash/./test
```

Pelican Software Suite

- Just like you configure a HTCondor client to talk to the OSPool, you need to configure Pelican to talk to the OSDF.
- To ease the lives of OSDF users, one can rename/symlink the binary to 'osdf' and it'll default to the OSDF instance:

```
$ osdf object get /osgconnect/public/dweitzel/blast/queries/query1 ./
```

- Similarly, if you are an avid user of stashcp or HTCSS's OSDF plugin, rename the binary to have compatible behavior:

```
$ stashcp /osgconnect/public/dweitzel/blast/queries/query1 ./
```

An enormous amount of work to do here – watch this area over the next year!

OSDF Status – A view from Operations

Open Science Data Federation

- 1 Cache (50TB for cache) and one origin (1.6PB for origin storage): San Diego Super Computer Center - San Diego - California.
- 1 Cache (42TB for cache) and one origin (1.2 PB for origin storage): University of Nebraska-Lincoln - Lincoln - Nebraska.
- 1 Cache (29TB for cache) and one origin (1.2PB for origin storage): Massachusetts Green High-Performance Computing Center - Holyoke - Massachusetts.
- 1 Cache: Internet2 - Boise - Idaho (42TB for cache).
- 1 Cache: Internet2 - Houston - Texas (in installation).
- 1 Cache: Internet2 - Jacksonville - Florida (42TB for cache).
- 1 Cache: Internet2 - Denver - Colorado (42TB for cache).
- 1 Cache: Northeastern University - Boston - Massachusetts (42TB for cache).

Open Science Data Federation

Providing data access and transfer services for Open Science

BYTES READ
29,570 TB

↑ 938 MB/s Last 1 Year



FILES READ
1,188,355,803

↑ 38/s Last 1 Year



<https://osdf.osg-htc.org>

OSDF Client

StashCP is the client tool to download or upload files on the Open Science Data Federation.

Download StashCP for your platform

rpm, deb, exe, static binary

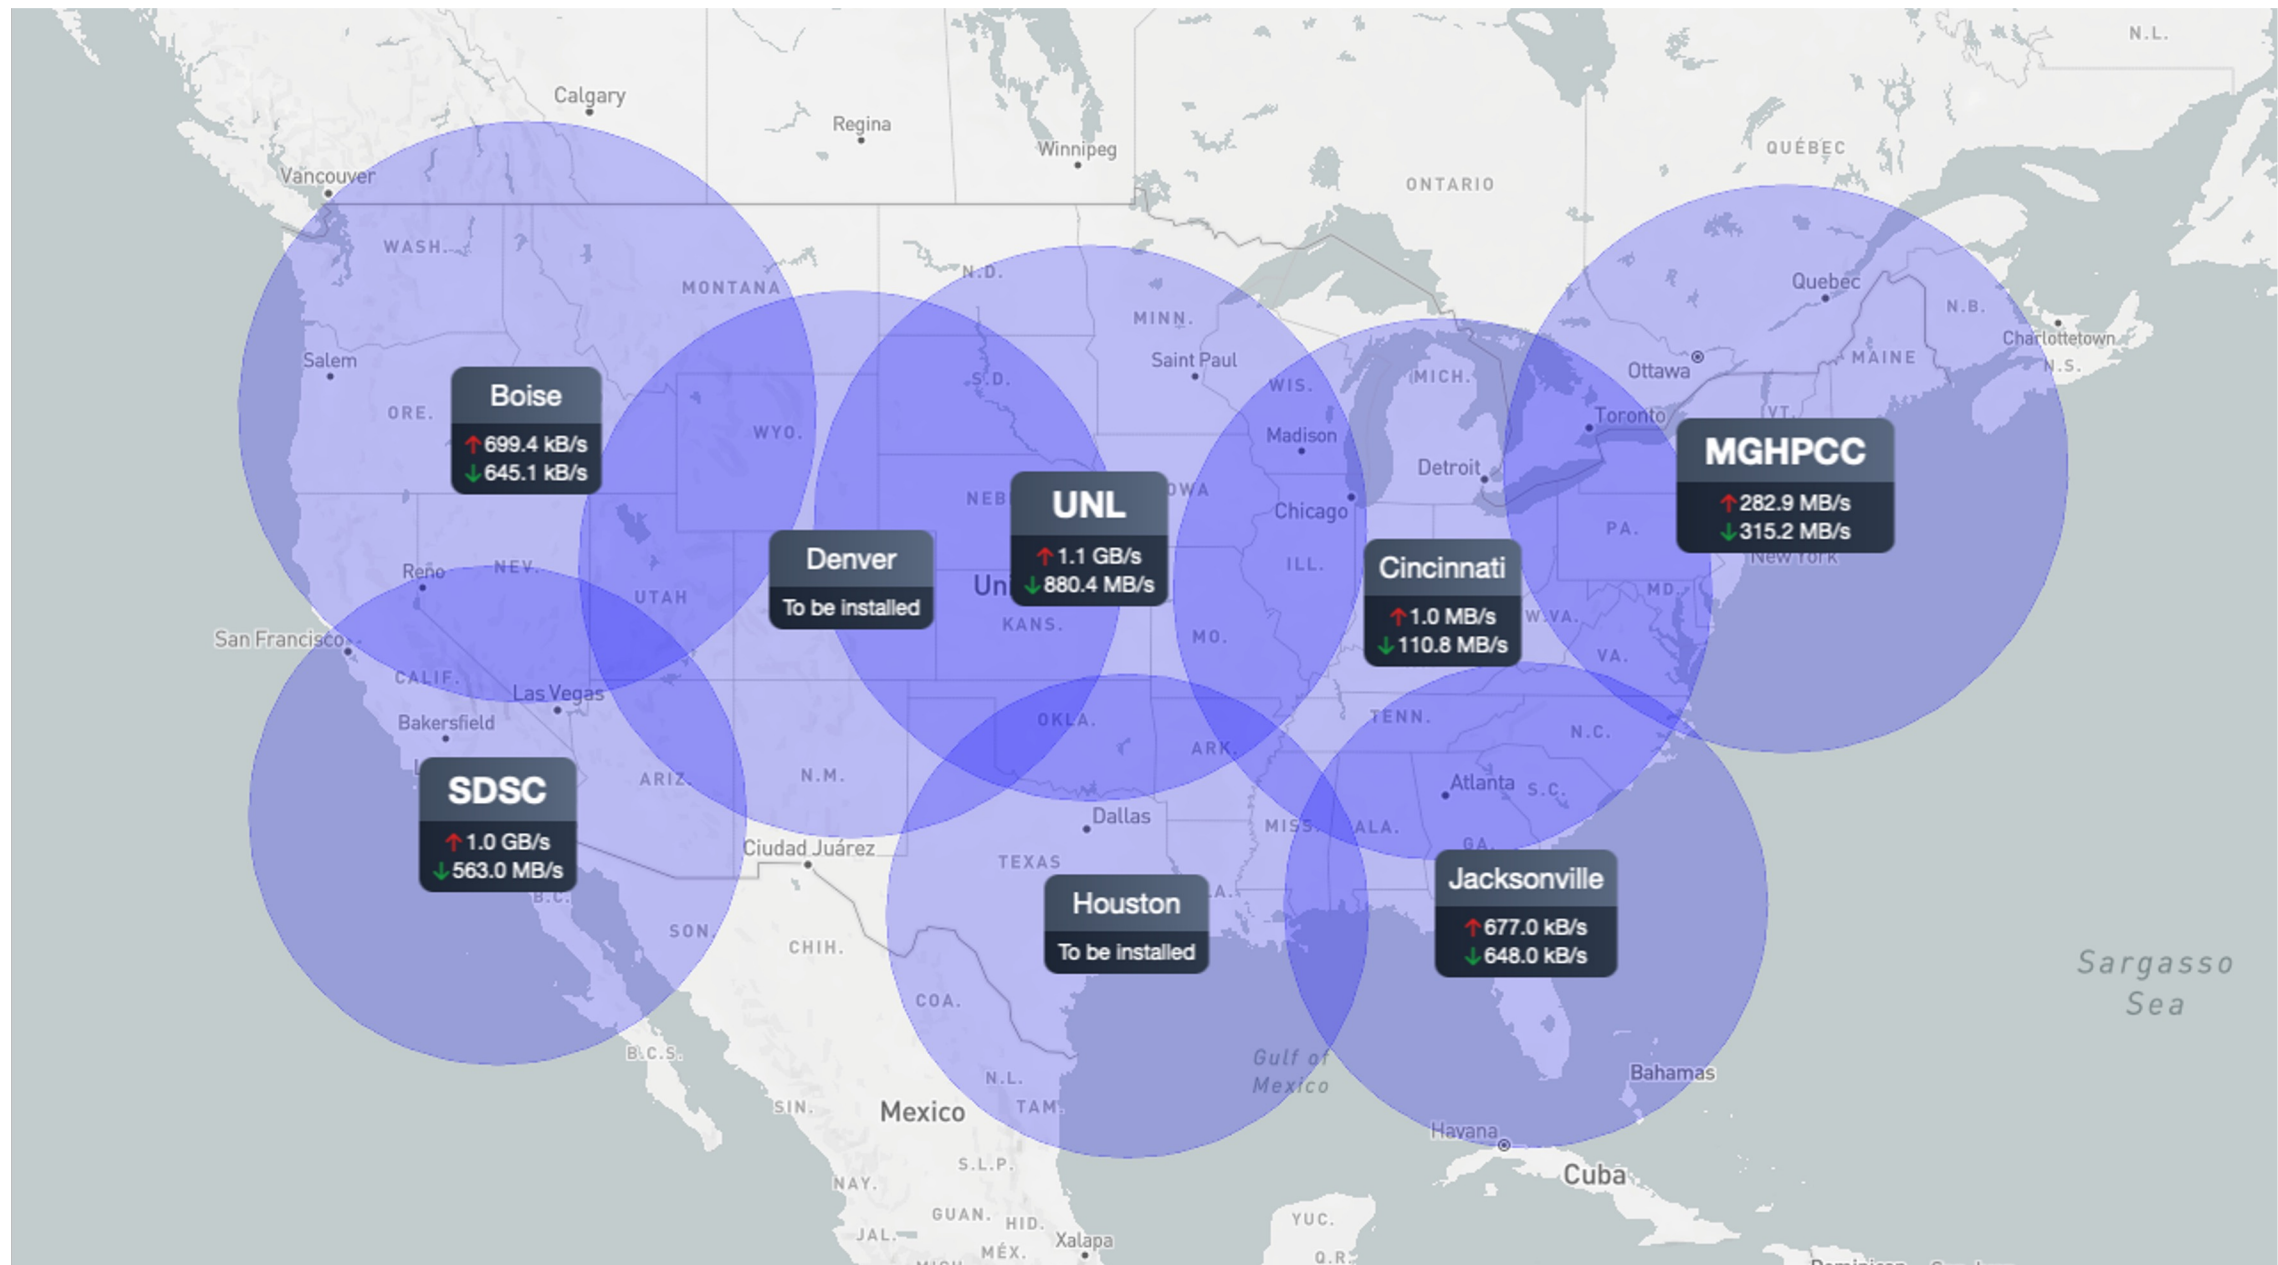


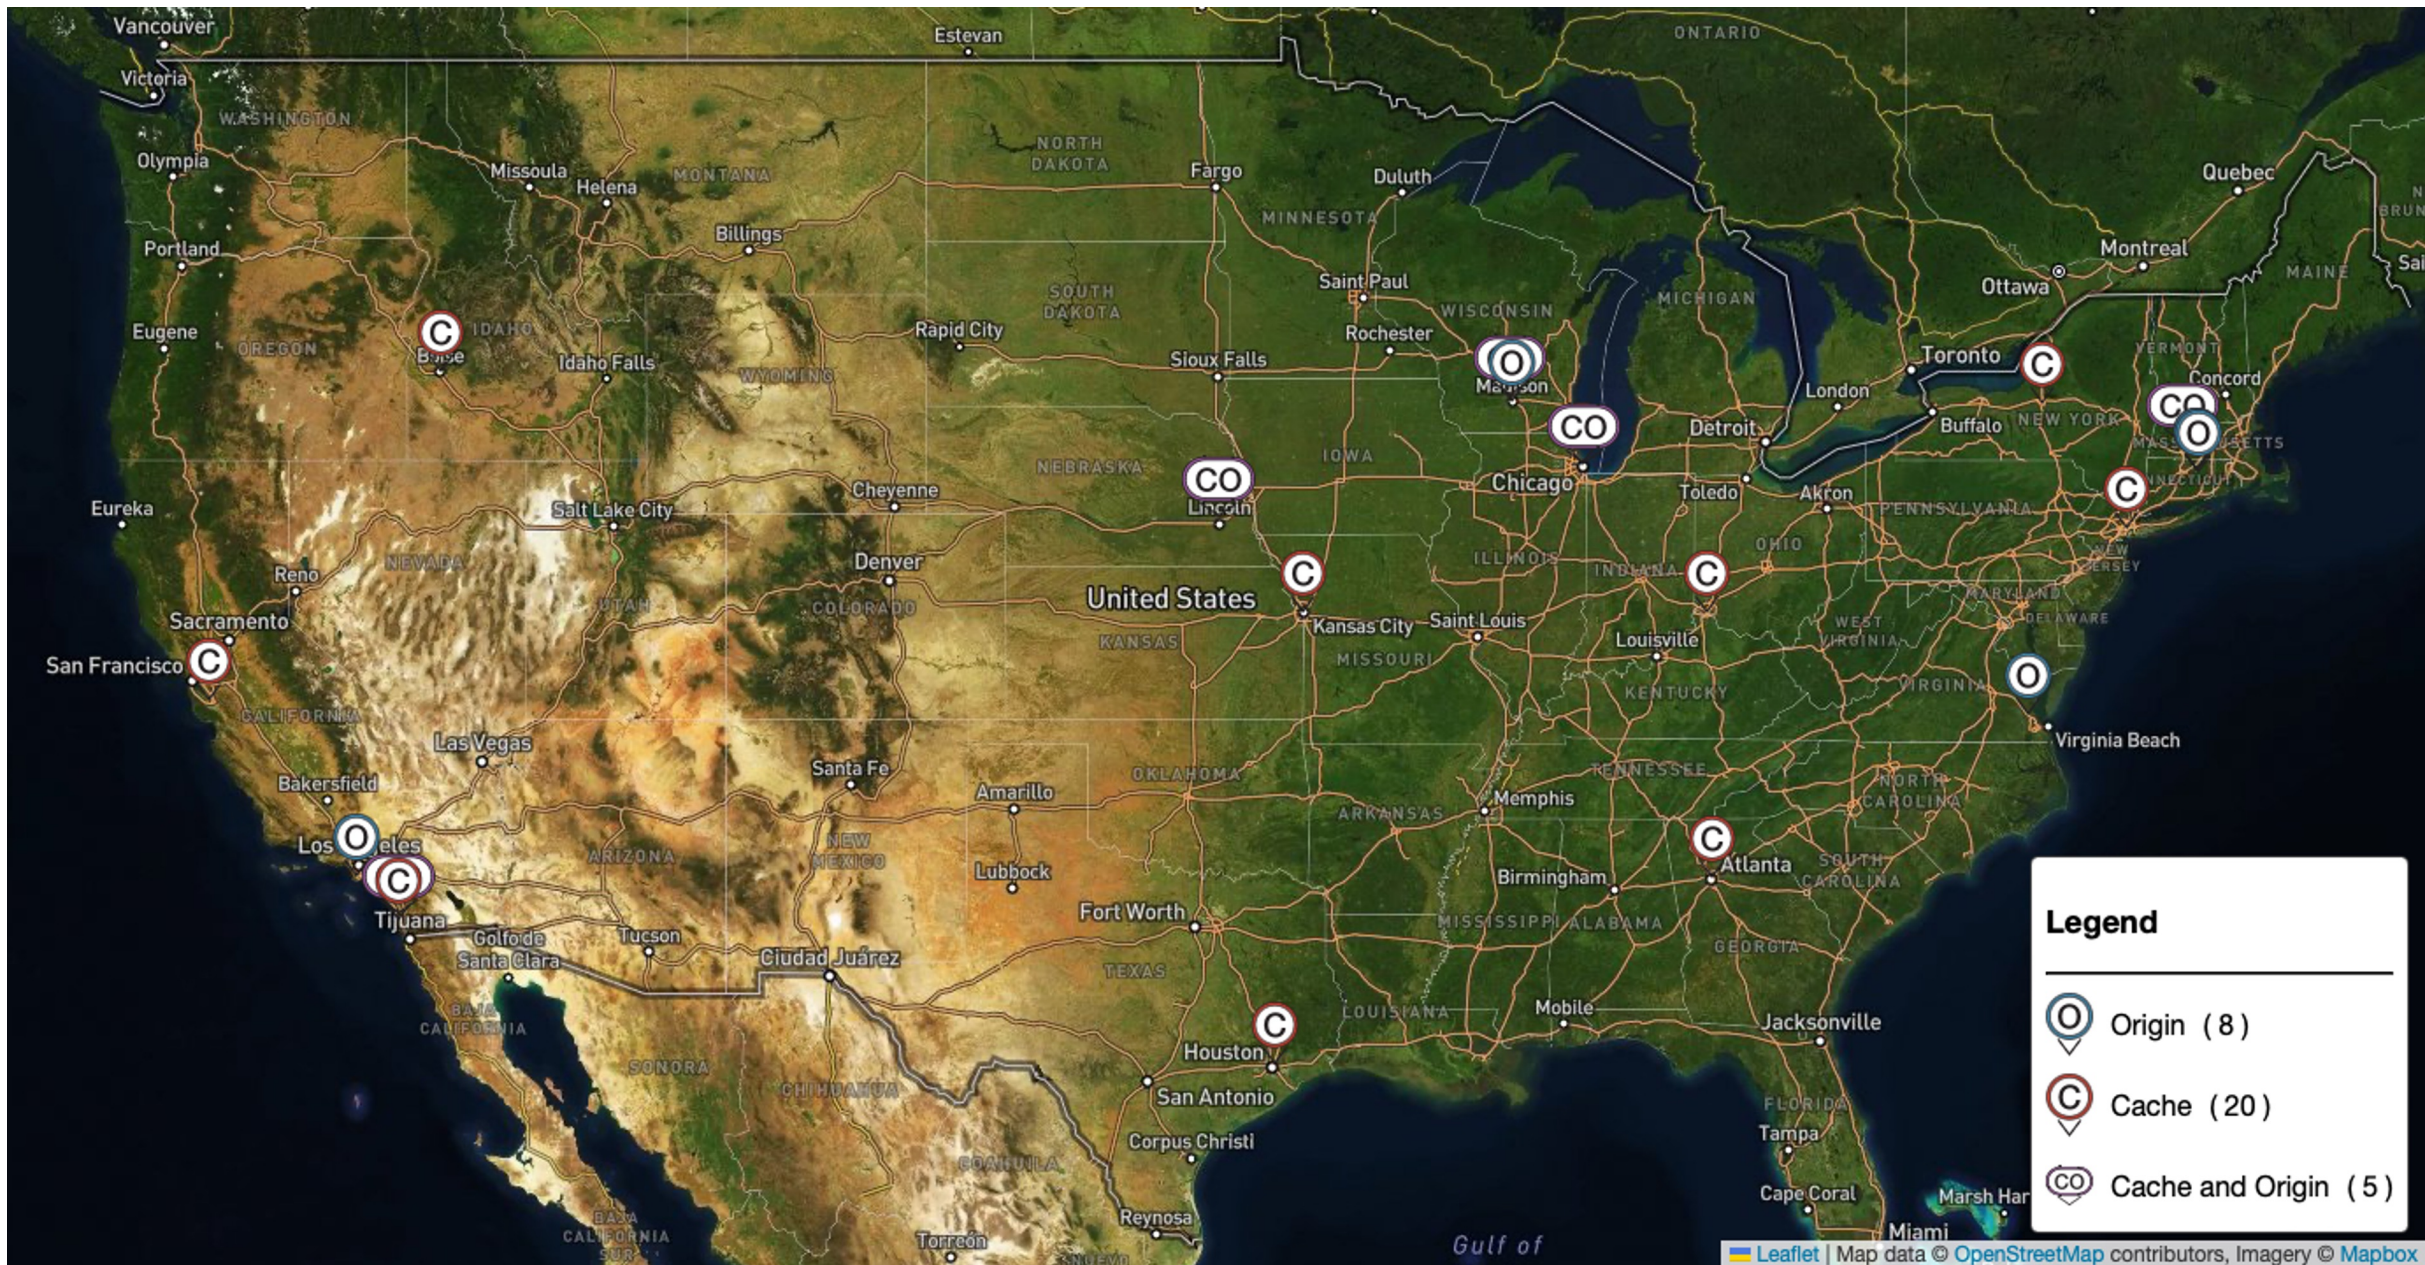
View the Source </>

Terminal

```
$ stashcp /osgconnect/public/dweitzel/blast/queries/query1 ./  
query1 done!
```

<https://osdf.osg-htc.org/client>







OSDF - Monitoring


832,169,349 hits

Sep 17, 2021 @ 01:06:40.442 - Mar 27, 2023 @ 00:00:00.000





Auto




[Hide chart](#)



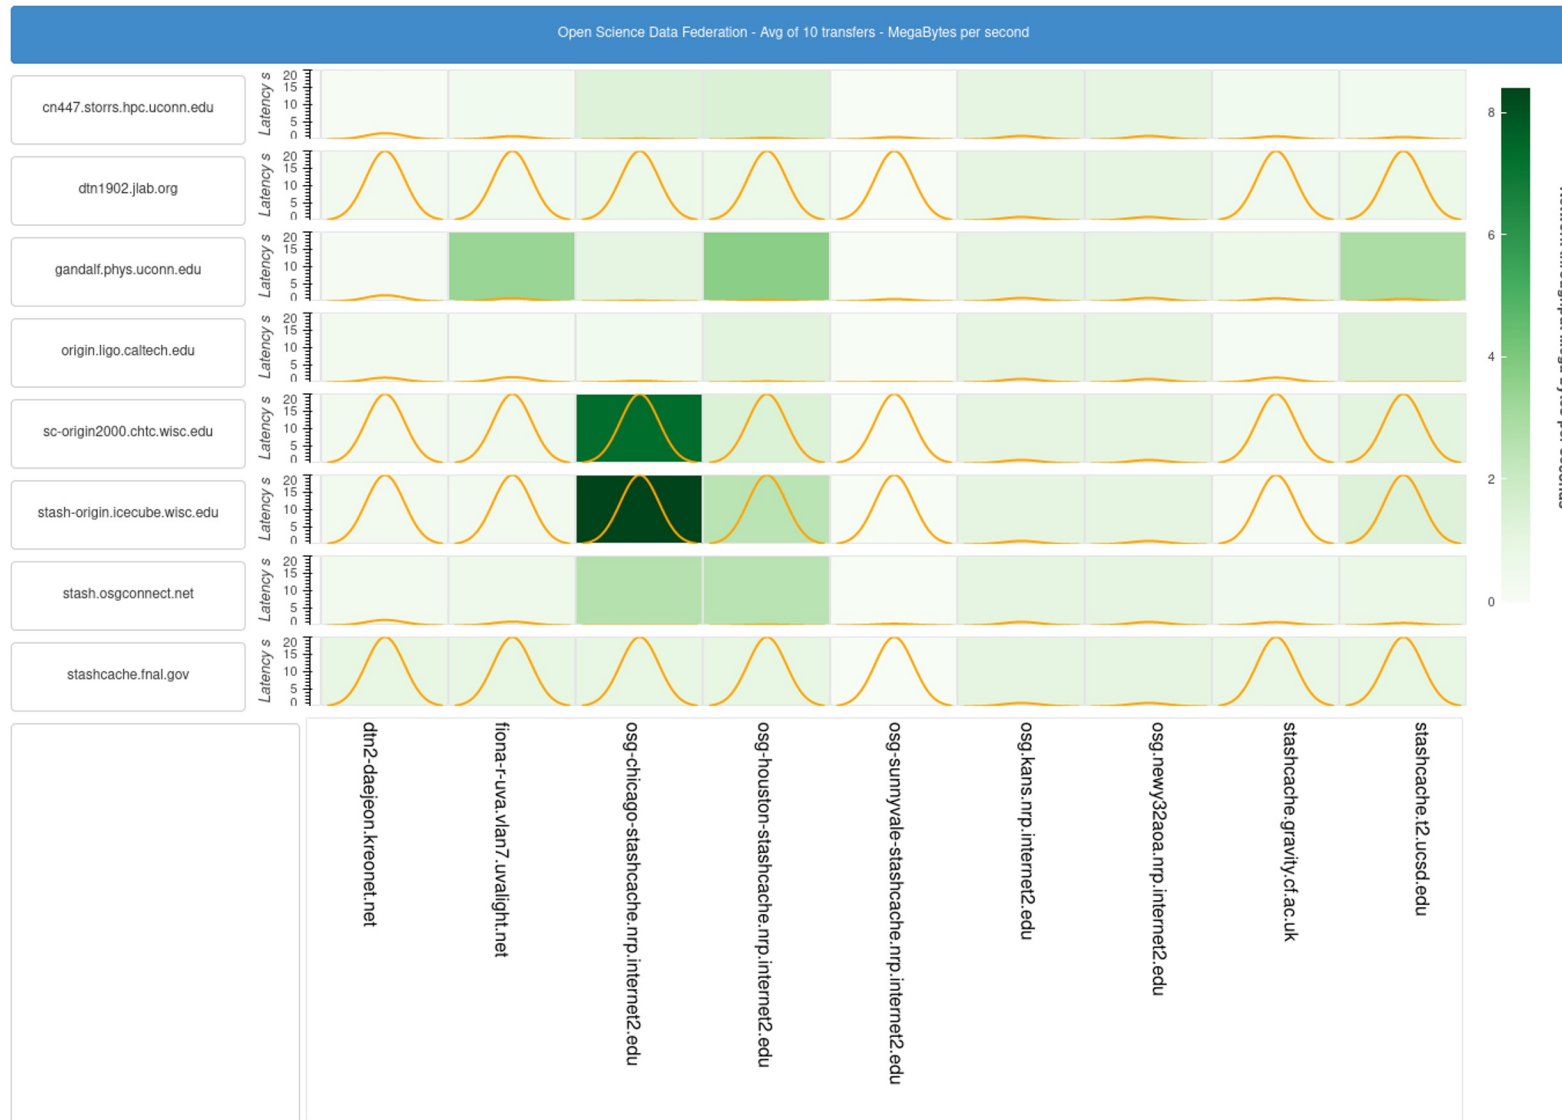
OSDF - Monitoring

OK	Check SSL certificate	≡	OK - Certificate 'stashcache.t2.ucsd.edu' will expire on Sat Nov 4 23:59:59 2023 +0000.	17 h	12.8 s	
OK	HTTP Check CVMFS Access	≡ 	HTTP OK: HTTP/1.1 200 OK - 94 bytes in 0.085 second response time	17 h	12.8 s	<div><div></div></div> 84.6 ms
OK	Read and Write check with Scitoken Authorization	≡	file on origin test0.2924439967576046: file on origin test0.2924439967576046: HTTP/1.1 200 OK	2023-03-24 22:36:39	40.8 s	
OK	Shoveler processed msg	≡	shoveler_packets_received = 7108974	25 h	52.8 s	
OK	Stash cache auth	≡ 		17 h	12.8 s	
OK	Stash cache copy check with authentication	≡	[' % Total % Received % Xferd Average Speed Time Time Time Current	17 h	7.78 s	
OK	Stash cache reports to GRACC in the past 4 days	≡ 	GRACC received 727843 records	2023-02-06 21:25:02	12.8 s	
OK	Stash cache shoveler status	≡	shovler_queue_size = 0	2023-03-23 16:31:51	12.8 s	
OK	XRootD copy	≡	[13B/13B][100%][=====] [13B/s] [13B/13B][100%] [=====][13B/s]	17 h	3.79 s	
OK	XRootD copy per host	≡	[8MB/10MB][80%][=====>][8MB/s] [10MB/10MB][100%][=====] [10MB/s] [10MB/10MB][100%] [=====][10MB/s]	17 h	3.79 s	
OK	XRootD StashCache cache Ping	≡ 	OK	17 h	52.8 s	

OSDF - Monitoring



OSDF - Monitoring



OSDF - Monitoring



Acknowledgements

This work is supported by NSF under Grant Nos. 2030508. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.