

OSG Storage Repositories- Serving more than software!

Derek Weitzel

*This project is supported by National Science Foundation
under Cooperative Agreements 2030508 and 1148698.*

Background


- The OSG maintains several CVMFS repos:
 - **oasis.opensciencegrid.org** - Traditional CVMFS repo for software
 - **singularity.opensciencegrid.org** - Container repository, similar to unpacked.cern.ch
 - ***.osgstorage.org** - Data repositories
- Today, I will only be talking about the *.osgstorage.org repos. For the others see previous [presentations](#) and [papers](#).

Motivation

- We want to present a global filesystem for researcher's data
- The users should be provided with a familiar interface to the data, a posix like interface
- CVMFS provides a filesystem to software (primarily), can it provide the same interface to data?


Motivation

- Our goal is to make using data on the OSG easy!

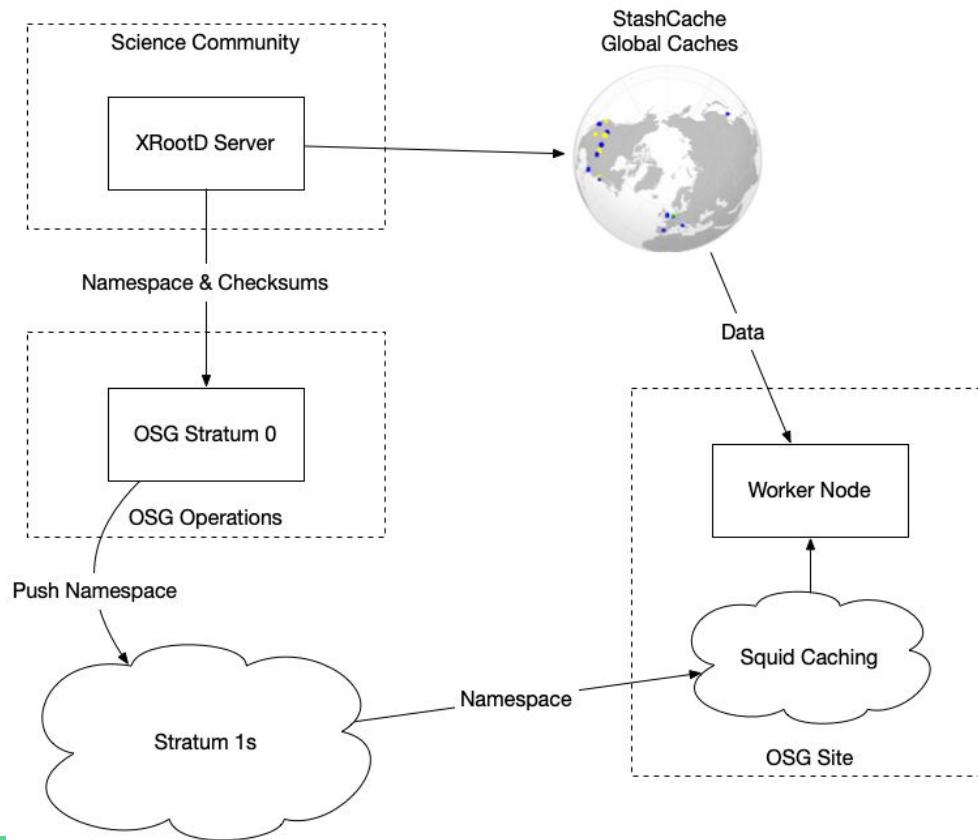


Motivation

- Our goal is to make using data on the OSG easy!

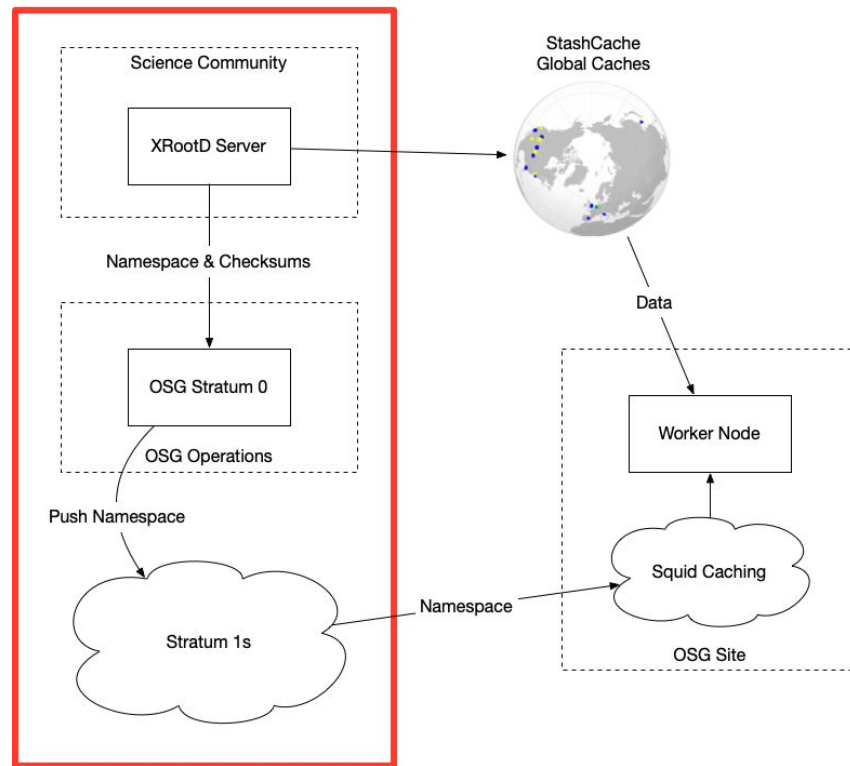


OSG Storage Repositories



Namespace Publishing

- A periodic process scans the XRootD server for new files
- Copies the file details and checksum to the CVMFS namespace
- Publishes namespace to Stratum 1's

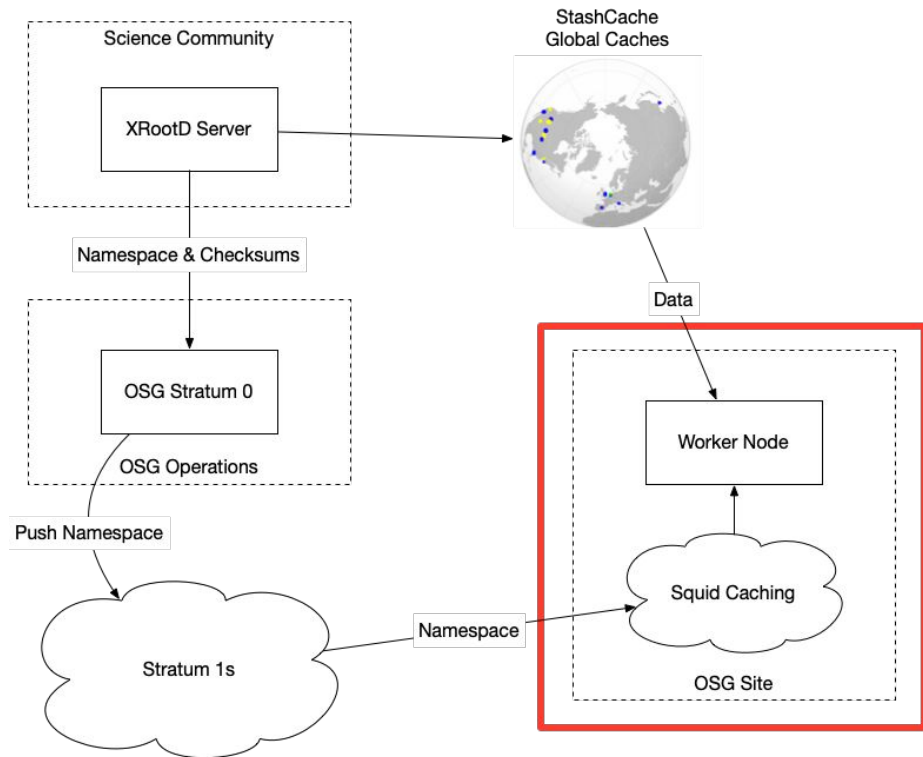


Namespace

- “Grafting” an external filesystem onto CVMFS using [cvmfs-sync](#)
- Creates a special file in the directory named: `.cvmfsgraft-<filename>`
- Use the `cvmfs_swissknife graft` CLI, piping in the contents of the file.
- The graft tool checksums the data and formats the graft file.
- We configured it to use 24MB chunks for checksumming.

Worker Node

- Job access data in *.osgstorage.org repo
- Namespace is available just like a normal CVMFS repo
- When the contents of the file are accessed, download from regional StashCache cache.



Worker Node

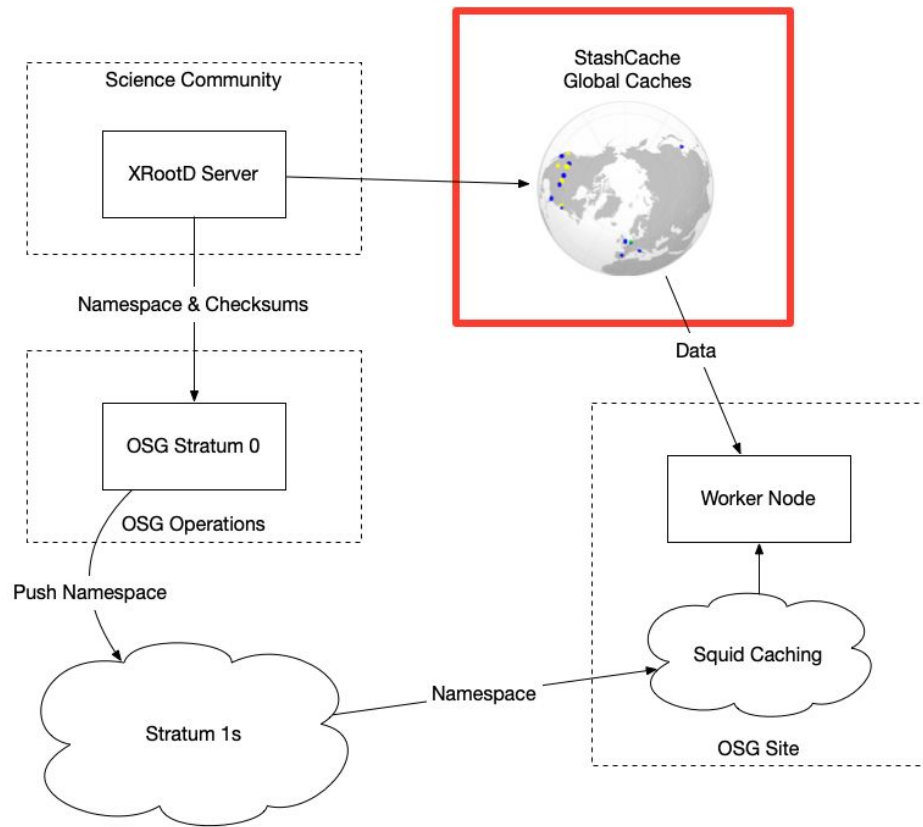
- Worker node has a list of StashCache servers to use:

```
CVMFS_EXTERNAL_URL="http://stashcache.t2.ucsd.edu:8000/;http://mwt2-stashcache.campuscluster.illinois.edu:8000/;..."
```

- We set the CVMFS cache on the node to only **1GB**
 - It is expected that the data size pulled from these repos will be less than 1GB, or too large to cache in CVMFS
 - The regional caches are designed to be the caching layer

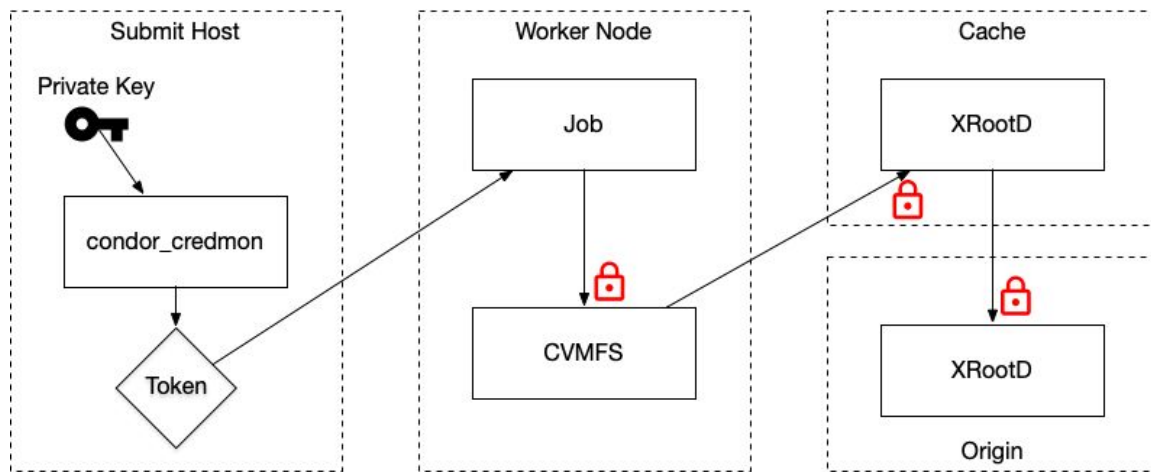
StashCache

- Global network of caches maintained by multiple organizations
- Caches data from the science community's server



Building on: Authenticated Repositories

- The OSG maintains one authenticated repository, **ligo.osgstorage.org**
- Access requires LIGO credentials, either x509 or a **SciToken** (new)



Usage - Last 1 year

	Directory	Data Read	Access
1	/pnfs/fnal.gov/usr/nova	7.4PB	CVMFS
2	/user/ligo	7.4PB	CVMFS
3	/gwdata/O2	1.4PB	CVMFS
4	/gwdata/O1	529.1TB	CVMFS
5	/osgconnect/public/cgomes02	492.5TB	
6	/pnfs/fnal.gov/usr/uboone	459.5TB	CVMFS
7	/pnfs/fnal.gov/usr/minerva	386.5TB	CVMFS
8	/osgconnect/public/odgerk	369.3TB	
9	/chtc	362.5TB	CVMFS
10	/osgconnect/public/gziegler	283.5TB	