NSDF-FUSE: A Testbed for Studying Object Storage via FUSE File Systems

Paula Olaya*, Jakob Luettgau*, Naweiluo Zhou*, Giorgio Scorzelli§,
Jay Lofstead‡, Valerio Pascucci§, Michela Taufer
*University of Tennessee Knoxville, ‡Sandia National Laboratories, §University of Utah

Legacy Applications using Object Storage

- Data is being generated at unprecedented rates across cloud platforms; managing this large amount of data impacts scalability and resilience in legacy applications.
- Object storage can provide scalable and resilient solutions, however, most legacy applications are optimized to run on local and HPC file systems; rewriting them is not an option.
- Filesystem userspace-based (FUSE) mapping packages serve as a bridge to enable legacy applications to access data in an object storage as though they were executed from local or HPC file systems, still…

Users need to understand merits and pitfalls of existing packages when mapping object storage to file systems.

NSDF-FUSE is a testbed for evaluating settings and performance of FUSE-based file systems on top of S3-compatible object storage

NSDF-FUSE supports multiple mapping packages, I/O jobs, and cloud platforms.

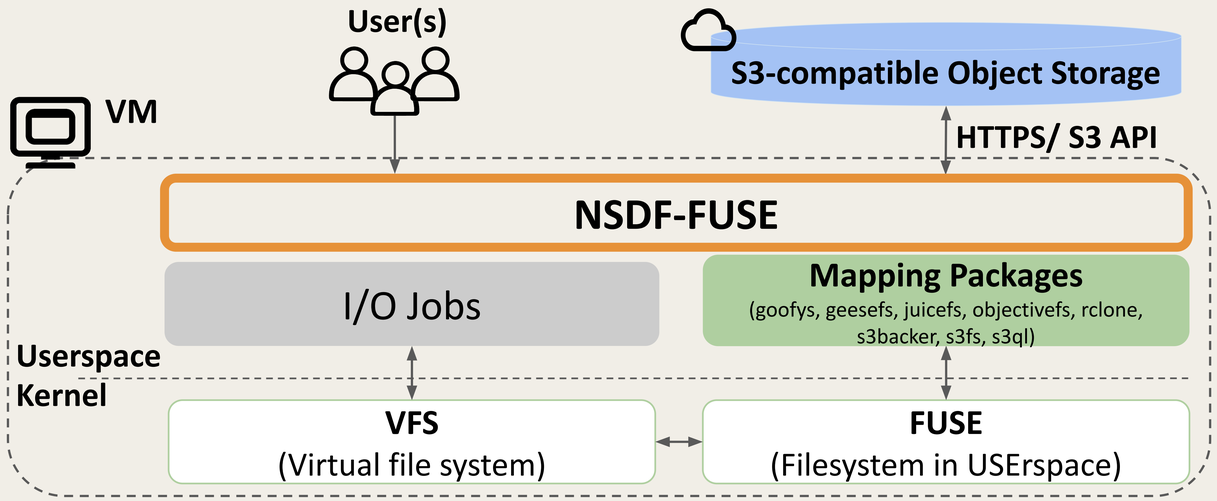

NSDF-FUSE Capabilities:

- Creation/deletion of buckets.
- Installation of mapping package.
- Mount/unmount buckets as FS.
- Evaluate I/O performance through I/O jobs.

Mapping Packages

NSDF-FUSE evaluates 8 common mapping packages.

Mapping Package	Open Source	POSIX	Data Mapping	Metadata Location	Compression
Goofys	✓	Partial	Direct	In name	✗
GeeseFS	✓	Partial	Direct	In name	✗
JuiceFS	✓	Full	Full	In bucket*	✓
ObjectiveFS	✗	Full	Full	In bucket	✓
rclone	✓	Partial	Direct	In bucket	✗
s3backer	✓	Full	Full	In bucket	✓
s3fs	✓	Partial	Direct	In name	✗
s3ql	✓	Full	Full	In bucket	✗

** JuiceFS offers a dedicated server for metadata*

POSIX: the full or partial support of POSIX standards. **Data Mapping:** the direct (i.e., file-object) or chunked (i.e., file-blocks-objects) transformation of data layout from file system to object storage.
Metadata location: Inferred from the name of the objects or as an independent object within the bucket

I/O Jobs

NSDF-FUSE includes 6 I/O jobs, representatives of different data access patterns.

Job 1	**Sequential write, large files, 1 writer**
Job 2	**Sequential read, large files, 1 reader**
Job 3	**Sequential write, large files, 8 writers**
Job 4	**Sequential read, large files, 8 readers**
Job 5	**Random write, small files, 16 writers**
Job 6	**Random read, small files, 16 readers**

Each pattern mimics possible I/O accesses in real applications on the cloud and at the edge.

I/O Performance Across Cloud Platforms Using NSDF-FUSE

We deploy NSDF-FUSE to measure peak I/O performance for six I/O jobs on two cloud platforms.

Mapping Package	Cloud A - Peak Throughput [MiB/s]						Cloud B - Peak Throughput [MiB/s]					
	Job 1	Job 2	Job 3	Job 4	Job 5	Job 6	Job 1	Job 2	Job 3	Job 4	Job 5	Job 6
Goofys	248	**546**	481	**1638**	9	28	136	**431**	356	**910**	15	**78**
GeeseFS	248	455	**910**	585	19	34	136	409	356	146	28	51
JuiceFS	**455**	327	744	431	13	25	**148**	47	327	43	11	15
ObjectiveFS	195	315	273	327	41	39	117	240	282	356	**62**	40
rclone	107	85	372	682	8	16	89	95	**372**	630	32	47
s3backer	84	81	102	91	**62**	**51**	39	130	42	126	29	34
s3fs	74	117	91	136	1	3	34	512	41	585	4	12
s3ql	44	64	56	117	32	9	13	46	6	31	12	9

Best I/O

- For Job 1, JuiceFS and for Job 2, Goofys enable the highest I/O performance for both cloud platforms.
- For Job 3 and Job 4, the highest performance is achieved in Cloud A using Goofys.
- For Job 5, ObjectiveFS and for Job 6 Goofys provide the optimal I/O in Cloud B.

Even when there is not an optimal mapping package and cloud platform that provides the highest I/O performance for all data patterns, **NSDF-FUSE allows the user to reach comprehensive conclusions about mapping packages given different data patterns and cloud platforms.**

Acknowledgments:
Work supported by National Science Foundation Grant No. 2138811

Check NSDF-FUSE on Github