

High Performance Storage System



Future of Tape 2026, May 2026

Jim Gerry
Architect and Consultant



Covered in 30 minutes...

- HPSS today.
- Our next releases.
- When HPSS developers talk.
- Moving to HPSS.



HPSS – a collaboration, storage software, services

- HPSS is a 33-year-old development collaboration.
- HPSS is storage software
 - Single namespace for extreme scale capacity, bandwidth and file count.
 - Tiers of tape , cloud, spinning disk, and flash.
- HPSS is an IBM Consulting service offering
 - Delivery, annual support and consulting services
 - **Not licensed by disk capacity or tape capacity**



Big savings because HPSS is not licensed by capacity!



		Capacity (PB)	M Files	Since
	European Centre for Medium-Range Weather Forecasts	1,450	801	2002
	Shared Services Canada	1,158	82	2017
	United Kingdom Met Office	845	2,011	2009
	Meteo France – French Weather and Climate	797	1,183	2015
	Brookhaven National Laboratory	587	312	1998
⋮				
	Los Alamos National Laboratory – Open	17	294	1997
	India's National Centre for Medium Range Weather Forecasting	15	375	2020
	NASA Atmospheric Science Data Center	11	51	2018
	National Center for Environmental Information	9	137	1998
	University of Tasmania Research Data Management System	6	100	2020

- Most start small and grow.
- HPSS scales with hardware...
 - More files and more transactions.
 - More bandwidth.
 - More capacity.



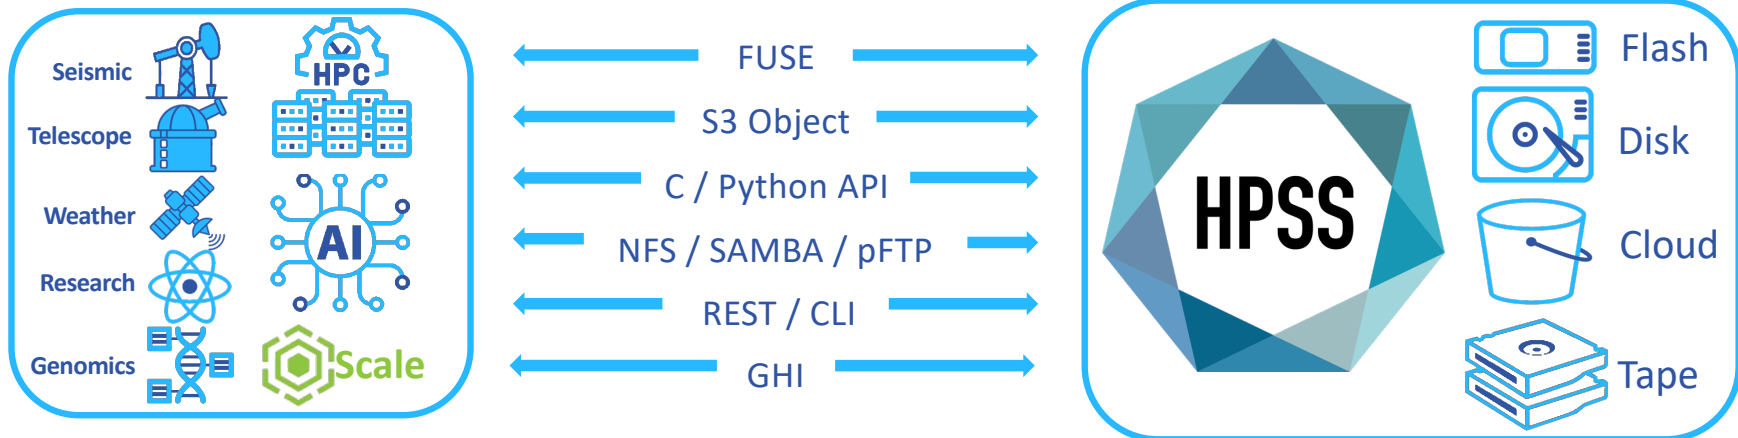
<http://www.hpss-collaboration.org>

Variety of interfaces for flexibility



Capacity (PB) M Files Since

- Traditional HPC workflows via FUSE to mount the HPSS namespace.
- AI and other hybrid-cloud workflows via S3, with S3-Glacier style tape recalls.
- C-language and Python client API for software integration.
- NFS, SAMABA and parallel FTP protocol support.
- Restful and CLI interfaces to interrogate and transfer files.
- Automated space management and backup/restore of IBM Storage Scale.




<http://www.hpss-collaboration.org>

HPSS at Future of Tape 2026

Differentiating features


- Best of class for tape data integrity
 - File checksums and block checksums.
 - Read-after-write verifies data on disk and tape.
 - Corrupted data never reaches tape.

**Battle Against
SILENT DATA
CORRUPTION**



- Tape stripes with rotating parity to cut redundant tape costs, and to improve data durability and tape library durability with RAIT and RAIL

- $N+P = 2+1$
 - Cut redundant tape costs by 50%
 - Native striped IO near 800 MB/s
- $N+P = 4+1$
 - Cuts redundant tape costs by 75%
 - Native striped IO near 1600 MB/s





Differentiating performance and efficiency features...

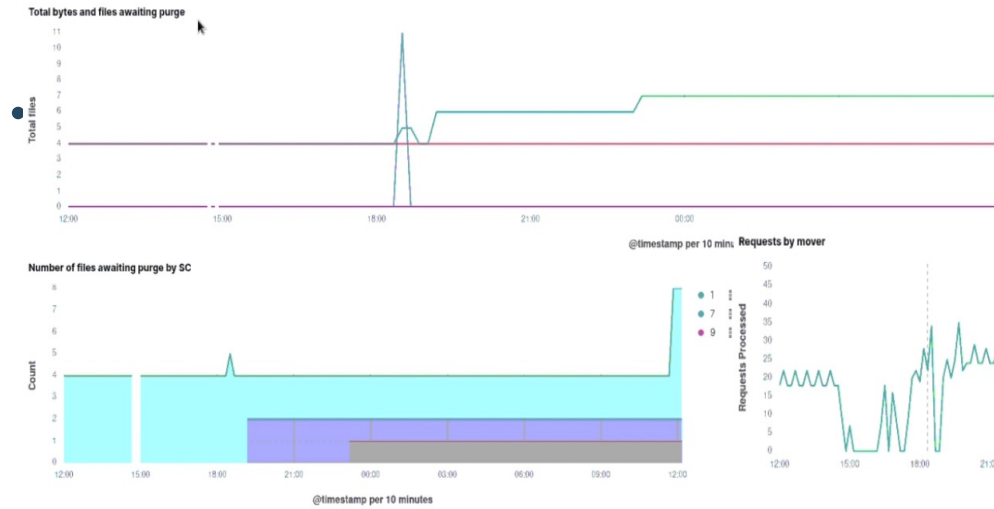
- HPSS achieves near-native tape transfers with all file sizes.
- Automatic grouping and files on tape by policy.
- Automatic grouping of tape recalls by tape.
- Fully integrated with RAO (recommended access order) to minimize tape seeks.
- Features improve tape cartridge longevity.
- Tape recall tools to efficiently recall files back to HPSS disk or recall files to user storage endpoint.
- Features enable near-native tape transfers on write and on read.

More
Data
Less
Hardware



Metrics & visualization...

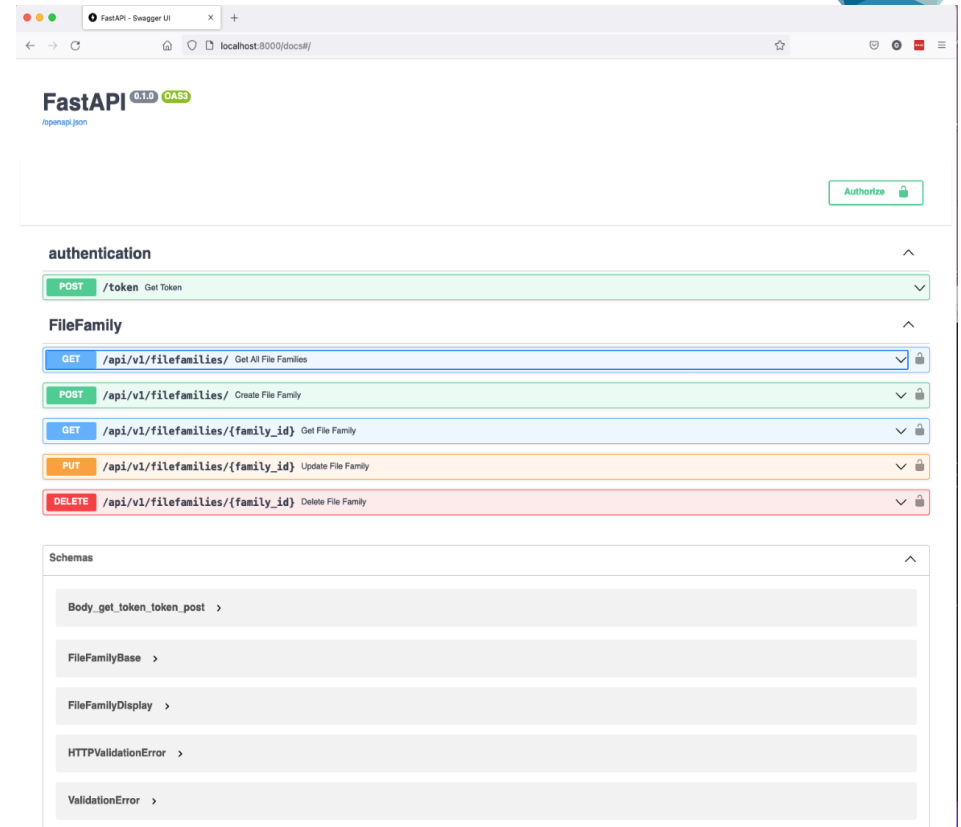
- HPSS tools output data in JSON
- Easily ingested into systems like Kibana, Victoria Metrics, Grafana, etc.





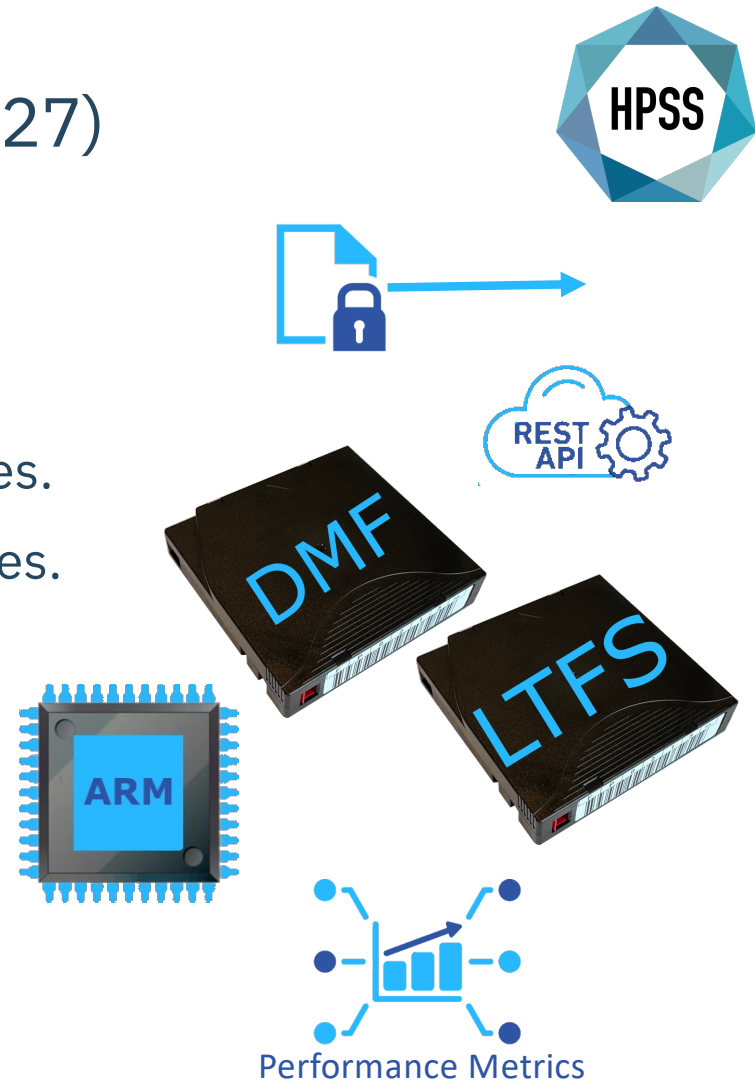
RESTful system management

- Improve ability for sites to programmatically access SSM
- Enable RESTful access with no client-side deployment
- Provide a REST interface to interact with SSM
- API Versioning



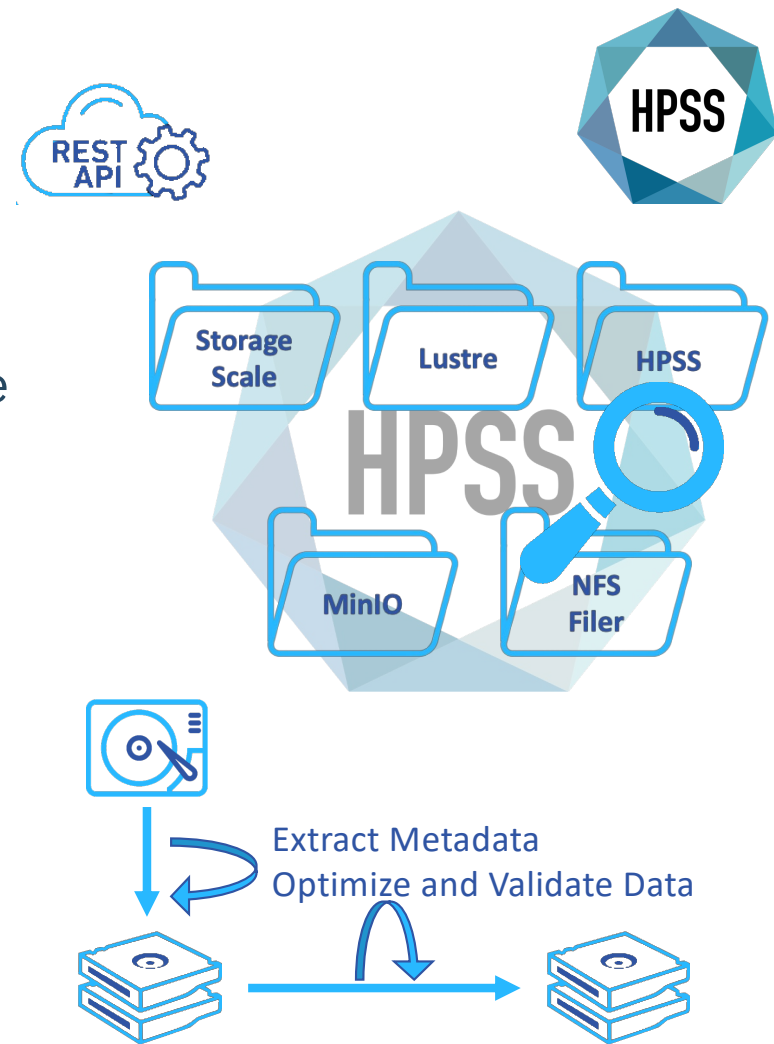
HPSS 12.1 (July 2026) and 12.2 (1Q2027)

- Data encryption in flight.
- More RESTful system management and browser interface enhancements.
- Support for metadata migration / reading DMF tapes.
- Support for metadata migration / reading LTFS tapes.
- ARM support for S3Proxy.
- Add byte counts and transfer rates to read/write/stage logs.
- More RESTful system management and browser interface enhancements.



When developer's talk ...

- More RESTful system management and browser interface enhancements.
- Unified Data Catalog to search and policy-manage data across storage endpoints – clustered file systems, NFS filers, S3 object stores, and HPSS.
- Metadata Extraction Workflow.
- Integrate tape drive and tape cartridge metrics.
- Log analysis service.
- Writing LTFS Tapes.



Migrate data to HPSS

- Export metadata from your legacy storage system(s).
- Transform your metadata.
- Load metadata into HPSS Db2.
- HPSS supports single partition and multi-partition tapes.
- Oracle SAM QFS, DMF, LTFS

