High Performance Storage System

Extreme-Scale Hierarchical Storage Management (HSM) Software

HPSS is disk and tape storage software designed to manage and access exabytes of data at high data rates. Files can be on disk, disk protected by tape, disk space-managed by tape, or stored directly to tape. For files on tape, files are recalled to disk or accessed directly from tape. Standalone, HPSS presents its own file system directly to the user. When coupled with IBM Storage Scale, users interact with Scale, and HPSS provides space management and disaster recovery services to one or more IBM Storage Scale file systems.

Cluster Architecture for Extreme Scalability

HPSS has a cluster design that combines the power of multiple computers, disk storage units, tape libraries, and tape drives, into a single, integrated storage system. HPSS is capable of managing billions of files, exabytes of data, and extreme data transfer rates. No matter how large, the storage system always appears to its clients as a single storage service with a unified common name space.

Migration Services

IBM offers a service to migrate TSM, LTFS, Oracle HSM, DMF, and DXUL files to HPSS. IBM has experience migrating large tape data stores, with or without physically copying the tapes.

End-to-End Data Integrity

HPSS fully integrated with SCSI T10 Logical Block Protection technology of modern tape drives and file checksums to ensure that your files are accurately written to tape. High-speed re-validation of data on tape is accomplished without recalling files back to disk.

High Data Rate

HPSS enables disks and tapes to be striped to create files that can be accessed using parallel I/O operations. A single instance of HPSS is capable of concurrently accessing hundreds of tapes for extremely high aggregate data transfers. HPSS provides automatic aggregation of small files when writing to tape, which increases tape performance and conserves tape space.

HPSS Native RAIT

A Redundant Array of Independent Tapes (RAIT) software feature for HPSS is available. One or more tape parity configurations can be defined to cut the cost of tape redundancy by 50% or more. RAIT has the same performance benefits as tape striping.

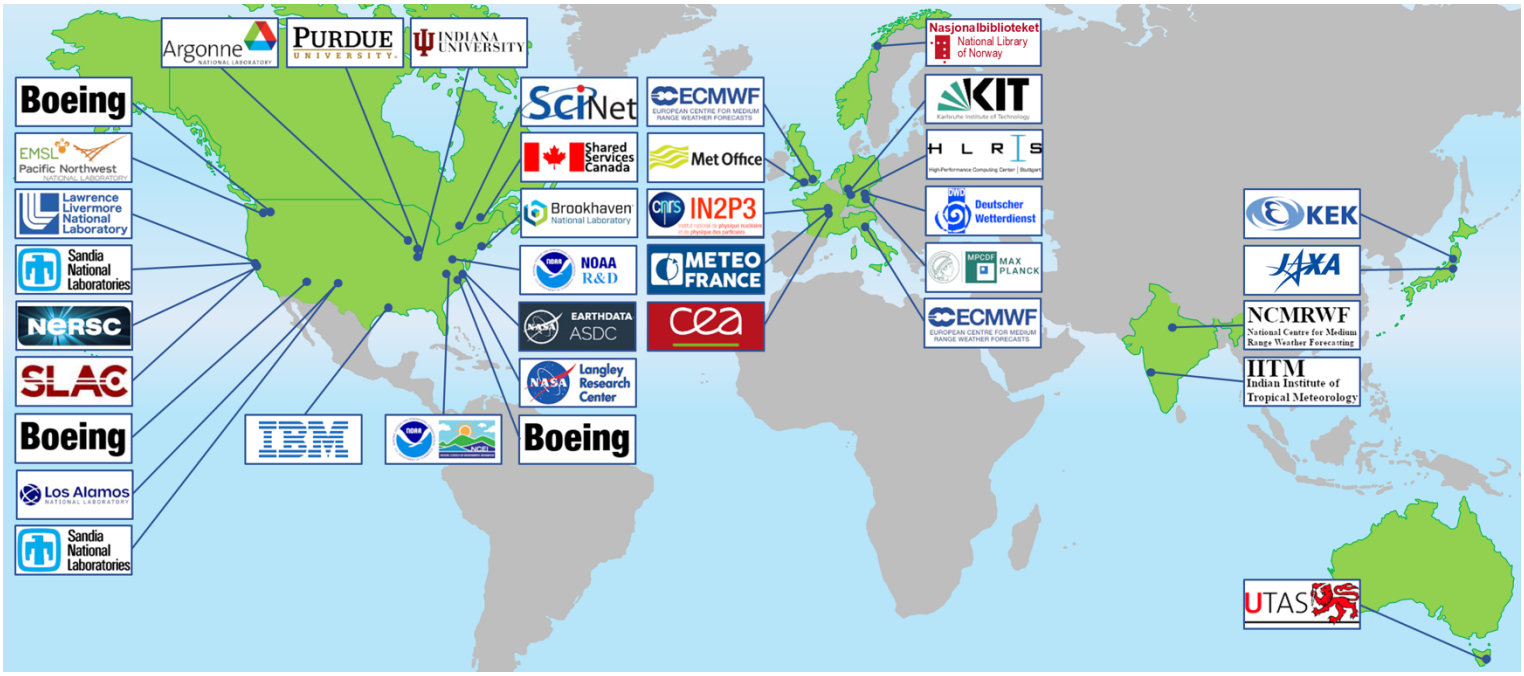

Publicly disclosed HPSS sites

The following organizations have reported data stores ranging from a few petabytes to over an exabyte in a single HPSS namespace:

- Argonne National Lab (ANL)
- Boeing
- Brookhaven National Lab (BNL)
- European Center for Medium-Range Weather Forecasts (ECMWF)
- French Atomic Energy Commission (CEA)
- French National Institute of Nuclear Physics (IN2P3)
- French National Meteorological Service (Météo-France)
- French National Space Center (CNES)
- German Weather Center (DWD)
- High Energy Accelerator Research Organization (KEK) in Japan
- Indian Institute of Tropical Meteorology (IITM)
- Indiana University (IU)
- Japan Aerospace Exploration Agency (JAXA)
- Karlsruhe Institute of Technology (KIT) in Germany
- Langley Research Center (NASA LaRC)
- Lawrence Berkeley National Lab (LBNL)
- Lawrence Livermore National Lab (LLNL)
- Los Alamos National Lab (LANL)
- Max Planck Computing and Data Facility (MPCDF)
- National Centers for Environmental Information (NOAA NCEI)
- National Centers for Environmental Prediction (NOAA NCEP)
- National Centre for Medium Range Weather Forecasting (NCMRWF)
- National Library of Norway
- Pacific Northwest National Lab (PNNL)
- Purdue University
- SciNet, Canada's largest supercomputer center
- Shared Services Canada (SSC)
- Stanford Linear Accelerator Center (SLAC)
- Sandia National Lab (SNL)
- United Kingdom Meteorological Office (UKMO)
- University of Stuttgart (HLRS)
- University of Tasmania

www.hpss-collaboration.org

Contacts:

Ramin Nosrat ramin@us.ibm.com
Jim Gerry jgerry@us.ibm.com

Transfer interfaces

HPSS C and Python Application Programming Interface: The most powerful interface in terms of control, performance, and functionality. HPSS API is the foundation of every HPSS interface, and customers have ported open-source applications including NFS-Ganesha, Globus GridFTP, iRODS, and SFTP to directly interact with HPSS.

Parallel File Transfer Protocol: The high performance Parallel FTP (PFTP) interface moves files in and out of HPSS at high data rates. Standard FTP and high-performance parallel FTP commands are both supported.

HSI & HTAR: The Hierarchical Storage Interface (HSI) provides a familiar UNIX shell-style interface for managing and transferring files. HTAR is a high-performance utility for transferring groups of files between file systems and HPSS.

HPSSFS FUSE Interface: Linux applications benefit from a near-POSIX standard read-write file system interface. This interface enables HPSS to be mounted as a Linux file system in user space (FUSE). Customers are using HPSSFS FUSE with DSpace, Bacula, Open SSL, SaMBa, NFS and Apache.

HPSS S3 Interface: An S3 interface for HPSS that supports automatic class of service (COS) selection, automatic HPSS end-to-end data integrity support with S3 md5 object checksums, and shared access of S3 objects by other HPSS interfaces.

IBM Storage Scale for HPSS Interface: HPSS can be coupled with IBM Storage Scale (previously named GPFS) to automatically: by policy, copy files from Scale to HPSS; purge inactive files from Scale when space thresholds are reached; automatically recall files from HPSS when accessed by users; and save a point-in-time SOBAR backup on a periodic basis for disaster recovery purposes.

HPSS Storage Broker: A high-performance interface to store, protect, and error correct project datasets as SNIA SIRF preservation containers. Demonstrating near-native tape transfer rates when storing datasets containing thousands of tiny files.

System Integrity

IBM Db2 Protected Metadata: All metadata is stored-in and protected-by an IBM Db2 partitioned database enabling, extreme-scale HPSS file-counts and file-transactions, with rapid restart and failure recovery.

HPSS High Availability: HA-HPSS optionally leverages redundant hardware, the Db2 log shipping feature, and failover scripts to minimize HPSS downtime.

Equipment Supported

The full suite of HPSS software runs on Intel and Power computers using Red Hat Enterprise Linux. Tape libraries from IBM, Oracle, Quantum and Spectra Logic are supported, as are all-current IBM, HP and Oracle tape drives for the HPSS tape tiers. HPSS supports IBM Storage Scale, and most enterprise HDD and SSD local attached or network attached block storage units for the HPSS disk cache tiers.

How HPSS is Offered

HPSS is licensed and supported by IBM under an agreement between IBM and the Department of Energy. HPSS is sold, installed and supported as a service offering of IBM Consulting. IBM also offers system engineering services, custom feature development services, and migration services. For U.S. Government organizations, HPSS is listed on NASA's SEWP government-wide acquisition web site, *http://swep.nasa.gov*.

HPSS is developed by an ongoing collaboration of five national laboratories of the United States Department of Energy and IBM. This collaboration has been working together since 1992, through seven major releases.