



High Performance Storage System



Extreme-Scale Hierarchical Storage Management (HSM) Software

HPSS is disk and tape storage software designed to manage and access an exabyte 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, an option determines if files are recalled to disk or accessed directly from tape. Standalone, HPSS presents its own file system directly to the user. When coupled with Spectrum Scale, users interact with Spectrum Scale, and HPSS provides space management and disaster recovery services to one or more Spectrum Scale.

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, SAM-FS, 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

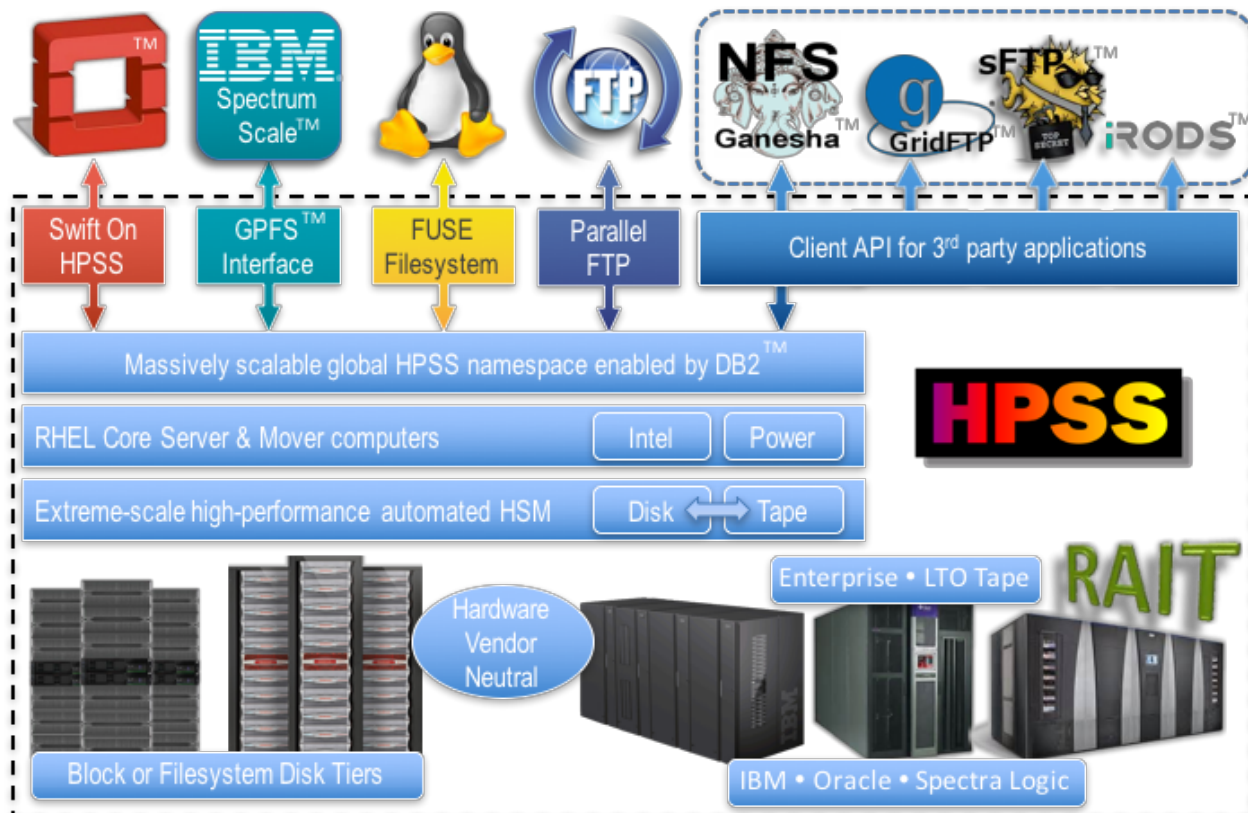
HPSS leverages SCSI T10 Logical Block Protection technology of modern tape drives and file checksums to ensure that your files are accurately written to tape. High-speed revalidation of data on tape is accomplished without recalling files back to disk.

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 now 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 two petabytes to over 270 petabytes in a single HPSS namespace:

- Argonne National Lab (ANL)
- Brookhaven National Lab (BNL)
- Atomic Energy Commission (CEA-DAM) in France
- Comprehensive Large Array-data Stewardship System (NOAA CLASS)
- European Center for Medium-Range Weather Forecasts (ECMWF)
- French National Institute of Nuclear Physics (IN2P3)
- French National Meteorological Service (Météo-France)
- French National Space Center (CNES)
- German Climate Research Center (DKRZ)
- German Weather Center (DWD)
- High Energy Accelerator Research Organization (KEK) in Japan
- Indiana University (IU)
- Japan Aerospace Exploration Agency (JAXA)
- Karlsruhe Institute of Technology (KIT) in Germany
- Langley Research Center (NASA LaRC)
- Lawrence Livermore National Lab (LLNL)
- Lawrence Berkeley National Lab (LBNL)
- Los Alamos National Lab (LANL)
- Max Planck Computing and Data Facility (MPCDF)
- National Center for Atmospheric Research (NOAA NCAR)
- (NCSA) National Center for Supercomputing Applications
- National Climatic Data Center (NOAA NCDC)
- National Centers for Environmental Prediction (NOAA NCEP)
- Oak Ridge National Lab (ORNL)
- Pacific Northwest National Lab (PNNL)
- Purdue University
- SciNet, Canada's largest supercomputer center
- Stanford Linear Accelerator Center (SLAC)
- Sandia National Lab (SNL)
- United Kingdom Meteorological Office (UKMO)
- University of Stuttgart (HLRS)

www.hpss-collaboration.org

Contacts:

Bob Coyne coyne@us.ibm.com
Jim Gerry jgerry@us.ibm.com

Transfer interfaces:

HPSS Application Programming Interface: The Client API is the most powerful interface in terms of control, performance, and rich functionality. The HPSS Client API is the foundation of every HPSS interface, and customers have ported other open source applications including NFS-Ganesha, 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.

Globus GridFTP: GridFTP is a high-performance, secure, reliable data transfer protocol optimized for high-bandwidth wide-area networks. The globus-url-copy tool provides a command-line client for requesting transfers to, from, or between servers - http://toolkit.globus.org/grid_software/data/globus-url-copy.php

HPSS Virtual File System (VFS) Interface: Linux applications benefit from a true 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 HPSS VFS with Open SSL, OpenStack, Samba, NFS and Apache.

OpenStack Swift Object Storage Interface: Highly scalable "SwiftOnHPSS" object storage interface allows OpenStack cloud users to PUT/GET objects to/from HPSS. Supports OpenStack Dashboard and Keystone, HPSS end-to-end data integrity, and the sharing of objects with other HPSS interfaces – <https://github.com/openstack/swiftonhpss>

GPFS-HPSS Interface (GHI): Intended for HPC use, GHI can be coupled with Spectrum Scale (previously named GPFS) to automatically: copy files from Spectrum Scale to HPSS; purge Spectrum Scale files that are not being used when space thresholds are reached; recall files from HPSS when accessed by Spectrum Scale users; and save a point-in-time snapshot of Spectrum Scale.

System Integrity

IBM DB2 Protected Metadata: All metadata is stored in and protected by an IBM DB2 database enabling very rapid restart and failure recovery.

HPSS High Availability: HPSS supports failover redundancy for core servers, movers, and gateways using optional functionality built upon Red Hat Enterprise Linux cluster services. Without HA services, HPSS can maintain operational integrity after partial loss of disks, tape drives, and movers.

Equipment Supported

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