

High Performance Storage System **HPSS**

An IBM Technology Offering

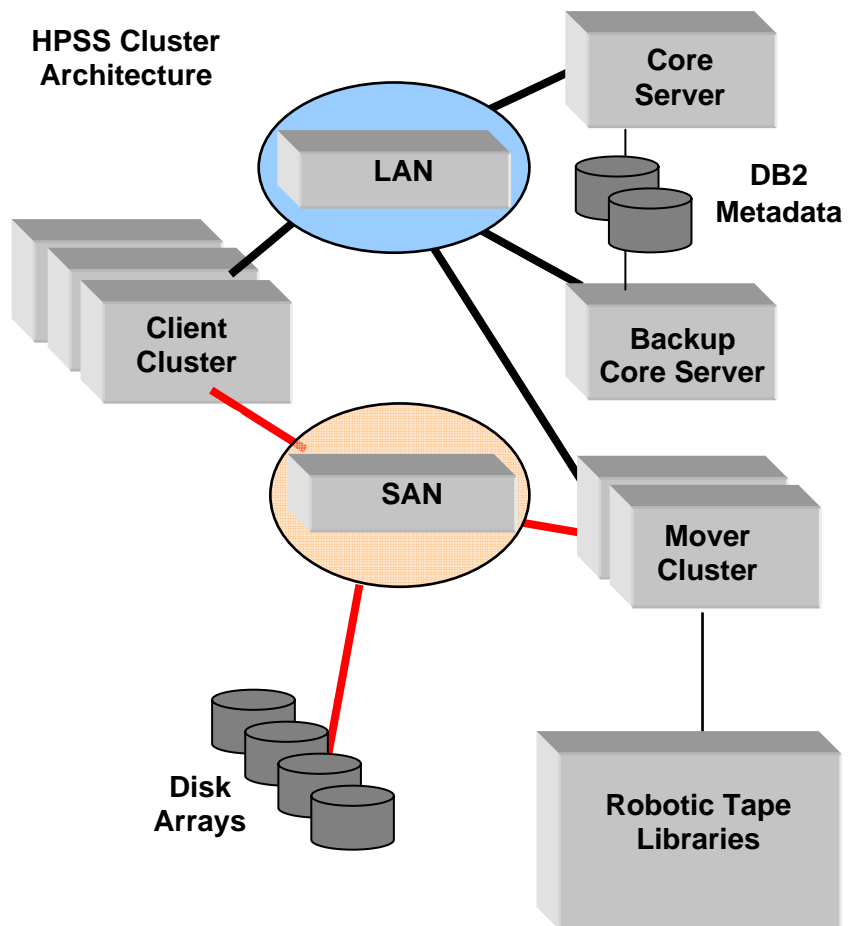
Highlights

- **Software for storing petabytes of data on disk arrays and tape libraries**
- **Enables and supports data life cycle management and when used with disk-only cluster file systems such as Lustre and GPFS**
- **Cluster architecture supports many concurrent data paths**
- **Core servers may be Linux or AIX; movers and clients may be Linux, AIX, or Solaris**
- **An internal IBM DB2 database provides a robust metadata library for storing the identity, location and status of all files and devices**
- **A Linux file system interface provides the syntax of a standard POSIX file system, even when writing to or reading from tape**
- **HPSS Basic, Standard, and Premium offerings deliver full HPSS functionality at a price that is affordable for many sub-petabyte systems**

High Performance Storage System (HPSS) is hierarchical file system software designed to manage and access petabytes of data at high data rates. While appearing to the user as a disk file system, HPSS manages the life cycle of data by moving inactive data to tape and retrieving it the next time it is referenced. HPSS has been used successfully for very large digital image libraries, scientific data repositories, university mass storage systems, and weather forecasting systems, as well as defense and national security applications.

The cluster aspect of HPSS combines the power of multiple computer nodes into a single, integrated storage system. The computers that comprise the HPSS platform may be of different makes and models, yet the storage system appears to its clients as a single storage service with a unified common name space.

HPSS provides both SAN access to disks and efficient network access to disks using HPSS Mover nodes to create SAN-like access over cost-effective TCP/IP networks.



The Petabyte Club

The following organizations have reported HPSS data stores of between one and eleven petabytes, each in a single name space:

Brookhaven National Laboratory (BNL)

Commissariat à l'Energie Atomique /Direction des Applications Militaires (CEA/DAM) Computer Center in France

European Centre for Medium-Range Weather Forecasts (ECMWF)

Indiana University

Lawrence Livermore National Laboratory (LLNL)

Los Alamos National Laboratory (LANL)

National Climatic Data Center (NCDC)

National Centers for Environmental Prediction (NCEP)

National Energy Research Scientific Computing Center (NERSC)

RIKEN in Japan

San Diego Supercomputer Center (SDSC)

Stanford Linear Accelerator Center (SLAC)

For more information

www.hpss-collaboration.org

Or contact:

Bob Coyne
281-992-7129
coyne@us.ibm.com

Harry Hulen
281-488-2473
hulen@us.ibm.com

High Data Rates

The cluster aspect of HPSS combines the power of multiple computer nodes into a single, integrated storage system. The computers that comprise the HPSS platform may be of different makes and models, yet the storage system appears to its clients as a single storage service with a unified common name space.

HPSS enables disks and tapes to be striped to create files that can be accessed at high data rates through parallel I/O operations. With 16-way striping, single file disk data rates of over two gigabytes per second have been achieved.

With adequate provisioning, HPSS supports horizontal scaling, by adding disks, tape libraries, movers, and core servers, to 10s of petabytes, 100s of millions of files, and gigabytes per second.

Interfaces

HPSS API. HPSS provides a Client API that is similar to a POSIX standard read-write interface with additional syntax for supporting striping and classes of service. This Client API is the most powerful interface in terms of control and rich functionality.

PFTP. Parallel FTP (PFTP) is a get-put interface similar to ftp but with parallel capabilities. For wide-area applications, this multi-threaded, TCP/IP-based service has achieved long-distance file transfers of 200 megabytes per second between Department of Energy national laboratories.

HPSS VFS interface for Linux. Linux applications benefit from a true POSIX standard read-write file system interface. This enables HPSS to be mounted as a Linux file system. This interface enables many standard commercial programs that include file I/O to use HPSS as file space, making them into hierarchical disk-tape applications.

DB2/HPSS Active Tablespace Archive. IBM DB2 has been demonstrated to run with its tablespace managed by HPSS. This is a powerful way to create active, online archival databases that can migrate to tape yet be accessed normally and transparently, except for tape mount delays. This feature can be supported as a special bid.

Equipment Supported

The full suite of HPSS software runs on IBM's System p computers with AIX and on IBM System x and System p computers with Red Hat Linux. Clients may also be Sun computers with Solaris.

HPSS supports generic fibre channel disk arrays and many disk arrays from IBM, Apple, DDN, Sun StorageTek, and LSI Logic. Tape libraries from IBM, ADIC, Spectra Logic, and Sun StorageTek are supported, as are most current IBM, HP, and Sun StorageTek tape formats and Sony AIT.

How HPSS is Offered

When you purchase HPSS, there is an end-user license agreement but no license fee, and you purchase a subscription for the services you need from IBM. Significantly, there are no capacity charges.

HPSS is developed by an ongoing collaboration of five national laboratories of the United States Department of Energy and IBM. HPSS is licensed and supported by IBM under an agreement between IBM and the Department of Energy.