Hierarchical Storage Solution

A complete, production ready solution consisting of
IBM's Power, PureFlex or System x servers,
pre-configured High Performance Storage System software,
and on site integration of IBM or other vendor tape systems
to lower total cost of ownership of your massive data store

Hierarchical Storage Management

IBM has reformulated an approach for customers needing to quickly address massive data storage requirements. Hierarchical storage management is recognized as the most durable, most compact, least expensive, and greenest solution for massive long-term data stores. The foundational architecture of hierarchical storage is a tape library with a disk cache front-end, enabling a massive tape system to appear to end users as a familiar disk file system. Recently stored or recently accessed data is automatically read from the disk cache, minimizing the number of tape reads.

While several product offerings exist, one solution – HPSS – has stood alone as the most scalable solution, and with its non-escalating price, by far the least expensive solution for massive data stores reaching out to the 100 petabyte capacity range.

Hierarchical Storage Solution (HSS)

The unique technology of the High Performance Storage System (HPSS) is now being made available in a smaller scale pre-packaged cluster offering that licenses only the essential software features of HPSS, combined with tape library deployment services in IBM's Hierarchical Storage Solution (HSS).

HSS offers the most widely used features of HPSS in a pre-configured offering. The HSS cluster consists of components of the IBM Power, PureFlex or System x server portfolio, selected by IBM to meet the functional and performance needs of your HSS installation. HPSS, the servers, and disk storage arrays are integrated and tested at your site and are configured for your pre-approved tape library and tape drives. The IBM Power, PureFlex or System x servers and HPSS come with a standard service contract, which may optionally be extended.

Together with a range of validated optional storage hardware upgrade offerings, HSS offers unparalleled value for a new hierarchical storage system facing a massive capacity future.

What's inside this new package? Core Server and Data Movers

HPSS is the only hierarchical storage manager designed from the beginning to make use of an extensible cluster platform. The IBM Power, PureFlex or System x basis for HSS uses a broad portfolio of Power and Intel servers, disk arrays and networking products to compose a robust and upgradeable platform. The HSS offering comes with a core server and can be initially configured with up to four data movers or later upgraded, at the customer's request, during the term of support to accommodate upgrades for dual disk caches and as many as 64 tape drives in multiple locally attached libraries.

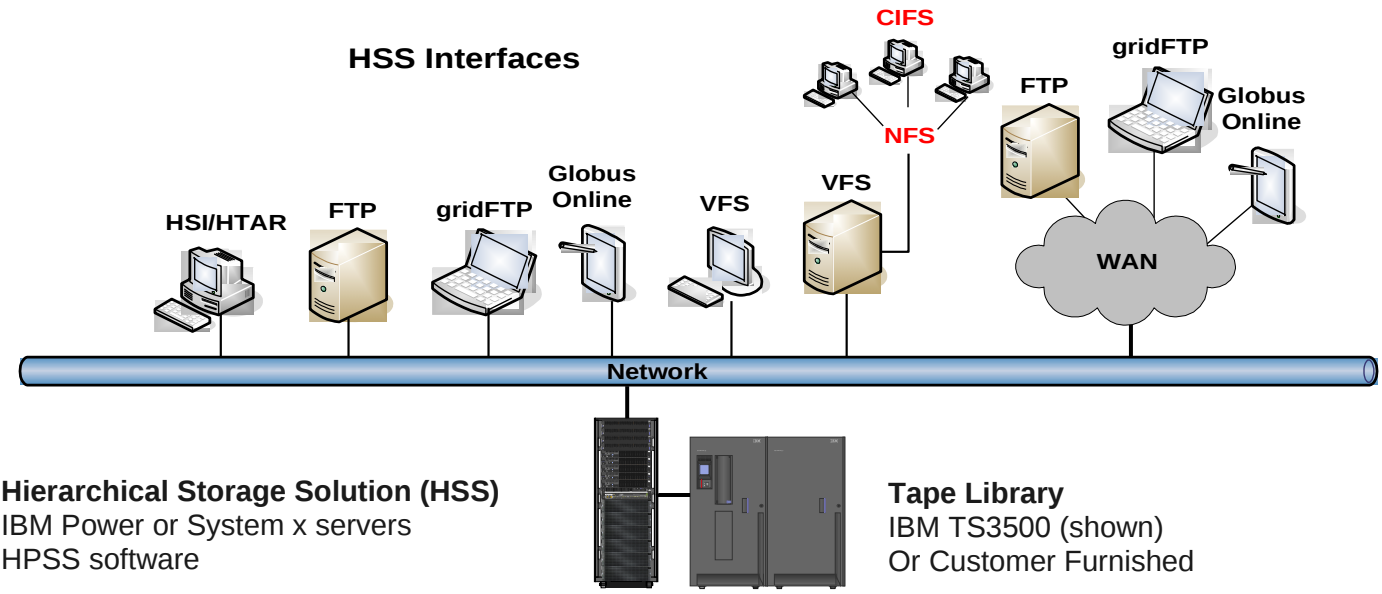

Hierarchical Storage Solution (HSS)
IBM Power or System x servers
HPSS software

Tape Library
IBM TS3500 (shown)
Or Customer Furnished

How does HSS work?

While the end-user simply sees what appears to be a massive unified single namespace mounted as a Linux Virtual File System (VFS), this masks the complexity of actual components and the automation that is occurring. The Core Server receives all read and write commands and keeps track of the locations and status of every file and directory. The HSS you put into production could have billions of files, far too many to be securely or efficiently managed by the "inodes" used in other storage systems. For this extremely demanding task there is an embedded copy of IBM's DB2 relational database on the Core Server. Upon receiving a command from a user, the Core Server uses DB2 to save or retrieve information about the file called "metadata" and then dispatches orders to one or more Data Movers to carry out the request.

The Core Server is thus the most indispensable component of this design. For this reason, a stand-by passive copy of the Core Server is also kept on one of the Data Mover nodes. An HSS operator can initiate a failover to the copy residing on the Mover while the primary Core Server is being serviced.

Data Movers are dedicated servers that perform the actual reads and writes to HSS disks and tapes. They receive data from and send data to client computers. The HSS offering can have from one to four Data Mover computers, each of which is capable of supporting up to 16 tape drives and hundreds of terabytes of disk cache. The maximum aggregated level of capacity for this release of HSS need not be part of the initial purchase, but can be expanded to this extent incrementally.

Reducing the complexity of this offering is an important feature of HSS. Keep in mind that the user, however, need not even be aware of the existence of the server components. The entire HSS system looks like a single point of service to the user. No matter how many files, tape cartridges, or tape libraries in the solution, the user does not have to track on which server the user's file is located. Thus, the user truly "sees" this packaged offering as a single, infinitely large disk file system. The various interfaces supported by HSS to this hierarchical store are depicted in the diagram *HSS Interfaces* on the first page.

HPSS is offered without limits

If capability or capacity beyond the Hierarchical Storage Solution's package maximums is ever required, HSS can quickly and easily be upgraded using existing hardware to the standard High Performance Storage System (HPSS) offering.

You can choose vendor and model of tape libraries and tape drives

Although it is true that the value of all data is not necessarily equal, the notion that older data is less valuable than recent data is incorrect. Tape archived data can offer new insights, new opportunities for analysis, or a convenient repository of infrequently accessed data. So long as the HSS tape based data remains easily available to users, the opportunities to exploit it still exist. The tape components supported by HSS include tape drives, tape cartridges, and one or more robotic tape libraries. Tape systems used may include a non-IBM tape system, as long as it is already supported by HPSS and is approved by the HPSS support team. At the time of this writing, most tape libraries from IBM, Oracle, Spectra Logic, QualStar, and Quantum are supported, as are all current IBM, HP, and Oracle tape drives. HSS can extend up to a maximum of four tape libraries and 64 tape drives with this offering.

You can choose RAIT or mirroring to protect your data

HSS offers two approaches for protecting tape-resident data from corruption and loss. For write-optimized data, a Redundant Array of Independent Tapes (RAIT) feature offers substantial savings in tape media. For read-optimized data, mirroring is the right choice, as there is usually only one mount to get to a file. HSS supports both concurrently, using different HPSS "classes of service."

HSS RAIT is configured with one equivalent parity cartridge for every three equivalent data cartridges (parity is actually rotated across the four cartridges). Redundancy with at least the strength of mirroring is therefore achieved with about one third the additional cartridges that mirroring requires.

You are not locked in!

At the end of your three year support term for HSS, you can renew your HSS contract. You can also "aspire" to the standard HPSS offering, all your HSS components can be re-used in a standard HPSS installation, or choose to "retire" your support for this hierarchical store.

You must always be sure you have the right to freely access and move your own data and storage system metadata. HPSS assures you that you will. HPSS tape formats are conventional and are not proprietary. IBM can provide services to identify and extract files from HSS created tapes without HPSS.

To purchase the Hierarchical Storage Solution,
have your IBM Business Partner or IBM sales representative contact
Patrick Schaefer, 713-320-3249, pschaef@us.ibm.com
Ross Aiken, 720-395-2973, rmaiken@us.ibm.com.