

What's going on next door?

The 2017 HPSS User Forum!

Jim A. Gerry

Senior IT Architect and Consultant

HPSS User Forum



High Performance Storage System



- A gathering of existing HPSS customers, future customers, developers, system engineers, support personnel, project managers, and HPSS leaders.
- Discussing software roadmaps, advanced technologies, burning issues, admin tips, lessons learned, techniques, and successes stories.

HPSS Collaboration



- The HPSS Collaboration is a development partnership between five Department of Energy labs and IBM
- Established in 1992 and going strong today
- Original goal was to develop high performance storage system software
- Today's top priorities of the HPSS collaboration include ...
 - Actively solving exascale storage requirements
 - Success of the HPSS customers around the globe
- There are over 50 people in the HPSS development collaboration



What is HPSS?

- HPSS is an IBM Software-as-a-Service offering that maximizes the value of tape within large-scale storage deployments
- The IBM services are focused on the success of big data storage deployments with demanding requirements such as ...
 - Managing tens of petabytes to exabytes of data within a single name space which may span billions to trillions of files
 - Supporting highly redundant solutions and fast failure recovery
 - Managing high data transfer rates with hundreds of tape drives
 - Minimizing the impact of silent data corruption
 - Delivering less hardware while meeting ingest and recall requirements
 - Supporting bulk data ingest of legacy data

HPSS is considered “best of breed for tape”

- No other storage software uses tape as efficiently as HPSS
- Hardware efficiency translates to lower total cost of ownership (TCO) of data
- HPSS can move more data with less hardware by organizing I/O, minimizing waste and delays, and maximizing transfer performance
- HPSS is sold as an IBM service offering
 - Proven delivery of complex storage solutions
 - Annual support service price NOT based on capacity
 - Proactive and personalized support
- An active HPSS community

Efficient



HPSS is considered “best of breed for tape”

- Striping to **move big files faster**
- Striping with parity (RAIT) to **reduce cost of redundant tape**
- Small file aggregation to **move small files faster**
- Recommended Access Order (RAO) for **faster tape recall**
- HPSS end-to-end data integrity to **identify silent data corruption**
- Extreme scale tape mount/dismount logic to **maximize the number of tape mounts per hour from each tape library**

Striping = FAST



Primary Tape

Redundant Tape



Thank you!



High Performance Storage System

Questions?

Jim Gerry jgerry@us.ibm.com