High Performance Storage System (HPSS)

DOE 40th Anniversary Technology Panel

Advanced Scientific Computing Advisory Committee Meeting

September 26-27, 2017

Buddy Bland
OLCF Program Director

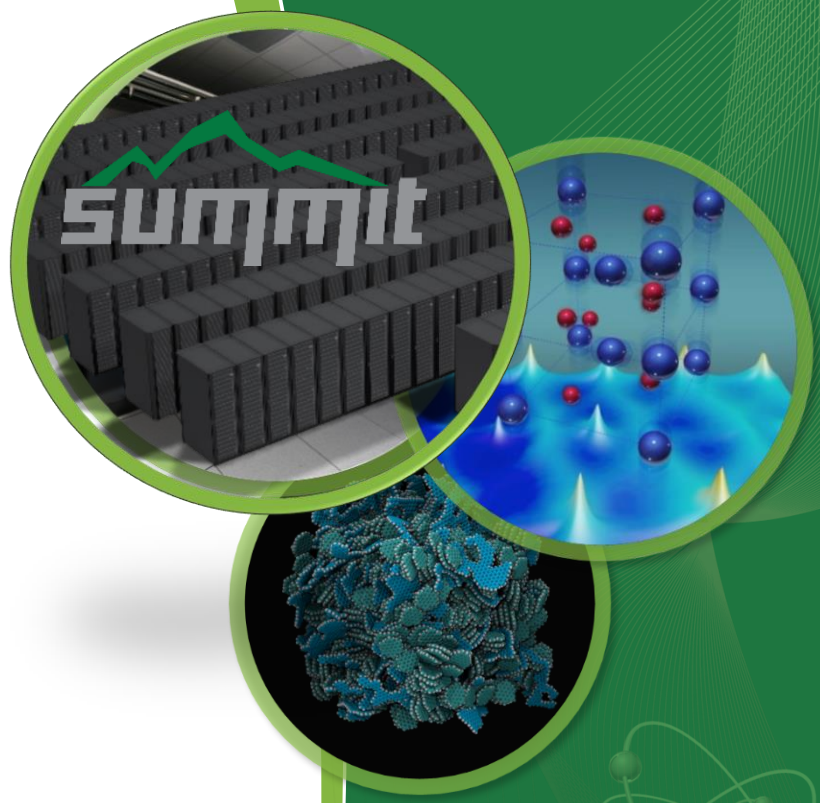

What is the High Performance Storage System?

HPSS is a scalable archival storage system developed and commercialized through a collaboration among 5 DOE Labs and IBM.

- Founded in 1992 – Performance Goals:
  - Bandwidth > 1 GB/s
  - File Size > 1 TB
  - Metadata > 1 Billion files
- Beta release in 1995, General availability 1996
- Development Partners:
  - Livermore, Los Alamos, Oak Ridge, Sandia, Berkeley, IBM
- Won an R&D100 award in 1997

Important features:

- Parallel data transfers
- Access control lists
- RAIT – Redundant Arrays of Independent Tape for data integrity
- Multiple classes of service and data storage devices
  - User configurable stripe width, number of copies
  - Tape, Disk, SSD
- Scalable metadata performance, bandwidth, and capacity
- Globus online and Cloud support

Berkeley Lab | Los Alamos National Laboratory | IBM | University of California Lawrence Livermore National Laboratory | Sandia National Laboratories | Oak Ridge National Laboratory | Leadership Computing Facility

HPSS Collaboration

- Over the years, the HPSS development team consisted of two to five people from each of the development labs and 15-20 people from IBM. Others have contributed to specific features.
- DOE Labs, IBM, and other users provide the requirements for new features and performance goals
- Technical and Executive committees prioritize features to meet the programmatic goals of the users, and the business needs of IBM
- The development team then develops, tests, and documents the changes and packages them into released versions

Why is HPSS important?

- DOE labs continue to push the leading edge of computing performance
- HPSS provides the scalable storage system that has and will keep up with the demands of tera-, peta-, and exascale computing
- This has required *sustained investment* from ASCR and NNSA/ASC
  - We can't wait for an exascale system to arrive to start working on supporting exascale systems
  - Industry finds it difficult to invest in these long-term projects that have limited commercial prospects, but with lab investment, this can be viable.

HPSS Well Positioned for Exascale

- HPSS has continuously evolved and remained world-class, demonstrating production scalability in I/O rates and amount of data stored by factors of 1,000s since it went into production about 21 years ago.
  - Continued scaling of I/O rates for Exascale era is primarily a matter of providing adequate and balanced network and storage device bandwidth.
- HPSS architecture pioneered separation of data from metadata and I/O from storage.
  - Extensibility of architecture evident as HPSS continues to scale
- HPSS supports striped files and parallel file I/O and is cloud enabled through OpenStack Swift
- Adding new classes of storage devices that may emerge is straightforward. Currently putting solid state disk into production for metadata storage.
  - Tape remains a viable archive media. Industry continues research and development to improve tape media bit density

HPSS Well Positioned for Exascale

- HPSS metadata engine is the scalable COTS relational database management system (IBM Db2).
- Small file performance supported through use of user explicit and system supported implicit small file aggregation.
- Because HPSS metadata is built on a RDBMS, capabilities exist for user level metadata storage (aka User Defined Attributes), which while not fully leveraged, are currently being used by various utilities and are expected to be key to enabling user optimization of data set management.
- Commercial multi-exabyte HPSS solution is currently being deployed

HPSS Collaboration

- After twenty-five years, the collaboration is going strong.
  - Software development and releases of new capabilities and fixes continue
  - HPSS is best of breed scalable tape data repository for HPC and big data
  - HPSS is well positioned to exploit new storage technology and configurations
- Most major DOE HPC facilities now rely upon HPSS as their primary archival storage solution. HPSS supports universities, weather and climate data centers, nuclear energy programs, defense systems and research labs in ~40 sites world-wide.
- DOE's Office of Science and NNSA continue to support the development of HPSS

Questions?

Publicly disclosed HPSS sites