

Kronos Nearline Storage System

Brenna Miller

HPC Storage and Archive Group

National Center for Computational
Sciences

ORNL is managed by UT-Battelle LLC
for the US Department of Energy



U.S. DEPARTMENT OF
ENERGY

What is Kronos?



134-petabyte (PB) multi-programmatic nearline storage system with disaster-recovery (DR) copy on tape



Nearline disk component uses IBM Storage Scale (also known as GPFS or Spectrum Scale)



DR tape leverages LTFS (Linear Tape File System)



The two components are tied together by IBM's Storage Archive application, a hierarchical storage manager



Data is replicated by automated policy according to programmatic requirements

Differences from HPSS

Provides resilience through dual-copy of all files

- One on disk, one on tape

Users will never need to interact with tape

- Time-to-first-byte is greatly reduced

Quotas are enforced

- Programmatic quota – the amount of space a program purchased in the system
 - OLCF – 109PB
 - Air Force Weather Wing – 26PB
- OLCF Project quota – currently 200 terabytes (TB), increases can be requested

Performance Benchmarks



Run during acceptance phase (April 2024)



Disk subsystem aggregate I/O bandwidth

450GB/s sequential write
407GB/s random read



DR subsystem capable of ingest of ~2PB/day

Implies files larger than 1GB



Globus performance

12GB/s, single client

User Interfaces – Current and Upcoming

POSIX mount on OLCF/NCCS Moderate DTNs

- /nl/kronos/{olcf | afw}


Globus Collection

- OLCF Kronos

Frontier Login Nodes via POSIX mount (upcoming)

Marble (NCCS Kubernetes, upcoming)

HPSS to Kronos Transition – Significant Dates



31 July: Kronos available to early users

- 8PB of data ingested in first ~30 days
- Please begin migrating data


30 August: HPSS read-only

- HPSS Globus Collection disabled
- Please begin migrating data sooner rather than later

31 January, 2025: HPSS decommission date

- Data remaining on HPSS will be unavailable for retrieval
- Please do not wait until the last minute to migrate data

HPSS to Kronos Transition – Best Practices



Movement of existing data should use the hsi_xfer tool

Provides an optimized recall of data from tape

Reduces time-to-transfer and load on the HPSS systems



For assistance, reach out to help@olcf.ornl.gov



Do not wait to start moving data

Questions?

