



# Tape's Not Dead

At LBNL/NERSC

Nick Balthaser  
MSST 2019  
May 21, 2019

# Storage @NERSC

- NERSC is a DOE HPC Center
  - Broad range of unclassified research, thousands of users
  - Tiered storage hierarchy:
    - Burst buffer
    - Lustre scratch
    - GPFS campaign (homes, project)
    - Large tape archive
- > 180PB stored in 2 tape archives running HPSS
  - > 3PB/month growth rate

# Ongoing Tape-related Projects

- Lots of ongoing work at NERSC related to archival storage
  - Locating 3rd library at Berkeley Data Center (CRT)
  - Data center move Oakland (OSF) to Berkeley
  - GHI testing
  - HPSS Upgrade

# Locating 3rd Library

- Berkeley site currently has 2 IBM TS4500 libraries with integrated cooling feature
  - 16 frames, 13k slots, 64 TS1155 drives per library
  - Libraries are self-contained for environmental stability in open-air cooled data center
- Currently working on delivery of 3rd TS4500 for future growth and data center move

# Data Center Move

- Moving 140PB from Oakland data center SL8500s to Berkeley
  - First 30PB IBM Enterprise media moved by truck over the course of 2 weeks
  - Copying 300TB/day out of 4 Oracle SL8500s via HPSS repack utility
  - Estimated 2 year time frame, then SL8500s will be retired

# GHI Testing

- GHI is GPFS/HPSS integration
  - Policy-based migration of files from GPFS to HPSS
  - Ongoing test effort with goal of limited production deployment after next HPSS upgrade

# HPSS Upgrade

- Current production HPSS version has been EOL for some time
  - Working on upgrading HPSS to supported version (7.5.1) for bug fixes (e.g. SCSI PVR) and new features (RAIT, GHI)

Thanks!

