NERSC

National Energy Research
Scientific Computing Center

Long-term Data Management In the NERSC Archive

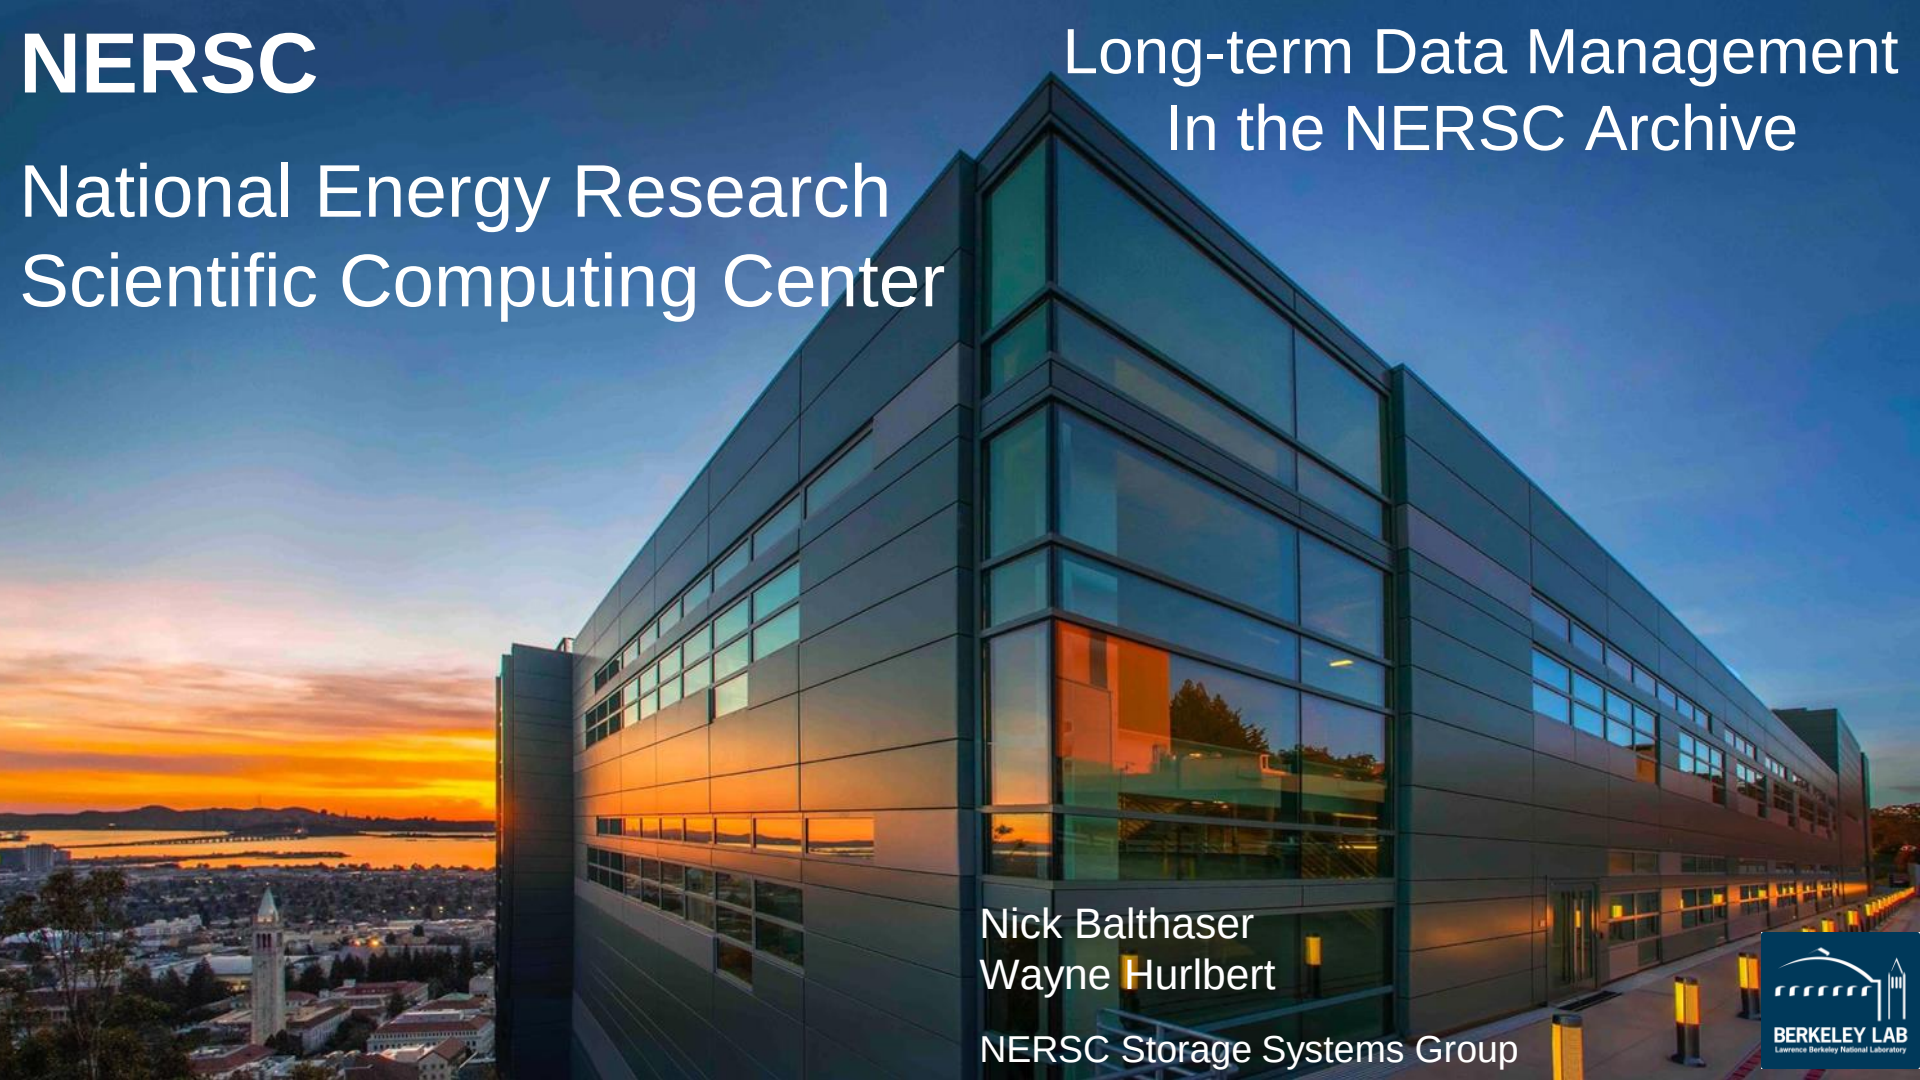

Nick Balthaser
Wayne Hurlbert

NERSC Storage Systems Group

BERKELEY LAB
Lawrence Berkeley National Laboratory

NERSC Data Archive

- 43 years of data archived by the scientific community
  - ~15,000 tape cartridges in 3 IBM TS4500 libraries
  - In some cases the NERSC archive is the only copy of the data
- HPSS software in production since 1998
- 2 systems
  - Archive – user-facing: 180PB
  - Regent – center backups: 30PB
- Data transfers via client - no direct file system interface

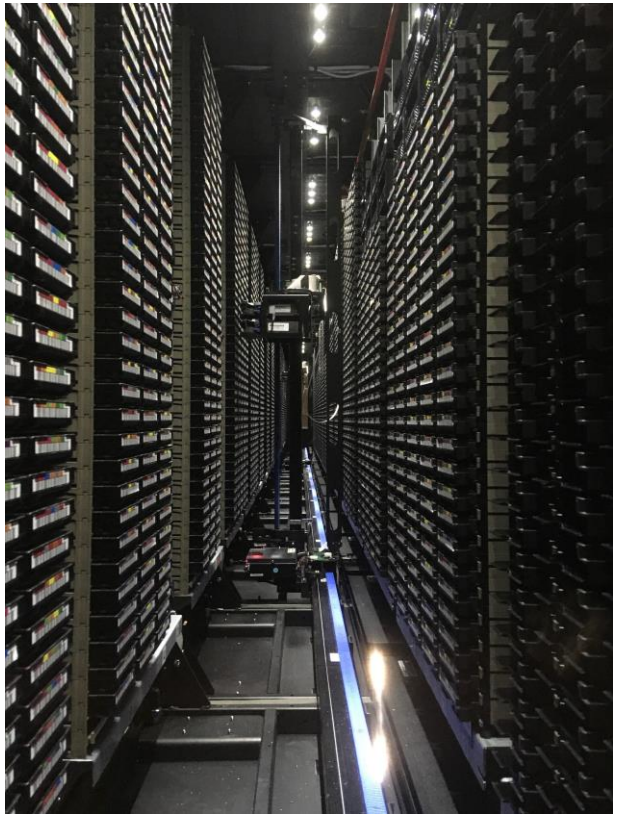

Archival Data Growth

- Historically data has grown in the archive at a rate of 1.4x/year

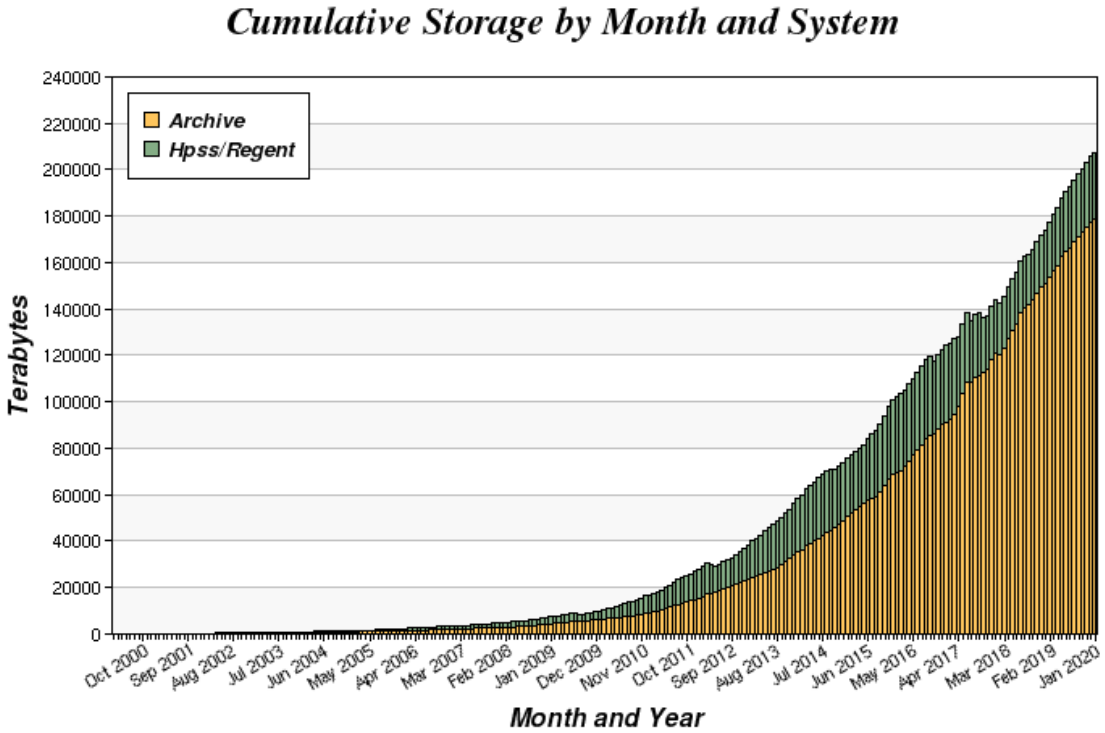

Preventing Data Loss
3-Tiered Approach

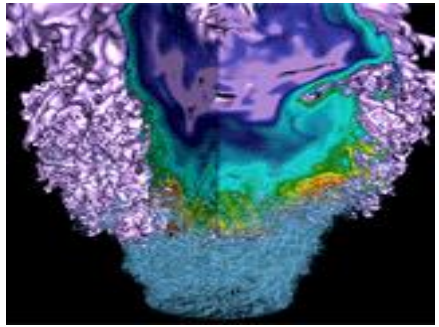

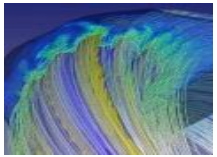
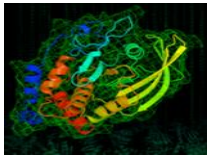
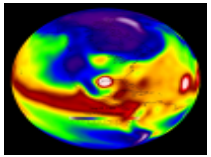
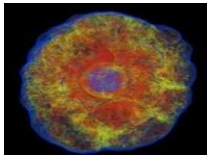
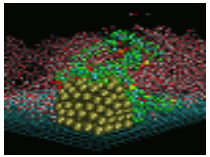

- Hardware and Software
- System and Data Management Practices
- Data Environment

1.) Hardware/Software

- Hardware
  - Vendor and media diversity are important
    - Issues affecting one media/drive type may not affect another – drive, firmware, faulty media, etc.
  - Enterprise Tape
    - Higher reliability, capacity and performance
    - Media reuse between drive generations – cost savings
  - Enterprise class servers and disk arrays
    - High reliability needed for data availability, also due to certain issues with HPSS – e.g. single metadata instance
    - Reliability and prompt vendor support important
- Software
  - HPSS – 20+ years of active development and support by IBM and 5 DOE Labs
  - Has the most tape features and supported HW
  - Scales to Exabytes
  - Roadmaps years into the future
- Integration engineering is performed in-house

2.) System and Data Management Practices

- Periodic Drive/media refresh
  - Media and drives refreshed every 3 - 4 years to contain data growth
    - Old media copied to new on continual basis – all data re-read periodically
  - 2017 - 2019 data center move:
    - Read 121PB in 29k cartridges, 230M files from old STK infrastructure to new IBM libraries
    - Loss of 1964 files from 90 cartridges, total of 0.0008%
      - The general mode of data loss is failure to read a very small portion of a tape, usually due to physical damage
  - Other hardware (servers, arrays) refreshed on basis of cost, performance, and vendor support availability
- Planned vs. Unplanned Maintenance
  - Periodic scheduled maintenance avoids unscheduled outages[1]
  - Planned incremental vs. large-scale changes – minimizes impact
  - Maintenance planned to level of individual commands

[1] Hurlbert, Wayne, *U.S. Department of Energy Best Practices Workshop on File Systems & Archives*, September 26 – 27, 2011

System and Data Management Practices, continued

- Careful and conservative system management
  - Testing – new hardware, config changes, upgrades
  - Monitoring – hardware and media health, environment
  - Problems addressed promptly as they occur
- Dual copy media
  - 2nd copy of files made to separate media where possible
- Other:
  - Daily metadata and quarterly offsite backups
  - configuration management, automated host provisioning

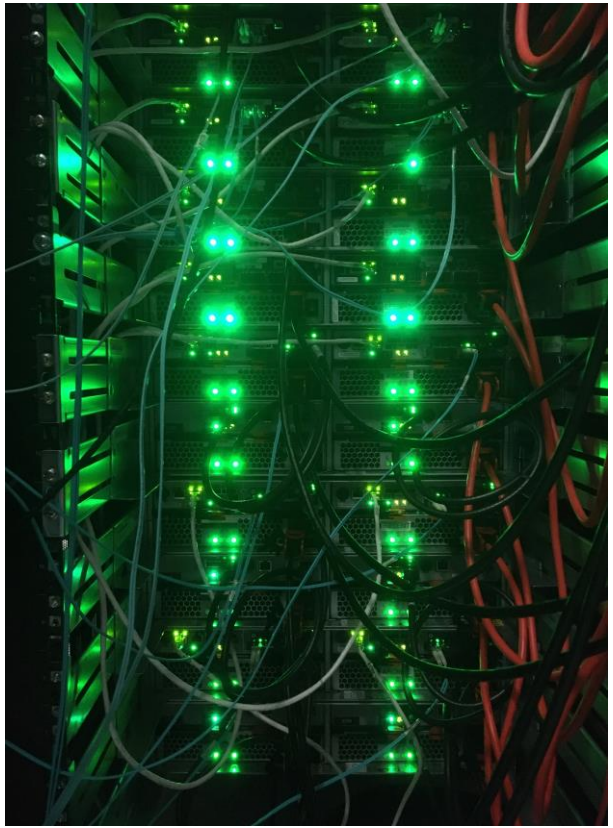

3.) Data Environment

- Airborne particulate monitoring
  - Oracle/STK and IBM tape storage equipment require ISO Class 8 cleanroom
  - NERSC OTG group monitors particulate levels in the data center
  - We shut down libraries if particulate exceeds thresholds
- Temperature and humidity control
  - Tape media reliability can decrease with rapid temperature and humidity fluctuations
  - NERSC data center uses open air cooling to conform to LEED Gold energy standards
    - Temperature and humidity can fluctuate rapidly
  - We use IBM TS4500 libraries with integrated cooling to keep tape environment stable
    - Recirculates internal air – during 2017 Northern CA wildfires library internal particulate levels stayed within Class 8 spec

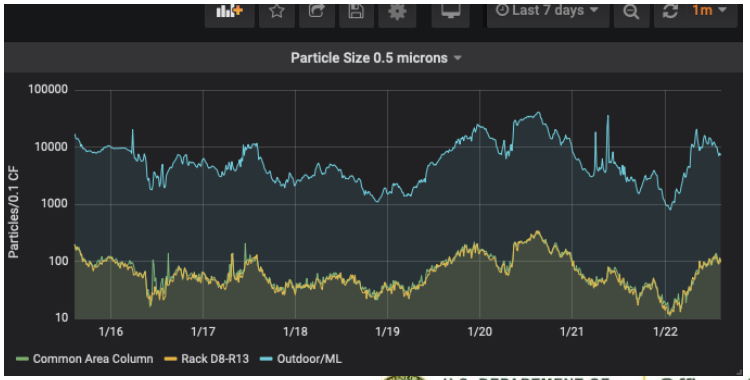

Incidents

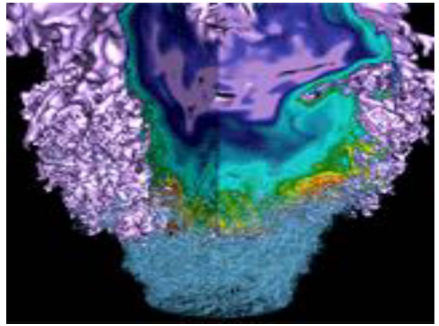

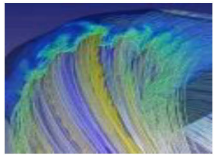

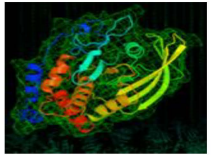

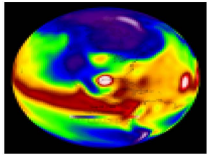

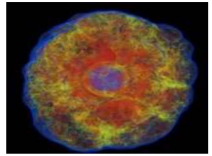

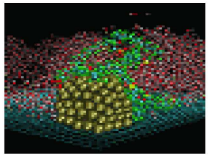

- Dust Incident
- 9840 “Dimple Syndrome”

Dust Incident

- In 2010, data center construction exposed Oracle/STK environment to a significant amount of drywall dust contamination
  - Observable dust layer coated library internal surfaces - robots, cartridges, drives
    - Unknown damage to media and library components at time of incident
  - HPSS Team worked with vendor and NERSC Facilities to mitigate equipment and media damage – at first we weren't sure if tape IO was safe
    - Oracle determined 9840 drives and media unaffected – HW diversity is important
    - Extensive cleaning of library internals and cartridge surfaces after-hours
    - Library containment/positive pressurization/HEPA filtration fabricated by NERSC Facilities
    - "Clean/Dirty" HPSS environment configured via HPSS PVR functionality– contaminated tapes copied to new media using vendor cleaned and recertified tape drives
  - Data loss
    - All files on 7,380 contaminated T10KB cartridges recovered – no data loss
    - Independent LBNL analysis[2] showed only the first 30ft of tape affected
  - Changes to data management practices
    - Established particulate monitoring/alerting, library complex shutdown (manual) on alert
    - Oracle/STK libraries continue to operate with in-house HEPA filtration system

[2] McKinney, Wayne R., Voronov, Dmitry L., Yashchuk, Valeriy V. *"Measurements Contaminated Magnetic Tapes for LBNL Library,"* December 20, 2010

“Dimple Syndrome”

- In 2011 9840 drive errors started to occur frequently
  - Difficulty reading files from 9840 media
  - HPSS Team notified vendor regarding potential drive or media issue
    - Vendor sent team on site to visually inspect 10,400 9840 cartridges in STK library complex
    - Determined drive firmware change caused some level of physical damage to media in ~3,000 9840 cartridges (“Dimple Syndrome”)
    - Mitigation involved special “recovery” firmware to more tightly pack affected media around internal cartridge spools to flatten damaged sections
  - Data loss
    - Lost ~2,000 files on ~500 damaged 9840 cartridges, a total of 200GB
    - Another argument for hardware diversity
  - Changes to data management practices
    - Now dual-copying small files to separate media – low overhead
    - Also try to keep all small files permanently resident in HPSS disk cache

Futures

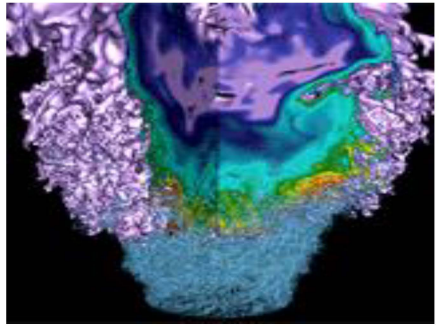

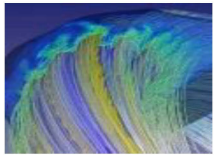

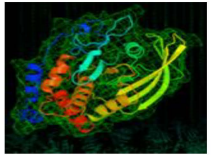

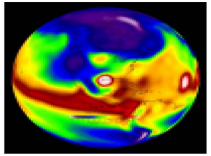

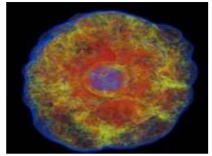

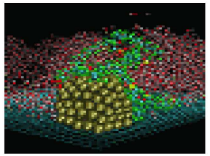

- Dual Copy
- RAIT
- Media verification
- Offsite Disaster Recovery

Future Enhancements

- Technologies we are planning and/or would like to implement to enhance data reliability
  - Full redundancy for all files in the archive
    - Permanent dual copy for small and medium files
      - Disk and/or alternate small-file compatible technology e.g. WORM DVD
  - RAIT for large files
    - Like RAID for disk, i.e. parity written to tape
    - Supported in current versions of HPSS
  - Continuous automated media verification
  - Offsite disaster recovery (DR) site
    - 2nd geographic site for important/irreplaceable data
  - No plans for HA systems at this time
    - Added cost of complexity and system management issues worth the benefit?

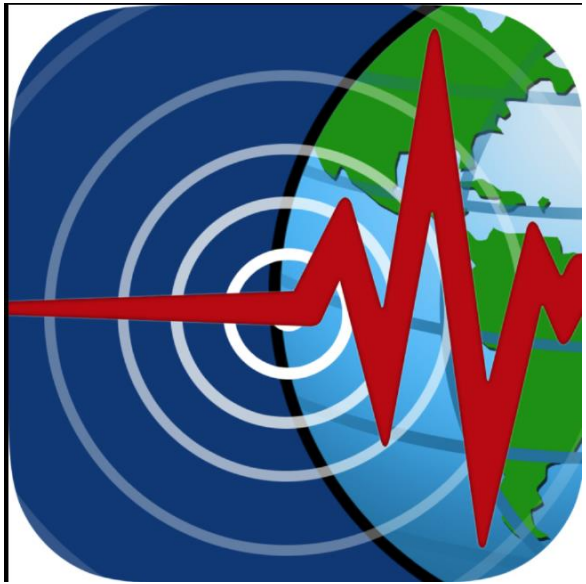

Thank You

"Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Networking and Information Technology Research and Development Program."

The Networking and Information Technology Research and Development (NITRD) Program

Mailing Address: NCO/NITRD, 2415 Eisenhower Avenue, Alexandria, VA 22314

Physical Address: 490 L'Enfant Plaza SW, Suite 8001, Washington, DC 20024, USA Tel: 202-459-9674, Fax: 202-459-9673, Email: nco@nitrd.gov, Website: https://www.nitrd.gov