NERSC HPSS
Site Update

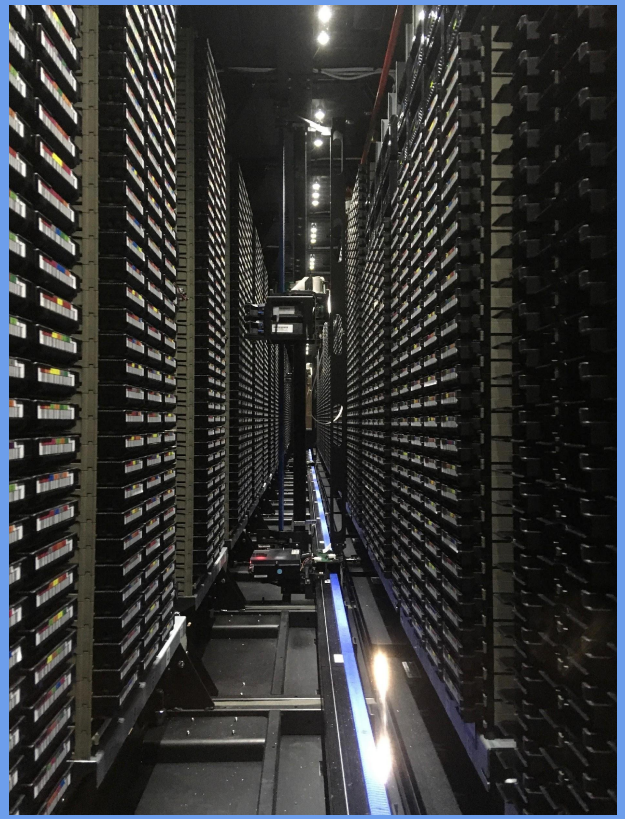

NERSC Storage Systems Group
October 20, 2021

Presenter
Nick Balthaser

NERSC is the mission computing facility

for the U.S. Department of Energy Office of Science

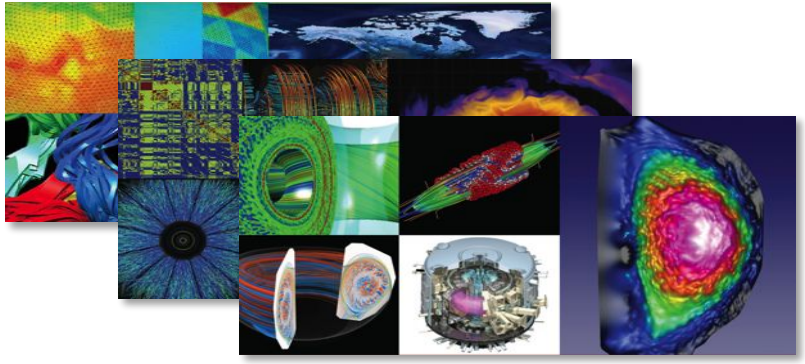

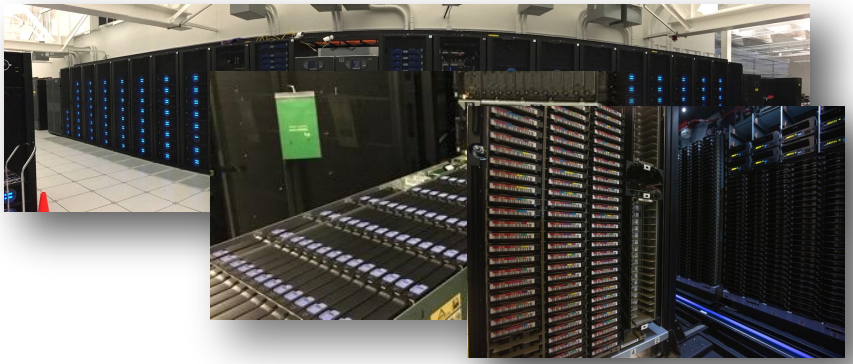

- **Largest funder of physical sciences research in the U.S.**
- **Diverse user community**
  - 8,000 active users, 900 projects
  - 700 applications (sim, data, AI)
- **We design for our workload**
  - Many jobs at many scales (40% of hours go to capability jobs)
  - Small, incoherent I/O
  - Not just checkpoint/restart!

2020 NERSC by the Numbers

9 BILLION CORE HOURS USED IN 2020

>1,800
Refereed Publications
Cited NERSC

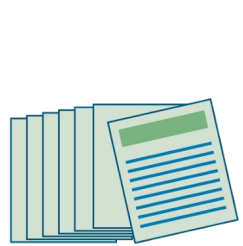

NERSC | BERKELEY LAB | U.S. DEPARTMENT OF ENERGY | Office of Science

Slide credit: Richard Gerber

NERSC's infrastructure for science

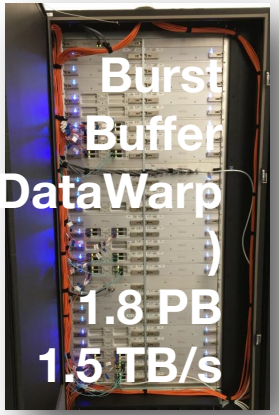

16x MDS + 274 OSS
1x AMD Epyc 7502P
2x Slingshot NICs
24x 15.36 TB NVMe

1,536 GPU nodes
1x AMD Epyc 7763
4x NVIDIA A100
4x Slingshot NICs

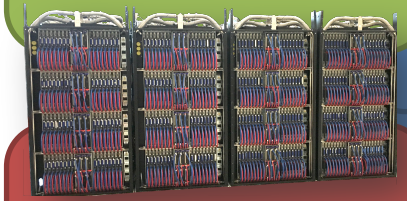

Slingshot
200 Gb/s
2-level dragonfly

3,072 CPU nodes
2x AMD Epyc 7763
1x Slingshot NIC

24x Gateway nodes
2x Slingshot NICs
2x 200G HCAs

SAN

SAN

2x Arista 7804 routers
400 Gb/s/port
> 10 Tb/s routing

WAN

NERSC Data Archive

45 years of data archived by the scientific community

- 20,000 cartridges in 3 IBM TS4500 libraries

HPSS software in production since 1998

- 2 systems:
  - Archive – user-facing: 240PB
  - Regent – center backups: 35PB

High utilization

- Active archive: > 30 - 40% retrieval rate
- 1.4 - 1.7x yearly growth

Unique environmental controls

- IBM integrated cooling libraries: enable operation in green data centers
  - 1st DOE site to use this technology
- Library particulate/AQI monitoring

2020 - 2021 Projects and Milestones

1. HPSS Upgrade
2. Modernized DevOps Processes, Tools & Automation

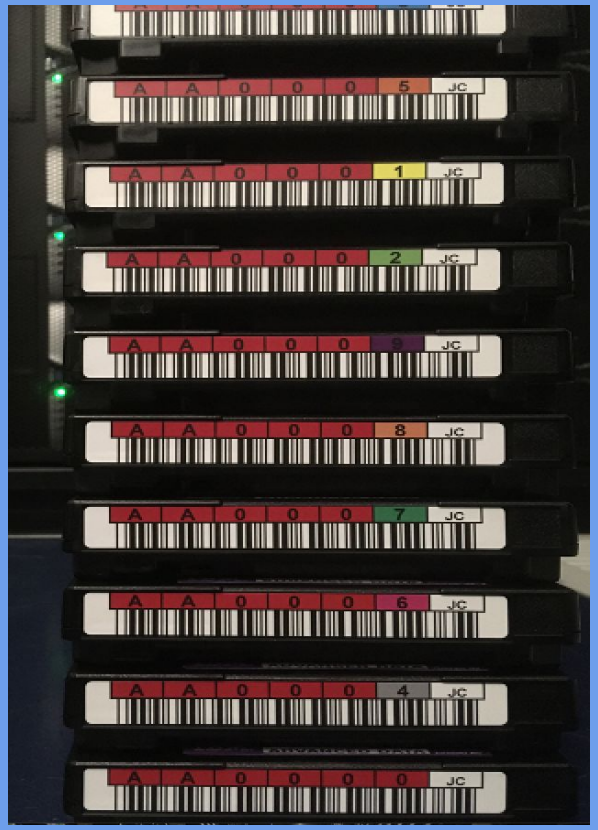

HPSS Upgrade

- 1st HPSS upgrade in 8 years
- Accomplished via remote work
- 6 day scheduled outage, 1 year planning
- Joint Dev/Ops effort
  - Integration of local mods and new HW
- Numerous process enhancements
  - Install automation, test & issue tracking
- SW Updates:
  - Metadata: DB2 9.7 ➡ 11.1 on flash
  - Latest Globus/HPSS interface (DSI)
  - Upgraded HSI/HTAR
- Retired:
  - 6 IBM P55A/AIX servers - 2008
  - 3 IBM DS3500 metadata arrays - 2009
  - 1 Team Lead

HPSS Upgrade: Agile Process Improvements

Adapting the model of Continuous Improvement, our team has taken steps to make local HPSS software maintenance more efficient, simplified, and automated.

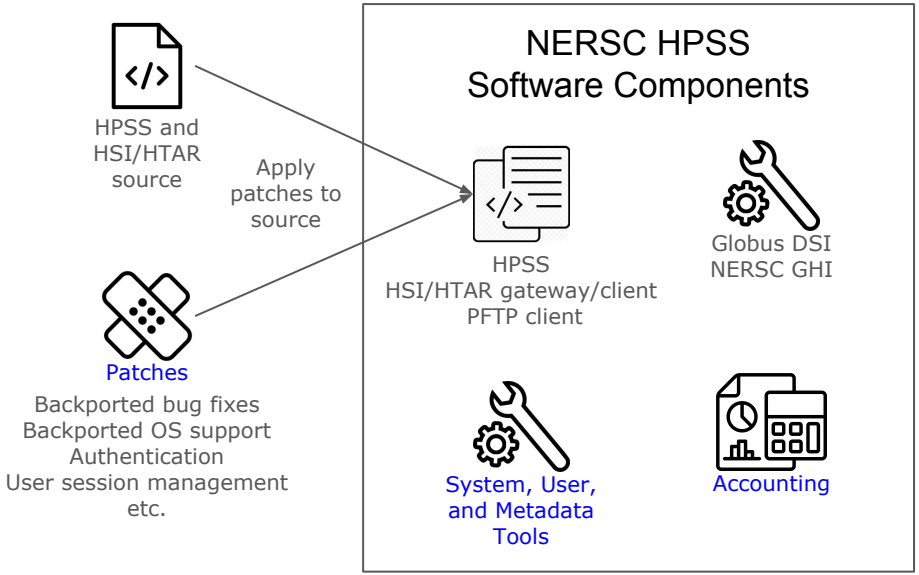

FOCUS AREAS

- Version control
- Issue tracking
- Local patches
- Automated build
- Automated deployment
- Test case management
- Documentation

TOOLS

GitLab, yum, rpm, TestRail, docker, xCAT, Ansible, MARKDOWN

Operational Challenges

1. Data Loss Incident
2. Drive & Accessor Issues
3. AQI

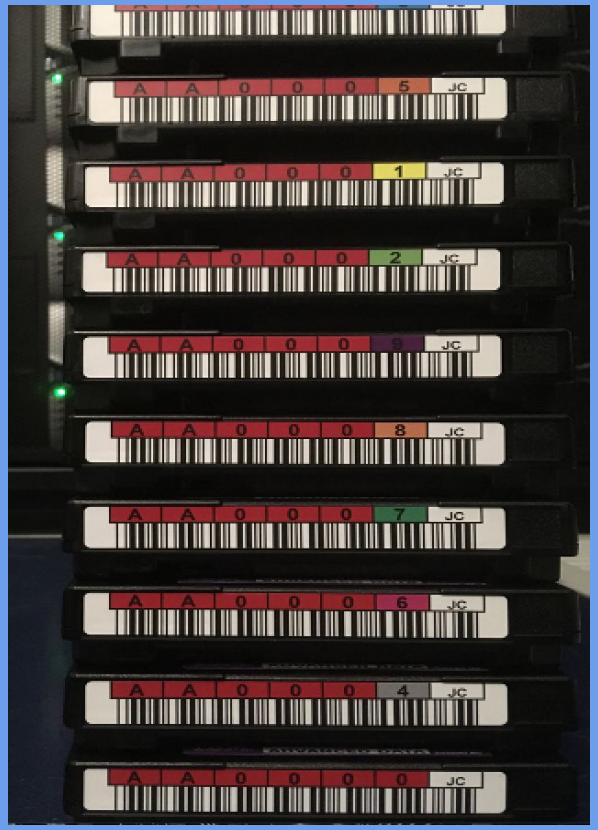

Data Loss

04/2021

- Mount issues with AG1126JD
  - EOM tape, 16TB, 1400 files

05/2021

- Attempts to read data via repack, force migrate, etc.
- Engaged IBM recovery after a month of read attempts - 68 files/1.7TB remaining
- Enabled Media Validation in TS4500s but eventually disabled it

09/2021

- Tape copy returned from IBM. Crease found at beginning of tape (BOT), files in that section unreadable
  - Tape copy has zero-padded files where there was damage
  - To do:
    - Restore files that are intact (non-padded)
    - Contact users to see if they're interested in zero-padded files

Drive & Accessor Issues

05/2021

- Upgrade 55F drive FW from **A14** to **B12** tensioning fix
- Upgrade TS4500 FW from **1.7.0.1** to **1.7.0.3**
  - **12 drive failures** after update (10%)
  - Increase* in tape & drive errors - ~2/day
    - * we enabled library reporting - previous rate(?)

08/2021

- **Dual Accessor failure** TS4500-1
  - Library down 4 hours
  - Accessor B failed due to worn pivot belt and severed X-track cable - blocked access to data cartridges
  - Accessor A failed twice for unknown reasons - no issues found

09/2021

- TS4500-2 Accessor A failed due to **track obstruction**
  - Accessor B unable to clear it
  - Accessor A failed again due to X-track cable short
  - Pursuing preventative library maintenance w/IBM
- Upgrade 55F FW from **B12** to **B8F**
- Upgrade TS4500 FW from **1.7.0.3** to **1.7.0.4**
  - **6 drive failures** post update

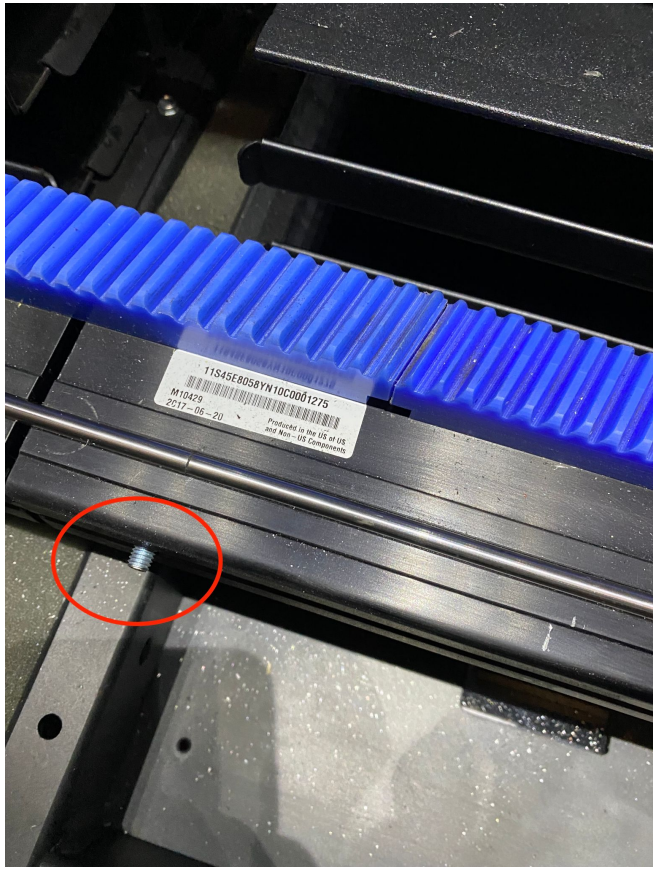

Photo credit: Kirill Lozinskiy

Accessor Obstruction Detail

Metal shavings

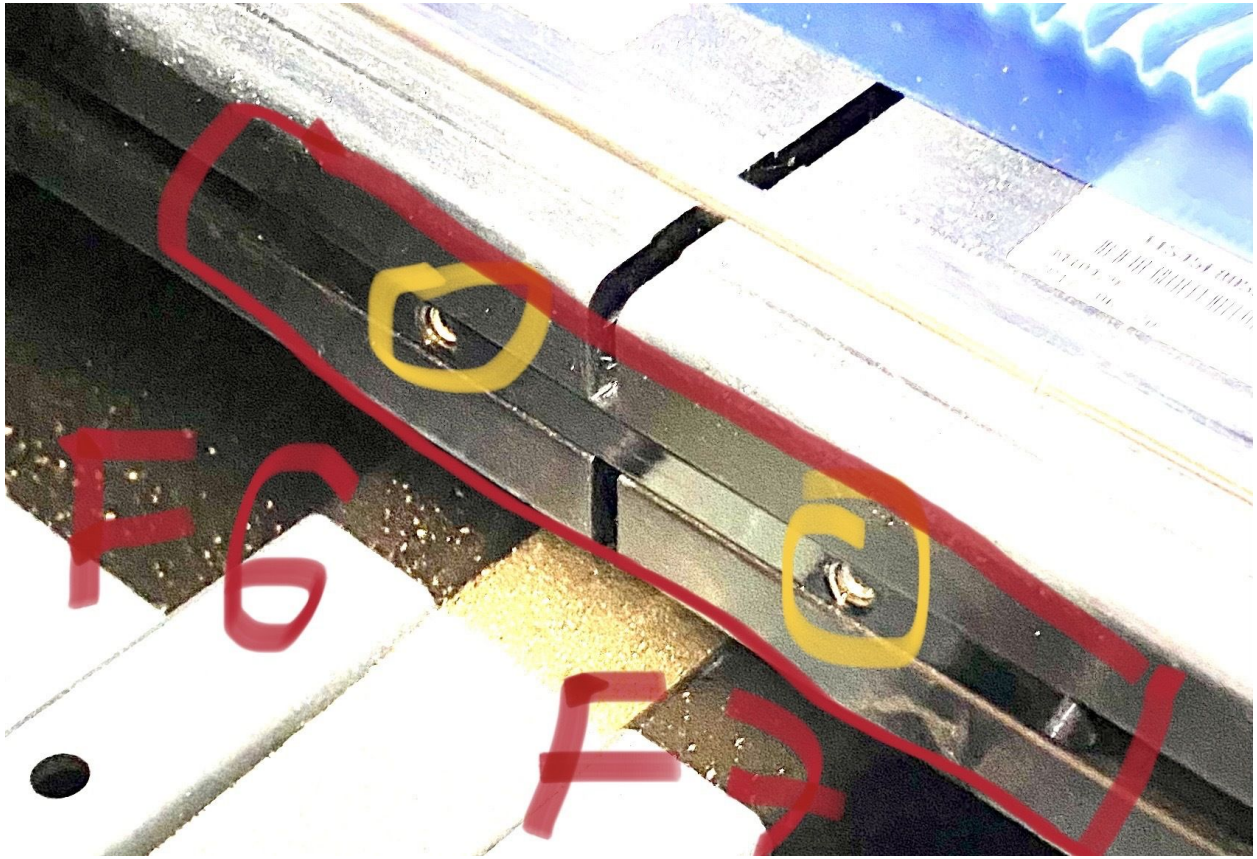

Photo credit: Kirill Lozinskiy

AQI and Particulate Monitoring

[IBM documentation] specifies ISO 14644-1 class 8 standard:

Data centers must meet the cleanliness level of ISO 14644-1 class 8. For data centers without

Class	maximum particles/m³					
	≥0.1 µm	≥0.2 µm	≥0.3 µm	≥0.5 µm	≥1 µm	≥5 µm
ISO 8	~~1.0×10⁸~~	~~2.37×10⁷~~	~~1.02×10⁷~~	3,520,000	832,000	29,300

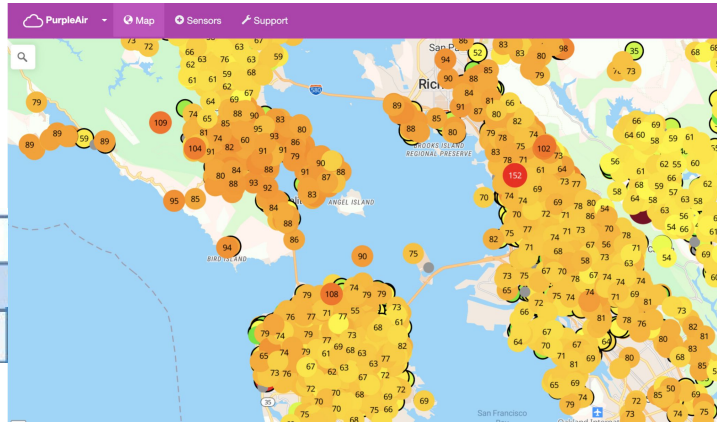

- AQI and airborne particulate monitoring are critical during CA fire season to maintain class 8 spec
- Monitoring: multiple particulate sensors within the data center, one in TS4500-2, [purpleair] site
- Recent Incidents:
  - **09/2020**: Tape IO disabled 5 consecutive days
  - **08/2021**: 2 AQI incidents of several hours each
- Data center AQI mitigation has improved
  - Still looking for a solution

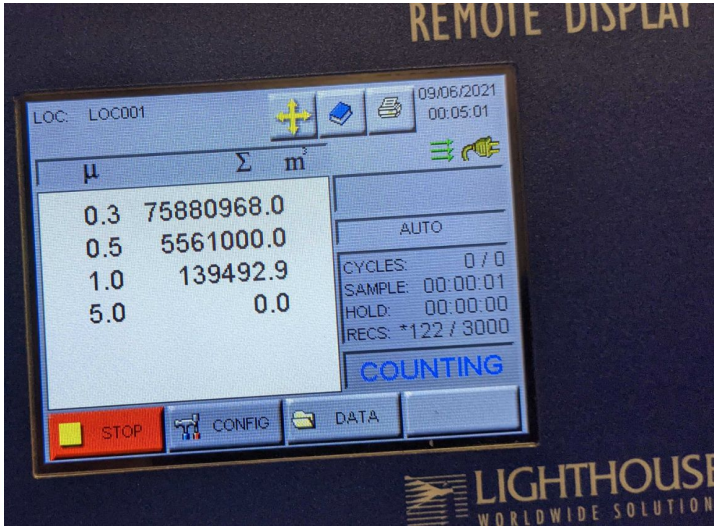

2021 Staffing Update

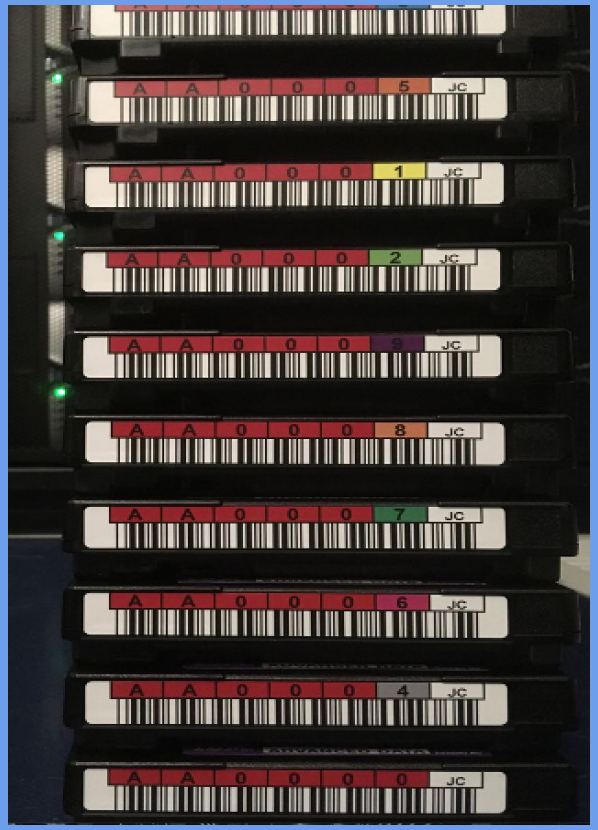

HPSS Staffing Changes

HPSS Team Lead Wayne Hurlbert retired after 30 years

- Wayne's legacy includes fundamental NERSC system design principles:
  - Use of Enterprise tape
  - "SuperMover" concept
  - Highly parallel disk and tape IO
  - Excellent system reliability and user ratings

SSG welcomes GL Kristy Kallback-Rose 🎉

- Glenn Lockwood returned to research in 2020
- Kristy became acting GL and accepted permanent position in 2021
- Kristy brings many years of storage systems and management experience at NERSC and IU

NERSC HPSS welcomes Francis Dequenne and Owen James 🎉

- Francis, formerly at a UK site, joined NERSC in 01/2021 bringing decades of storage experience
- Owen joins HPSS from NERSC OTG. Owen has been our advocate since 2016, leading several high visibility HPSS projects in his former role
- Leo Saavedra from NRAO to join soon (remote)

Special Thanks

Rosario Martinez, HPSS Software Intern

- Joined NERSC HPSS Dev Team through the Community College Internship program then returned as a limited term employee, 01/2021 - 06/2021
- Rosario's work includes containerized HPSS client setup and porting NERSC HPSS authentication utilities to Python 3
- Now completing CS degree at Georgia Tech

NERSC HPSS Team

Rocko
Group Lead & TC Rep
[Kristy Kallback-Rose]

Parker
HPSS Development
[Melinda Jacobsen]

Rosie
DevOps & Systems Deployment
[Kirill Lozinskiy]

Cleo
DevOps & Systems Deployment
[Francis Dequenne]

Noni
DevOps & Systems Deployment
[Owen James]

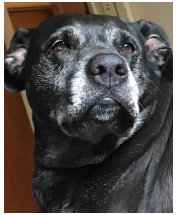

Zoe
DevOps & Systems Deployment
[Nick Balthaser]

Upcoming Projects & Works in Progress

1. Operational Projects
2. HPSS Development
3. Center Initiatives

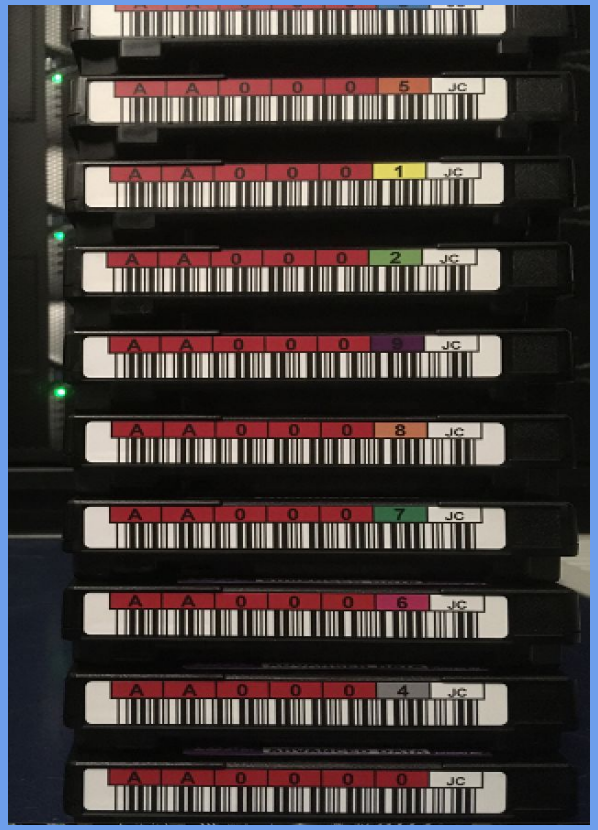

Upcoming Operational Projects

Clearing Long Term Backlog
Multi-year building move and year long HPSS upgrade left significant HW and operational backlog

TS1160 drive/JE media update for Large COS
- Drives in boxes since 06/2020

Fibre Channel Refresh 💰 💰 💰
- Cisco 9513s > 10YO, EOL 04/2022
- Over 400 ports needed for disk and tape

TS4500-4 purchase/installation
- Rapidly filling up 3 TS4500s
- "Integrated cooling" feature - approaching last sale date

HPSS Upgrade - again
- HPSS 7.4.3 EOL 2016
- Target 8.3

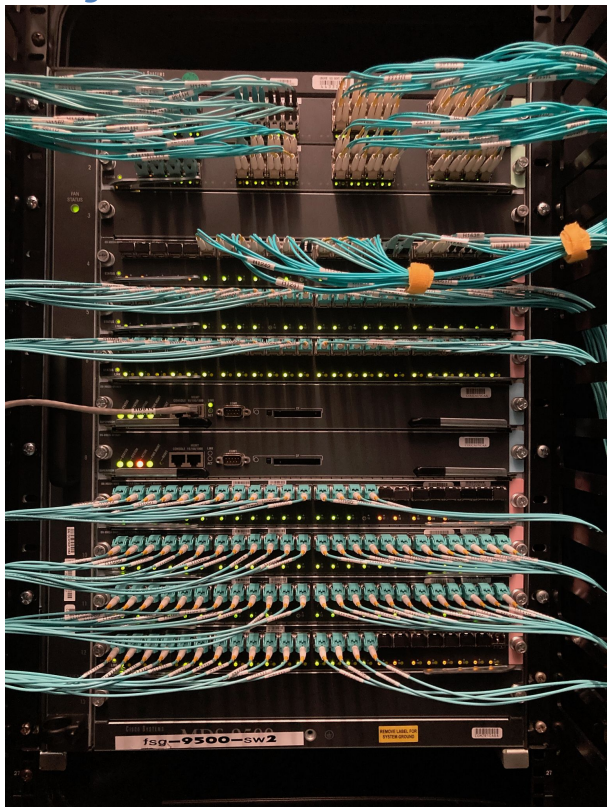

HPSS development at NERSC

Containerized Clients

On Perlmutter users work within containers orchestrated by Kubernetes

- NERSC HPSS clients including FUSE will need to be able to operate and transfer data from within a containerized environment
  - Ongoing work at NERSC to containerize HSI/HTAR and Globus
  - Melinda covered this topic in detail on 10/13

SuperFacility API

- NERSC effort to automate common HPC tasks & operations, e.g. job status queries, data transfers, etc. via REST API
- Ongoing work at NERSC to determine how this will be applied to HPSS clients and data transfers

HTAR Enhancements

- Size limitation, UID/GID, pathname issues deferred over 10 years
- Will feature newer TAR format (PAX) & Index file
- Planned for HPSS release 10.1

NERSC Center Initiatives

FAIR

Making data searchable across NERSC, particularly within HPSS

- **Possible use of UDAs** for data tagging
- Rich metadata for describing and locating files - probably external to HPSS

Resilience

Interest at NERSC in increasing system reliability and uptime

- Designing systems resistant to:
  - Component failure
  - Environmental (e.g. power, earthquake, AQI) issues
  - Cybersecurity attack
  - Data loss/corruption
- Near term HPSS solutions:
  - RAIT, E2EDI
- Long term:
  - Power, network, environmental controls
  - Cross-site cooperation for DR
  - RAIL?

Future Directions

Goals

- Ensure Archive is *Exascale-ready*
  - **Automated data movement** between tiers* - file system integration is key
    - Ongoing GHI and FUSE evals
  - Extreme speed and capacity
    - Parallelism in network, disk, and tape systems

- Continuing DevOps process improvements:
  - Labor-intensive task and deployment automation
    - Continued use and evaluation of automated deployment and system management tools
    - Remote system management

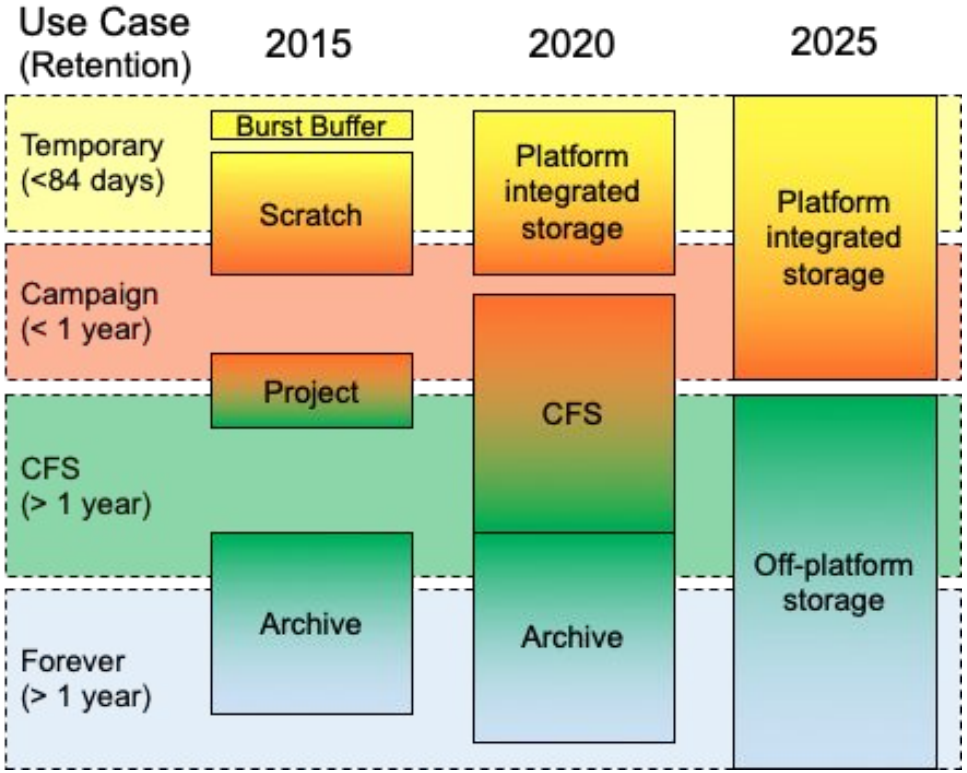

* More info: G. K. Lockwood *et al.*, "Storage 2020: A Vision for the Future of HPC Storage," Berkeley, CA, 2017.

Thank you!

Questions?

Backup and Supporting Slides

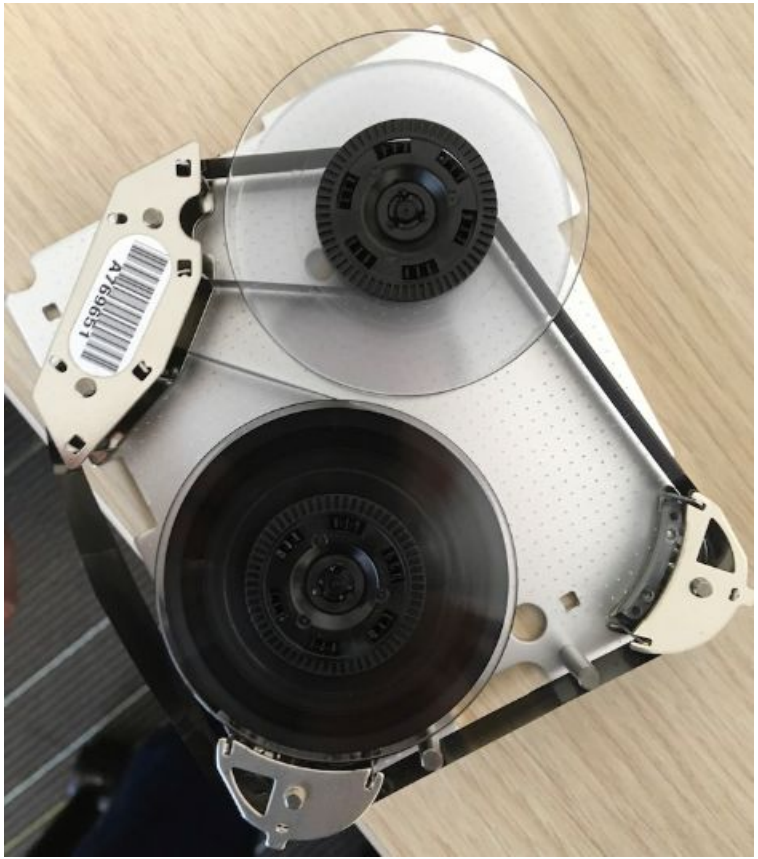

26

2020 NERSC by the Numbers

NERSC BERKELEY LAB U.S. DEPARTMENT OF ENERGY Office of Science

2020 NERSC by the Numbers

7,887 ANNUAL USERS FROM **~1,750** Institutions + National Labs

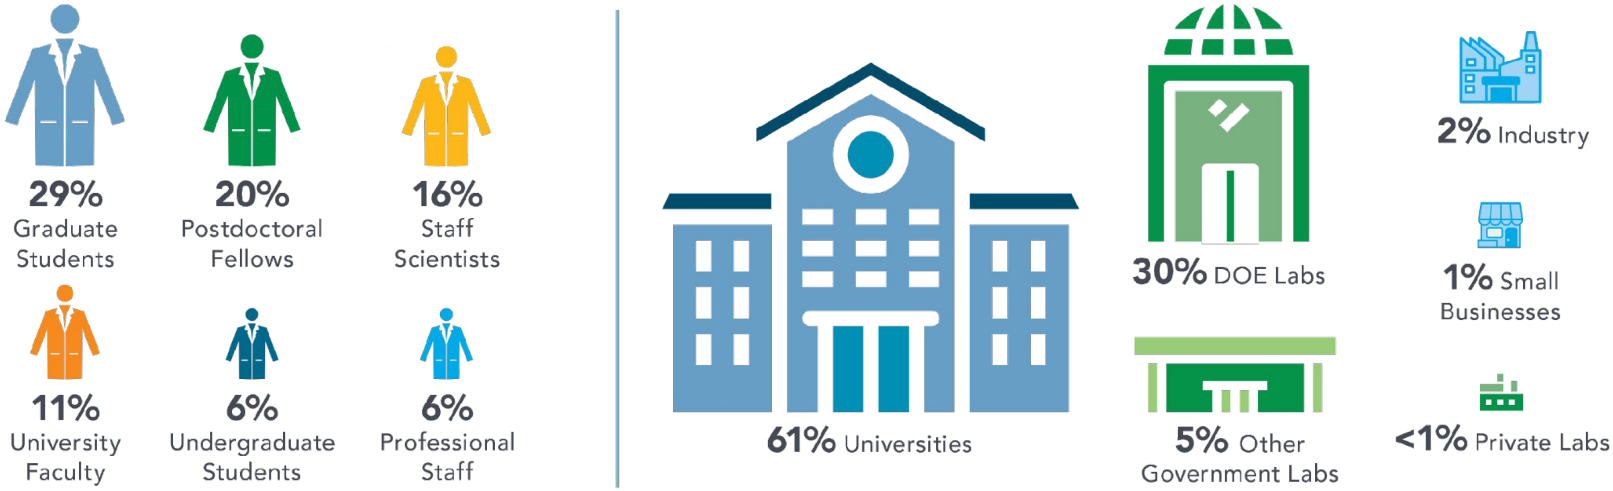

NERSC | BERKELEY LAB | U.S. DEPARTMENT OF ENERGY | Office of Science