


Visual Data Management at NERSC




**Annette Greiner and
Lisa Gerhardt**

**Supercomputing State of Practice,
November 17, 2020**

NERSC supports many users and projects from DOE SC's experimental and observational facilities




planck



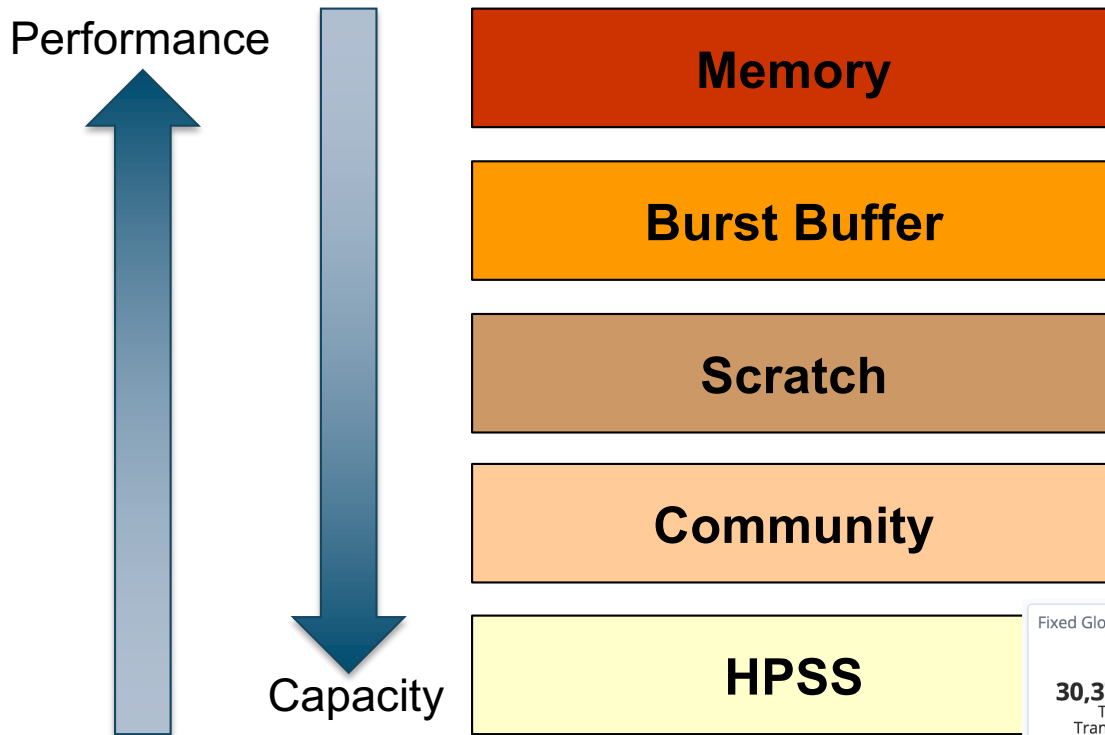
Experiments
operating now

Future
experiments



See NERSC's talk about Superfacilities: **“Cross-Facility Science: The Superfacility Model at Lawrence Berkeley National Laboratory”** also in this session

Data At NERSC



Community File System: Intended for sharing scientific data between groups of NERSC users.

27 PB and 3 billion inodes used

Medium-term (1 - 3 years) active data

Very useful, but has some issues

- Group quotas that fill up
- Permissions drift
- Data migration between the tiers

30% of all data movement is purely internal

Fixed Globus Metrics	Globus Internal Trans...	Globus Internal Data Tra...	Globus Collaboration En...
30,377,622.13 Total Data Transferred [GB]	13,965,183.829 Total Data Transferred Internally [GB]	11,926,864.919 Total Data Transferred Internally to/from Cori Scratch [GB]	816,667.051 Total Data Transferred Using Collab Endpoint [GB]
785 Total Users	145 Total Users	99 Total Users	10 Total Users
31.95 Average Unique Users Per Day			

Data Dashboard: Managing Space



My NERSC [Email] [Menu] [User]

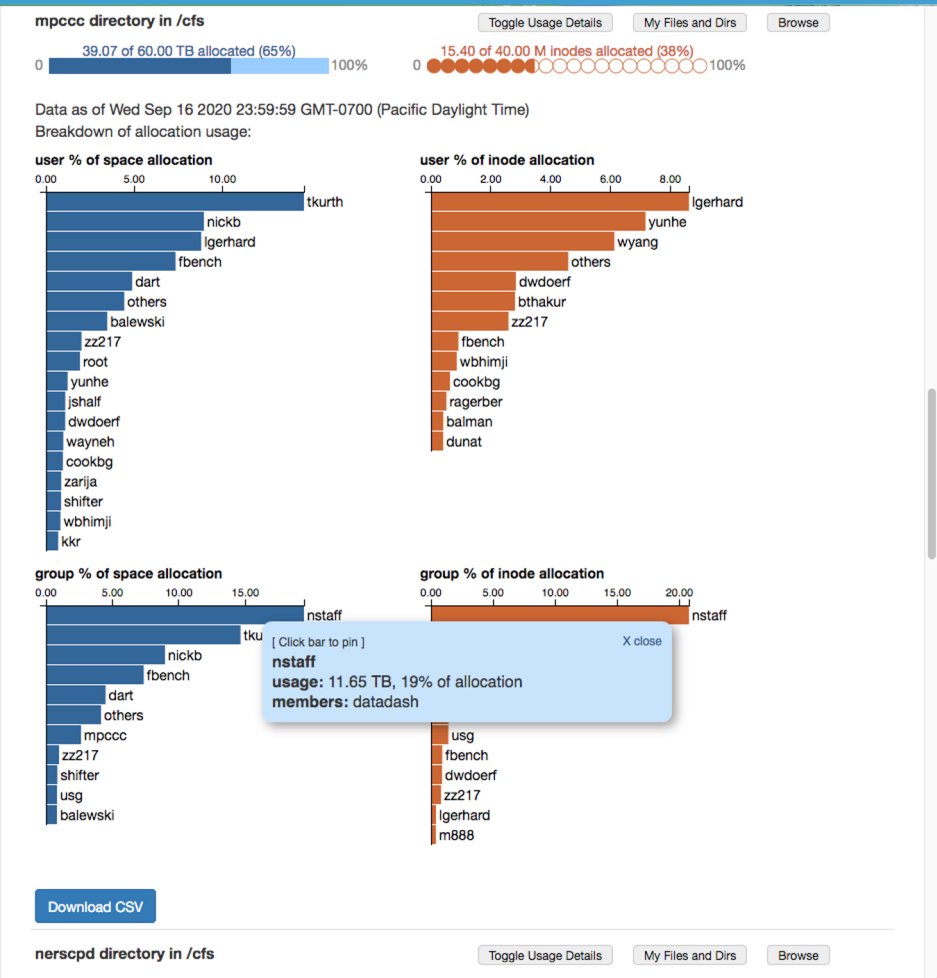
- Dashboard
- Jobs
- Center Status
- File Browser
- Service Tickets
- Data Dashboard**
- NX Desktop
- Jupyter Hub
- NERSC Homepage
- Documentation Portal
- Accounts Portal

Data Dashboard

Showing disk space and inode usage for global directories at NERSC to which you have access as PI, PI proxy, or user (includes /cfs, /projecta, and /projectb/sandbox)

Directory	Disk Space Usage	Inode Usage	Actions
m1541 directory in /cfs	0.42 of 1.00 TB allocated (42%)	0.25 of 20.00 M inodes allocated (1%)	Toggle Usage Details My Files and Dirs Browse
m2002 directory in /cfs	0.00 of 20.00 TB allocated (0%)	0.01 of 20.00 M inodes allocated (0%)	Toggle Usage Details My Files and Dirs Browse
m2043 directory in /cfs	151.41 of 250.00 TB allocated (61%)	3.26 of 20.00 M inodes allocated (16%)	Toggle Usage Details My Files and Dirs Browse
m888 directory in /cfs	15.44 of 26.00 TB allocated (59%)	4.36 of 12.00 M inodes allocated (36%)	Toggle Usage Details My Files and Dirs Browse

Data Dashboard



Data Dashboard



My biggest /cfs/m2002 files and dirs

as of Wed Sep 16 2020 23:59:59 GMT-0700 (Pacific Daylight Time)

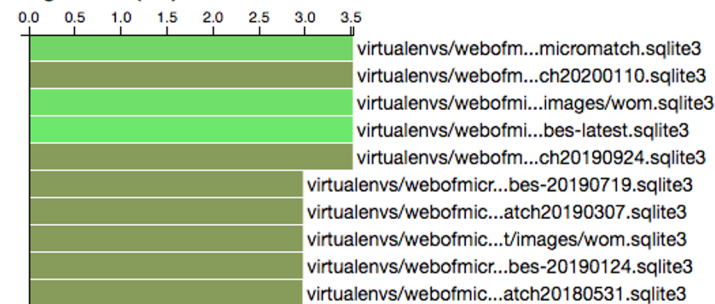
Color by: Access Time Change Time

New ≥1yr

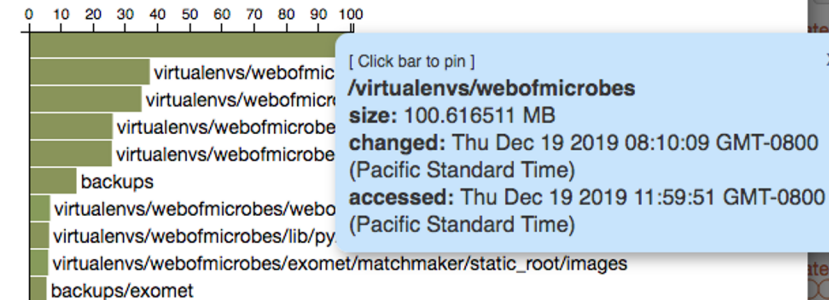
Show parents of top child directories*

Show the top

Largest Files (MB)



Largest Dirs* (MB)



[Click bar to pin] X close

/virtualenvs/webofmicrobes
size: 100.616511 MB
changed: Thu Dec 19 2019 08:10:09 GMT-0800 (Pacific Standard Time)
accessed: Thu Dec 19 2019 11:59:51 GMT-0800 (Pacific Standard Time)

Most Inodes* (thousand)

Data Dashboard




My /cfs/mpccc files and dirs

as of Wed Sep 16 2020 23:59:59 GMT-0700 (Pacific Daylight Time)

sunburst icicle

copy path

mpccc/www/agreiner/mynerc/live/js/plugins/ace-builds/.git 3.75%



Data Dashboard Front End



- HTML, Javascript, D3
- Built into existing MyNERSC user site
- Thin PHP shell
- NEWT web API for calls to microservices
- Microservices read data files or query Postgres DB



Data Dashboard Backend



- ILM policy scans (Spectrum Scale specific tool) do a complete scan of the whole file system in ~3 hours
- Generates a ~600GB text file with 3 billion lines that contains all the HW metadata for each file and directory (uid, gid, mtime, ctime, size)
- File is analyzed in Spark to generate aggregated information for each project organized by both user and directory
- Produces jsons for web page within 10 minutes of scan finishing
- Uploads cleaned data to a postgres DB for most advanced web features (takes ~4 hours)



PostgreSQL

PI Toolbox: Permission Wrangling



My NERSC

agreiner

- Dashboard
- Queues
- Center Status
- File Browser
- Completed Jobs
- My Tickets
- Data Dashboard
- Changelog
- NERSC Homepage

PI Toolbox

Jump to: datamap directory in cfs-dev Make entire project dir group-readable

Path: /global/cfs-dev/cdirs/datamap/datasets Show

Select All

Select	Name	Owner	Group	Size	Date	Permissions
	Parent Directory					
<input type="checkbox"/>	This directory	agreiner	agreiner	4096	Sep 17 14:52	drwxrwx---
<input type="checkbox"/>	dataset15	agreiner	datamap	0	Sep 17 14:52	-rw-rw----
<input type="checkbox"/>	dataset16	agreiner	datamap	0	Sep 17 14:52	-rw-rw----
<input type="checkbox"/>	dataset17	agreiner	datamap	0	Sep 17 14:51	-rw-rw----
<input type="checkbox"/>	dataset18	agreiner	datamap	0	Sep 17 14:51	-rw-rw----
<input type="checkbox"/>	dataset19	agreiner	datamap	0	Sep 17 14:51	-rw-rw----
<input type="checkbox"/>	dataset20	agreiner	datamap	0	Sep 17 14:50	-rw-rw----
<input type="checkbox"/>	dataset21	agreiner	datamap	0	Sep 17 14:50	-rw-rw----
<input type="checkbox"/>	dataset22	agreiner	datamap	0	Sep 17 14:50	-rw-rw----
<input type="checkbox"/>	dataset23	agreiner	datamap	0	Sep 17 14:50	-rw-rw----
<input type="checkbox"/>	dataset24	agreiner	datamap	0	Sep 17 14:50	-rw-rw----

Action on Selected: Change Group Change Permissions

PI Toolbox



My NERSC

agreiner

- Dashboard
- Queues
- Center Status
- File Browser
- Completed Jobs
- My Tickets
- Data Dashboard
- Changelog
- NERSC Homepage

PI Toolbox

Owning group can:

- Read (r)
- Write (w)
- Execute file, enter directory
- Execute binary file normally, as member of user's default group (x)
- Execute binary file as member of owning group and force new items in directory to be owned by the group (s)
- Just enter directory, no file execution (X)
- I want to apply these permissions recursively

Cancel Submit request

	Date	Permissions
<input type="checkbox"/>	Sep 17 14:52	drwxrwx---
<input checked="" type="checkbox"/>	Sep 17 14:52	-rw-rw----
<input checked="" type="checkbox"/>	Sep 17 14:52	-rw-rw----
<input checked="" type="checkbox"/>	Sep 17 14:51	-rw-rw----
<input checked="" type="checkbox"/>	Sep 17 14:51	-rw-rw----
<input checked="" type="checkbox"/>	Sep 17 14:51	-rw-rw----
<input checked="" type="checkbox"/>	Sep 17 14:50	-rw-rw----
<input checked="" type="checkbox"/>	Sep 17 14:50	-rw-rw----
<input checked="" type="checkbox"/>	Sep 17 14:50	-rw-rw----
<input checked="" type="checkbox"/>	Sep 17 14:50	-rw-rw----
<input checked="" type="checkbox"/>	Sep 17 14:50	-rw-rw----
<input checked="" type="checkbox"/>	Sep 17 14:50	-rw-rw----
<input checked="" type="checkbox"/>	Sep 17 14:50	-rw-rw----

Action on Selected: Change Group Change Permissions

PI Toolbox



<input checked="" type="checkbox"/>		dataset22	agreiner	datamap	0	Sep 17 14:50	-rw-rw----
<input checked="" type="checkbox"/>		dataset23	agreiner	datamap	0	Sep 17 14:50	-rw-rw----
<input checked="" type="checkbox"/>		dataset24	agreiner	datamap	0	Sep 17 14:50	-rw-rw----
<input checked="" type="checkbox"/>		dataset25	agreiner	datamap	0	Sep 17 14:50	-rw-rw----

Action on Selected:

▼ Pending Requests

Command	Requested	Status
chmod in /global/cfs-dev/cdirs/datamap/datasets directory	9/17/2020, 5:14:54 PM	pending

▼ Recently Completed Requests

Command	Requested	Status
chmod in /global/cfs-dev/cdirs/datamap/datasets directory	9/17/2020, 4:09:17 PM	completed
chgrp in /global/cfs-dev/cdirs/datamap directory	9/17/2020, 3:59:32 PM	completed
chgrp in /global/cfs-dev/cdirs/datamap/datasets directory	9/17/2020, 3:11:26 PM	completed
chmod in /dirs/das/www/agreiner directory	9/4/2020, 5:24:50 PM	error

PI Toolbox Front End



- HTML, Javascript, React
- Built into existing MyNERSC user site
- Thin PHP shell (for now)
- NEWT web API for calls to middleware
- Requests saved to Postgres DB
- Middleware scripts and DB run as containers in secure Rancher environment



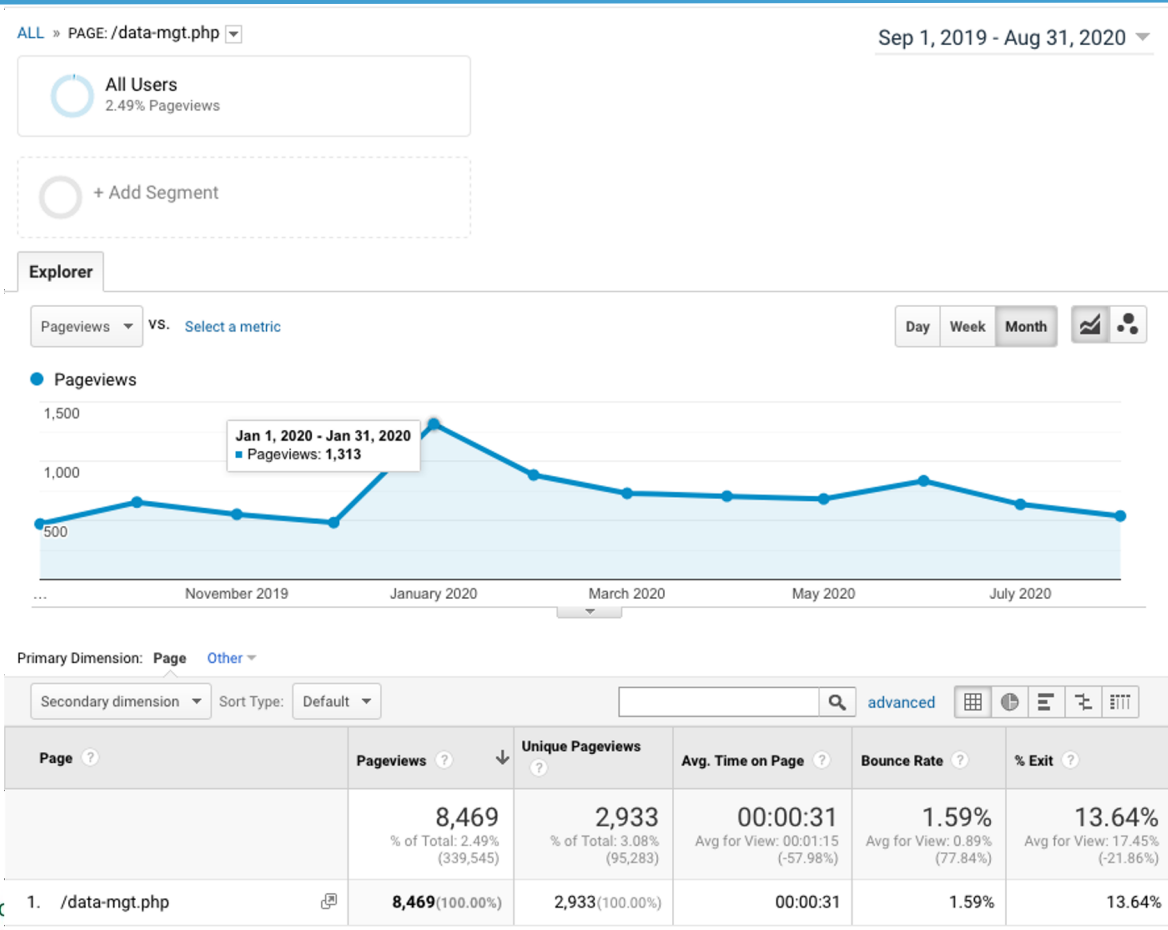
PI Toolbox Backend



- Python backend that reads commands from DB, sanitizes command string, and checks whether the user has permission to access the paths
- If all checks pass, the backend executes the commands as root and updates the DB with either "SUCCESS" or the error



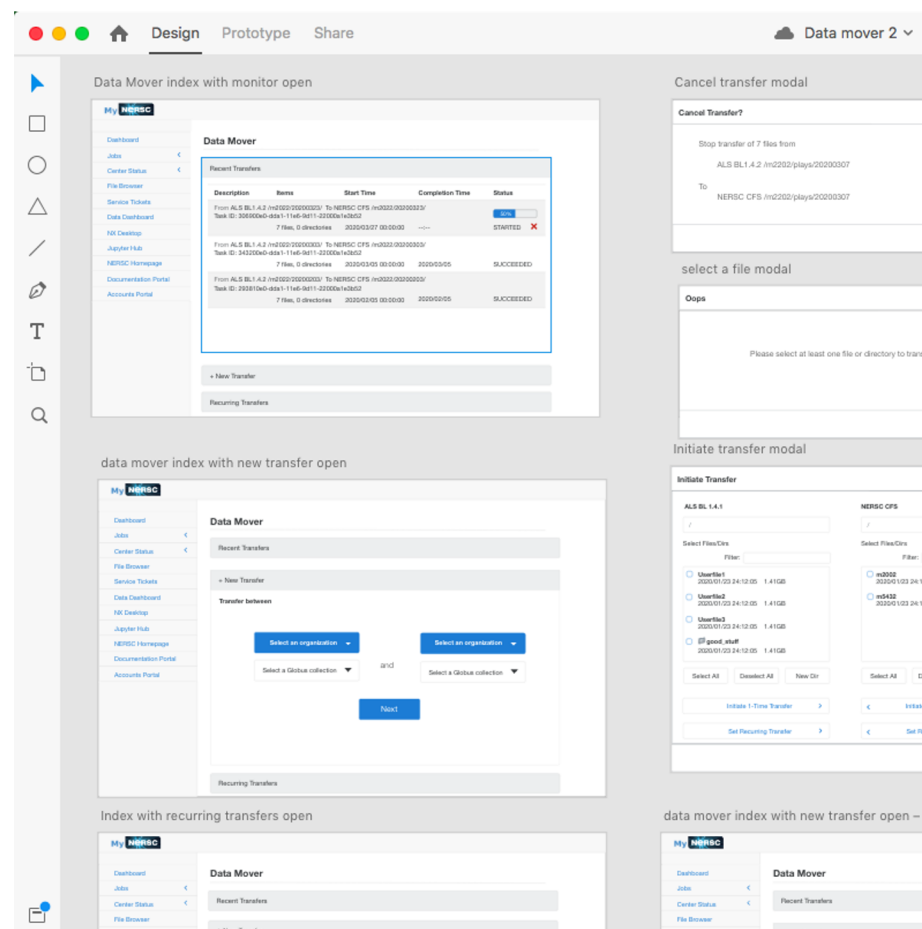
Usage Stats



Future Work and Conclusions



- PB Data Portal being designed
 - Adobe XD high-fidelity prototype for usability testing
 - Will be built with HTML, JavaScript, React
 - Will connect to Superfacility API to call microservices/Globus
- Giving users intuitive tools for data management improves scientific efficiency and saves staff time





Thank You