



HPC I/O Data Management Tools: What to do with your important data

ATPESC 2019

Philip Carns
Mathematics and Computer Science Division
Argonne National Laboratory


Q Center, St. Charles, IL (USA)
July 28 – August 9, 2019

The HPC Data Management Toolkit

- Getting your application to efficiently access data is only part of the battle
 - (To be fair, it is the hardest part. That's why we will spend most of the day on it!)
- What about stewardship of your data once it has been generated?
- We'll discuss two specific tasks:
 - Archiving data
 - Protecting data or saving it for long-term use
 - Keeping data when you have used up your quota
 - Transferring data across machines or across facilities
 - Using data at other facilities or sharing with collaborators

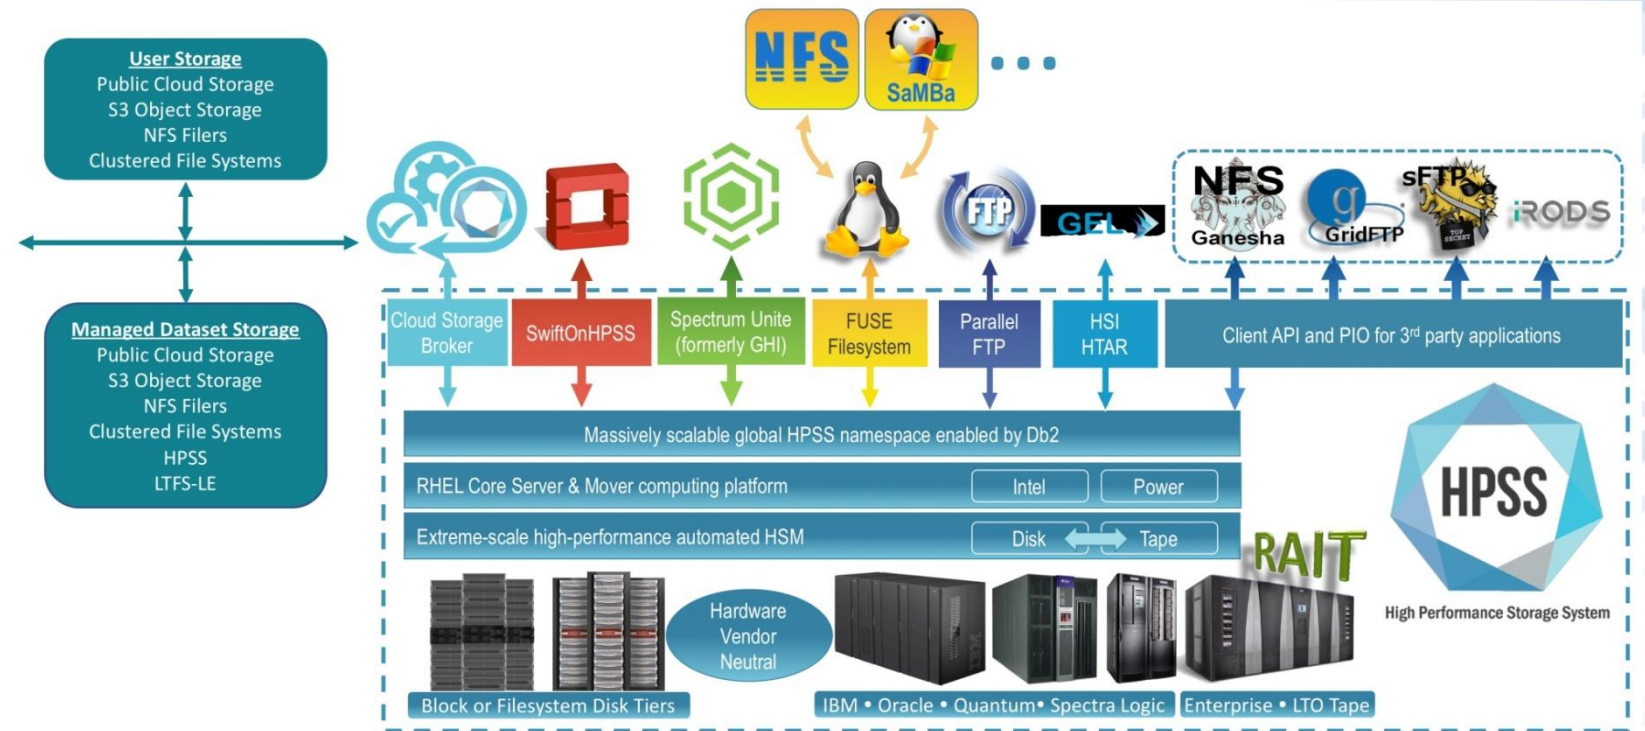
Archiving data: it's a big job

- If you have run out of disk space or need to protect or retain your data in the long term, then it's time to use an archival system.
- All of the DOE compute facilities have (very) large tape backup systems available for use from all of the major file systems
 - Dozens of tape controllers, hundreds of PiB of capacity, and **robots** 🤖
- A tape backup system like this could be difficult to use.
- Fortunately for us, it is all managed by a relatively easy-to-use system called HPSS.



HPSS overview

Image from <http://www.hpss-collaboration.org/>



- HPSS was developed as a collaboration between IBM and the DOE laboratories
- Organizes large collections of tape devices
- Designed to meet the needs of HPC facilities (especially scalability)
- Database to index metadata, disk array to cache data
- Can be accessed directly from Globus as well

Using HPSS

- <https://www.alcf.anl.gov/user-guides/using-hpss/>
- <http://www.nersc.gov/users/storage-and-file-systems/hpss/>
- There are a few different tools for interacting with HPSS, but “hsi” is the easiest
- “module load hsi”
- You can run it in a script (“hsi <command>”) or just launch it with no arguments to start an interactive session

HPSS example

```
carns@thetalogin4:~> module load hsi
carns@thetalogin4:~> hsi
*****
*           ANL HPSS 7.4 System
*****
Username: carns  UID: 4279  Acct: 4279(4279)  Copies: 1  Firewall: off [hsi.5.0.2.
p7 Thu Nov 30 16:42:05 UTC 2017]
[HSI]/home/carns->put notes.txt
put 'notes.txt' : '/home/carns/notes.txt' ( 4965 bytes, 935.9 KBS (cos=411))
[HSI]/home/carns->ls
/home/carns:
notes.txt
[HSI]/home/carns->quit
carns@thetalogin4:~>
```

- Load the hsi module
- Start an hsi session
- Use “put” command to place a file in the archive system.
- Use “ls” to browse

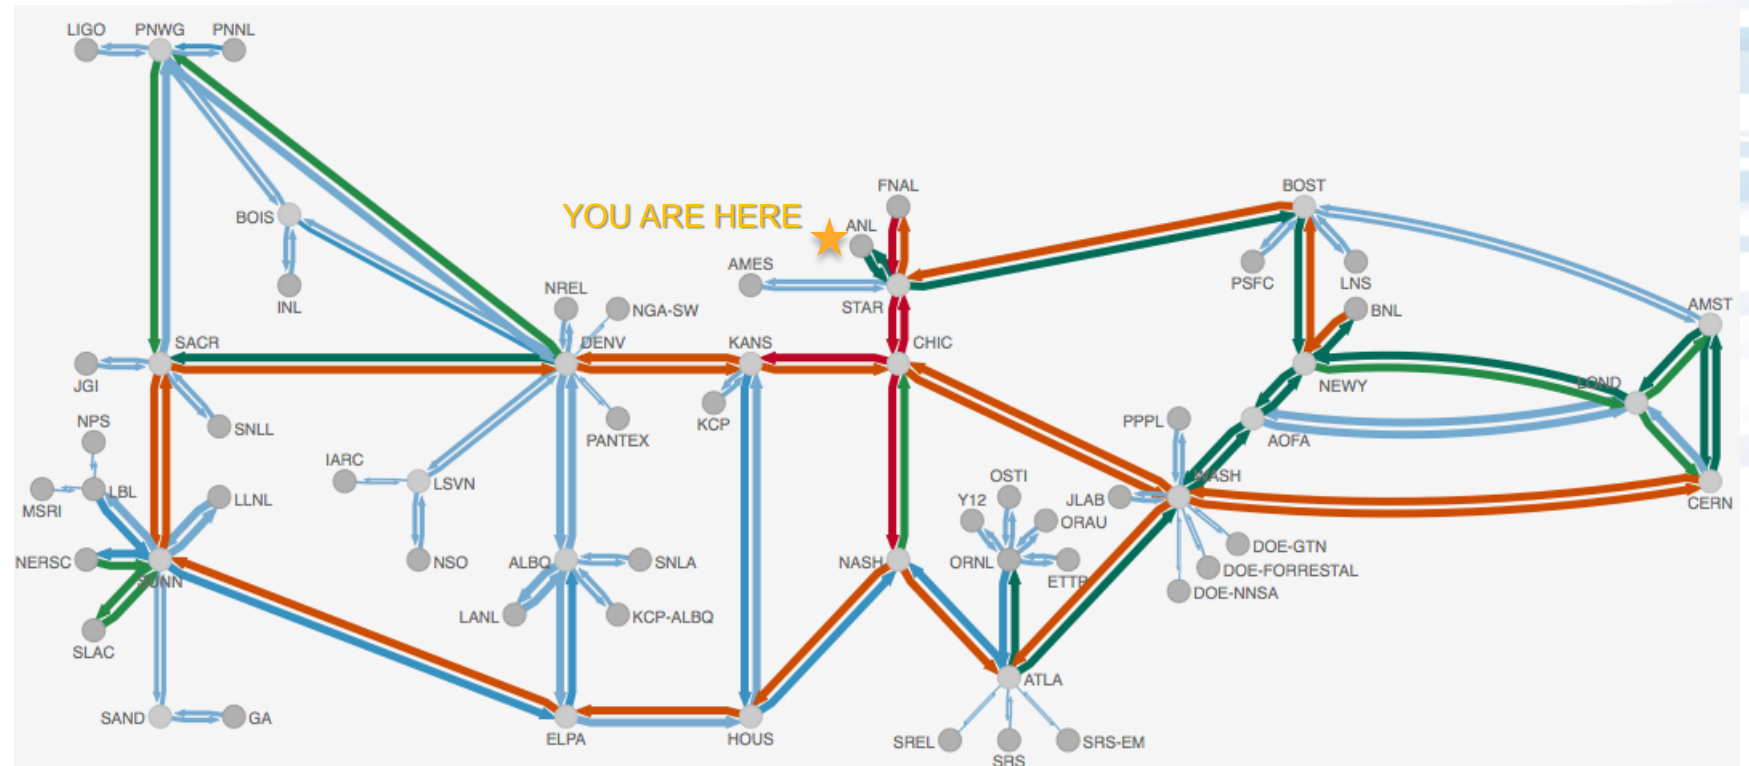
HPSS example

- That's about it!
 - “put” stores a file
 - “put -R” stores a directory
 - “get” brings them back
- If hsi doesn't work for you, contact support to have it activated
- *HPSS will not work for ALCF guest accounts during today's hands-on!*

Pro tip: there is an alternative command line tool available called “htar”. Avoid htar if possible, though: it has some subtle limitations that can cause problems.

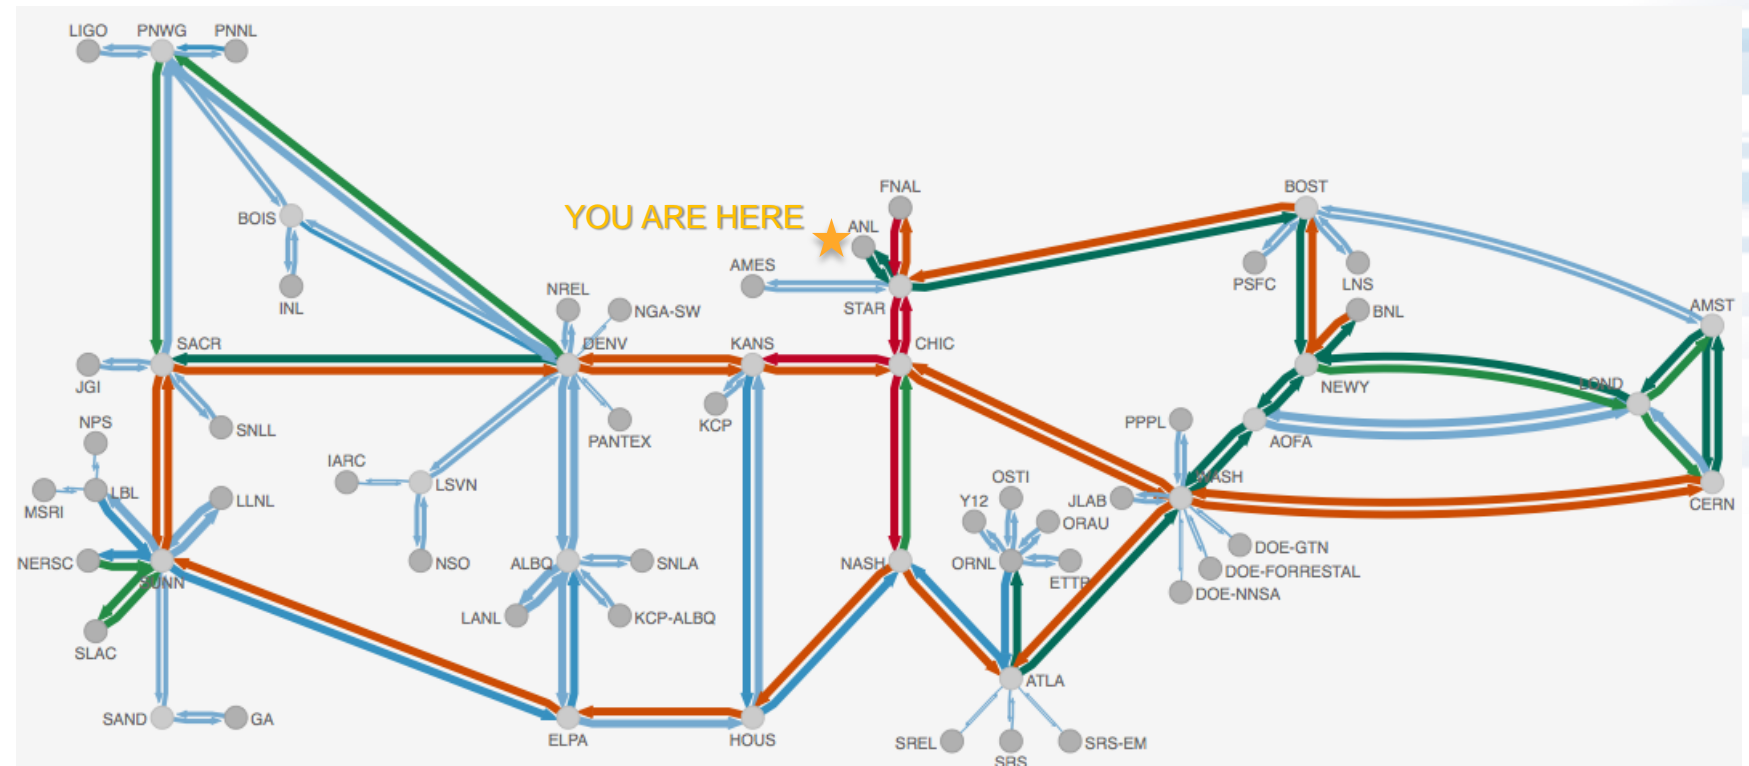
Transferring data sets: it's also a big job

- Data sets are much more valuable if you can share them for collaboration or to leverage resources at other facilities
- The DOE's Energy Sciences Network (ESNet) is here to help with dedicated links to a wide variety of facilities.



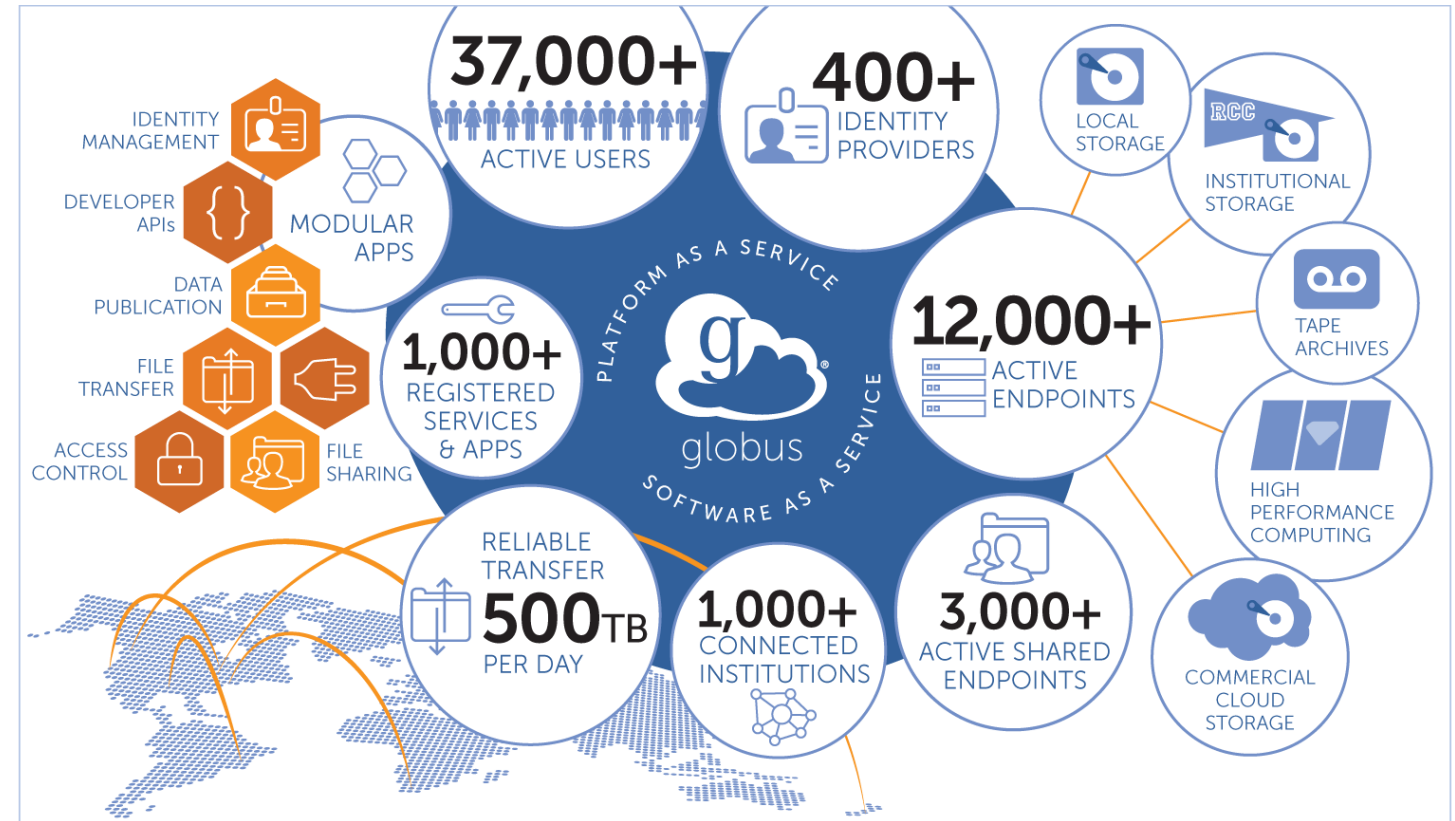
Transferring data sets: it's also a big job

- You can copy files over this network using any tool you like, but conventional desktop tools (scp, ftp, etc.) are incredibly inefficient for transferring big data over fast links
- We recommend Globus.



Globus overview

- Globus is a data management service that can transfer files, share files, publish files, etc.
- Available at most HPC facilities
- You can use it as a web service or through scriptable APIs
- In a nutshell: sets up efficient 3rd party transfers between sites
 - Restarts transfers if interrupted
 - Validates data
 - Spans authentication systems



Using Globus

- <https://www.alcf.anl.gov/user-guides/using-globus>
- <http://www.nersc.gov/users/storage-and-file-systems/transferring-data/globus-online/>
- There are several ways to use Globus; the easiest way to learn is to log in to <https://www.globus.org> and try the web interface

Logging in to Globus

globus Globus Account Log In

Log in to use Globus Web App

Use your existing organizational login
e.g., university, national lab, facility, project

Argonne LCF

Didn't find your organization? Then use [Globus ID to sign in.](#) ([What's this?](#))

Continue

Or

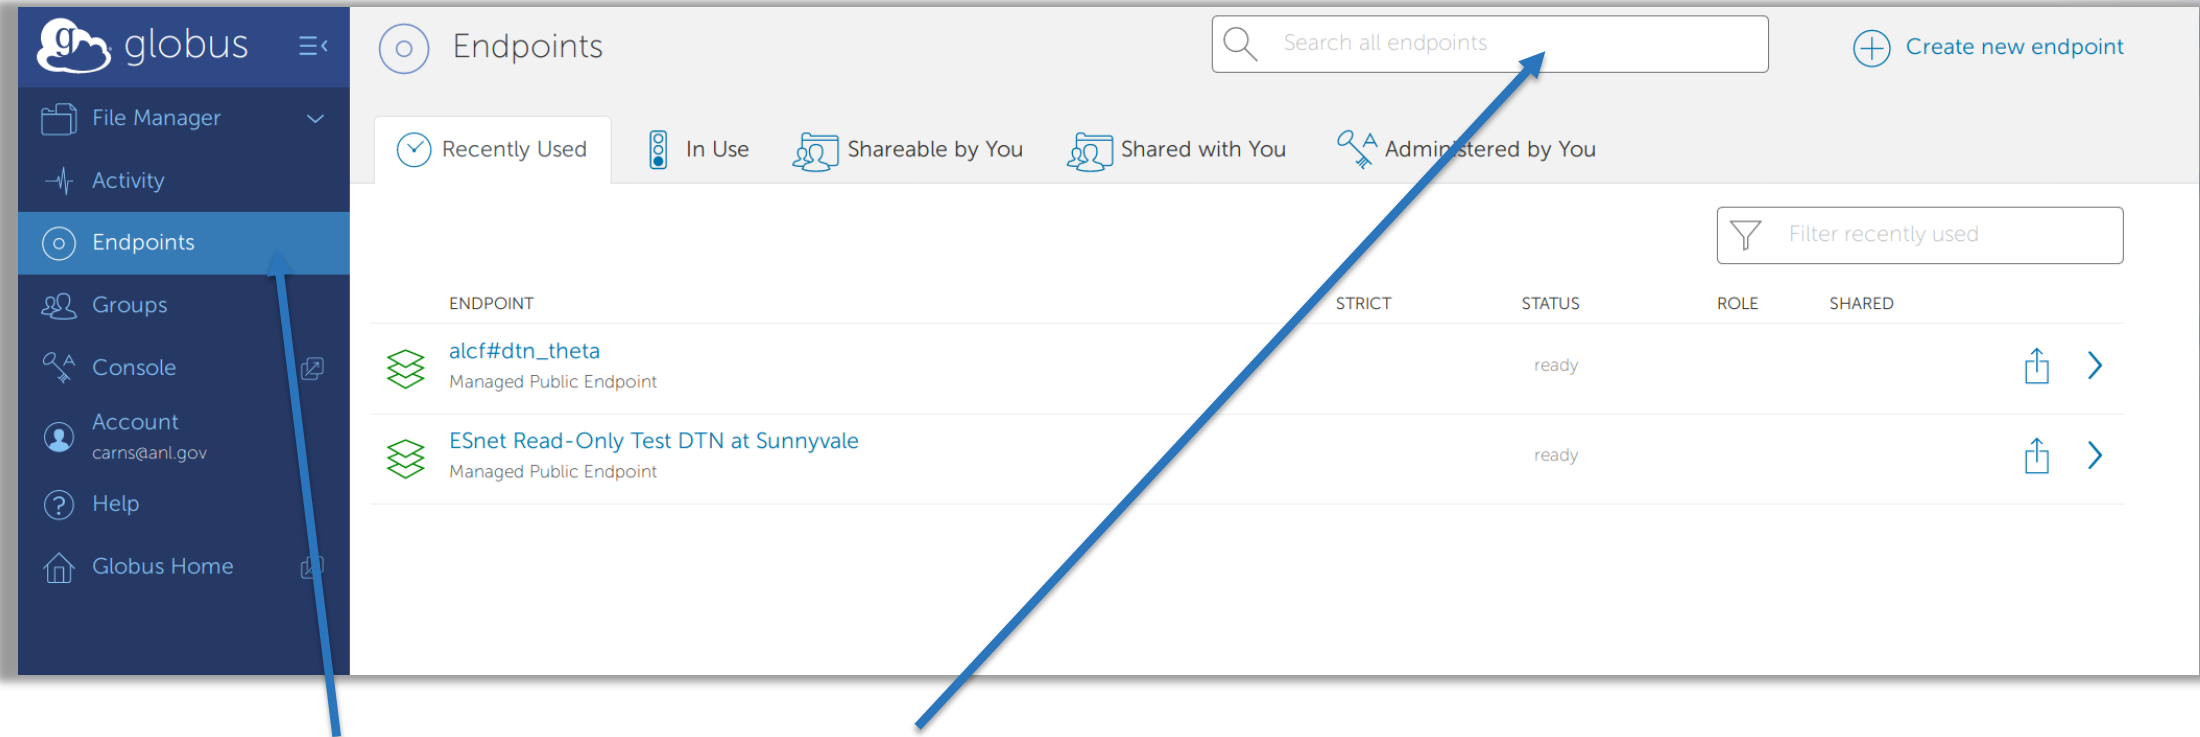
Sign in with Google Sign in with ORCID iD

You can log into Globus using a variety of credentials:







- Facility accounts (like the ALCF, using a token)
- Google account
- ORCID

You can link multiple accounts- this is important for transferring data between sites (like the ALCF and NERSC) that don't share a common account.

Finding endpoints



The screenshot shows the Globus Endpoints interface. The left sidebar is dark blue with the Globus logo and navigation options: File Manager, Activity, Endpoints (highlighted), Groups, Console, Account (carns@anl.gov), Help, and Globus Home. The main content area is titled 'Endpoints' and features a search bar 'Search all endpoints' and a 'Create new endpoint' button. Below the search bar are filter tabs: 'Recently Used' (selected), 'In Use', 'Shareable by You', 'Shared with You', and 'Administered by You'. A 'Filter recently used' dropdown is also present. The main area displays a table of endpoints:

ENDPOINT	STRICT	STATUS	ROLE	SHARED
 alcf#dtn_theta Managed Public Endpoint		ready		 
 ESnet Read-Only Test DTN at Sunnyvale Managed Public Endpoint		ready		 

Two blue arrows are overlaid on the image: one points from the 'Endpoints' tab in the sidebar to the main content area, and the other points from the search bar to the 'Endpoints' tab in the sidebar.

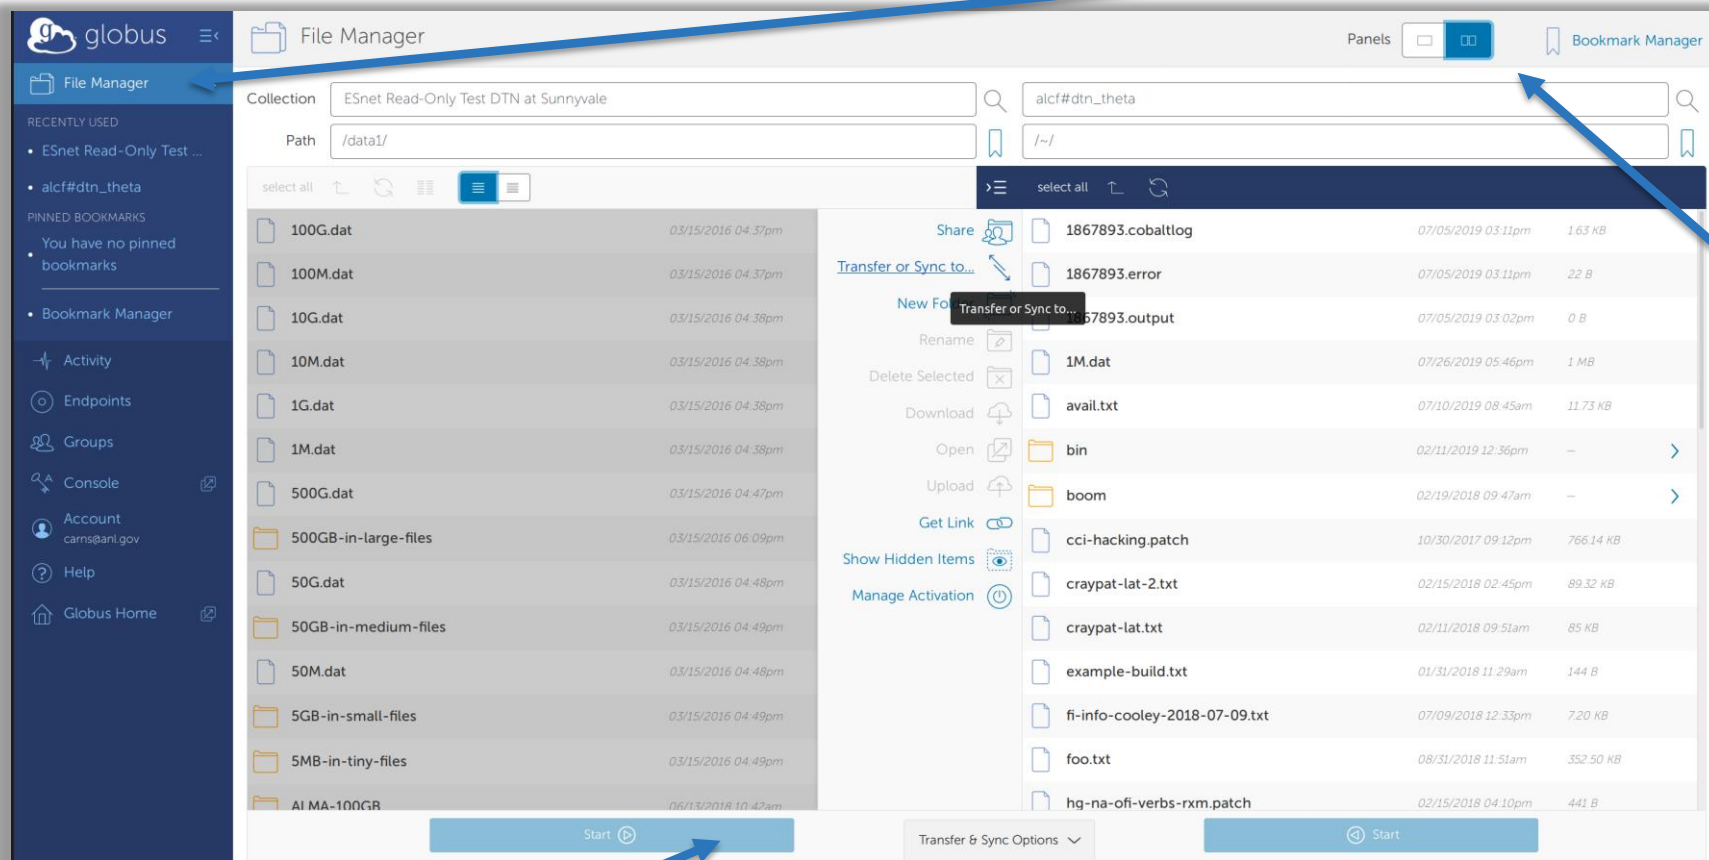
Use the “endpoints” tab to search for data repositories. In this case:

- **alcf#dtn_theta**: transfer data to and from Theta’s primary file system
- **Esnet Read-Only Test DTN**: available to anyone for test transfers

Transferring files

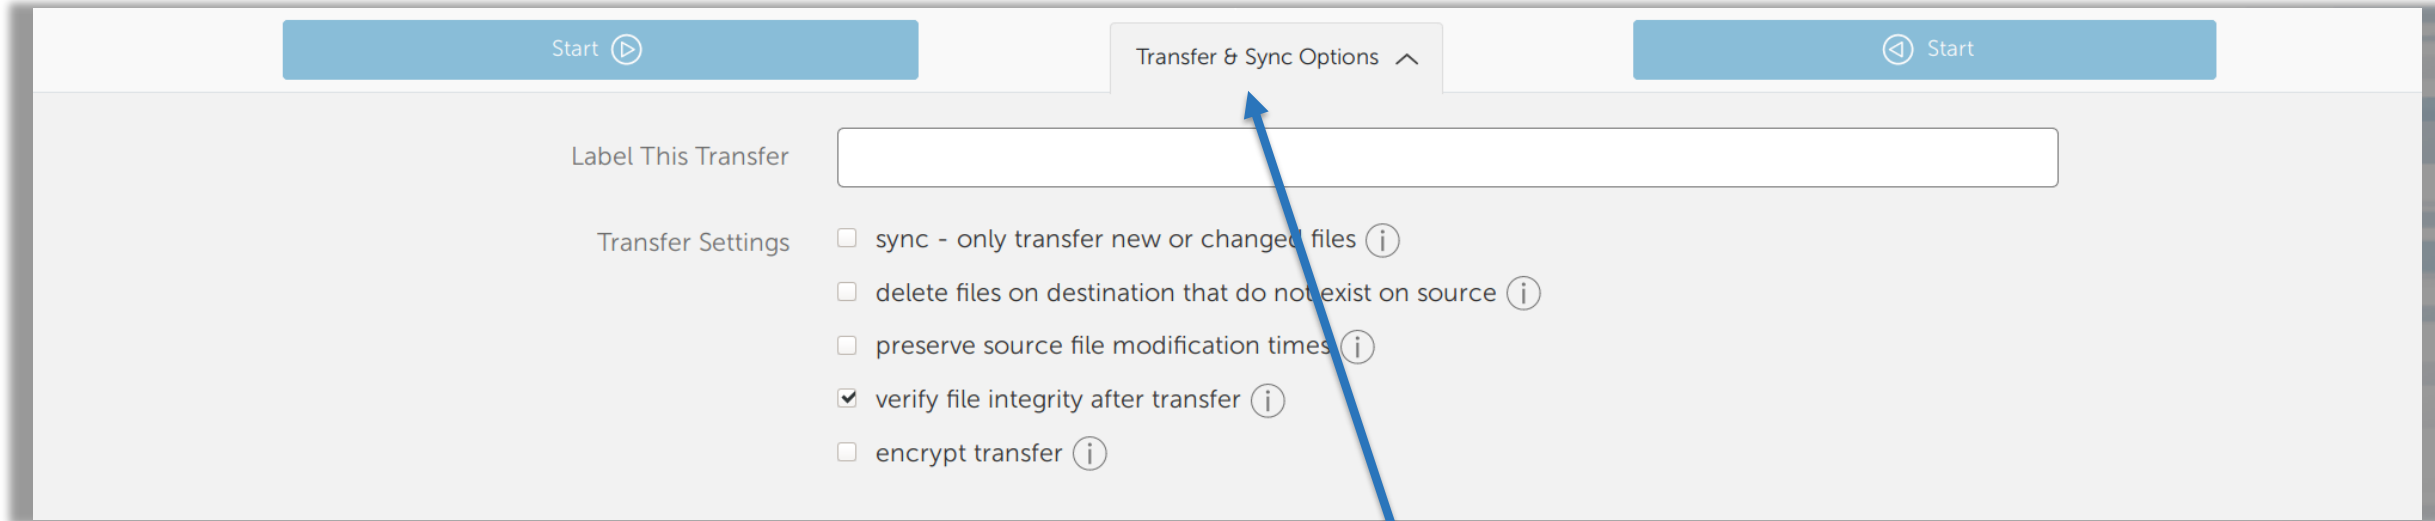
The “file manager” tab can be used to browse files and move them between endpoints.

Use the “panels” button to toggle to a view that shows 2 endpoints at once.



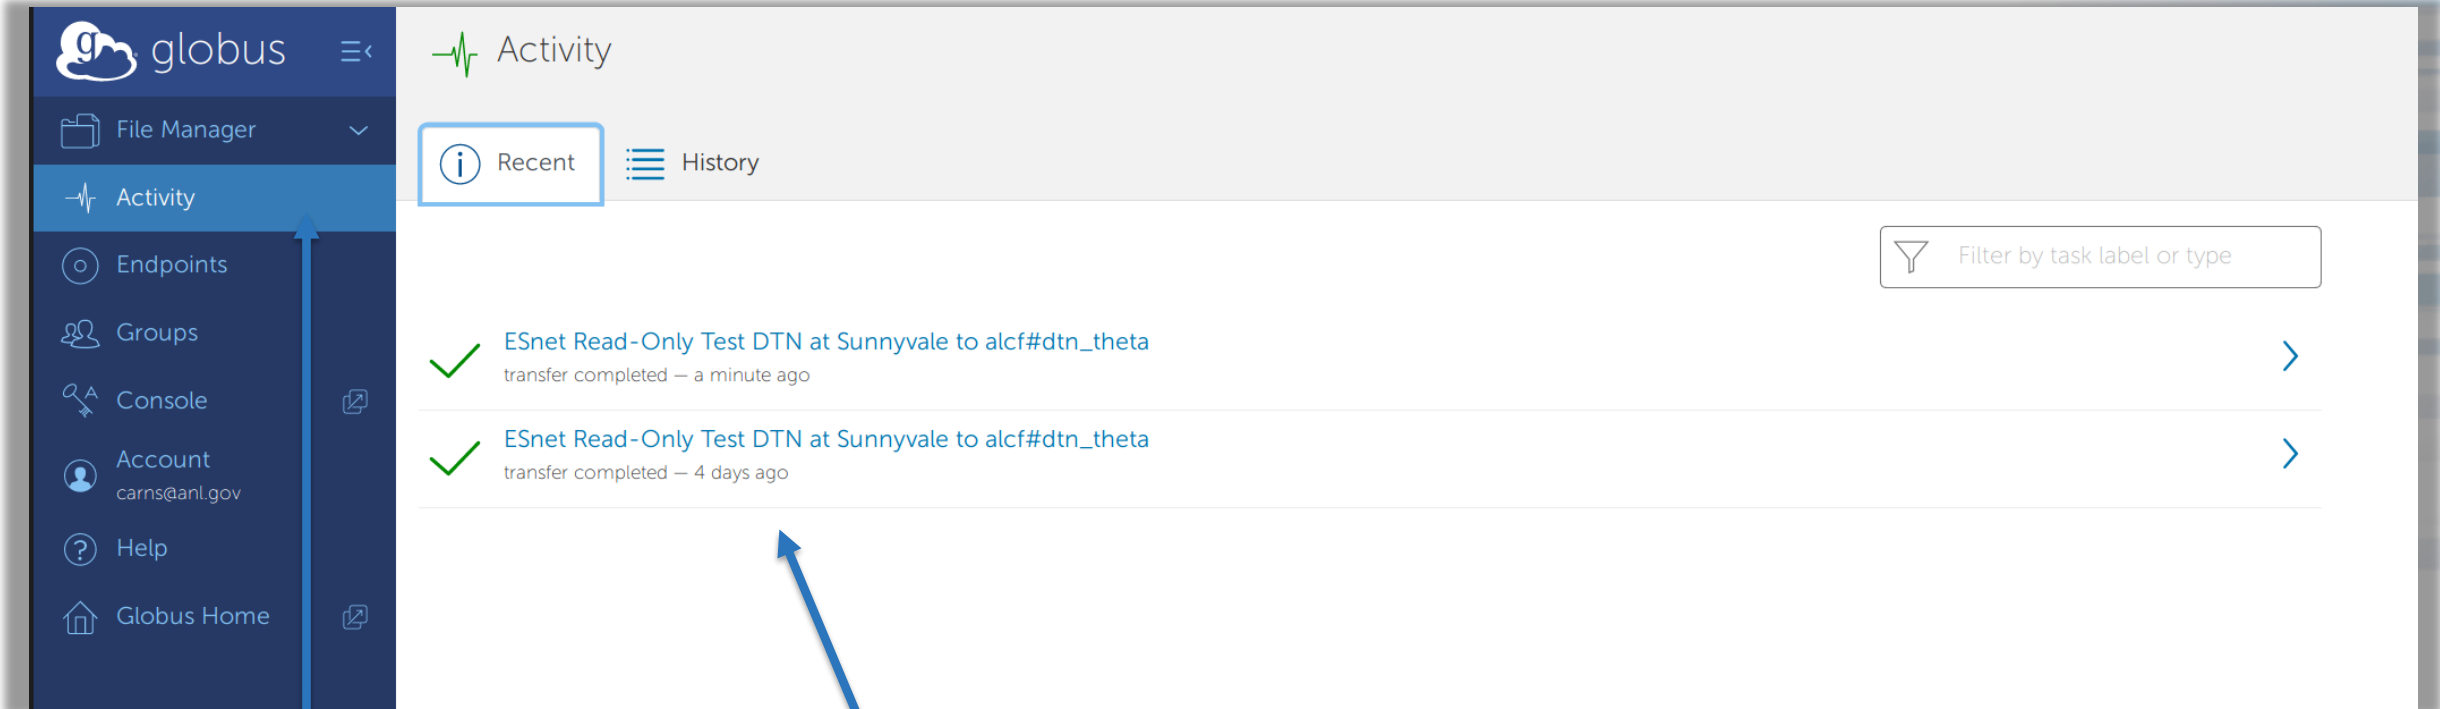
Highlight a set of files on one panel and click the “Start” button to move them to the other endpoint.

File transfer options



Clicking on the “Options” tab at the bottom of the screen pulls up a set of options that control the behavior of the file transfer.

Transfer notifications



The “Activity” tab displays a list of in-progress or completed transfers.

Globus example

- There are several options for transfer notifications, email by default
- No need to stay logged into the web site, your transfer will continue on it's own!
- You can also install a transfer agent on your laptop to move files to and from your laptop.
- You can also publish data sets so that they are persistently available to collaborators or the public.

Globus example

Hands on exercise: <https://xgitlab.cels.anl.gov/ATPESC-IO/hands-on#globus>

- Instructions are in the README.md file in the hands-on repository
- Can you transfer the example file and decode the message?

Agenda (revisited)

- Morning:
 - Introductory concepts and tools ✓
 - MPI-IO and PnetCDF
- Afternoon
 - HDF5
 - Architectures and tuning
- Evening
 - Hands-on exercises



Building up more detail
as the day goes on

ATPESC attendees have access to a dedicated reservation on Theta throughout the day today. See the link at the top of each slide for details.



Thank you!

