



HPC I/O Data Management Tools: So you have some data. Now what?

ATPESC 2018

Phil Carns
Mathematics and Computer Science Division
Argonne National Laboratory


Q Center, St. Charles, IL (USA)
July 29 – August 10, 2018

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 spent 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

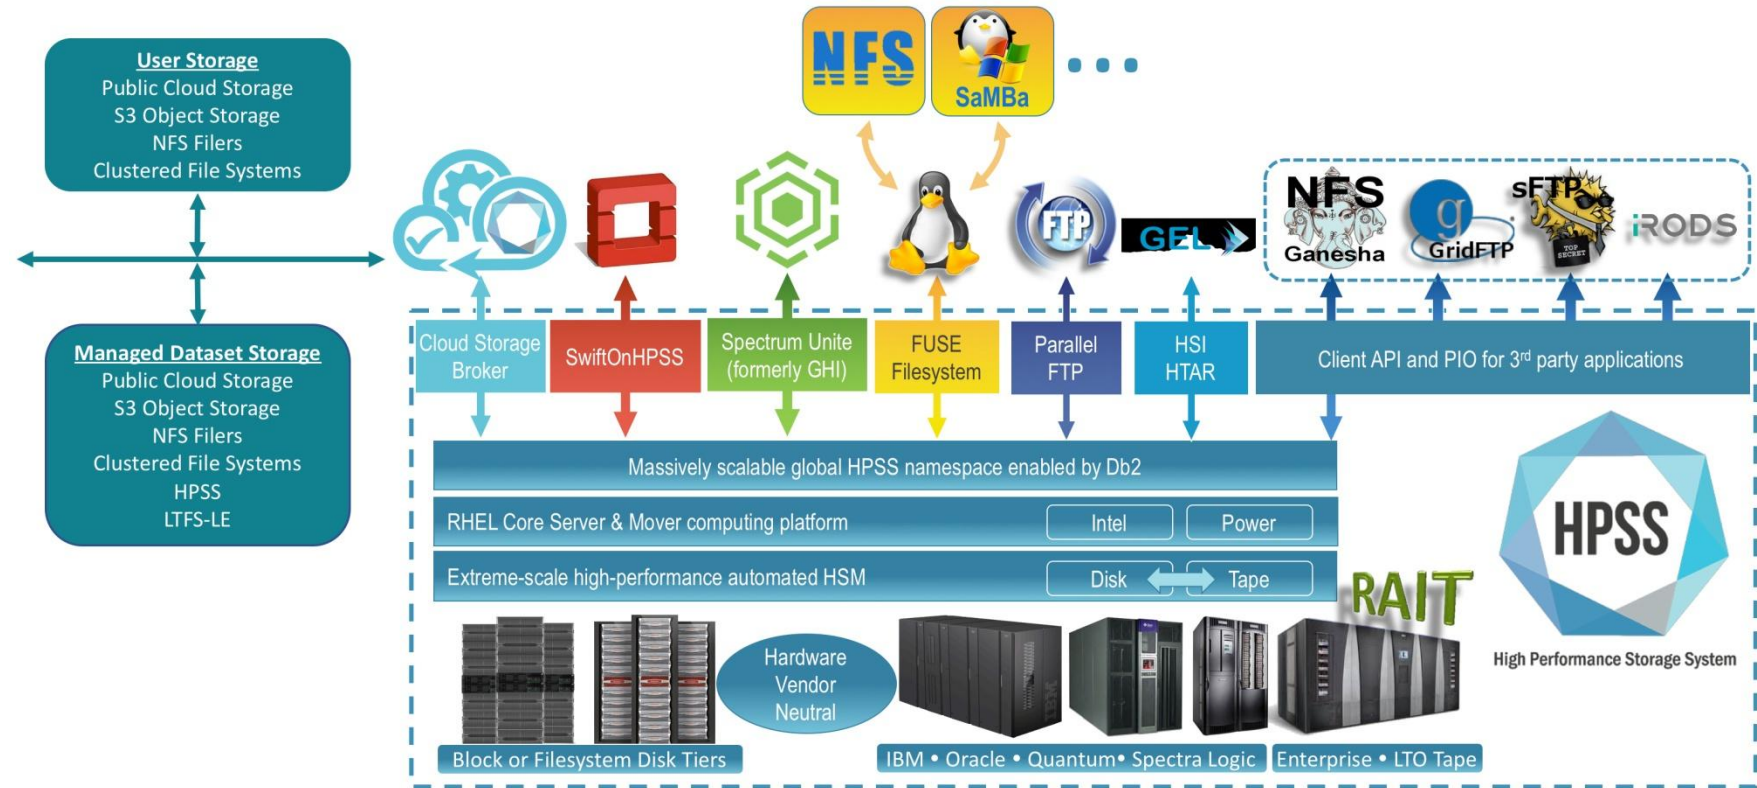
Archiving data: it's a big job

- 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

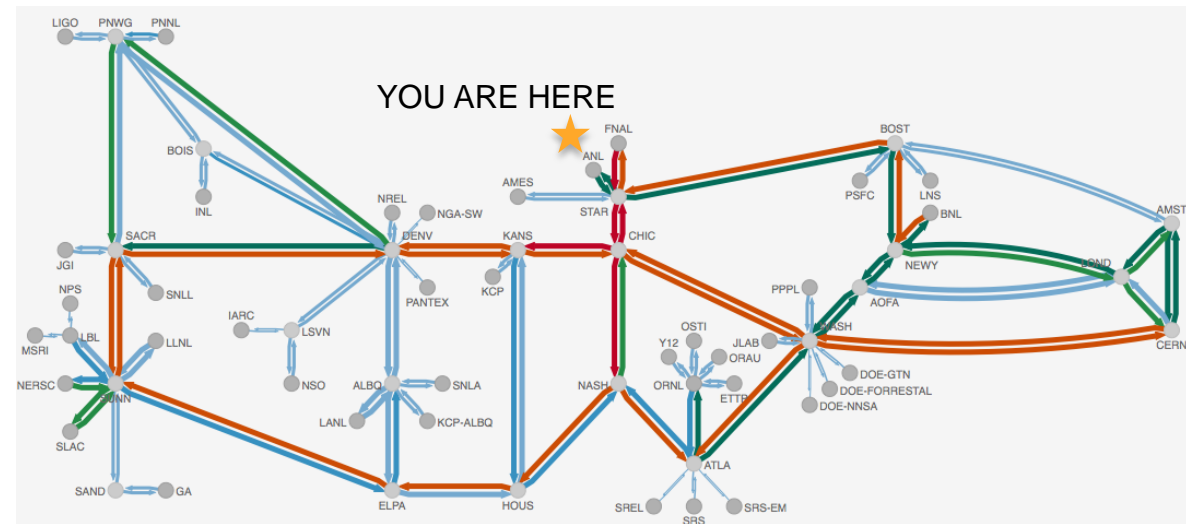
Saving a file called notes.txt in HPSS

```
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:~>
```

- That's 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!*

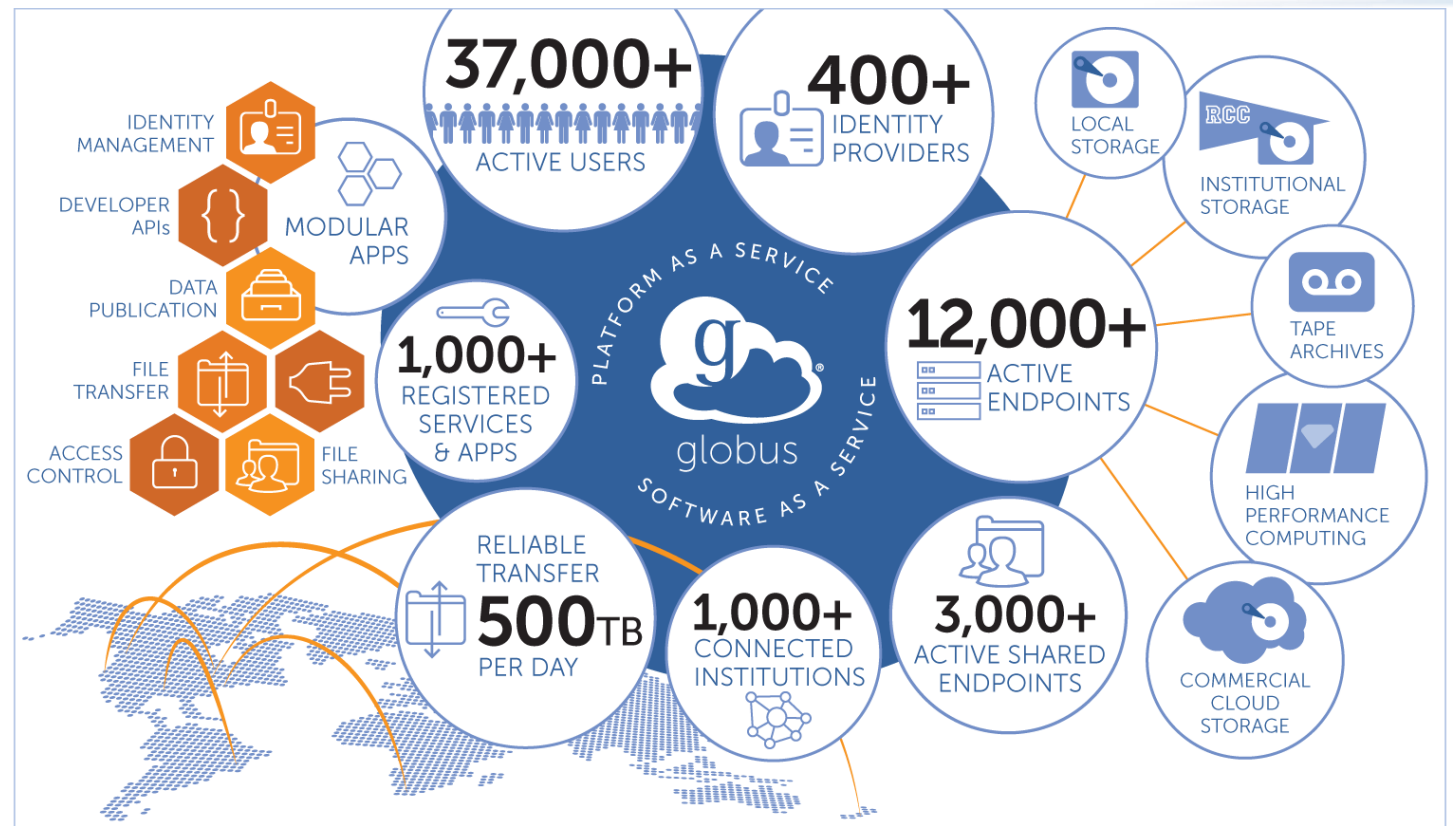
Transferring data sets: it's also a big job

- If data falls in a forest when no one is there to hear it...
- Data sets are much more valuable if you can share them for collaboration or to leverage resources at other facilities
- The DOE's Earth Sciences Network (ESNet) is here to help, with dedicated links to a wide variety of facilities
- 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



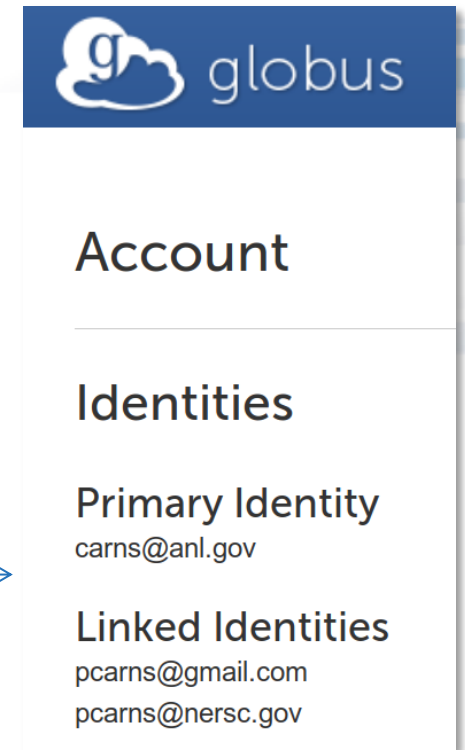
Globus overview

- Globus is a data management service that can transfer files, share files, publish files, etc.
- Supported by a non-profit business from Uchicago
- Available at most HPC facilities
- You can use it as a web service or through scriptable APIs
- In a nutshell: sets up 3rd party transfers between two sites
 - Restarts transfers if interrupted
 - Validates data
 - Uses GridFTP for fast, parallel data transfers
 - 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
- Use an organization account (uses cryptocard if needed for places like the ALCF), Google account, or ORCiD
- Important: you can link multiple accounts! →
- Enables you to transfer files across systems even though they have different usernames and accounts



Globus example

The screenshot shows the Globus web interface for managing data transfers. At the top, there's a navigation bar with 'Manage Data', 'Publish', 'Groups', 'Support', and 'Account'. Below that, there are links for 'Transfer Files', 'Activity', 'Endpoints', 'Bookmarks', and 'Console'. The main area is titled 'Transfer Files' and features a 'Get Globus Connect Personal' button and 'RECENT ACTIVITY' indicators. Two endpoint panels are visible: the left one is for 'alcf#dtn_theta' and the right one is for 'NERSC DTN'. Both panels show a file browser with folders and files, and a 'Go' button. Below the file browsers, there's a 'Label This Transfer' field and 'Transfer Settings' with checkboxes for 'sync - only transfer new or changed files', 'delete files on destination that do not exist on source', 'preserve source file modification times', 'verify file integrity after transfer' (checked), and 'encrypt transfer'.

- The most fundamental thing you can do with Globus is transfer files across sites
- Web GUI lets you open two “endpoints” and view them side by side
- In this case: ALCF and NERSC directories
- See web sites for Globus endpoint names
- Use right and left arrows to transfer files back and forth

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.
- You can also publish data sets so that they are persistently available to collaborators or the public.

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?

Next up

- This presentation gave a brief overview of **Data Management Tools**.
- After a break for dinner we will reconvene for **Hands-On Exercises**.

Thank You!