

File Systems/ Data Transfers

George S. Markomanolis

HPC Engineer

Oak Ridge National Laboratory

Summit New User Training

3 June 2020

ORNL is managed by UT-Battelle, LLC for the US Department of Energy




U.S. DEPARTMENT OF
ENERGY

Outline

- File Systems
 - Available file systems and options for archiving
- Data Transfer
 - Transfer your files

Summit and filesystems



NFS

- User home: /ccs/home/\$USER
- Project home: /ccs/proj/[projid]
- **Long-term** storage for your general data under home or related to project under proj
- **Build** your code in /tmp/\$USER it is faster and **install** in /ccs/proj/[projid]
- There is provided a **backup**
- User home and project home are accessible read-only from the Summit compute nodes
- **Not purged**
- **Quota** of 50GB
- User home is user-centric

NFS (cont.)

- Check quota on user home

```
> quota -Qs
```

Disk quotas for user gmarkoma (uid XXXX):

Filesystem	space	quota	limit	grace	files	quota	limit	grace
------------	-------	-------	-------	-------	-------	-------	-------	-------

XXXXXXXXXX:/nccs/home2								
------------------------	--	--	--	--	--	--	--	--

51031M	51200M	51200M						
--------	--------	--------	--	--	--	--	--	--

598K	4295m	4295m						
------	-------	-------	--	--	--	--	--	--

NFS (continue)

- I deleted a file from my NFS, how to recover it?
- Answer: snapshots
 - Go to the .snapshot folder (ls will not show this folder):
 - `cd .snapshot`

```
ls -l
```

```
drwx----- 27 gmarkoma gmarkoma 4096 Nov 21 16:51 daily.2018-11-23_0010
```

```
drwx----- 27 gmarkoma gmarkoma 4096 Nov 21 16:51 daily.2018-11-24_0010
```

```
...
```

HPSS

- Project archive: /hpss/prod/[projid]/
- **Long-term** storage for large amount of general data under related to project under projid.
- **Quota** of 100 TB for project archive. If any of the used files during htar is bigger than 68 GB size, then it will fail, similar if there are more than 1 million files per archive
- **Not purged**
- Tools to be used from terminal: **hsi, htar**

HPSS – Details (I)

```
> hsi ls
```

```
...
```

```
/home/gmarkoma:
```

```
gen110@ stf007@ stf106@
```

```
> hsi
```

```
...
```

```
O:[/home/gmarkoma]:
```

```
O:[/home/gmarkoma]: ls -l
```

```
/home/gmarkoma:
```

```
lrwxrwxrwx 1 root root
```

```
lrwxrwxrwx 1 root root
```

```
lrwxrwxrwx 1 root root
```

```
17 Mar 4 11:03 gen110@ -> /hpss/prod/gen110
```

```
17 Mar 4 11:03 stf007@ -> /hpss/prod/stf007
```

```
17 Mar 4 11:03 stf106@ -> /hpss/prod/stf106
```


HPSS – Details (II)

```
O:[/home/gmarkoma]: cd /hpss/prod/stf007
```

```
O:[/hpss/prod/stf007]: ls
```

```
/hpss/prod/stf007:
```

```
proj-shared/  users/      world-shared/
```

```
O:[/hpss/prod/stf007]: ls proj-shared/gmarkoma/
```

```
proj-shared/gmarkoma/:
```

```
all1_lot/      archive.tar.gz  data/          data.txt       data10/        small/
```

- 100 TB storage for /hpss/prod/stf007
- Maybe you need to create your username folder in world-shared
- Documentation: <https://docs.olcf.ornl.gov/data/archiving.html>

HPSS – Details (III)

```
> hsi put file_list : /hpss/prod/stf007/proj-shared/gmarkoma/file_list
...
put 'file_list' : '/hpss/prod/stf007/proj-shared/gmarkoma/file_list' ( 181700 bytes, 49107.9 KBS
(cos=11))
```

```
> hsi ls -l /hpss/prod/stf007/proj-shared/gmarkoma/
...
-rw-----  1 gmarkoma gmarkoma    181700 Jun  3 00:09 file_list
...
```

HPSS commands

Command	Function
cd	Change current directory
get, mget	Copy one or more HPSS-resident files to local files
cget	Conditional get - get the file only if it doesn't already exist
cp	Copy a file within HPSS
rm mdelete	Remove one or more files from HPSS
ls	List a directory
put, mput	Copy one or more local files to HPSS
cput	Conditional put - copy the file into HPSS unless it is already there
pwd	Print current directory
mv	Rename an HPSS file
mkdir	Create an HPSS directory
rmdir	Delete an HPSS directory

HTAR example: `htar -cvf /hpss/prod/[projid]/users/[userid]/allfiles.tar dir1/*`

HPSS (cont.)

- Check HPSS quota:

```
> showusage -s hpss
```

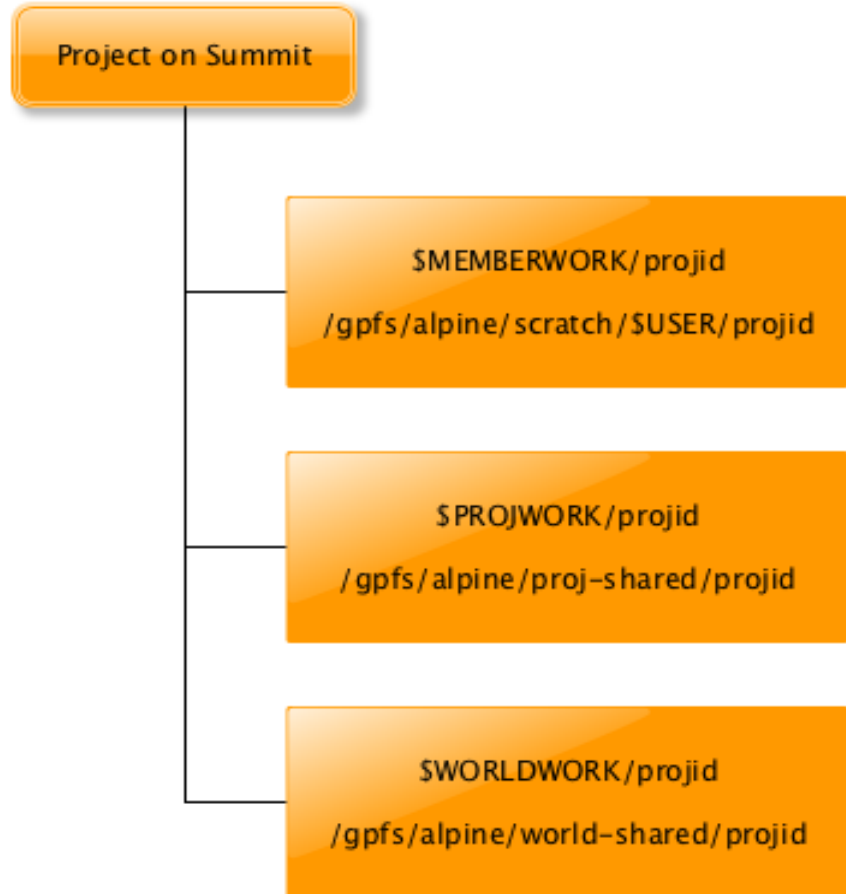
HPSS Storage in GB:

Project	Project Totals Storage	gmarkoma Storage
stf007	47501.84	363.50

Spider III - Alpine

- Alpine, is a Spectrum Scale (ex-GPFS) file system of 250 PB of used space, which is mounted on Summit and Data Transfer Nodes (DTN) with maximum performance of 2.5 TB/s for sequential I/O and 2.2 TB/s for random I/O
- Largest GPFS file system installation
- Up to 2.6 million accesses per second of 32 KB small files
- It is constituted by 154 Network Shared Disk (NSD) servers
- It is a shared resource among users, supporting File Per Process (FPP), Single Shared File (SSF) and any of their combination

Alpine (cont.)



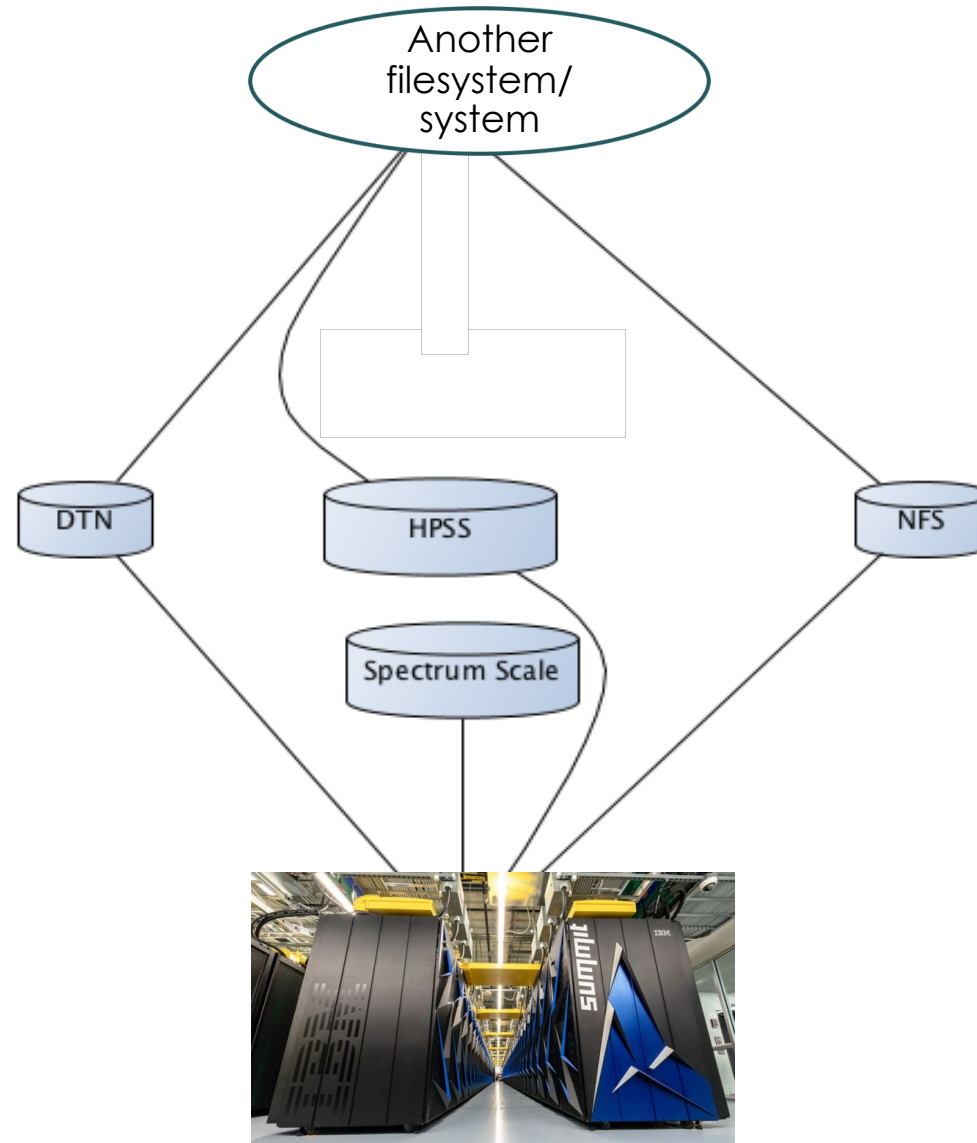
- Memberwork:
 - Short-term storage of user data related to the project but not shared
- Projwork:
 - Short-term storage of project data shared among the members of the project
- Worldwork:
 - Short-term storage of project data shared with OLCF users outside the project
- **No backup**
- **Quota 50 TB**
- **Purged after 90 days if not accessing a file**

Storage policy

Area	Path	Type	Permissions	Quota	Backups	Purged	Retention	On Compute Nodes
User Home	<code>/ccs/home/[userid]</code>	NFS	User set	50 GB	Yes	No	90 days	Read-only
User Archive [1]	<code>/home/[userid]</code>	HPSS	User set	2TB	No	No	90 days	No
User Archive [2]	<code>/home/[userid]</code>	HPSS	700	N/A	N/A	N/A	N/A	No
Project Home	<code>/ccs/proj/[projid]</code>	NFS	770	50 GB	Yes	No	90 days	Read-only
Member Work	<code>/gpfs/alpine/[projid]/scratch/[userid]</code>	Spectrum Scale	700 [3]	50 TB	No	90 days	N/A [4]	Yes
Project Work	<code>/gpfs/alpine/[projid]/proj-shared</code>	Spectrum Scale	770	50 TB	No	90 days	N/A [4]	Yes
World Work	<code>/gpfs/alpine/[projid]/world-shared</code>	Spectrum Scale	775	50 TB	No	90 days	N/A [4]	Yes
Member Archive	<code>/hpss/prod/[projid]/users/\$USER</code>	HPSS	700	100 TB	No	No	90 days	No
Project Archive	<code>/hpss/prod/[projid]/proj-shared</code>	HPSS	770	100 TB	No	No	90 days	No
World Archive	<code>/hpss/prod/[projid]/world-shared</code>	HPSS	775	100 TB	No	No	90 days	No

- Documentation: <https://docs.olcf.ornl.gov/data/policies.html#policy>

Data Transfer



Data Transfer Nodes (DTN) improve the performance by reducing the load on the login and service nodes of the HPC facilities. Moreover, transfer data outside the HPC facility.

Data Transfer (cont. I)

- When you log-in to Summit maybe you would like to transfer some files to start working or even extract your simulation output to other filesystems
- There are many ways to transfer files but in many cases we propose Globus
- We will mention all the approaches and some performance results from some older systems

Data Transfer (cont. II)

- Using home NFS
- If the data size is less than 50 GB and there is enough free space in your home directory you can use tools such as scp, rsync

```
outside> scp -r data username@summit.olcf.ornl.gov:/path/
```

- It is simple, but is it fast? Transferring 100 MB (from home) to Alpine took ~**16** minutes. Rsync was a bit faster

Data Transfer (cont. III)

- Using home NFS

```
outside> scp -r data username@dtn.olcf.ornl.gov:/path/
```

- Transferring 100 MB (from home) to Alpine took **97** seconds
- DTN nodes are dedicated for data transfer

Globus

- Globus transfers fast, parallel and reliable files between two endpoints
- Endpoints are different locations where data can be moved using the Globus transfer
- Visit www.globus.org to register and/or login
- You can find the **OLCF DTN** endpoint.
- If you want to save data on HPSS, you could use the **OLCF HPSS** endpoint.
- Using Globus with HPSS:
<https://docs.olcf.ornl.gov/data/archiving.html#using-globus>
- Using Globus from your local computer:
<https://docs.olcf.ornl.gov/data/transferring.html#using-globus-from-your-local-machine>

Globus(cont.)



The screenshot shows the Globus website homepage. The browser's address bar displays 'globus.org'. The navigation bar includes links for 'I Want To...', 'Pricing', 'Resources', 'Support', 'About', and a blue 'Log In' button, which is highlighted by a red arrow. The main content area features a large red hexagon with a white virus icon on the left. To its right, the text reads 'Doing COVID-19 research?' followed by 'We have a box of goodies that we'd like to share with you!' and a link 'LEARN MORE ABOUT THIS FREE ACCESS OFFER'. On the right side, there is a 3D illustration of an open box labeled 'globus Research Data Management Platform' containing several orange hexagons with icons representing various research and data management functions. The background is a dark blue network of nodes and lines. At the bottom, a dark blue banner contains the text 'Research data management simplified.' and three icons with labels: 'TRANSFER', 'SHARE', and 'BUILD'.

globus
a **uchicago** non-profit service

I Want To... Pricing Resources Support About Log In

Doing COVID-19 research?

We have a box of goodies that we'd like to share with you!

LEARN MORE ABOUT THIS FREE ACCESS OFFER

Research Data Management Platform

Research data management simplified.

TRANSFER SHARE BUILD

Globus(cont.)

Use Globus and select the organization that **YOU** belong

or

Use Globus from personal computer or if your organization not listed above

globus

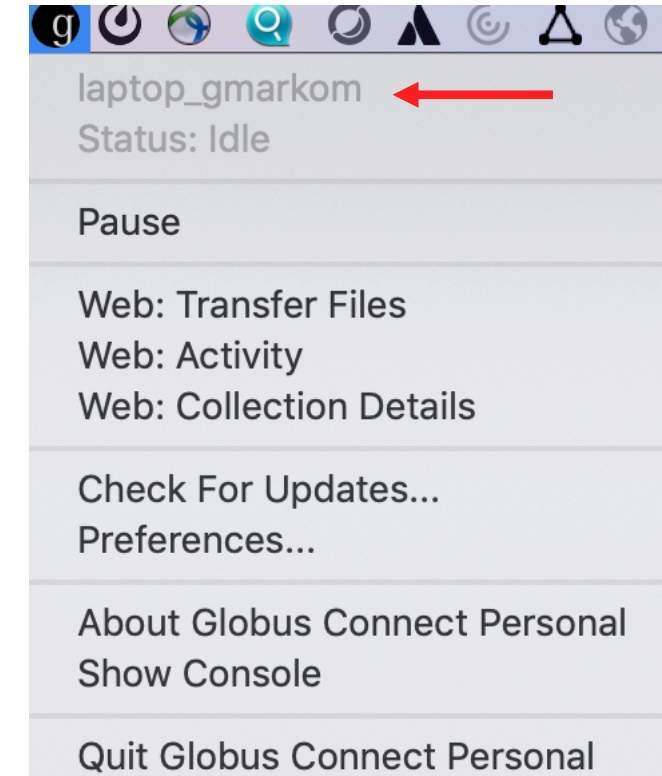
Log in to use Globus Web App

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

Oak Ridge National Laboratory|
Northern Illinois University
Northwestern University
North-West University
NTB - Hochschule für Technik Buchs
NTU - Nanyang Technological University
Nursing School of Coimbra
NUS - National University of Singapore
NYU Langone Health
Oak Ridge National Laboratory


Or



Sign in with Google Sign in with ORCID ID







Personal Globus does not work with the ORNL VPN (cisco)



Globus(cont. personal computer)

 File Manager


Panels  


Collection  


Path  














Bookmark  Bookmark 


select all


 up one folder


 refresh list


 view


NAME 	LAST MODIFIED	SIZE	
 data	06/02/2020 05:56pm	—	
 data_10_files	06/02/2020 06:04pm	—	
 data_2_files	06/02/2020 06:08pm	—	
 data_4_files	06/02/2020 06:08pm	—	
 data.txt	06/02/2020 02:18pm	1.04 GB	
 test.img	06/02/2020 02:23pm	1.04 GB	
 test2.img	06/02/2020 02:23pm	104.85 MB	
 test3.img	06/02/2020 06:08pm	52.42 MB	


 Share


 Transfer or Sync to...


 New Folder


 Rename


 Delete Selected


 Download

 Open

 Upload

 Get Link

 Show Hidden Items

 Manage Activation

Globus – Personal computer and Alpine

The screenshot displays the Globus File Manager interface. On the left is a dark blue sidebar with navigation icons for File Manager, Bookmarks, Activity, Endpoints, Groups, Console, Account, Logout, and Help. The main area is titled 'File Manager' and shows a 'Collection' of 'laptop_gmarkom' and a 'Path' of '/~/Downloads/last/'. A table lists files and folders with columns for Name, Last Modified, and Size. A right-hand menu offers actions like Share, Transfer or Sync to..., New Folder, Rename, Delete Selected, Download, Open, Upload, Get Link, Show Hidden Items, and Manage Activation. At the bottom, there are 'Start' buttons and a 'Transfer & Sync Options' dropdown. Red arrows highlight the 'Panels' toggle, the search bars, the path input, the 'Start' button, and the 'Transfer & Sync Options' dropdown.

File Manager

Collection: laptop_gmarkom

Path: /~/Downloads/last/

NAME	LAST MODIFIED	SIZE
data	06/02/2020 05:56pm	—
data_10_files	06/02/2020 06:04pm	—
data_2_files	06/02/2020 06:08pm	—
data_4_files	06/02/2020 06:08pm	—
data.txt	06/02/2020 02:18pm	1.04 GB
test.img	06/02/2020 02:23pm	1.04 GB
test2.img	06/02/2020 02:23pm	104.85 MB
test3.img	06/02/2020 06:08pm	52.42 MB

Start

Transfer & Sync Options

Start

Globus - Activity

FILE MANAGER

BOOKMARKS

ACTIVITY

ENDPOINTS

GROUPS

CONSOLE

ACCOUNT

LOGOUT

?

Activity List

laptop_gmarkom to OLCF DTN
transfer completed

Overview

Event Log

Task Label

Source

Destination

Task ID

Owner

Condition

Requested

Completed

Transfer Settings

laptop_gmarkom to OLCF DTN

laptop_gmarkom ⓘ

OLCF DTN ⓘ

6f49ff2c-a52a-11ea-8f07-0a21f750d19b

George Markomanolis (george@markomanolis.com@accounts.google.com) ⓘ

SUCCEEDED

2020-06-02 07:40 pm

2020-06-02 07:41 pm

- verify file integrity after transfer
- transfer is not encrypted
- overwriting all files on destination

1 Files

1 Directories

104.85 MB Bytes Transferred

1.39 MB/s Effective Speed

0 Skipped

[View debug data](#)

Globus – From Alpine to HPSS (1 file)

The screenshot displays the Globus File Manager interface with two panels side-by-side. The left panel is titled 'OLCF HPSS' and shows a file list with columns for NAME, LAST MODIFIED, and SIZE. The right panel is titled 'OLCF DTN' and shows a single folder named 'data'. A red arrow points to the 'Start' button in the bottom right corner of the interface.

OLCF HPSS Panel:

NAME	LAST MODIFIED	SIZE
all1_lot	03/07/2019 03:20pm	—
archive.tar.gz	09/12/2019 09:20am	3.01 MB
data.txt	09/12/2019 09:55am	22.02 GB
data10	03/06/2019 03:32pm	—
small	02/28/2019 04:36pm	—
transfer	06/02/2020 06:28pm	—

OLCF DTN Panel:

NAME	LAST MODIFIED	SIZE
data	06/02/2020 07:40p...	—

Globus - Activity

FILE MANAGER

BOOKMARKS

ACTIVITY

ENDPOINTS

GROUPS

CONSOLE

ACCOUNT

LOGOUT

HELP

Activity List

OLCF DTN to OLCF HPSS

transfer completed

Overview

Event Log

Task Label

Source

Destination

Task ID

Owner

Condition

Requested

Completed

Transfer Settings

OLCF DTN to OLCF HPSS

OLCF DTN ⓘ

OLCF HPSS ⓘ

c83f9ad4-a52e-11ea-9a3a-0255d23c44ef

George Markomanolis (george@markomanolis.com@accounts.google.com)

SUCCEEDED

2020-06-02 08:11 pm

2020-06-02 08:11 pm

- verify file integrity after transfer
- transfer is not encrypted
- overwriting all files on destination

1 Files

1 Directories

104.85 MB Bytes Transferred

33.16 MB/s Effective Speed

0 Skipped

[View debug data](#)

Globus - From Alpine to HPSS (1 larger file)

The screenshot displays the Globus File Manager interface with two panels side-by-side. The left panel is titled 'OLCF HPSS' and shows a collection of files with the path '/~/stf007/proj-shared/gmarkoma/transfer/'. It contains two folders: 'data' (last modified 06/02/2020 06:28pm) and 'data_10_files' (last modified 06/02/2020 08:51pm). The right panel is titled 'OLCF DTN' and shows a collection of files with the path '/gpfs/alpine/stf007/scratch/gmarkoma/transfer/data1/'. It contains a single file 'data.txt' (last modified 06/02/2020 08:58p..., size 23.06 GB), which is highlighted with a red arrow. A red arrow also points to the 'Start' button in the bottom right corner of the interface.

NAME	LAST MODIFIED	SIZE
data	06/02/2020 06:28pm	—
data_10_files	06/02/2020 08:51pm	—

NAME	LAST MODIFIED	SIZE
data.txt	06/02/2020 08:58p...	23.06 GB

Globus - Activity

FILE MANAGER

BOOKMARKS

ACTIVITY

ENDPOINTS

GROUPS

CONSOLE

ACCOUNT

LOGOUT

HELP

Activity List

OLCF DTN to OLCF HPSS

transfer completed

Overview

Event Log

Task Label

Source

Destination

Task ID

Owner

Condition

Requested

Completed

Transfer Settings

OLCF DTN to OLCF HPSS

OLCF DTN

OLCF HPSS

92aa8b02-a535-11ea-bee4-0e716405a293

George Markomanolis (george@markomanolis.com@accounts.google.com)

SUCCEEDED

2020-06-02 09:00 pm

2020-06-02 09:01 pm

- verify file integrity after transfer
- transfer is not encrypted
- overwriting all files on destination

1

Files

0

Directories

23.06 GB

Bytes Transferred

272.54 MB/s

Effective Speed

0

Skipped

View debug data

Globus - From Alpine to HPSS (22 files)

FILE MANAGER

BOOKMARKS

ACTIVITY

ENDPOINTS

GROUPS

CONSOLE

ACCOUNT

LOGOUT

HELP

File Manager

Panels

Collection

Path

select all

view

NAME	LAST MODIFIED	SIZE
data	06/02/2020 06:28pm	—
data_10_files	06/02/2020 08:51pm	—
data.txt	06/02/2020 09:00pm	23.06 GB

Start

Transfer & Sync Options

Start

OLCF DTN

Path

select none

view

NAME	LAST MODIFIED	SIZE
data1	06/02/2020 08:58p...	—
data22	06/02/2020 08:58p...	—

Start

Globus - Activity

g

FILE MANAGER

BOOKMARKS

ACTIVITY

ENDPOINTS

GROUPS

CONSOLE

ACCOUNT

LOGOUT

HELP

Activity List

OLCF DTN to OLCF HPSS
transfer completed

Overview

Event Log

Task Label

Source

Destination

Task ID

Owner

Condition

Requested

Completed

Transfer Settings

OLCF DTN to OLCF HPSS

OLCF DTN

OLCF HPSS

e52096e2-a535-11ea-bee4-0e716405a293

George Markomanolis (george@markomanolis.com@accounts.google.com)

SUCCEEDED

2020-06-02 09:02 pm

2020-06-02 09:02 pm

- verify file integrity after transfer
- transfer is not encrypted
- overwriting all files on destination

22

Files

1

Directories

23.06 GB

Bytes Transferred

1.43 GB/s

Effective Speed

0

Skipped

View debug data

Globus – From outside system to HPSS

File Manager

Collection: **laptop_gmarkom** Path: **/~/Downloads/last/**

Collection: **OLCF HPSS** Path: **/~/stf007/proj-shared/gmarkoma/transfer/**

NAME	LAST MODIFIED	SIZE
data	06/02/2020 05:56pm	—
data_10_files	06/02/2020 06:04pm	—
data_2_files	06/02/2020 06:08pm	—
data_4_files	06/02/2020 06:08pm	—
data.txt	06/02/2020 02:18pm	1.04 GB
test.img	06/02/2020 02:23pm	1.04 GB
test2.img	06/02/2020 02:23pm	104.85 MB
test3.img	06/02/2020 06:08pm	52.42 MB

NAME	LAST MODIFIED	SIZE
data_10_files	06/02/2020 08:51pm	—
data.txt	06/02/2020 09:00p...	23.06 GB
data22	06/02/2020 09:02p...	—

Start

Transfer & Sync Options

Start

Globus - Activity

FILE MANAGER

BOOKMARKS

ACTIVITY

ENDPOINTS

GROUPS

CONSOLE

ACCOUNT

LOGOUT

HELP

Activity List

laptop_gmarkom to OLCF HPSS

transfer completed

Overview

Event Log

Task Label

Source

Destination

Task ID

Owner

Condition

Requested

Completed

Transfer Settings

laptop_gmarkom to OLCF HPSS

laptop_gmarkom ⓘ

OLCF HPSS ⓘ

b890ca84-a59e-11ea-bee4-0e716405a293

George Markomanolis (george@markomanolis.com@accounts.google.com)

SUCCEEDED

2020-06-03 09:32 am

2020-06-03 09:34 am

- verify file integrity after transfer
- transfer is not encrypted
- overwriting all files on destination

1 Files

1 Directories

104.85 MB Bytes Transferred

1.32 MB/s Effective Speed

0 Skipped

View debug data

Performance Results

- Study case: Transfer data from an older filesystem Atlas to Alpine with 3 approaches. Copy the files through NFS, HPSS, or Globus

Type	Home NFS*	HPSS*	Globus*	HPSS**	HPSS** from DTN
	Time in seconds to finish the transfer				
Transfer 22 files of 1 GB each	323	270	10	227	20
Transfer 1 file of 22 GB	308	301	80	345	26
Transfer 4 files of 1 GB each	69	53	9	39.2	4.6

- Globus is the most efficient approach to transfer files when you have many of them, however, HPSS tools from DTN, can be more efficient in some cases. There are available some traditional tools such as scp, rsync

* 29 November 2018 ** 16 May 2019

Conclusions – File systems/Data transfer

- Use NFS for installing your libraries (long-term storage)
- There are many approaches to transfer files, some tools are more efficient depending on the number of files and the file size.
- Use HPSS for large files that you don't plan to use soon and to backup soon to expire projects with important data
- Do not forget the storage policy!

Thank you!
Questions?