

# Storage Areas / Data Transfers

George S. Markomanolis,  
HPC Engineer  
Oak Ridge National Laboratory  
Summit Training Workshop  
3 December 2018

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




U.S. DEPARTMENT OF  
**ENERGY**

# Outline

- Storage Areas
  - Available file systems and options for archiving
- Data Transfer
  - Transfer your files between Titan and Summit

# 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 is not accessible from compute nodes
- Project home is accessible only for reading from the 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 14850):

Filesystem	blocks	quota	limit	grace	files	quota	limit	grace
------------	--------	-------	-------	-------	-------	-------	-------	-------

nccs-svm1.lb.ccs.ornl.gov:/nccs/home2								
---------------------------------------	--	--	--	--	--	--	--	--

3237M	51200M	51200M			49161	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 $HOME/.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

- User archive: /home/\$USER
- Project archive: /proj/[projid]
- **Long-term** storage for large amount of general data under home or related to project under proj.
- **Quota** of 2 TB and 100 TB for user and project archive respectively. Maximum 68 GB size per file and max 1 million files per archive
- **Not purged**
- User archive is user-centric

## HPSS (cont.)

- Check HPSS quota:

```
> showusage -s hpss
```

HPSS Storage in GB:

### Project Totals


Project	Storage	Storage
stf007	46868.90	0.00



# 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
- 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 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 150 days**

# Storage policy

<i>Name</i>	<i>Path</i>	<i>Type</i>	<i>Permissions</i>	<i>Backups</i>	<i>Purged</i>	<i>Quota</i>
<i>User Home</i>	<code>\$HOME</code>	NFS	User Set	yes	no	50GB
<i>User Archive</i>	<code>/home/\$USER</code>	HPSS	User Set	no	no	2TB
<i>Project Home</i>	<code>/ccs/proj/[projid]</code>	NFS	770	yes	no	50GB
<i>Member Work</i>	<code>/gpfs/alpine/scratch/[userid]/[projid]/</code>	Spectrum Scale	700	no	150 days	50TB
<i>Project Work</i>	<code>/gpfs/alpine/proj-shared/[projid]</code>	Spectrum Scale	770	no	150 days	50TB
<i>World Work</i>	<code>/gpfs/alpine/world-shared/[projid]</code>	Spectrum Scale	775	no	150 days	50TB
<i>Project Archive</i>	<code>/proj/[projid]</code>	HPSS	770	no	no	100TB

# 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.)

- When you log-in to Summit you would like to have access to your old files (if you are already user of OLCF HPC facilities)
- There are many ways to transfer files but in general we propose Globus
- We will mention all the approaches and some performance results.

# Data Transfer (cont.)

- Using home NFS
- If the data size is less than 50 GB and there is enough free space in your home directory is through home.

```
titan> cp -r data $HOME  
summit> cp -r $HOME/data .
```

- It is simple, but is it fast?

# Data Transfer (cont.)

- Using HPSS
- Send one folder to HPSS and retrieve it from the destination. There is significant higher data size limit

```
titan> htar -cvf transfer_test.tar transfer_test/*
```

```
HTAR: a  transfer_test/data0.txt
```

```
HTAR: a  transfer_test/data10.txt
```

```
...
```

```
HTAR: a  /tmp/HTAR_CF_CHK_8183_1543522594
```

```
HTAR Create complete for transfer_test.tar. 23,068,684,800 bytes  
written for 22 member files, max threads: 3 Transfer time: 186.324  
seconds (123.809 MB/s) wallclock/user/sys: 186.521 30.654 105.275  
seconds
```

```
HTAR: HTAR SUCCESSFUL
```

```
summit> htar -xvf transfer_test.tar
```

# Transferring files through NFS and HPSS

```
titan>ls -l
total 1824016
-rw-r--r-- 1 gmarkona gmarkona 1048576000 Dec 1 09:52 data.txt
titan>htar -cvf small_transfer.tar data.txt
HTAR: a data.txt
HTAR: a /tmp/HTAR_CF_CHK_52054_1543676913
HTAR Create complete for small_transfer.tar. 1,048,578,048 bytes written for 1 member
Files, max threads: 3 Transfer time: 8.409 seconds (123,811 MB/s) wallclock/user/sys
: 8.562 1.308 4.796 seconds
HTAR: HTAR SUCCESSFUL
titan>
```

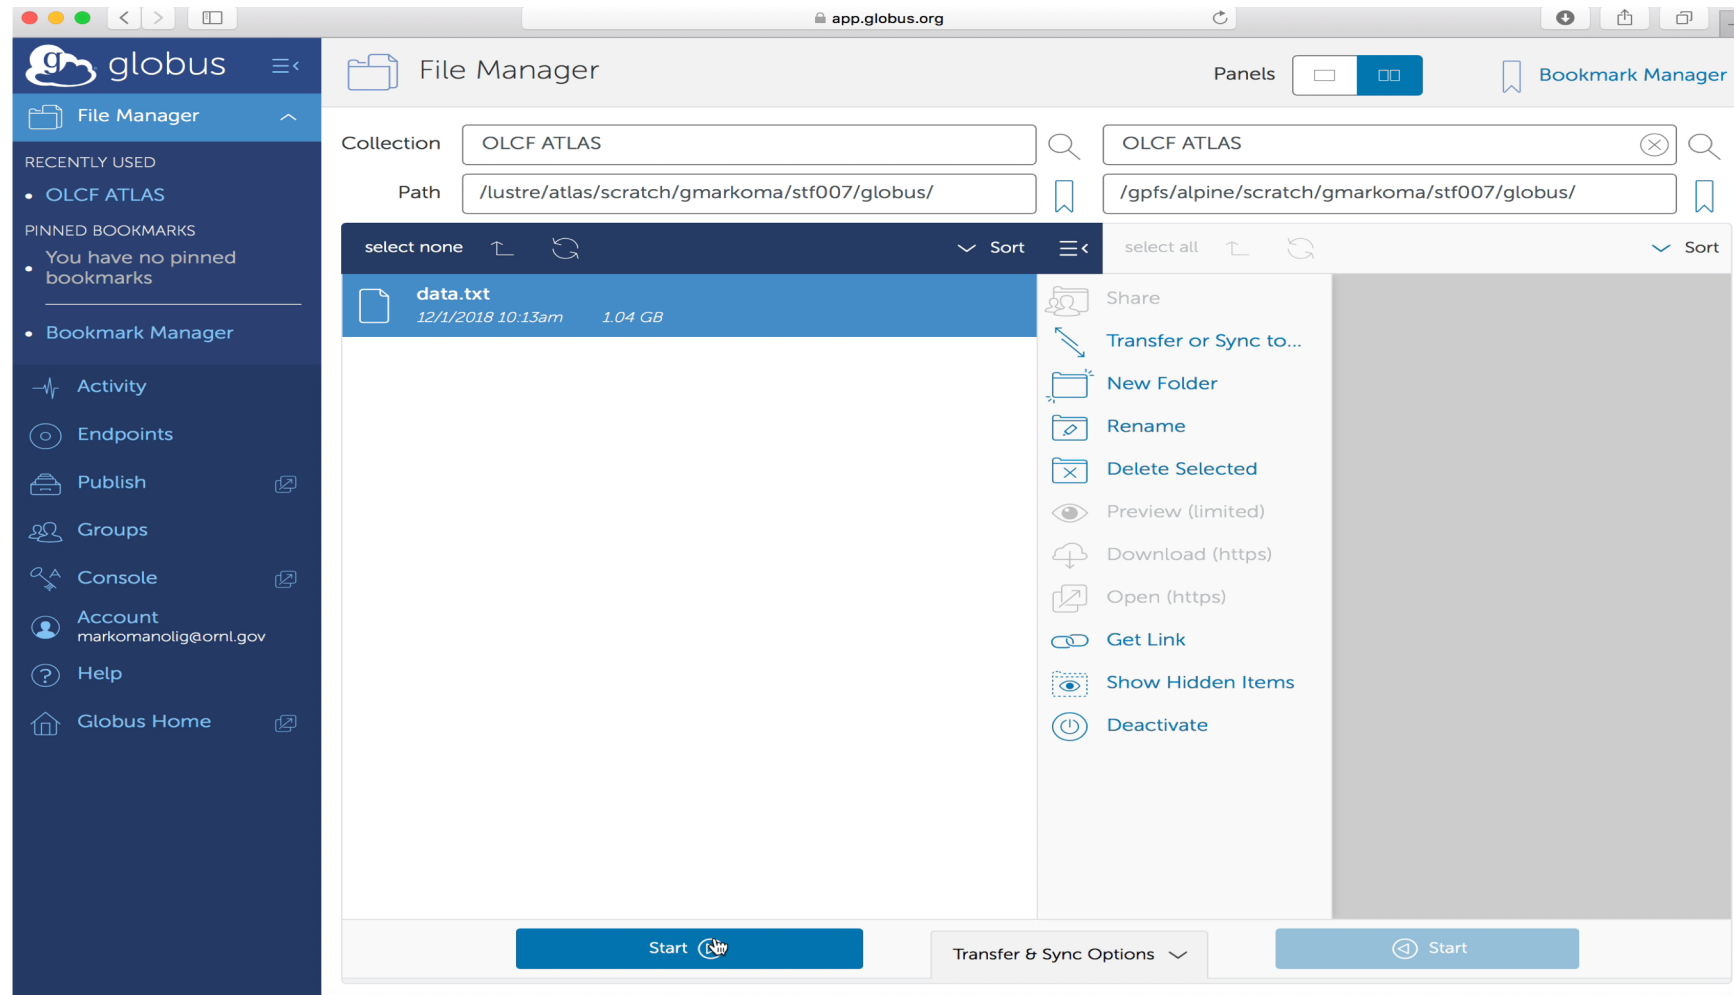
Submit>



# 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](http://www.globus.org) to register and/or login
- You can find few OLCF endpoints such as OLCF Atlas. However, on 11<sup>th</sup> December we'll define the **OLCF DTN** globus endpoint where both Lustre and Spectrum Scale will be mounted and possible to transfer files. You will receive the official announcement soon.

# Globus(cont.)



**The globus endpoint will be OLCF DTN on 11<sup>th</sup> December, it is not available yet!**

# Globus demo, transfer from Titan to Summit



# Performance Results

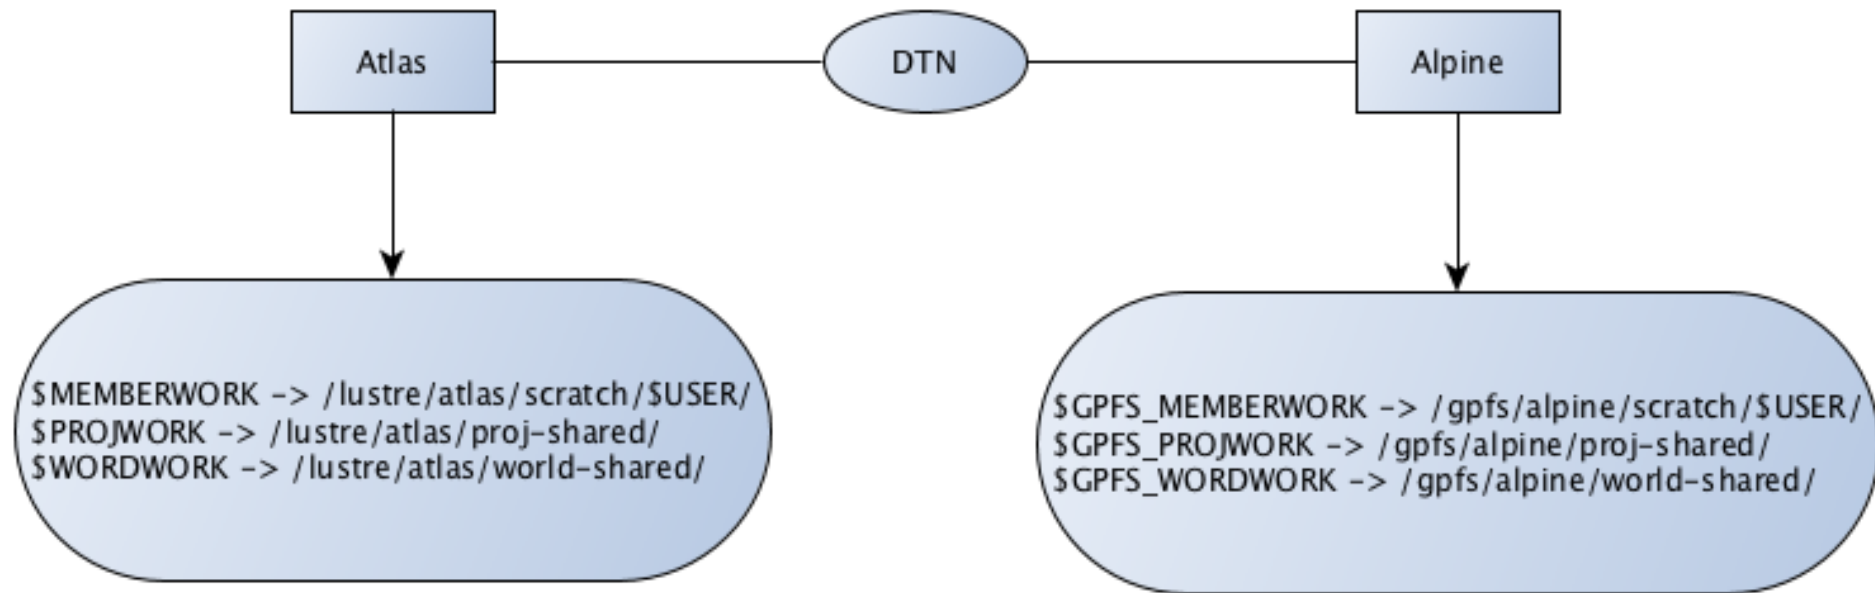
- Study case: Transfer data from Atlas to Alpine with 3 approaches. Copy the files through NFS, use HPSS, or use Globus

Type	Home NFS	HPSS	Globus
	Time in seconds to finish the transfer		
Transfer 22 files of 1GB each	323	270	10
Transfer 1 file of 22 GB	308	301	80
Transfer 4 files of 1GB each	69	53	9

- Globus is the most efficient approach to transfer files for all the evaluated cases, for small files though, transferring through NFS should be efficient.
- There are available some traditional tools such as scp, rsync
- The tests took place on 29<sup>th</sup> November

# DTN

- As long as we have both Atlas and Alpine on DTN, we use the following variables (GPFS variables are not active yet)



# Acknowledgement

*This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.*

Thank you!  
Questions?