

High Performance Computing in GW

Shuyue (Frank) Guan

The Medical Imaging & Image Analysis (MIA) Laboratory

Nov. 2020

Outline

- Introduce to the new HPC: PEGASUS
- How to get an account
- How to upload files
- How to log in the HPC system
- To run jobs by SLURM

New HPC in GW



- One CPU node
 - Dual 20-Core 3.70GHz Intel Xeon processors
 - 192GB RAM
 - 800 GB SSD
- One GPU node
 - 2 NVIDIA **Tesla V100** GPU (4 for large nodes)
 - Dual 20-Core 3.70GHz Intel Xeon (18-Core Xeon for large nodes)
 - 192GB RAM (384GB for large nodes)
 - 800 GB SSD
- High throughput node, High memory node (3TB RAM!),...

Tesla V100 highlight
640 **Tensor Cores**
5120 **CUDA Cores**
16 GB **Memory**

First: request for an account

<https://colonialone.gwu.edu/getting-access/>

mia



NetID *
[Redacted]
Before the '@' symbol in your email.

Full Name *
[Redacted]

Email *
[Redacted]

Principal Investigator (PI) or Faculty Sponsor email *
loew@gwu.edu
If none, please enter your email address.

Modify my existing group membership(s) ?

- Add a group
- Change my primary group
- Remove me from a group

For users that already have a Pegasus account and need group associations changed.

Research Group
mia
Research Group that you would like to be added to.

Please enter your school (or affiliate organization) *
BME, SEAS

Message
[Redacted]

For any questions or concerns, please contact hpchelp@gwu.edu.

Submit

File Transfer with Globus

Globus is used literally everywhere by everyone transferring data in the HPC world. AWS, National Labs, Universities ,..., and GWU.

<https://www.globus.org/>

globus
a ucicago non-profit service

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

237,823 years of non-stop video calls

100x ALL PRINTED MATERIAL of the Library of Congress

400 HUMAN BRAINS worth of memory storage

137 years of observational data from the Rubin Observatory in Chile

1,000,000,000,000,000,000,000
1 EXABYTE - A QUINTILLION BYTES TRANSFERRED BY GLOBUS

266 BILLION human genomes worth of sequence information

66.7 YEARS of the Large Hadron Collider's experimental data

Research data management simplified.

LEARN MORE ABOUT THIS HUGE MILESTONE

Research data management simplified.

TRANSFER SHARE BUILD

When log in



Log in to use Globus Web App

Use your existing organizational login

e.g., university, national lab, facility, project

The George Washington University

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

Continue



Globus uses CILogon to enable you to Log In from this organization. By clicking Continue, you agree to the [CILogon privacy policy](#) and you agree to share your username, email address, and affiliation with CILogon and Globus. You also agree for CILogon to issue a certificate that allows Globus to act on your behalf.

Log in using GW NetID



The George Washington University Web Single Sign-on

Login to **CILogon**

NetID

Password

- Don't Remember Login
- Clear prior granting of permission for release of your information to this service.

> [Forgot your password?](#)

Login

Please do not bookmark this page!

The new single sign-on system works by authenticating your GW NetID and password and then redirecting you back to the system to which you are logging in. If you bookmark this page, the system will not receive all needed information and will display an error when you try to log in.

Additionally, to protect your personal information, please close all open browser windows and tabs when you log out. For Mac, please quit the browser application.

GW students, faculty, staff and alumni can access this service by using their GW NetIDs (the part of the e-mail address that precedes @gwu.edu) and corresponding passwords.

If you are a member of the GW community and do not have a NetID, you can get one by following the instructions on the IT Support Center website.

For further assistance accessing this service, please contact the IT Support Center at 202-994-GWIT (4948), ithelp@gwu.edu or IT.GWU.EDU. The IT Support Center is available Monday - Friday, 7:00am - 10:00pm.

Reminder: GW will never ask for your password via e-mail.

We have upgraded. Please [Check Here](#) for more information.



The George Washington University Web Single Signon Terms of Use

The George Washington University Web Single Signon service: Terms of Use (ToU)

A. Data Protection

"The End User notes that personal data about the End User is compiled from generally available sources and from communications received from the End User and other Universities as well as from off-site sources. Such data will be used, inter alia, to authenticate and authorize the access to and use of various resources within the University and on other sites ("Approved Uses"). The End User hereby consents to the collection, processing, use and release of such data to the extent reasonably necessary for the Approved Uses. Such consent includes, but is not limited to, the release of personal data to other institutions by employing cookies and electronically exchanging, caching and storing personal authorization attributes."

B. Limitation of Liability

"To the extent permitted by the applicable law, the End User hereby waives all and any claims for cost and damages, whether direct or indirect, incidental, or consequential (including, inter alia, loss of use and lost profits), both in contract and in tort, arising from the use or in any way related to the inter-organizational authentication and authorization services which allow the End User to access certain resources of other organizations. This waiver of claims shall be valid and effective in relation to all participants of the inter-organizational authentication and authorization services including the GWU Service Providers and its affiliates, officers, employees and agents."

Refuse



I accept the terms of use

Submit

The screenshot shows the Globus interface. On the left is a dark blue sidebar with icons for FILE MANAGER, BOOKMARKS, ACTIVITY, ENDPOINTS, GROUPS, CONSOLE, ACCOUNT, LOGOUT, and HELP. The FILE MANAGER icon is highlighted with a red checkmark. The top header area is light gray and contains a folder icon and the text 'Collection Search'. Below this is a search input field with the text 'Pegasus HPC Cluster' highlighted in yellow. A red arrow points from the search input field to a search result. The search result is for 'Pegasus HPC Cluster' and includes the owner information 'Owner: 066de31f-e0f6-452f-a797-15abf551fdc7@clients.auth.globus.org' and the description 'Collection of home, group and lustre directories'. A red arrow also points to the search result title.

Collection Search

Collection

 Pegasus HPC Cluster 

Owner: 066de31f-e0f6-452f-a797-15abf551fdc7@clients.auth.globus.org

Collection of home, group and lustre directories

After login, choose [File Manager]
Search collection: **“Pegasus HPC Cluster”**




FILE MANAGER



BOOKMARKS



ACTIVITY



ENDPOINTS



GROUPS




CONSOLE *ip*



ACCOUNT



LOGOUT



HELP

Collection

Pegasus HPC Cluster

Path

/

select all up one folder refresh list

NAME ▾

CCAS

CHILDRENS

GWSPH

home

lustre

SEAS

SMHS

File Manager

Collection: Pegasus HPC Cluster

Path: /SEAS/groups/

select none up one folder refresh list

NAME

- geomech
- huanglab
- iasl
- mia
- mpnl
- parkgrp

Folders

- Personal path: **/SEAS/home/\$your_netID\$**
- For group: **/SEAS/groups/mia**

At this point, you should have the right access (HPC account)




FILE MANAGER



BOOKMARKS




ACTIVITY



ENDPOINTS



GROUPS



CONSOLE



ACCOUNT



LOGOUT



HELP

Collection Pegasus HPC Cluster

Path /SEAS/groups/mia/

File Manager interface showing a large empty area with a dashed border. A blue button at the bottom center reads "Select Files to Upload".



FILE MANAGER



BOOKMARKS



ACTIVITY



ENDPOINTS



GROUPS



CONSOLE



ACCOUNT



LOGOUT



HELP

Collection Pegasus HPC Cluster

Path /SEAS/groups/mia/

Bookmark

Authentication/Consent is required for the Globus web app to perform HTTPS uploads on your behalf.


[Continue](#)

- Share
- Transfer or Sync to...
- New Folder
- Rename
- Delete Selected
- Download
- Open
- Upload**
- Get Link
- Show Hidden Items
- Manage Consent

Upload

- Get Link
- Show Hidden Items
- Manage Consent

Upload files




FILE MANAGER



BOOKMARKS



ACTIVITY



ENDPOINTS



GROUPS




CONSOLE



ACCOUNT



LOGOUT



HELP

Collection

Path

Bookmark ▼

File management toolbar with icons for selection, refresh, view settings, and a menu icon.

NAME	LAST MODIFIED	SIZE
<input checked="" type="checkbox"/> blank.png		2.07 KB

- Transfer or Sync to...
- Rename
- Delete Selected
- Download
- Open
- Get Link

- Share
- Transfer or Sync to...
- New Folder
- Rename
- Delete Selected
- Download
- Open
- Upload
- Get Link
- Show Hidden Items
- Manage Consent

Files modification

Or... using the “Sync model”

The screenshot shows the File Manager interface with a context menu open over an empty folder. The menu item 'Transfer or Sync to...' is highlighted in yellow. The interface includes a sidebar with navigation options like FILE MANAGER, BOOKMARKS, ACTIVITY, ENDPOINTS, GROUPS, CONSOLE, ACCOUNT, LOGOUT, and HELP. The main area shows the folder path '/SEAS/groups/mia/' and a search bar containing 'trying'. A table on the right shows file details with columns for LAST MODIFIED and SIZE.

File Manager

Collection: Pegasus HPC Cluster

Path: /SEAS/groups/mia/

Search: trying

Path: /~/C/Users/

This folder is empty.

Transfer or Sync to...

Share

New Folder

Rename

Delete Selected

Download

Open

Upload

Get Link

Show Hidden Items

LAST MODIFIED SIZE

09/22/2020 0...

Start

Transfer & Sync Options

Start

Uploads


13

Endpoints [Create a personal endpoint](#)

Recently Used In Use Shareable By You Shared With You Administered By You


Create a personal endpoint

- FILE MANAGER
- BOOKMARKS
- ACTIVITY
- ENDPOINTS**
- GROUPS
- CONSOLE
- ACCOUNT
- LOGOUT
- HELP

 Download Globus Connect Personal

Globus Connect Personal allows you to share and transfer files to and from your laptop or desktop computer. Download Globus Connect Personal below and follow the install steps to turn your personal computer into a Globus collection. [Learn more about Globus Connect Personal](#)



 [Download Globus Connect Personal for Windows](#)

[Show me other supported operating systems](#)

Download, install and set the server app to create a personal endpoint...

Review 1

- Request for an account
 - Research group name “**mia**”
- Endpoint name “**Pegasus HPC Cluster**”
- Path
 - Personal path: **/SEAS/home/\$your_netID\$**
 - For group: **/SEAS/groups/mia**

- ✓ Introduce to the new HPC: PEGASUS
- ✓ How to get an account
- ✓ How to upload files

Next How to log in the HPC system
To run jobs by SLURM

VPN Access

It needs GW network environment.

- Off-campus access (at home): requires GW's VPN.

<https://seascf.seas.gwu.edu/vpn-access>

(1) Download the **Cisco AnyConnect VPN Client**

(2) Connect to “**go.vpn.gwu.edu**”

(3) Use NetID and password to connect

- On-campus access: *skip this step.*





SSH client

GW HPC system requires the **SSH** connection.

“SSH (Secure Shell) is a cryptographic network protocol for operating network services securely over an unsecured network.”

- MobaXterm <https://mobaxterm.mobatek.net/>
- PuTTY <https://www.putty.org/>
- ...

Quick connect...

 User sessions
 WSL-Legacy

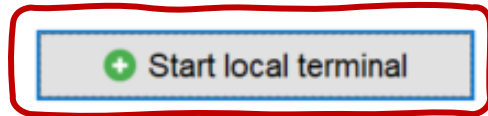
Sessions

Tools

Macros



MobaXterm

 Start local terminal

Find existing session or server name...

Welcome to MobaXterm

Press <return> to start a local terminal

MobaXterm

ssh NetID@pegasus.colonialone.gwu

Password of NetID

Welcome info

Running jobs on the login nodes is prohibited

Command line

Login is successful when see this.

```
2. frank @log001 ~
? MobaXterm Personal Edition v20.3 ?
(X server, SSH client and network tools)

> Your computer drives are accessible through the /drives path
> Your DISPLAY is set to 192.168.1.198:0.0
> When using SSH, your remote DISPLAY is automatically forwarded
> Each command status is specified by a special symbol (✓ or ✗)

? Important:
This is MobaXterm Personal Edition. The Professional edition
allows you to customize MobaXterm for your company: you can add
your own logo, your parameters, your welcome message and generate
either an MSI installation package or a portable executable.
We can also modify MobaXterm or develop the plugins you need.
For more information: https://mobaxterm.mobatek.net/download.html

10/11/2020 23:08.13 /home/mobaxterm ssh frank @pegasus.colonialone.gwu
frank @pegasus.colonialone.gwu.edu's password:

** GW Pegasus Login Node **
** Unauthorized use is prohibited, authorized users must adhere to GW and
** Colonial One policies and procedures at all times. Report usage
** violations to hpchelp@gwu.edu.
**
** Online documentation can be found at http://colonialone.gwu.edu/
**
** PLEASE NOTE:
** Lustre is scratch storage only and not for archiving purposes. Any data
** left on Lustre will be automatically deleted after 30 days or if the
** Lustre partition fills up, whichever comes first. Do not expect data
** to be retained: please move any data you want to preserve to either
** your home directory, group directory or an external storage service.
**
** For support, email hpchelp@gwu.edu
**
[frank @log001 ~]$
```

Linux commends

- Current directory: **pwd**

```
$ pwd  
/SEAS/groups/mia
```

- Go to user's main directory: **cd ~**
- Go to root: **cd /**
- Parent directory: **cd ..**
- Go to a directory: **cd /xxx/xxxx/xx**

Review 2

- ✓ Introduce to the new HPC: PEGASUS
- ✓ How to get an account
- ✓ How to upload files
- ✓ How to log in the HPC system


Next To run jobs by SLURM

- VPN connection to “**go.vpn.gwu.edu**” (off-campus)
- SSH client (like “MobaXterm”)
 - `ssh NetID@pegasus.colonialone.gwu`

Do not run jobs on the login nodes.

Do use the SLURM system.

SLURM is a cluster management and job scheduling system for Linux clusters.



sinfo command

Number of nodes

Node state

Types of nodes

```
$ sinfo
PARTITION AVAIL  TIMELIMIT  NODES  STATE NODELIST
defq*      up    14-00:00:0    1  drain*  cpu153
defq*      up    14-00:00:0    1  down*   cpu066
defq*      up    14-00:00:0    1  drain   cpu004
defq*      up    14-00:00:0   158  alloc   cpu[005-065,067-152,154-164]
short      up    1-00:00:00    19  alloc   cpu[101-119]
tiny       up    8:00:00       1  down*   cpu066
tiny       up    8:00:00       1  drain   cpu004
tiny       up    8:00:00      95  alloc   cpu[005-065,067-100]
highMem    up    14-00:00:0    2  alloc   hmm[001-002]
highThru   up    7-00:00:00    6  alloc   hth[001-006]
debug      up    4:00:00       3  idle    cpu[001-002],gpu013
debug-cpu  up    4:00:00       2  idle    cpu[001-002]
debug-gpu  up    4:00:00       1  idle    gpu013
large-gpu  up    7-00:00:00   22  alloc   gpu[001-012,021-023,032-038]
small-gpu  up    7-00:00:00    1  maint   gpu015
small-gpu  up    7-00:00:00   12  alloc   gpu[016-020,024-030]
small-gpu  up    7-00:00:00    1  idle    gpu031
```

Time limit for running
day – h : m : s

“idle” is ready to use

Node type (Non-GPU)

- **defq*** – default compute nodes, CPU only
- **highThru** – Nodes in this category have **high CPU frequency nodes with a low core count**. These nodes are used for single threaded jobs that cannot take advantage of multiple cores.
- **highMem** – Nodes in this category have **large memory (3TB)** and are for jobs that require more memory intensive jobs.
- **short** – Designed for **quicker** turnaround of shorter running jobs (1 day).
- **tiny** – Designed for **even quicker** turnaround of shorter running jobs (8 hrs).
- **debug-cpu** – See more details later

Node type (GPU)

- **small-gpu** – Each has **two** NVIDIA Tesla V100 SXM2 16GB GPUs.
- **large-gpu** – Each has **four** NVIDIA Tesla V100 SXM2 16GB GPUs with NVLink enabled.
- **debug-gpu** – See more details later

Debug partition


The debug partition is intended for running **interactive jobs, or for short-duration test cases** (e.g., making sure your job scripts work correctly), and is not intended for general use.

The hardware is the same as the cpu/gpu partition nodes, but time limit is **4 hours**.

- **debug-cpu**
- **debug-gpu**

You must specify a time limit for your job.

Workflow

1. Upload data and code files 
2. Create a SLURM script to assign the job
3. Run the script
4. Outputs will be saved
5. Receive states and notifications by Email / or check by commands

A Python example

Find the primes in [2,100].

Code file name: **prime.py**; no data file.

```
if __name__ == '__main__':  
    for num in range(2, 101):  
        prime = True  
        for i in range(2, num):  
            if (num % i == 0):  
                prime = False  
        if prime:  
            print(num)
```


SLURM script for Python

submit.sh


```
#!/bin/sh
#SBATCH --time 1:00 ← To limit running time; format: day – h : m : s
#SBATCH -o testing%j.out File for outputs; “%j” is task number
#SBATCH -e testing%j.err File for error/warning report
#SBATCH -p defq Point the node
#SBATCH --mail-user=user@email.com Email address for receiving states
#SBATCH --mail-type=ALL Email type
module load python/2.7.16 Load modules for running code
python /SEAS/home/$user_name$/prime.py Run code
```

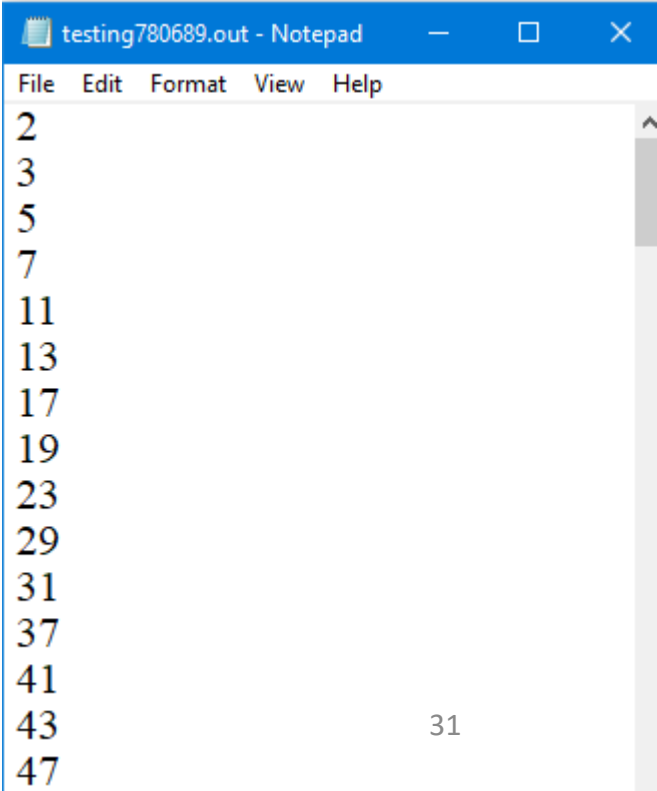
Submit job -- sbatch

```
$ sbatch submit.sh  
Submitted batch job xxxxxxxx
```

If your job starts to run, the job is done or other **status changes**, you will get an email:

<input type="checkbox"/>	☆	slurm	Slurm Job_id=820457 Name=submit.sh Ended, Run time 00:00:00, CANCELLED, ExitCode 0
<input type="checkbox"/>	☆	slurm	Slurm Job_id=820456 Name=submit.sh Ended, Run time 00:00:01, COMPLETED, ExitCode 0
<input type="checkbox"/>	☆	slurm	Slurm Job_id=820456 Name=submit.sh Began, Queued time 00:00:54

Running (printed) results are in the **testingxxxxxxx.out** file. 
Error/warning reports are in the **testingxxxxxxx.err** file.
Open by *nano* or *Notepad* in Win.



```
testing780689.out - Notepad  
File Edit Format View Help  
2  
3  
5  
7  
11  
13  
17  
19  
23  
29  
31  
37  
41  
43  
47
```

Job status

If the job is in queue, to check job's status:

```
$ squeue -u <username>
```

JOBID	PARTITION	NAME	USER	ST	TIME	NODES	NODELIST (REASON)
820457	defq	submit.s	frank	PD	0:00	1	(Resources)

Statuses:

CA = CANCELLED	CD = COMPLETED	CG = COMPLETING
F = FAILED	NF = NODE_FAIL	PD = PENDING
R = RUNNING	S = SUSPENDED	TO = TIMEOUT

Estimate when a job will start:

```
$ squeue -u <username> --start
```

Cancel job:

```
$ scancel <jobid>
```


Code file name: **matlab_simple.m**

A Matlab example

Submit_mat.sh

```
#!/bin/sh
#SBATCH --time 1:00
#SBATCH -o testing%j.out
#SBATCH -e testing%j.err
#SBATCH -p defq
#SBATCH --mail-user=user@email.com
#SBATCH --mail-type=ALL
```

```
module load matlab/2018b
matlab -nodisplay < /SEAS/home/$user_name$/matlab_simple.m
```

```
% Creates a 10x10 Magic square
M = magic(10);
disp(M)
exit
```

*testing781182.out - Notepad

File Edit Format View Help

>>>>>>

M=

92	99	1	8	15	67	74	51	58	40
98	80	7	14	16	73	55	57	64	41
4	81	88	20	22	54	56	63	70	47
85	87	19	21	3	60	62	69	71	28
86	93	25	2	9	61	68	75	52	34
17	24	76	83	90	42	49	26	33	65
23	5	82	89	91	48	30	32	39	66
79	6	13	95	97	29	31	38	45	72
10	12	94	96	78	35	37	44	46	53
11	18	100	77	84	36	43	50	27	59

>>

Code file name: **c_sample.c**

A C example

submit_c.sh

```
#!/bin/sh
#SBATCH --time 1:00
#SBATCH -o testing%j.out
#SBATCH -e testing%j.err
#SBATCH -p defq
#SBATCH --mail-user=user@email.com
#SBATCH --mail-type=ALL
```


```
module load gcc/6.5.0
gcc -Wall /SEAS/home/$user_name$/c_sample.c -o c_sample
```

```
#include <stdio.h>
int main (void)
{
    printf ("Hello, world!\n");
    return 0;
}
```

testing781188.out - Notepad
File Edit Format View Help

Hello, world!

Helpful resources

- **GWU HPC documents:**
<https://colonialone.gwu.edu/about/presentations/>
Highly recommend 
- GWU HPC support email: hpchelp@gwu.edu **← For help**
- SLURM system user guide:
<https://slurm.schedmd.com/>
- The Linux command line for beginner:
<https://ubuntu.com/tutorials/command-line-for-beginners#1-overview>