Introduction to HPC at UD

Created By: IT-RCI

Prerequisites

- HPC Account on a UD Cluster
  - Caviness
  - DARWIN
- Basic Linux Knowledge
  - Navigating the CLI
  - Basic File Operations
  - File Editor (Nano/vim)
- Basic Bash Scripting
- Basic Programming Experience (C, Python, R, etc.)

Caviness Cluster

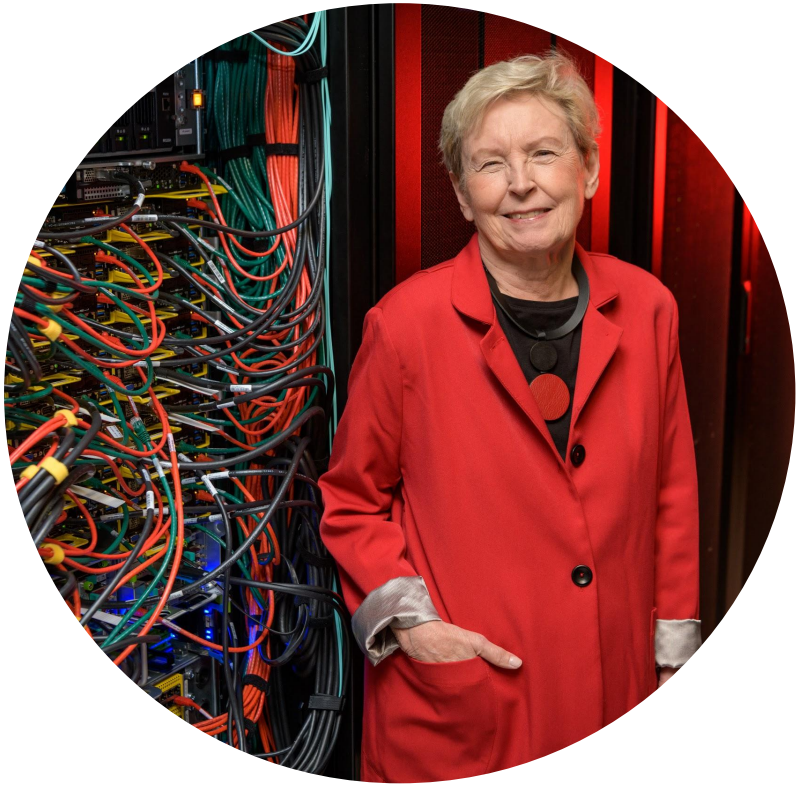

- Community Cluster
  - Research groups purchase node(s) and add'l storage
  - PI of research group must sponsor accounts
  - Workgroup partitions based on resources purchased
  - Any idle nodes can be used altruistically, but can be preempted

DARWIN Cluster

- Allocation Cluster
  - Research groups apply for allocations based on resources needed
    - Start up
    - Research
    - Education
  - PI of allocation requests accounts to be added
  - SU's are required to run jobs
  - CPU, GPU and Storage
  - Free to use

Workshop Overview

SSH Connections & Keys

Workgroups

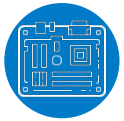

HPC Components

File Systems

Data Management

Data Transfer Methods

Software and VALET

Running Jobs with Slurm

SSH Connections & Keys

How to connect with SSH

- Windows
  - PuTTY
  - WSL
  - PowerShell
  - Terminal Emulator
- Linux/Mac OS
  - Terminal

SSH Connections & Keys

How to connect with SSH

- Windows
  - PuTTY
  - WSL
  - PowerShell
  - Terminal Emulator
- Linux/Mac OS
  - Terminal

```
mkyle@ITRC-G92D633:~$ hostname
ITRC-G92D633
mkyle@ITRC-G92D633:~$
```

SSH Connections & Keys

How to connect with SSH

- Windows
  - PuTTY
  - WSL
  - PowerShell
  - Terminal Emulator
- Linux/Mac OS
  - Terminal

```
mkyle@ITRC-G92D633:~$ hostname
ITRC-G92D633
mkyle@ITRC-G92D633:~$ ssh <UDelNetID>@<Remote_Server>
```

SSH Connections & Keys

How to connect with SSH

- Windows
  - PuTTY
  - WSL
  - PowerShell
  - Terminal Emulator
- Linux/Mac OS
  - Terminal

```
mkyle@ITRC-G92D633:~$ hostname
ITRC-G92D633
mkyle@ITRC-G92D633:~$ ssh mkyle@darwin.hpc.udel.edu
```

SSH Connections & Keys

How to connect with SSH

- Windows
  - PuTTY
  - WSL
  - PowerShell
  - Terminal Emulator
- Linux/Mac OS
  - Terminal

```
mkyle@ITRC-G92D633:~$ hostname
ITRC-G92D633
mkyle@ITRC-G92D633:~$ ssh mkyle@darwin.hpc.udel.edu
Last login: Fri Feb 23 14:09:14 2024 from net10-7-181-136.host.udel.edu
[mkyle@login00.darwin ~]$
```

SSH Connections & Keys

SSH Keys:
What are they?

How are they helpful?

How to set them up?

- ➔ Windows
  - PuTTY
  - WSL/Terminal Emulator
- ➔ Linux/Mac OS

SSH Keys:
What are they?

- ➔ A pair of keys to initiate a secure handshake between remote parties. The key pair contains a public and private key.
  - Private Key ⟶ Key
  - Public Key ⟶ Lock

SSH Connections & Keys

SSH Keys:
What are they?

How are they helpful?

How to set them up?

- ➔ Windows
  - ◆ PuTTY
  - ◆ WSL/Terminal Emulator
- ➔ Linux/Mac OS

SSH Keys:
How are they helpful?

- ➔ Eliminates entering your password for each remote connections - only need to remember a passphrase of your choice*
- ➔ More convenient and efficient especially with other applications such as scp and sftp

* if you set a passphrase when you generate your SSH keys

SSH Connections & Keys

SSH Keys:
What are they?

How are they helpful?

How to set them up?
- ➔ Windows
  - ◆ PuTTY
  - ◆ WSL/Terminal Emulator
- ➔ Linux/Mac OS

SSH Keys:
How to set them up?

Windows:
- ➔ PuTTY
- ➔ WSL/Terminal Emulators
  - ◆ Should be similar to setting up SSH Keys on a Linux or Mac OS system.

SSH Connections & Keys

SSH Keys:
What are they?

How are they helpful?

How to set them up?

- ➔ Windows
  - ◆ PuTTY
  - ◆ WSL/Terminal Emulator
- ➔ Linux/Mac OS

SSH Keys:
How to set them up?

Linux/Mac OS:

- ➔ Generate an SSH Key Pair
- ➔ Upload The Public Key

SSH Connections & Keys

SSH Keys:
What are they?

How are they helpful?

How to set them up?

- ➔ Windows
  - ◆ PuTTY
  - ◆ WSL/Terminal Emulator
- ➔ Linux/Mac OS

SSH Keys:
How to set them up?

Linux/Mac OS:

- ➔ Generate an SSH Key Pair
- ➔ Upload The Public Key
  - ◆ Use caution when uploading to not overwrite existing SSH Keys. You might want to create a copy of ~/.ssh/authorized_keys on the remote machine.

WORKGROUP

Workgroup(s) are assigned when accounts are created and are disseminated to users in welcome emails.

- How to check your assigned workgroups
  - **`workgroup --menu`**
  - `workgroup -q workgroups`
- How to join a workgroup
  - `workgroup -g <workgroup_name>`

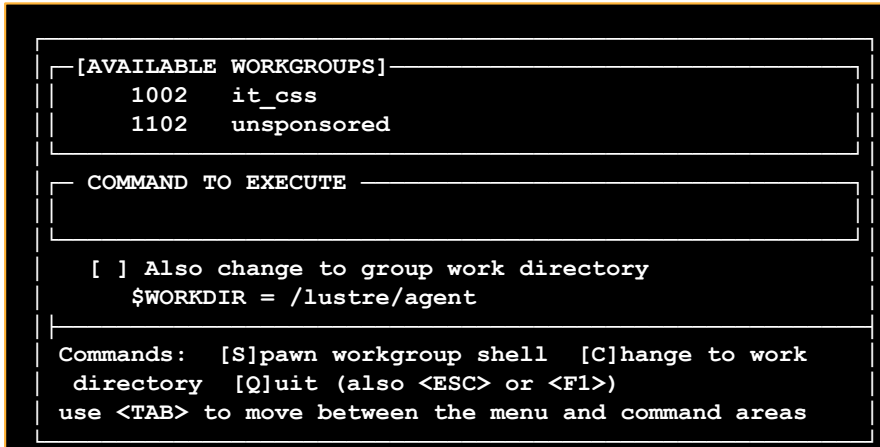

WORKGROUP

Workgroup(s) are assigned when accounts are created and are disseminated to users in welcome emails.

- How to check your assigned workgroups
  - `workgroup --menu`
  - **`workgroup -q workgroups`**
- How to join a workgroup
  - `workgroup -g <workgroup_name>`

```
[mkyle@login01.darwin ~]$ workgroup --menu
[mkyle@login01.darwin ~]$ workgroup -q workgroups
 1002  it_css
 1102  unsponsored
[mkyle@login01.darwin ~]$
```

WORKGROUP

Workgroup(s) are assigned when accounts are created and are disseminated to users in welcome emails.

- How to check your assigned workgroups
  - `workgroup --menu`
  - `workgroup -q workgroups`
- How to join a workgroup
  - **`workgroup -g <workgroup_name>`**

```
[mkyle@login01.darwin ~]$ workgroup --menu
[mkyle@login01.darwin ~]$ workgroup -q workgroups
 1002  it_css
 1102  unsponsored
[mkyle@login01.darwin ~]$ workgroup -g it_css
[(it_css:mkyle)@login01.darwin ~]$
```

HPC Components

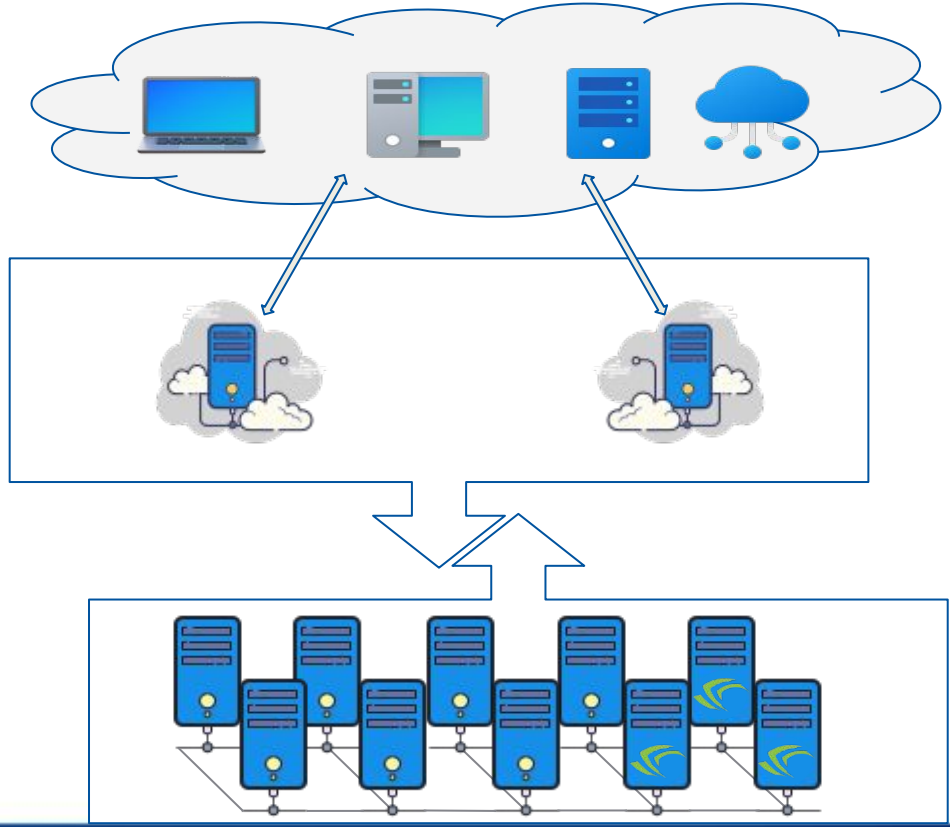

Internet at Large

Login Nodes

- Accessible from the internet
- Can connect to compute nodes

Compute Nodes

- Private Network
- Only accessible from login/other compute nodes

HPC Components

Login Nodes

- Caviness & DARWIN each have 2 login nodes

Do’s

- Remote connection
- Transfer data (<=30 minutes)
- Compile code
- Quickly test code
- Use Slurm to submit jobs

Do Not's

- Run applications/executable code
- Transfer data (>30 minutes)
- Run interactive sessions

HPC Components

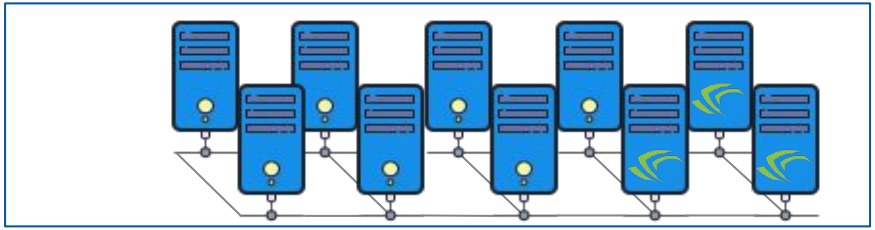

Compute Nodes

- CPU and GPU Nodes Available
- Amounts vary

Do’s

- Run long computational jobs
- Transfer >30 minutes
- Run interactive sessions
- Scale code
- Large scale test/debug code
- Compile code (only available on devel partition on Caviness)

Do Not's

- Run applications/executable code directly without submitting via Slurm

File Systems

Caviness

Home Storage

- 20GB
- Personal use files
- /[illegible]<[illegible]d>
- `df -h $HOME`

Backed Up

Lustre Storage

- 403 TB
- Omni-path Infiniband
- Temporary/Scratch storage
- `/lustre/scratch`

NOT BACKED UP

General Information

- Head 10 GBit/s
- Compute 1 GBit/s
- Home & Workgroup Backups: `.zfs/snapshots`

Workgroup Storage

- >= 1TB
- Shared within workgroup
- /[illegible]<[illegible]e>
- `df -h $WORKDIR`

Backed Up

Compute Node

NOT BACKED UP

- Vary in size by compute node
- Temporary job storage
- `$TMPDIR`

File Systems

DARWIN

Home Storage

- 20GB
- Personal use files
- /home/<user_id>
- $HOME

Inform

- No Back Ups!!

Storage

- >= TB
- red within workgroup
- /lustre/<workgroup_name>/
- $WORKDIR

L

- 1.1PB
- Omni-pa nfiniband

Compute Node

- >= 2TB
- Temporary job storage
- $TMPDIR

NOT BACKED UP

Data Management

HPC Focused Checklist

- Data Production
  - Types of Data
  - Reproducibility
  - Tools & Software
- **Data Size**
  - **How Much**
  - Rate of Growth
  - Frequency of Change
- **Data Usage**
  - **Who Needs Access**

- Privacy & Security
  - Data Classification Level
- Data Documentation
  - Organization & Naming
- Storage and Backup

Data Management
Data Size

Quotas on Caviness

- **df -h $HOME**
- df -h $WORKDIR
- df -h /lustre/scratch
- my_quotas

```
[(it_css:mkyle)@login01.caviness ~]$ df -h $HOME
Filesystem                      Size  Used Avail Use% Mounted on
r03nfs0-10Gb:/fs/r03nfs0/home/2179   20G   15G  5.9G  71% /home/2179
[(it_css:mkyle)@login01.caviness ~]$
```

Data Management
Data Size

Quotas on Caviness

- df -h $HOME
- **df -h $WORKDIR**
- df -h /lustre/scratch
- my_quotas

```
[(it_css:mkyle)@login01.caviness ~]$ df -h $HOME
Filesystem                      Size  Used Avail Use% Mounted on
r03nfs0-10Gb:/fs/r03nfs0/home/2179   20G   15G  5.9G  71% /home/2179
[(it_css:mkyle)@login01.caviness ~]$ df -h $WORKDIR
Filesystem                           Size  Used Avail Use% Mounted on
r03nfs0-10Gb:/fs/r03nfs0/work/it_css  1.1T  761G  265G  75% /work/it_css
[(it_css:mkyle)@login01.caviness ~]$
```

Data Management
Data Size

Quotas on Caviness

- df -h $HOME
- df -h $WORKDIR
- **df -h /lustre/scratch**
- my_quotas

```
[(it_css:mkyle)@login01.caviness ~]$ df -h $HOME
Filesystem                          Size  Used Avail Use% Mounted on
r03nfs0-10Gb:/fs/r03nfs0/home/2179   20G   15G  5.9G  71% /home/2179
[(it_css:mkyle)@login01.caviness ~]$ df -h $WORKDIR
Filesystem                             Size  Used Avail Use% Mounted on
r03nfs0-10Gb:/fs/r03nfs0/work/it_css  1.1T  761G  265G  75% /work/it_css
[(it_css:mkyle)@login01.caviness ~]$ df -h /lustre/scratch
Filesystem                                      Size  Used Avail Use%
Mounted on
10.65.32.18@o2ib:10.65.32.19@o2ib:/scratch  367T  290T   77T  80%
/lustre/scratch
[(it_css:mkyle)@login01.caviness ~]$
```

Data Management
Data Size

Quotas on Caviness

- `df -h $HOME`
- `df -h $WORKDIR`
- `df -h /lustre/scratch`
- **`my_quotas`**

```
[(it_css:mkyle)@login01.caviness ~]$ df -h /lustre/scratch
Filesystem                                      Size  Used Avail Use%
Mounted on
10.65.32.18@o2ib:10.65.32.19@o2ib:/scratch  367T  290T   77T  80%
/lustre/scratch
[mkyle@login00.caviness ~]$ my_quotas
Type  Path                  In-use / kiB   Available / kiB  Pct
----- -------------------- ------------ ----------------- ----
user  /home/2179                14814208          20971520  71%
group /work/it_css             477801472        1075068928  44%
[(it_css:mkyle)@login01.caviness ~]$
```

Data Management
Data Size

Quotas on DARWIN

- **df -h $HOME**
- my_quotas

```
[mkyle@login01.darwin ~]$ df -h $HOME
Filesystem                Size  Used Avail Use% Mounted on
nfs0-ib:/beagle/home/2179   20G  6.7G   14G  34% /home/2179
[mkyle@login01.darwin ~]$
```

Data Management
Data Size

Quotas on DARWIN

- `df -h $HOME`
- **`my_quotas`**

```
[mkyle@login01.darwin ~]$ df -h $HOME
Filesystem                Size  Used Avail Use% Mounted on
nfs0-ib:/beagle/home/2179  20G  6.7G   14G  34% /home/2179
[mkyle@login01.darwin ~]$ my_quotas
Type  Path                In-use / kiB Available / kiB  Pct
----- ------------------- ------------ ----------- ---- ----
user  /home/2179              6950912        20971520  33%
group /lustre/it_css       1483539500      1383539500  93%
[mkyle@login01.darwin ~]$
```

Data Management
Data Usage

chgrp

- `chgrp [options] <group_name> <files/dirs>`
- Change group ownership
- Sharing is caring

chmod

- `chmod [options] <permissions> <files/dirs>`
- Change file mode bits
- AKA permissions

```
[(it_css:mkyle)@login01.darwin ~]$ cp -a group_project /lustre/it_css
[(it_css:mkyle)@login01.darwin it_css]$
```

Data Management
Data Usage

chgrp

- `chgrp [options] <files/dirs>`
- Change group ownership
- Sharing is caring

chmod

- `chmod [options] <files/dirs>`
- Change file mode bits
- AKA permissions

```
[(it_css:mkyle)@login01.darwin ~]$ cd $WORKDIR
[(it_css:mkyle)@login01.darwin it_css]$ ls
anita  group_project  jnhuffma  mkyle  sw  system_status  thuachen  traine
users  vnc
```

Data Management
Data Usage

chgrp

- `chgrp [options] <files/dirs>`
- Change group ownership
- Sharing is caring

chmod

- `chmod [options] <files/dirs>`
- Change file mode bits
- AKA permissions

```
[(it_css:mkyle)@login01.darwin it_css]$ cd $WORKDIR
[(it_css:mkyle)@login01.darwin it_css]$ ls
anita  group_project  jnhuffma  mkyle  sw  system_status  thuachen  traine
users  vnc
[(it_css:mkyle)@login01.darwin it_css]$ ls -ld group_project
drwxr-sr-x 2 mkyle everyone 41472 Apr 15 10:57 group_project
```

Data Management
Data Usage

chgrp

- `chgrp [options] <files/dirs>`
- Change group ownership
- Sharing is caring

chmod

- `chmod [options] <files/dirs>`
- Change file mode bits
- AKA permissions

```
[(it_css:mkyle)@login01.darwin it_css]$ cd $WORKDIR
[(it_css:mkyle)@login01.darwin it_css]$ ls
anita  group_project  jnhuffma  mkyle  sw  system_status  thuachen  traine
users  vnc
[(it_css:mkyle)@login01.darwin it_css]$ ls -ld group_project
drwxr-sr-x 2 mkyle everyone 41472 Apr 15 10:57 group_project
[(it_css:mkyle)@login01.darwin it_css]$ ls -l group_project | head -n 5
total 50
-rw-r--r-- 1 mkyle everyone 0 Apr 15 10:57 file1.csv
-rw-r--r-- 1 mkyle everyone 0 Apr 15 10:57 file10.csv
-rw-r--r-- 1 mkyle everyone 0 Apr 15 10:57 file100.csv
-rw-r--r-- 1 mkyle everyone 0 Apr 15 10:57 file11.csv
[(it_css:mkyle)@login01.darwin it_css]$
```

Data Management
Data Usage

chgrp

- `chgrp [options] <files/dirs>`
- Change group ownership
- Sharing is caring

chmod

- `chmod [options] <files/dirs>`
- Change file mode bits
- AKA permissions

```
[(it_css:mkyle)@login01.darwin it_css]$ cd $WORKDIR
[(it_css:mkyle)@login01.darwin it_css]$ ls
anita  group_project  jnhuffma  mkyle  sw  system_status  thuachen  traine
users  vnc
[(it_css:mkyle)@login01.darwin it_css]$ ls -ld group_project
drwxr-sr-x 2 mkyle everyone 41472 Apr 15 10:57 group_project
[(it_css:mkyle)@login01.darwin it_css]$ ls -l group_project | head -n 5
total 50
-rw-r--r-- 1 mkyle everyone 0 Apr 15 10:57 file1.csv
-rw-r--r-- 1 mkyle everyone 0 Apr 15 10:57 file10.csv
-rw-r--r-- 1 mkyle everyone 0 Apr 15 10:57 file100.csv
-rw-r--r-- 1 mkyle everyone 0 Apr 15 10:57 file11.csv
[(it_css:mkyle)@login01.darwin it_css]$ chgrp -R it_css group_project
[(it_css:mkyle)@login01.darwin it_css]$
```

Data Management
Data Usage

chgrp

- `chgrp [options] <files/dirs>`
- Change group ownership
- Sharing is caring

chmod

- `chmod [options] <files/dirs>`
- Change file mode bits
- AKA permissions

```
[(it_css:mkyle)@login01.darwin it_css]$ cd $WORKDIR
[(it_css:mkyle)@login01.darwin it_css]$ ls
anita  group_project  jnhuffma  mkyle  sw  system_status  thuachen  traine
users  vnc
[(it_css:mkyle)@login01.darwin it_css]$ ls -ld group_project
drwxr-sr-x 2 mkyle everyone 41472 Apr 15 10:57 group_project
[(it_css:mkyle)@login01.darwin it_css]$ ls -l group_project | head -n 5
total 50
-rw-r--r-- 1 mkyle everyone 0 Apr 15 10:57 file1.csv
-rw-r--r-- 1 mkyle everyone 0 Apr 15 10:57 file10.csv
-rw-r--r-- 1 mkyle everyone 0 Apr 15 10:57 file100.csv
-rw-r--r-- 1 mkyle everyone 0 Apr 15 10:57 file11.csv
[(it_css:mkyle)@login01.darwin it_css]$ chgrp -R it_css group_project
[(it_css:mkyle)@login01.darwin it_css]$ ls -ld group_project
drwxr-sr-x 2 mkyle it_css 41472 Apr 15 10:57 group_project
[(it_css:mkyle)@login01.darwin it_css]$ ls -l group_project | head -n 5
total 50
-rw-r--r-- 1 mkyle it_css 0 Apr 15 10:57 file1.csv
-rw-r--r-- 1 mkyle it_css 0 Apr 15 10:57 file10.csv
-rw-r--r-- 1 mkyle it_css 0 Apr 15 10:57 file100.csv
-rw-r--r-- 1 mkyle it_css 0 Apr 15 10:57 file11.csv
```

Data Management
Data Usage

chgrp

- `chgrp [options] <files/dirs>`
- Change group ownership
- Sharing is caring

chmod

- `chmod [options] <files/dirs>`
- Change file mode bits
- AKA permissions

```
[(it_css:mkyle)@login01.darwin it_css]$ ls -ld group_project
drwxr-sr-x 2 mkyle it_css 41472 Apr 15 10:57 group_project
[(it_css:mkyle)@login01.darwin it_css]$ ls -l group_project | head -n 5
total 50
-rw-r--r-- 1 mkyle it_css 0 Apr 15 10:57 file1.csv
-rw-r--r-- 1 mkyle it_css 0 Apr 15 10:57 file10.csv
-rw-r--r-- 1 mkyle it_css 0 Apr 15 10:57 file100.csv
-rw-r--r-- 1 mkyle it_css 0 Apr 15 10:57 file11.csv
[(it_css:mkyle)@login01.darwin it_css]$ chmod -R 760 group_project
[(it_css:mkyle)@login01.darwin it_css]$
```

Data Management
Data Usage

chgrp

- `chgrp [options] <files/dirs>`
- Change group ownership
- Sharing is caring

chmod

- `chmod [options] <files/dirs>`
- Change file mode bits
- AKA permissions

```
[(it_css:mkyle)@login01.darwin it_css]$ ls -ld group_project
drwxr-sr-x 2 mkyle it_css 41472 Apr 15 10:57 group_project
[(it_css:mkyle)@login01.darwin it_css]$ ls -l group_project | head -n 5
total 50
-rw-r--r-- 1 mkyle it_css 0 Apr 15 10:57 file1.csv
-rw-r--r-- 1 mkyle it_css 0 Apr 15 10:57 file10.csv
-rw-r--r-- 1 mkyle it_css 0 Apr 15 10:57 file100.csv
-rw-r--r-- 1 mkyle it_css 0 Apr 15 10:57 file11.csv
[(it_css:mkyle)@login01.darwin it_css]$ chmod -R 760 group_project
[(it_css:mkyle)@login01.darwin it_css]$ ls -ld group_project
drwxrws--- 2 mkyle it_css 41472 Apr 15 10:57 group_project
[(it_css:mkyle)@login01.darwin it_css]$ ls -l group_project | head -n 5
total 50
-rwxrw---- 1 mkyle it_css 0 Apr 15 10:57 file1.csv
-rwxrw---- 1 mkyle it_css 0 Apr 15 10:57 file10.csv
-rwxrw---- 1 mkyle it_css 0 Apr 15 10:57 file100.csv
-rwxrw---- 1 mkyle it_css 0 Apr 15 10:57 file11.csv
```

Data Transfer Methods

File Transfer Methods

- Basic
  - **sftp**
  - scp
  - rsync
- Advanced
  - rclone
  - Globus

```
[mkyle@ITRC-G92D633 ~]$ sftp mkyle@darwin.hpc.udel.edu
Connected to darwin.hpc.udel.edu.
sftp>
```

Data Transfer Methods

File Transfer Methods

- Basic
  - **sftp**
  - scp
  - rsync
- Advanced
  - rclone
  - Globus

```
[mkyle@ITRC-G92D633 ~]$ sftp mkyle@darwin.hpc.udel.edu
Connected to darwin.hpc.udel.edu.
sftp> lpwd
Local working directory: /home/mkyle/new_data
sftp>
```

Data Transfer Methods

File Transfer Methods

- Basic
  - **sftp**
  - scp
  - rsync
- Advanced
  - rclone
  - Globus

```
[mkyle@ITRC-G92D633 ~]$ sftp mkyle@darwin.hpc.udel.edu
Connected to darwin.hpc.udel.edu.
sftp> lpwd
Local working directory: /home/mkyle/new_data
sftp> pwd
Remote working directory: /home/2179
sftp>
```

Data Transfer Methods

File Transfer Methods

- Basic
  - **sftp**
  - scp
  - rsync
- Advanced
  - rclone
  - Globus

```
[mkyle@ITRC-G92D633 ~]$ sftp mkyle@darwin.hpc.udel.edu
Connected to darwin.hpc.udel.edu.
sftp> lpwd
Local working directory: /home/mkyle/new_data
sftp> pwd
Remote working directory: /home/2179
sftp> cd /lustre/it_css/group_project
sftp>
```

Data Transfer Methods

File Transfer Methods

- Basic
  - **sftp**
  - scp
  - rsync
- Advanced
  - rclone
  - Globus

```
[mkyle@ITRC-G92D633 ~]$ sftp mkyle@darwin.hpc.udel.edu
Connected to darwin.hpc.udel.edu.
sftp> lpwd
Local working directory: /home/mkyle/new_data
sftp> pwd
Remote working directory: /home/2179
sftp> cd /lustre/it_css/group_project
sftp> get file1.csv
sftp>
```

Data Transfer Methods

File Transfer Methods

- Basic
  - **sftp**
  - scp
  - rsync
- Advanced
  - rclone
  - Globus

```
[mkyle@ITRC-G92D633 ~]$ sftp mkyle@darwin.hpc.udel.edu
Connected to darwin.hpc.udel.edu.
sftp> lpwd
Local working directory: /home/mkyle/new_data
sftp> pwd
Remote working directory: /home/2179
sftp> cd /lustre/it_css/group_project
sftp> get file1.csv
sftp> lls file1.csv
sftp>
```

Data Transfer Methods

File Transfer Methods

- Basic
  - **sftp**
  - scp
  - rsync
- Advanced
  - rclone
  - Globus

```
[mkyle@ITRC-G92D633 ~]$ sftp mkyle@darwin.hpc.udel.edu
Connected to darwin.hpc.udel.edu.
sftp> lpwd
Local working directory: /home/mkyle/new_data
sftp> pwd
Remote working directory: /home/2179
sftp> cd /lustre/it_css/group_project
sftp> get file1.csv
sftp> lls file1.csv
sftp> put file101.csv
Uploading file101.csv to /lustre/it_css/group_project/file101.csv
file101.csv                                   100%    0     0.0KB/s   00:00
sftp>
```

Data Transfer Methods

File Transfer Methods

- Basic
  - **sftp**
  - scp
  - rsync
- Advanced
  - rclone
  - Globus

```
[mkyle@ITRC-G92D633 ~]$ sftp mkyle@darwin.hpc.udel.edu
Connected to darwin.hpc.udel.edu.
sftp> lpwd
Local working directory: /home/mkyle/new_data
sftp> pwd
Remote working directory: /home/2179
sftp> cd /lustre/it_css/group_project
sftp> get file1.csv
sftp> lls file1.csv
sftp> put file101.csv
Uploading file101.csv to /lustre/it_css/group_project/file101.csv
file101.csv                                   100%   0      0.0KB/s   00:00
sftp> ls file1??.csv
file100.csv         file101.csv
sftp>
```

Data Transfer Methods

File Transfer Methods

- Basic
  - **sftp**
  - scp
  - rsync
- Advanced
  - rclone
  - Globus

```
[mkyle@ITRC-G92D633 ~]$ sftp mkyle@darwin.hpc.udel.edu
Connected to darwin.hpc.udel.edu.
sftp> lpwd
Local working directory: /home/mkyle/new_data
sftp> pwd
Remote working directory: /home/2179
sftp> cd /lustre/it_css/group_project
sftp> get file1.csv
sftp> lls file1.csv
sftp> put file101.csv
Uploading file101.csv to /lustre/it_css/group_project/file101.csv
file101.csv                     100%   0      0.0KB/s   00:00
sftp> ls file1??.csv
file100.csv         file101.csv
sftp> quit
mkyle@ITRC-G92D633:~/new_data$
```

Data Transfer Methods

File Transfer Methods

- Basic
  - sftp
  - **scp**
  - rsync
- Advanced
  - rclone
  - Globus

```
mkyle@ITRC-G92D633:~/new_data$ pwd
/home/mkyle/new_data
mkyle@ITRC-G92D633:~/new_data$
```

Data Transfer Methods

File Transfer Methods

- Basic
  - sftp
  - **scp**
  - rsync
- Advanced
  - rclone
  - Globus

```
mkyle@ITRC-G92D633:~/new_data$ pwd
/home/mkyle/new_data
mkyle@ITRC-G92D633:~/new_data$ ls
file1.csv    file102.csv  file104.csv  file106.csv  file108.csv
file110.csv file101.csv  file103.csv  file105.csv  file107.csv
file109.csv
mkyle@ITRC-G92D633:~/new_data$
```

Data Transfer Methods

File Transfer Methods

- Basic
  - sftp
  - **scp**
  - rsync
- Advanced
  - rclone
  - Globus

```
mkyle@ITRC-G92D633:~/new_data$ pwd
/home/mkyle/new_data
mkyle@ITRC-G92D633:~/new_data$ ls
file1.csv    file102.csv  file104.csv  file106.csv  file108.csv
file110.csv file101.csv  file103.csv  file105.csv  file107.csv
file109.csv
mkyle@ITRC-G92D633:~/new_data$ scp [options] <source_file_or_directory>
<destination>
```

Data Transfer Methods

File Transfer Methods

- Basic
  - sftp
  - **scp**
  - rsync
- Advanced
  - rclone
  - Globus

```
mkyle@ITRC-G92D633:~/new_data$ pwd
/home/mkyle/new_data
mkyle@ITRC-G92D633:~/new_data$ ls
file1.csv    file102.csv  file104.csv  file106.csv  file108.csv
file110.csv file101.csv  file103.csv  file105.csv  file107.csv
file109.csv
mkyle@ITRC-G92D633:~/new_data$ scp file1??.csv
mkyle@darwin.hpc.udel.edu:/lustre/it_css/group_project/
file101.csv                                   100%   0      0.0KB/s   00:00
file102.csv                                   100%   0      0.0KB/s   00:00
file103.csv                                   100%   0      0.0KB/s   00:00
file104.csv                                   100%   0      0.0KB/s   00:00
file105.csv                                   100%   0      0.0KB/s   00:00
file106.csv                                   100%   0      0.0KB/s   00:00
file107.csv                                   100%   0      0.0KB/s   00:00
file108.csv                                   100%   0      0.0KB/s   00:00
file109.csv                                   100%   0      0.0KB/s   00:00
file110.csv                                   100%   0      0.0KB/s   00:00
mkyle@ITRC-G92D633:~/new_data$
```

Data Transfer Methods

File Transfer Methods

- Basic
  - sftp
  - **scp**
  - rsync
- Advanced
  - rclone
  - Globus

```
file106.csv                                   100%    0       0.0KB/s   00:00
file107.csv                                   100%    0       0.0KB/s   00:00
file108.csv                                   100%    0       0.0KB/s   00:00
file109.csv                                   100%    0       0.0KB/s   00:00
file110.csv                                   100%    0       0.0KB/s   00:00
mkyle@ITRC-G92D633:~/new_data$ scp
mkyle@darwin.hpc.udel.edu:/lustre/it_css/group_project/file9?.csv .
file90.csv                                    100%    0       0.0KB/s   00:00
file91.csv                                    100%    0       0.0KB/s   00:00
file92.csv                                    100%    0       0.0KB/s   00:00
file93.csv                                    100%    0       0.0KB/s   00:00
file94.csv                                    100%    0       0.0KB/s   00:00
file95.csv                                    100%    0       0.0KB/s   00:00
file96.csv                                    100%    0       0.0KB/s   00:00
file97.csv                                    100%    0       0.0KB/s   00:00
file98.csv                                    100%    0       0.0KB/s   00:00
file99.csv                                    100%    0       0.0KB/s   00:00
mkyle@ITRC-G92D633:~/new_data$
```

Data Transfer Methods

File Transfer Methods

- Basic
  - sftp
  - **scp**
  - rsync
- Advanced
  - rclone
  - Globus

```
file106.csv                                   100%    0       0.0KB/s   00:00
file107.csv                                   100%    0       0.0KB/s   00:00
file108.csv                                   100%    0       0.0KB/s   00:00
file109.csv                                   100%    0       0.0KB/s   00:00
file110.csv                                   100%    0       0.0KB/s   00:00
mkyle@ITRC-G92D633:~/new_data$ scp
mkyle@darwin.hpc.udel.edu:/lustre/it_css/group_project/file9?.csv .
file90.csv                                    100%    0       0.0KB/s   00:00
file91.csv                                    100%    0       0.0KB/s   00:00
file92.csv                                    100%    0       0.0KB/s   00:00
file93.csv                                    100%    0       0.0KB/s   00:00
file94.csv                                    100%    0       0.0KB/s   00:00
file95.csv                                    100%    0       0.0KB/s   00:00
file96.csv                                    100%    0       0.0KB/s   00:00
file97.csv                                    100%    0       0.0KB/s   00:00
file98.csv                                    100%    0       0.0KB/s   00:00
file99.csv                                    100%    0       0.0KB/s   00:00
mkyle@ITRC-G92D633:~/new_data$ ls file9*
file90.csv  file92.csv  file94.csv  file96.csv  file98.csv
file91.csv  file93.csv  file95.csv  file97.csv  file99.csv
mkyle@ITRC-G92D633:~/new_data$
```

Data Transfer Methods

File Transfer Methods

- Basic
  - `sftp`
  - `scp`
  - **`rsync`**
- Advanced
  - `rclone`
  - `Globus`

```
mkyle@ITRC-G92D633:~/new_data$ cd ..
mkyle@ITRC-G92D633:~$
```

Data Transfer Methods

File Transfer Methods

- Basic
  - sftp
  - scp
  - **rsync**
- Advanced
  - rclone
  - Globus

```
mkyle@ITRC-G92D633:~/new_data$ cd ..
mkyle@ITRC-G92D633:~$ mkdir data_backup
mkyle@ITRC-G92D633:~$
```

Data Transfer Methods

File Transfer Methods

- Basic
  - sftp
  - scp
  - **rsync**
- Advanced
  - rclone
  - Globus

```
mkyle@ITRC-G92D633:~/new_data$ cd ..
mkyle@ITRC-G92D633:~$ mkdir data_backup
mkyle@ITRC-G92D633:~$ cd data_backup
mkyle@ITRC-G92D633:~/data_backup$ rsync -azP
mkyle@darwin.hpc.udel.edu:/lustre/it_css/group_project
/home/mkyle/data_backup
```

Data Transfer Methods

File Transfer Methods

- Basic
  - sftp
  - scp
  - **rsync**
- Advanced
  - rclone
  - Globus

```
mkyle@ITRC-G92D633:~/new_data$ cd ..
mkyle@ITRC-G92D633:~$ mkdir data_backup
mkyle@ITRC-G92D633:~$ cd data_backup
mkyle@ITRC-G92D633:~/data_backup$ rsync -azP
mkyle@darwin.hpc.udel.edu:/lustre/it_css/group_project
/home/mkyle/data_backup
receiving incremental file list
group_project/
group_project/file1.csv   0 100%  0.00kB/s     0:00:00 (xfr#1,
to-chk=109/111)
group_project/file10.csv  0 100%  0.00kB/s     0:00:00 (xfr#2,
to-chk=108/111)
group_project/file100.csv 0 100%  0.00kB/s     0:00:00 (xfr#3,
to-chk=107/111)
group_project/file101.csv 0 100%  0.00kB/s     0:00:00 (xfr#4,
to-chk=106/111)
group_project/file102.csv 0 100%  0.00kB/s     0:00:00 (xfr#5,
to-chk=105/111)
group_project/file103.csv 0 100%  0.00kB/s     0:00:00 (xfr#6,
to-chk=104/111)
... dozens of additional files ...
group_project/file99.csv 0 100%   0.00kB/s     0:00:00 (xfr#110,
to-chk=0/111)
mkyle@ITRC-G92D633:~/data_backup$
```

Data Transfer Methods

File Transfer Methods

- Basic
  - sftp
  - scp
  - **rsync**
- Advanced
  - rclone
  - Globus

```
mkyle@ITRC-G92D633:~/data_backup$ rsync -azP
mkyle@darwin.hpc.udel.edu:/lustre/it_css/group_project
/home/mkyle/data_backup
receiving incremental file list
group_project/file1.csv    9 100% 8.79kB/s     0:00:00 (xfr#1,
to-chk=109/111)
group_project/file10.csv   9 100% 0.80kB/s     0:00:00 (xfr#2,
to-chk=108/111)
group_project/file100.csv  9 100% 0.68kB/s     0:00:00 (xfr#3,
to-chk=107/111)
group_project/file110.csv  9 100% 0.59kB/s     0:00:00 (xfr#4,
to-chk=96/111)
mkyle@ITRC-G92D633:~/data_backup$
```

Data Transfer Methods

File Transfer Methods

- Basic
  - `sftp`
  - `scp`
  - `rsync`
- Advanced
  - **`rclone`**
  - `Globus`

Popular Rclone Storage Systems

- Amazon S3
- Box
- Dropbox
- FTP
- Google Cloud/Drive*
- Azure Blob/File Storage
- OneDrive
- SFTP

* Currently there is issue with Google Drive connections

Data Transfer Methods

File Transfer Methods

- Basic
  - `sftp`
  - `scp`
  - `rsync`
- Advanced
  - `rclone`
  - **`Globus`**

Globus

Globus is a web based file transfer application that allows resilient, unattended file transfers between two Globus endpoints.

- Good for small and large transfers
- Web browser and or CLI
- Parallel transfers
- Encryption
- Performance tuning automatically
- Failed transfers automatically restart
- UD IT covers the subscription
- Globus is on Caviness and DARWIN
  - https://docs.hpc.udel.edu/software/globus/globus

Software and VALET

VALET

UD-developed software to help configure your environment for all IT-installed software packages.

- Similar to Modules used on other clusters, however VALET offers other features such as preventing conflicts of software versions being loaded
- Changes environment variables such as PATH, LD_LIBRARY_PATH and MANPATH
- Changes software development environment variables such as LDFLAGS and CPPFLAGS

```
[mkyle@login00.darwin ~]$ man valet
VALET(7)
       VALET User Manual
       VALET(7)

NAME
       valet - VALET Automates Linux Environment Tasks

DESCRIPTION
       One  particularly annoying aspect of cluster computing for most
users is getting the environment setup properly given a set of requisite
libraries and applications. Quite often in a cluster environment there
are multiple versions of a library available at any one time, and knowing
which to use when running a  program  can  often  mean  the  difference
       between  predictable, correct results and a flawed or even failed
execution.  A piece of software called (quite simply) "modules" has been
around for quite some time to
Manual page valet(7) line 1 (press h for help or q to quit)
```

Software and VALET

VALET commands

- **vpkg_list**
- vpkg_versions <package_name>
- vpkg_info <package_name>
- vpkg_require <package_name>
- vpkg_devrequire package_name>
- vpkg_history
- vpkg_rollback [all/#]
- vpkg_help

```
[mkyle@login00.darwin ~]$ vpkg_list
Available packages:
  in /opt/shared/valet/2.1/etc
       alphafold
       amd-rocm
       amd-uprof
       anaconda
       aocl
       arpack
       atlas
       autoconf
       binutils
       blacs
       blis
       boost
       charm
       cmake
       cryosparc
       cuda
       cudnn
       ... many more ...
```

Software and VALET

VALET commands

- vpkg_list
- vpkg_versions <package_name>
- **vpkg_info <package_name>**
- vpkg_require <package_name>
- vpkg_devrequire package_name>
- vpkg_history
- vpkg_rollback [all/#]
- vpkg_help

```
[mkyle@login01.darwin ~]$ vpkg_info gcc
[gcc] {
  contexts: all
  flags: no-development-env
  ...SKIPPING LINES..
  http://gcc.gnu.org/
  GCC Compiler Suite
  prefix: /opt/shared/gcc
  source file: /opt/shared/valet/2.1/etc/gcc.vpkg_yaml
  default version: gcc/system
  versions: {
       [gcc/4.8] {
       contexts: all
       alias-to: gcc/4.8.5
       }
       [gcc/4.8.5] {
       contexts: all
       CentOS system GCC with C, C++, Obj-C, Obj-C++, and Fortran
       prefix: /opt/shared/gcc/4.8.5
  ...SKIPPING LINES..
```

Software and VALET

VALET commands

- vpkg_list
- vpkg_versions <package_name>
- vpkg_info <package_name>
- **vpkg_require <package_name>**
- vpkg_devrequire package_name>
- vpkg_history
- vpkg_rollback [all/#]
- vpkg_help

```
[mkyle@login01.darwin ~]$ gcc --version
gcc (GCC) 4.8.5 20150623 (Red Hat 4.8.5-39)
Copyright (C) 2015 Free Software Foundation, Inc.
This is free software; see the source for copying conditions.  There is
NO
warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR
PURPOSE.
[mkyle@login01.darwin ~]$
```

Software and VALET

VALET commands

- vpkg_list
- vpkg_versions <package_name>
- vpkg_info <package_name>
- **vpkg_require <package_name>**
- vpkg_devrequire package_name>
- vpkg_history
- vpkg_rollback [all/#]
- vpkg_help

```
[mkyle@login01.darwin ~]$ gcc --version
gcc (GCC) 4.8.5 20150623 (Red Hat 4.8.5-39)
Copyright (C) 2015 Free Software Foundation, Inc.
This is free software; see the source for copying conditions.  There is
NO
warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR
PURPOSE.
[mkyle@login01.darwin ~]$[(it_css:mkyle)@login01.darwin ~]$ vpkg_require
gcc/12.2.0:nvidia,openacc
Adding dependency `binutils/2.35.1` to your environment
Adding dependency `cuda/10.2.89-440.33.01` to your environment
Adding package `gcc/12.2.0:openacc,nvidia` to your environment
[mkyle@login01.darwin ~]$
```

Software and VALET

VALET commands

- vpkg_list
- vpkg_versions <package_name>
- vpkg_info <package_name>
- **vpkg_require <package_name>**
- vpkg_devrequire package_name>
- vpkg_history
- vpkg_rollback [all/#]
- vpkg_help

```
[mkyle@login01.darwin ~]$ gcc --version
gcc (GCC) 4.8.5 20150623 (Red Hat 4.8.5-39)
Copyright (C) 2015 Free Software Foundation, Inc.
This is free software; see the source for copying conditions.  There is
NO
warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR
PURPOSE.
[mkyle@login01.darwin ~]$[(it_css:mkyle)@login01.darwin ~]$ vpkg_require
gcc/12.2.0:nvidia,openacc
Adding dependency `binutils/2.35.1` to your environment
Adding dependency `cuda/10.2.89-440.33.01` to your environment
Adding package `gcc/12.2.0:openacc,nvidia` to your environment
[mkyle@login01.darwin ~]$ gcc --version
gcc (GCC) 12.2.0
Copyright (C) 2022 Free Software Foundation, Inc.
This is free software; see the source for copying conditions.  There is
NO
warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR
PURPOSE
[mkyle@login01.darwin ~]
```

Software and VALET

VALET commands

- vpkg_list
- vpkg_versions <package_name>
- vpkg_info <package_name>
- vpkg_require <package_name>
- vpkg_devrequire package_name>
- **vpkg_history**
- vpkg_rollback [all/#]
- vpkg_help

```
[mkyle@login01.darwin ~]$ vpkg_history
[standard]
  binutils/2.35.1
  cuda/10.2.89-440.33.01
  gcc/12.2.0:openacc,nvidia
[mkyle@login01.darwin ~]$
```

Software and VALET

VALET commands

- vpkg_list
- vpkg_versions <package_name>
- vpkg_info <package_name>
- vpkg_require <package_name>
- vpkg_devrequire package_name>
- vpkg_history
- **vpkg_rollback [all/#]**
- vpkg_help

```
[mkyle@login01.darwin ~]$ vpkg_history
[standard]
  binutils/2.35.1
  cuda/10.2.89-440.33.01
  gcc/12.2.0:openacc,nvidia
[mkyle@login01.darwin ~]$ vpkg_rollback 1
[mkyle@login01.darwin ~]$
```

Software and VALET

VALET commands

- vpkg_list
- vpkg_versions <package_name>
- vpkg_info <package_name>
- vpkg_require <package_name>
- vpkg_devrequire package_name>
- vpkg_history
- **vpkg_rollback [all/#]**
- vpkg_help

```
[mkyle@login01.darwin ~]$ vpkg_history
[standard]
  binutils/2.35.1
  cuda/10.2.89-440.33.01
  gcc/12.2.0:openacc,nvidia
[mkyle@login01.darwin ~]$ vpkg_rollback all
[mkyle@login01.darwin ~]$
```

Software and VALET

VALET commands

- vpkg_list
- vpkg_versions <package_name>
- vpkg_info <package_name>
- vpkg_require <package_name>
- vpkg_devrequire package_name>
- vpkg_history
- **vpkg_rollback [all/#]**
- vpkg_help

```
[mkyle@login01.darwin ~]$ vpkg_history
[standard]
  binutils/2.35.1
  cuda/10.2.89-440.33.01
  gcc/12.2.0:openacc,nvidia
[mkyle@login01.darwin ~]$ vpkg_rollback all
[mkyle@login01.darwin ~]$ vpkg_history

[mkyle@login01.darwin ~]$
```

Software and VALET

VALET commands

- vpkg_list
- vpkg_versions <package_name>
- vpkg_info <package_name>
- vpkg_require <package_name>
- vpkg_devrequire package_name>
- vpkg_history
- vpkg_rollback [all/#]
- **vpkg_help**

```
[mkyle@login00.darwin ~]$ vpkg_help
VALET(7)
VALET User Manual
      VALET(7)
NAME
      valet - VALET Automates Linux Environment Tasks
DESCRIPTION
      One particularly annoying aspect of cluster computing for most users is getting the environment setup properly given a set of requisite libraries and
      applications. Quite often in a cluster environment there are multiple versions of a library available at any one time, and knowing which to use when running a
      program can often mean the difference between predictable, correct results and a flawed or even failed execution. A piece of software called (quite simply)
      "modules" has been around for quite some time to address the complexities of environment configuration for users. The tool itself is fairly complex, both in
      the commands it offers the user and the manner by which environment modifications are specified (a modulefile). valet -- a recursive acronym for VALET
... More Lines Below...
```

Software and VALET

`vpkg_require`

- Configure environment for one or more VALET packages.

VS

`vpkg_devrequire`

- Configure environment for one or more VALET packages.
- Includes software development environment variables such as CPPFLAGS and LDFLAGS.

Running Jobs with Slurm

What is Slurm?

Slurm is a workload manager user to manage and control the computing resources for all jobs submitted to a cluster. This includes load balancing, reconciling requests for memory and processor cores with availability of those resources, suspending and restarting jobs, and managing jobs with different priorities.

In order to schedule any job (interactively or batch) on a cluster, you must set your workgroup to define your allocation workgroup and explicitly request a single partition.

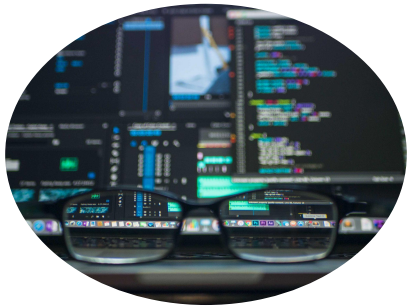

Caviness Partitions

Standard	Default partition if no `--partition` submission flag is specified; jobs can be preempted (killed)	• Maximum runtime of 7 days (default is 30 minutes) • Maximum number of CPUs per job = 360 • Maximum CPUs per user = 720
Devel	A partition with very short runtime limits and small resource limits; important to use for any development using compilers	• Maximum runtime of 2 hours (default is 30 minutes) • Each user can submit up to 2 jobs • Each job can use up to 4 cores on a single node

Workgroup	Partitions associated with specific kinds of compute equipment in the cluster purchased by a research `group «investing-entity»` (workgroup)	• Maximum runtime of 7 days (default is 30 minutes)
		• Per-workgroup resource limits (QOS) based on
		  ○ how many nodes your research group (workgroup) purchased (node=#)
		  ○ how many cores your research group (workgroup) purchased (cpu=#)
		  ○ how many GPUs your research group (workgroup) purchased (gres/gpu:<kind>=#)

DARWIN Partitions

standard	Contains all 48 standard memory nodes (64 cores, 512 GiB memory per node)	• Maximum run time of 7 days • Maximum of 400 jobs per user per partition • Default no preemption
large-mem	Contains all 32 large memory nodes (64 cores, 1024 GiB memory per node)	• Maximum run time of 7 days • Maximum of 400 jobs per user per partition • Default no preemption
xlarge-mem	Contains all 11 extra-large memory nodes (64 cores, 2048 GiB memory per node)	• Maximum run time of 7 days • Maximum of 400 jobs per user per partition • Default no preemption
extended-mem	Contains the single extended memory node (64 cores, 1024 GiB memory + 2.73 TiB NVMe swap)	• Maximum run time of 7 days • Maximum of 400 jobs per user per partition • Default no preemption

DARWIN Partitions Part II

gpu-t4	Contains all 9 NVIDIA Tesla T4 GPU nodes (64 cores, 512 GiB memory, 1 T4 GPU per node)	• Maximum run time of 7 days • Maximum of 400 jobs per user per partition • Default no preemption
gpu-v100	Contains all 3 NVIDIA Tesla V100 GPU nodes (48 cores, 768 GiB memory, 4 V100 GPUs per node)	• Maximum run time of 7 days • Maximum of 400 jobs per user per partition • Default no preemption
gpu-mi50	Contains the single AMD Radeon Instinct MI50 GPU node (64 cores, 512 GiB memory, 1 MI50 GPU)	• Maximum run time of 7 days • Maximum of 400 jobs per user per partition • Default no preemption
gpu-mi100	Contains the single AMD Radeon Instinct MI100 GPU node (64 cores, 512 GiB memory, 1 MI100 GPU)	• Maximum run time of 7 days • Maximum of 400 jobs per user per partition • Default no preemption
idle	**Contains all nodes in the cluster, jobs on this partition can be preempted but are not charged against your allocation**	• **Preemption is enabled for all jobs** • **Maximum of 320 jobs per user** • **Maximum of 640 CPUs per user (across all jobs in the partition)**

Running Jobs with Slurm

Running a Job?

Required to join workgroup

- **`workgroup -g <workgroup_name>`**

Job types

- Batch Jobs
  - `sbatch <job_script>`
- Interactive Jobs
  - `salloc`

```
[mkyle@login00.darwin ~]$ workgroup -g it_css
[(it_css:mkyle)@login00.darwin ~]$
```

Running Jobs with Slurm

Job types

- **Batch Jobs**

  sbatch <job_script>
  - **/opt/templates/slurm/generic/**
    - **serial.qs**
    - threads.qs
    - mpi/
  - /applications
    - RStudio-Server.qs
    - comsol.qs
    - gromacs.qs
    - matlab-mcr.qs
    - Several More…
- Interactive Jobs
  - salloc

```
[(it_css:mkyle)@login01.darwin ~]$ cp
/opt/templates/slurm/generic/serial.qs .
[(it_css:mkyle)@login01.darwin ~]$
```

Running Jobs with Slurm

Job types

- **Batch Jobs**

  sbatch <job_script>
  - **/opt/templates/slurm/generic/**
    - **serial.qs**
    - threads.qs
    - mpi/
  - /applications
    - RStudio-Server.qs
    - comsol.qs
    - gromacs.qs
    - matlab-mcr.qs
- Interactive Jobs
  - salloc

```
[(it_css:mkyle)@login01.darwin ~]$ cp
/opt/templates/slurm/generic/serial.qs  .
[(it_css:mkyle)@login01.darwin ~]$ cp serial.qs hpc_workshop/matlab_job.qs
[(it_css:mkyle)@login01.darwin ~]$ cd hpc_workshop
[(it_css:mkyle)@login01.darwin hpc_workshop]$
```

Running Jobs with Slurm

```
[(it_css:mkyle)@login01.darwin hpc_workshop]$ cat matlab_job.qs
#!/bin/bash -l
#
# DARWIN job script template, generated 2022-05-05T10:14:37-0400
#
# Sections of this script that can/should be edited are delimited by a
# [EDIT] tag.  All Slurm job options are denoted by a line that starts
# with "#SBATCH " followed by flags that would otherwise be passed on
# the command line.  Slurm job options can easily be disabled in a
# script by inserting a space in the prefix, e.g. "# SLURM " and
# reenabled by deleting that space.
#
# This is a batch job template for a program using a single processor
# core/thread (a serial job).
#
#SBATCH --nodes=1 --ntasks=1 --cpus-per-task=4
#
# [EDIT] All jobs have memory limits imposed.  The default is 1 GB per
#      CPU allocated to the job.  The default can be overridden either
#      with a per-node value (--mem) or a per-CPU value (--mem-per-cpu)
#      with unitless values in MB and the suffixes K|M|G|T denoting
#      kibi, mebi, gibi, and tebibyte units.  Delete the space between
#      the "#" and the word SBATCH to enable one of them:
#
# SBATCH --mem=8G
#SBATCH --mem-per-cpu=512M
#
# [EDIT] Each node in the cluster has local scratch disk of some sort
#      that is always mounted as /tmp.  Per-job temporary directories
#      are automatically created and destroyed by Slurm and can be
#      referenced via the $TMPDIR environment variable.  To ensure a
#      minimum amount of free space when your job is scheduled, the
```

Running Jobs with Slurm

```
#      --tmp option can be used; it has the same behavior unit-wise as
#      --mem and --mem-per-cpu.  Delete the space between the "#" and the
#      word SBATCH to enable:
#
#SBATCH --tmp=2G
#
# [EDIT] It can be helpful to provide a descriptive (terse) name for
#      the job (be sure to use quotes if there's whitespace in the
#      name):
#
#SBATCH --job-name=matlab_test
#
# [EDIT] The partition determines which nodes can be used and with what
#      maximum runtime limits, etc.  Partition limits can be displayed
#      with the "sinfo --summarize" command.
#
#      PLEASE NOTE:  On DARWIN every job is **required** to include the
#                  --partition flag in its submission!
#
#SBATCH --partition=idle
# [EDIT] Jobs that will run in one of the GPU partitions can request GPU
#      resources using ONE of the following flags:
#
#            --gpus=<count>
#            <count> GPUs total for the job, regardless of node count
#            --gpus-per-node=<count>
#            <count> GPUs are required on each node allocated to the job
#            --gpus-per-socket=<count>
#            <count> GPUs are required on each socket allocated to the
#            job
#            --gpus-per-task=<count>
#            <count> GPUs are required for each task in the job
```

Running Jobs with Slurm

```
#      PLEASE NOTE:  On DARWIN the --gres flag should NOT be used to
#                    request GPU resources.  The GPU type will be
#                    inferred from the partition to which the job is
#                    submitted if not specified.
#
# SBATCH --gpus=1
# SBATCH --gpus-per-task=1
# SBATCH --gpus-per-node=1
# SBATCH --gpus-per-socket=2
#
# [EDIT] The maximum runtime for the job; a single integer is interpreted
#      as a number of minutes, otherwise use the format
#
#             d-hh:mm:ss
#
#      Jobs default to the default runtime limit of the chosen partition
#      if this option is omitted.
#
#SBATCH --time=0-00:10:00
#
#      You can also provide a minimum acceptable runtime so the scheduler
#      may be able to run your job sooner.  If you do not provide a
#      value, it will be set to match the maximum runtime limit (discussed
#      above).
#
#SBATCH --time-min=0-00:01:00
#
# [EDIT] By default SLURM sends the job's stdout to the file "slurm-<jobid>.out"
#      and the job's stderr to the file "slurm-<jobid>.err" in the working
#      directory.  Override by deleting the space between the "#" and the
#      word SBATCH on the following lines; see the man page for sbatch for
#      special tokens that can be used in the filenames:
```

Running Jobs with Slurm

```
#SBATCH --output=%x-%j.out
#SBATCH --error=%x-%j.out
#
# [EDIT] Slurm can send emails to you when a job transitions through various
#      states: NONE, BEGIN, END, FAIL, REQUEUE, ALL, TIME_LIMIT,
#      TIME_LIMIT_50, TIME_LIMIT_80, TIME_LIMIT_90, ARRAY_TASKS.  One or more
#      of these flags (separated by commas) are permissible for the
#      --mail-type flag.  You MUST set your mail address using --mail-user
#      for messages to get off the cluster.
#
#SBATCH --mail-user='mkyle@udel.edu'
#SBATCH --mail-type=END,FAIL,TIME_LIMIT_90
#
# [EDIT] By default we DO NOT want to send the job submission environment
#      to the compute node when the job runs.
#
#SBATCH --export=NONE
#
#
# [EDIT] If you're not interested in how the job environment gets setup,
#      uncomment the following.
#
#UD_QUIET_JOB_SETUP=YES

#
# [EDIT] Define a Bash function and set this variable to its
#      name if you want to have the function called when the
#      job terminates (time limit reached or job preempted).
#
#      PLEASE NOTE:  when using a signal-handling Bash
#      function, any long-running commands should be prefixed
#      with UD_EXEC, e.g.
```

Running Jobs with Slurm

```
#
# [EDIT] Slurm only sets SLURM_MEM_PER_CPU when the --mem-per-cpu option is
#      used.  The job template system will attempt to set the missing
#      SLURM_MEM_PER_CPU when --mem was used and the job has a uniform number
#      of tasks per node (the only case when per-node memory yields a
#      uniform memory per task/cpu) if this variable is set:
#UD_PREFER_MEM_PER_CPU=YES
#
#      Uncomment the following variable if the job mandates a per-CPU memory
#      limit to be present or calculable when UD_PREFER_MEM_PER_CPU is set:
#UD_REQUIRE_MEM_PER_CPU=YES

#
# If you have VALET packages to load into the job environment,
# uncomment and edit the following line:
#
#vpkg_require intel/2019
vpkg_require matlab/2023b
#
# Do general job environment setup:
#
. /opt/shared/slurm/templates/libexec/common.sh

#
# [EDIT] Add your script statements hereafter, or execute a script or program
#      using the srun command.
#
#srun date
matlab -nodisplay -r "run('pendulum.m');exit;"
```

Running Jobs with Slurm

Job types

- **Batch Jobs**

  sbatch <job_script>
  - **/opt/templates/slurm/generic/**
    - **serial.qs**
    - threads.qs
    - mpi/
  - /applications
    - RStudio-Server.qs
    - comsol.qs
    - gromacs.qs
    - matlab-mcr.qs
- Interactive Jobs
  - salloc

```
[(it_css:mkyle)@login01.darwin hpc_workshop]$ ls
matlab_job.qs   pendulum.m   serial.qs
[(it_css:mkyle)@login01.darwin hpc_workshop]$
```

Running Jobs with Slurm

Job types

- **Batch Jobs**

  sbatch <job_script>
  - **/opt/templates/slurm/generic/**
    - **serial.qs**
    - threads.qs
    - mpi/
  - /applications
    - RStudio-Server.qs
    - comsol.qs
    - gromacs.qs
    - matlab-mcr.qs
- Interactive Jobs
  - salloc

```
[(it_css:mkyle)@login01.darwin hpc_workshop]$ ls
matlab_job.qs  pendulum.m  serial.qs
[(it_css:mkyle)@login01.darwin hpc_workshop]$ sbatch matlab_job.qs
Submitted batch job 5118101
[(it_css:mkyle)@login01.darwin hpc_workshop]$
```

Running Jobs with Slurm

Job types

- **Batch Jobs**

  sbatch <job_script>
  - **/opt/templates/slurm/generic/**
    - **serial.qs**
    - threads.qs
    - mpi/
  - /applications
    - RStudio-Server.qs
    - comsol.qs
    - gromacs.qs
    - matlab-mcr.qs
- Interactive Jobs
  - salloc

```
[(it_css:mkyle)@login01.darwin hpc_workshop]$ ls
matlab_job.qs  pendulum.m  serial.qs
[(it_css:mkyle)@login01.darwin hpc_workshop]$ sbatch matlab_job.qs
Submitted batch job 5118101
[(it_css:mkyle)@login01.darwin hpc_workshop]$ ls
matlab_job.qs  matlab_test-5118101.out  pendulum.m
pendulum_motion_plot.png  serial.qs
[(it_css:mkyle)@login01.darwin hpc_workshop]$
```

Running Jobs with Slurm

Job types

- **Batch Jobs**

  `sbatch <job_script>`
  - **`/opt/templates/slurm/generic/`**
    - **`serial.qs`**
    - `threads.qs`
    - `mpi/`
  - `/applications`
    - `RStudio-Server.qs`
    - `comsol.qs`
    - `gromacs.qs`
    - `matlab-mcr.qs`
- Interactive Jobs
  - `salloc`

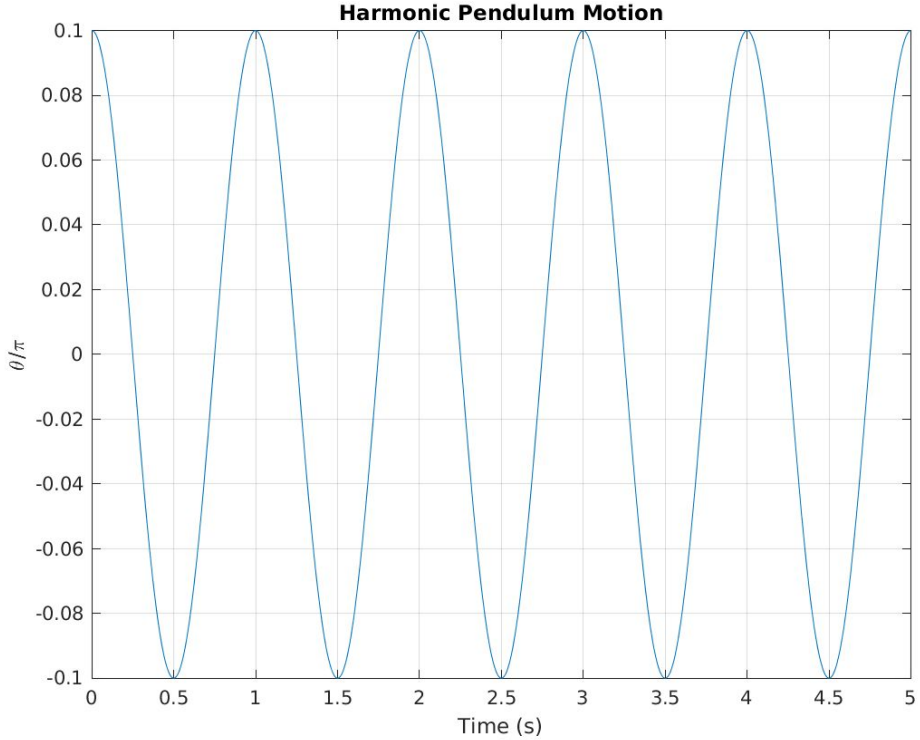

Running Jobs with Slurm

Run Job

Required to join workgroup

- `workgroup -g <workgroup_name>`

Job types

- Batch Jobs
  - `sbatch <job_script>`
- **Interactive Jobs**
  - **`salloc`**

```
[mkyle@login00.darwin ~]$ workgroup -g it_css
[(it_css:mkyle)@login01.darwin ~]$ salloc --partition=idle
```

Running Jobs with Slurm

Run Job

Required to join workgroup

- `workgroup -g <workgroup_name>`

Job types

- Batch Jobs
  - `sbatch <job_script>`
- **Interactive Jobs**
  - **`salloc`**

```
[mkyle@login00.darwin ~]$ workgroup -g it_css
[(it_css:mkyle)@login01.darwin ~]$ salloc --partition=idle
salloc: Granted job allocation 5089919
salloc: Waiting for resource configuration
salloc: Nodes r1n02 are ready for job
[(it_css:mkyle)@r1n02 ~]$
```

Running Jobs with Slurm

Run Job

Required to join workgroup

- `workgroup -g <workgroup_name>`

Job types

- Batch Jobs
  - `sbatch <job_script>`
- **Interactive Jobs**
  - **`salloc`**

```
[mkyle@login00.darwin ~]$ workgroup -g it_css
[(it_css:mkyle)@login01.darwin ~]$ salloc --partition=idle
salloc: Granted job allocation 5089919
salloc: Waiting for resource configuration
salloc: Nodes r1n02 are ready for job
[(it_css:mkyle)@r1n02 ~]$ sacct
       JobID  JobName  Partition  Account  AllocCPUS          State
ExitCode
------------ ---------- ---------- ---------- ---------- ----------
--------
5089919      interacti+    idle   it_css       1      RUNNING      0:0
5089919.int+ interacti+           it_css       1      RUNNING      0:0
5089919.ext+ extern               it_css       1      RUNNING      0:0
[(it_css:mkyle)@r1n02 ~]$
```

Running Jobs with Slurm

Run Job

Required to join workgroup

- `workgroup -g <workgroup_name>`

Job types

- Batch Jobs
  - `sbatch <job_script>`
- **Interactive Jobs**
  - **`salloc`**

```
[mkyle@login00.darwin ~]$ workgroup -g it_css
[(it_css:mkyle)@login01.darwin ~]$ salloc --partition=idle
salloc: Granted job allocation 5089919
salloc: Waiting for resource configuration
salloc: Nodes r1n02 are ready for job
[(it_css:mkyle)@r1n02 ~]$ sacct
       JobID  JobName  Partition  Account  AllocCPUS        State
ExitCode
------------ ---------- ---------- ---------- ---------- ----------
--------
5089919      interacti+    idle   it_css       1      RUNNING      0:0
5089919.int+ interacti+           it_css       1      RUNNING      0:0
5089919.ext+ extern               it_css       1      RUNNING      0:0
[(it_css:mkyle)@r1n02 ~]$exit
```

Running Jobs with Slurm

Run Job

Required to join workgroup

- `workgroup -g <workgroup_name>`

Job types

- Batch Jobs
  - `sbatch <job_script>`
- **Interactive Jobs**
  - **`salloc`**

```
[mkyle@login00.darwin ~]$ workgroup -g it_css
[(it_css:mkyle)@login01.darwin ~]$ salloc --partition=idle
salloc: Granted job allocation 5089919
salloc: Waiting for resource configuration
salloc: Nodes r1n02 are ready for job
[(it_css:mkyle)@r1n02 ~]$ sacct
       JobID  JobName  Partition  Account  AllocCPUS        State
ExitCode
------------ ---------- ---------- ---------- ---------- ----------
--------
5089919      interacti+    idle    it_css       1      RUNNING      0:0
5089919.int+ interacti+             it_css       1      RUNNING      0:0
5089919.ext+ extern                 it_css       1      RUNNING      0:0
[(it_css:mkyle)@r1n02 ~]$exit
logout
salloc: Relinquishing job allocation 5089919
salloc: Job allocation 5089919 has been revoked.
[(it_css:mkyle)@login01.darwin ~]$
```

Running Jobs with Slurm

Run Job

Required to join workgroup

- `workgroup -g <workgroup_name>`

Job types

- Batch Jobs
  - `sbatch <job_script>`
- **Interactive Jobs**
  - **`salloc`**

```
[mkyle@login00.darwin ~]$ workgroup -g it css
[(it css:mkyle)@login01.darwin ~]$ salloc --nodes=1 --ntasks-per-node=1
--cpus-per-task=4 --time=1:00:00 --partition=idle
salloc: Pending job allocation 5117900
salloc: job 5117900 queued and waiting for resources
salloc: job 5117900 has been allocated resources
salloc: Granted job allocation 5117900
salloc: Waiting for resource configuration
salloc: Nodes r1n03 are ready for job
[(it_css:mkyle)@r1n03 ~]$
```

Running Jobs with Slurm

Run Job

Required to join workgroup

- `workgroup -g <workgroup_name>`

Job types

- Batch Jobs
  - `sbatch <job_script>`
- **Interactive Jobs**
  - **`salloc`**

```
[mkyle@login00.darwin ~]$ workgroup -g it css
[(it css:mkyle)@login01.darwin ~]$ salloc --nodes=1 --ntasks-per-node=1
--cpus-per-task=4 --time=1:00:00 --partition=idle
salloc: Pending job allocation 5117900
salloc: job 5117900 queued and waiting for resources
salloc: job 5117900 has been allocated resources
salloc: Granted job allocation 5117900
salloc: Waiting for resource configuration
salloc: Nodes r1n03 are ready for job
[(it_css:mkyle)@r1n03 ~]$ vpgk_require python/2
Adding package `python/2.7.18` to your environment
[(it_css:mkyle)@r1n03 ~]$
```

Running Jobs with Slurm

Run Job

Required to join workgroup

- `workgroup -g <workgroup_name>`

Job types

- Batch Jobs
  - `sbatch <job_script>`
- **Interactive Jobs**
  - **`salloc`**

```
[(it_css:mkyle)@r1n03 ~]$ python threading_example.py
Thread Thread 1: 1Thread Thread 2: 1

Thread Thread 3: 1
Thread Thread 4: 1
Thread Thread 1: 2
 Thread Thread 3: 2
Thread Thread 2: 2
Thread Thread 4: 2
Thread Thread 3: 3Thread Thread 4: 3

 Thread Thread 2: 3
Thread Thread 1: 3
Thread Thread 3: 4
Thread Thread 4: 4
 Thread Thread 1: 4
Thread Thread 2: 4
Thread Thread 3: 5
 Thread Thread 1: 5
Thread Thread 2: 5
Thread Thread 4: 5
All 4 threads have finished execution.
[(it_css:mkyle)@r1n03 ~]$
```

Running Jobs with Slurm

Run Job

Required to join workgroup

- `workgroup -g <workgroup_name>`

Job types

- Batch Jobs
  - `sbatch <job_script>`
- **Interactive Jobs**
  - **`salloc`**

```
Thread Thread 4: 4
 Thread Thread 1: 4
Thread Thread 2: 4
Thread Thread 3: 5
 Thread Thread 1: 5
Thread Thread 2: 5
Thread Thread 4: 5
All 4 threads have finished execution.
[(it_css:mkyle)@r1n03 ~]$ sacct -j 5117900
JobID           JobName    Partition    Account    AllocCPUS      State  ExitCode
------------ ---------- --------- ---------- ---------- ------- --------
5117900      interacti+    idle    it_css       4      RUNNING      0:0
5117900.int+ interacti+            it_css       4      RUNNING      0:0
5117900.ext+ extern                it_css       4      RUNNING      0:0
[(it_css:mkyle)@r1n03 ~]$
```

Running Jobs with Slurm

Run Job

Required to join workgroup

- `workgroup -g <workgroup_name>`

Job types

- Batch Jobs
  - `sbatch <job_script>`
- **Interactive Jobs**
  - **`salloc`**

```
Thread Thread 4: 4
 Thread Thread 1: 4
Thread Thread 2: 4
Thread Thread 3: 5
 Thread Thread 1: 5
Thread Thread 2: 5
Thread Thread 4: 5
All 4 threads have finished execution.
[(it_css:mkyle)@r1n03 ~]$ sacct -j 5117900
JobID           JobName    Partition    Account    AllocCPUS      State  ExitCode
------------ ---------- ---------- ---------- ---------- ------- --------
5117900      interacti+    idle    it_css       4      RUNNING      0:0
5117900.int+ interacti+            it_css       4      RUNNING      0:0
5117900.ext+ extern                it_css       4      RUNNING      0:0
[(it_css:mkyle)@r1n03 ~]$ exit
logout
salloc: Relinquishing job allocation 5117900
[(it_css:mkyle)@login00.darwin ~]$
```

Running Jobs with Slurm

Manage Job

Slurm Commands:

- **squeue**
- ssqueue
- sacct
- ssacct

Command Details:

- Monitors a job while it is still active in Slurm.
- Provides information about:
  - Jobs status: queued, executing
  - Run time
  - CPU/GPU consummation of job

```
[(it_css:mkyle)@login00.darwin ~]$ squeue
```

Running Jobs with Slurm

Manage J

Slurm Com

- **sque**
- ssqu
- sacc
- ssac

Command

- Monit
  Slurm.
- Provid
  -
  -
  -

```
[(it_css:mkyle)@login00.darwin ~]$ squeue
             JOBID     PARTITION    NAME       USER    ST    TIME      NODES NODELIST(REASON)
           5091220       gpu-t4   interact  svalipou  R     5:20:31     1     r1t06
           5087532       gpu-t4   interact   bakerh   R   1-18:25:22    1     r1t05
           5090117       gpu-t4   interact   bakerh   R   20:57:14      1     r1t01
           5087531      gpu-v100  interact   bakerh   R   1-18:29:01    1     r2v01
           5092258      gpu-v100  interact  aanishp   R     10:24       1     r2v01
           5077716_267   idle       SCS     fortino   PD    0:00        1  (Resources)
           5077716_237   idle       SCS     fortino   PD    0:00        1  (Resources)
           5077716_271   idle       SCS     fortino   PD    0:00        1  (Resources)
           5077716_269   idle       SCS     fortino   PD    0:00        1  (Resources)
           5077716_268   idle       SCS     fortino   PD    0:00        1  (Resources)
           5077716_273   idle       SCS     fortino   PD    0:00        1  (Resources)
           5077716_272   idle       SCS     fortino   PD    0:00        1  (Resources)
           5077716_241   idle       SCS     fortino   PD    0:00        1  (Resources)
           5077716_277   idle       SCS     fortino   PD    0:00        1  (Resources)
           5077716_274   idle       SCS     fortino   PD    0:00        1  (Resources)
           5077716_242   idle       SCS     fortino   PD    0:00        1  (Resources)
```

Running Jobs with Slurm

Manage Jo

Slurm Comm

- **squeu**
- ssque
- sacct
- ssacc

Command D

- Monito
  Slurm.
- Provide
  - J
  - R
  - C

```
[(it_css:mkyle)@login00.darwin ~]$ squeue -j 5092258
             JOBID PARTITION     NAME     USER ST       TIME  NODES NODELIST(REASON)
           5092258  gpu-v100 interact    mkyle R      15:46      1 r2v01
[(it_css:mkyle)@login01.darwin ~]$
```

Running Jobs with Slurm

Manage Jo

Slurm Comm

- **squeu**
- ssque
- sacct
- ssacc

Command D

- Monito
  Slurm.
- Provide
  - J
  - R
  - C

```
[(it_css:mkyle)@login00.darwin ~]$ squeue -j 5092258
             JOBID PARTITION     NAME     USER ST       TIME  NODES NODELIST(REASON)
           5092258  gpu-v100 interact    mkyle R      15:46      1 r2v01
[(it_css:mkyle)@login01.darwin ~]$ squeue -j 5092258
--Format="jobid,name,state,partition,account,username,numnodes,numtasks,numcpus,gres"
JOBID       NAME        STATE     PARTITION  ACCOUNT    USER     NODES    TASKS    CPUS TRES_PER_NODE
5092258  interactive  RUNNING    gpu-v100    it_css    mkyle      1        1        4        N/A
```

Running Jobs with Slurm

Manage Job

Slurm Commands:

- `squeue`
- **`ssqueue`**
- `sacct`
- `ssacct`

Command Details:

- Monitors a job while it is still active in Slurm.
- Provides information about:
  - Jobs status: queued, executing
  - Run time
  - CPU/GPU consummation of job
  - **Column and Table Format**

```
[(it_css:mkyle)@login00.darwin ~]$ ssqueue -j 5092258
```

Running Jobs with Slurm

Manage J

Slurm Com

- sque
- **ssqu**
- sacc
- ssac

Command

- Monit Slurm.
- Provid
  -
  -
  -
  -

JOBID	PARTITION	NAME	USER	ST	TIME	NODES	NODELIST(REASON)
5092258	gpu-v100	interactive	mkyle	R	43:39	1	r2v01

[Q]uit [P]rev/[N]ext page Page [L]eft/[R]ight [E]nd/[B]eginning of list

Running Jobs with Slurm

Manage Job

Slurm Commands:

- `squeue`
- **`ssqueue`**
- `sacct`
- `ssacct`

Command Details:

- Monitors a job while it is still active in Slurm.
- Provides information about:
  - Jobs status: queued, executing
  - Run time
  - CPU/GPU consummation of job
  - **Column and Table Format**

```
[(it_css:mkyle)@login00.darwin ~]$ ssqueue -j 5092258
[(it_css:mkyle)@login01.darwin ~]$ ssqueue -j 5092258
--Format="jobid,name,state,partition,account,username,numnodes,numtasks,n
umcpus,gres"
```

Running Jobs with Slurm

JOBID	NAME	STATE	PARTITION	ACCOUNT	USER	NODES	TASKS	CPUS	TRES_PER_NODE
5092258	interactive	RUNNING	gpu-v100	jayaraman_lab	aanishp	1	1	4	N/A

[Q]uit [P]rev/[N]ext page Page [L]eft/[R]ight [E]nd/[B]eginning of list

Running Jobs with Slurm

Manage Job

Slurm Commands:

- `squeue`
- `ssqueue`
- **`sacct`**
- `ssacct`

Command Details:

- Information about a job from history
  - Job ID required
- Provides information about:
  - Job state: canceled, completed, preempted
  - Partition
  - CPU/GPU/Memory used

```
[(it_css:mkyle)@login00.darwin ~]$ sacct -j 5091988
JobID   JobName Partition  Account AllocCPUS  State  ExitCode
------- ------- --------- -------- ---------- ------ --------
5091988 serial_te+  idle   it_css       1    COMPLETED  0:0
5091988 .bat+       batch  it_css       1    COMPLETED  0:0
5091988 .ext+       extern it_css       1    COMPLETED  0:0
5091988 .0          date   it_css       1    COMPLETED  0:0
[(it_css:mkyle)@login00.darwin ~]$
```

Running Jobs with Slurm

Manage Job

Slurm Commands:

- `squeue`
- `ssqueue`
- **`sacct`**
- `ssacct`

Command Details:

- Information about a job from history
  - Job ID required
- Provides information about:
  - Job state: canceled, completed, preempted
  - Partition
  - CPU/GPU/Memory used

```
[(it_css:mkyle)@login00.darwin ~]$ sacct -j 5091988
--format="user,jobid,jobname,partition,state,time,start,end,elap
sed,MaxRss,MaxVMSize,nnodes,ncpus,nodelist"
```

Manag
Slurm C
- sc
- ss
- **sa**
- ss

Comma
- Int
- Pr

```
[(it_css:mkyle)@login00.darwin ~]$ sacct -j 5091988
--format="user,jobid,jobname,partition,state,time,start,end,elap
sed,MaxRss,MaxVMSize,nnodes,ncpus,nodelist"
User    JobID  JobName  Partition  State  Timelimit    Start   End   Elapsed MaxRSS  MaxVMSize  NNodes NCPUS NodeList
------ ----- ------- --------- ------ ---------- ----- ---- -------- ------- ---------- ------ ---- --------
mkyle 5091988 serial_te+ idle  COMPLETED  00:01:00  11:15:22 11:15:24 00:00:02  1    1                      r1n07
      5091988.bat+       batch COMPLETED           11:15:22 11:15:24 00:00:02  0    169808K   1      1    r1n07
      5091988.ext+     extern  COMPLETED           11:15:22 11:15:24 00:00:02  0    4364K     1      1    r1n07
```

Running Jobs with Slurm

Manage Job

Slurm Commands:

- `squeue`
- `ssqueue`
- `sacct`
- **`ssacct`**

Command Details:

- Information about a job from history
  - Job ID required
- Provides information about:
  - Job state: canceled, completed, preempted
  - Partition
  - CPU/GPU/Memory used
  - **Column and Table Format**

```
[(it_css:mkyle)@login00.darwin ~]$ ssacct -j 5091988
```

Running Jobs with Slurm

Manage Job

Slurm Commands:

- `squeue`
- `ssqueue`
- `sacct`
- **`ssacct`**

Command Details:

- Information about a job from history
  - Job ID required
- Provides information about:
  - Job state: canceled, completed, preempted
  - Partition
  - CPU/GPU/Memory used
  - **Column and Table Format**

```
JobID          | JobName     | Partition | Account | AllocC
---------------|-------------|-----------|---------|-------
5091988        | serial_test | idle      | it_css  |
5091988.batch  | batch       |           | it_css  |
5091988.extern | extern      |           | it_css  |
5091988.0      | date        |           | it_css  |

 [Q]uit   [P]rev/[N]ext page   Page [L]eft/[R]ight [E]nd
```

Running Jobs with Slurm

Manage Job

Slurm Commands:

- `squeue`
- `ssqueue`
- `sacct`
- **`ssacct`**

Command Details:

- Information about a job from history
  - Job ID required
- Provides information about:
  - Job state: canceled, completed, preempted
  - Partition
  - CPU/GPU/Memory used
  - **Column and Table Format**

```
[mkyle@login00.darwin ~]$ ssacct -j 5091988
[mkyle@login00.darwin ~]$ ssacct -j 5091988
--format="user,jobid,jobname,partition,state,time,start,end,elapsed,nnode
s,ncpus,nodelist"
```

Running Jobs with Slurm

User	JobID	JobName	Partition	State	Timelimit	Start	End	Elapsed	NNodes	NCPUS	NodeLis
mkyle	5091988	test	idle	COMPLETED	00:01:00	2024-04-02T11:15:22	2024-04-02T11:15:24	00:00:02	1	1	r1n07
	.batch	batch		COMPLETED		2024-04-02T11:15:22	2024-04-02T11:15:24	00:00:02	1	1	r1n07
	.extern	extern		COMPLETED		2024-04-02T11:15:22	2024-04-02T11:15:24	00:00:02	1	1	r1n07
	.0	date		COMPLETED		2024-04-02T11:15:24	2024-04-02T11:15:24	00:00:00	1	1	r1n07

[Q]uit   [P]rev/[N]ext page   Page [L]eft/[R]ight   [E]nd/[B]eginning of list

Thank you!

For more information, contact:

askit@udel.edu

Helpful HPC Resources

Scan 4 url's

hpc-ask

caviness-users

darwin-users

hpc-wiki

Help Desk
askit@udel.edu

Id faucibus nisl tincidunt eget.

- Id faucibus nisl tincidunt eget.
- Id faucibus nisl tincidunt eget.
- Id faucibus nisl tincidunt eget.
- Id faucibus nisl tincidunt eget.

Data Management

HPC Focused Checklist

- Data Production
  - Types of Data
  - Reproducibility
  - Tools & Software
- Data Size
  - How Much
  - Rate of Growth
  - Frequency of Change
- Data Usage
  - Who Needs Access

- Privacy & Security
  - Data Classification Level
- Data Documentation
  - Organization & Naming
- Storage and Backup