

Getting Started on DARWIN

IT-RCI
Fall 2024

Getting Started on DARWIN

Required Prerequisites

- Introduction to Linux Workshop
- Active DARWIN Account
 - UD or ACCESS
- SSH Client installed on your system (and know how to use it)
 - i.e. PuTTY for Windows

Suggested Prerequisites

- Familiarity with Coding
 - python
 - C / C++
 - Fortran
- Introduction to HPC at UD

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




Connecting with SSH



Software and VALET



Joining Workgroups



Setting up Jobs



File Systems & Storage



SUs & Accounting



Collecting Data



Running and Monitoring
Jobs with Slurm

Connecting with SSH

How to connect with SSH

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

```
mky1e@ITRC-G92D633:~$ ssh <user_name>@darwin.hpc.udel.edu
```

Connecting with SSH

How to connect with SSH

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

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

Connecting with SSH

How to connect with SSH

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

```
mkyle@ITRC-G92D633:~$ ssh mkyle@darwin.hpc.udel.edu
Last login: Fri Sep 20 10:20:29 2024 from net10-7-191-11.host.udel.edu
[mkyle@login00.darwin ~]$
```

Joining Your Workgroup

How you join your workgroup?

- Finding workgroup(s)
 - ◆ `workgroup --menu`
 - ◆ `workgroup -q workgroups`
- Joining Workgroups
 - ◆ `workgroup -g`
`[workgroup_name]`

```
[mkyle@login00.darwin ~]$ workgroup --menu
```


Joining Your Workgroup

How you join your workgroup?

- Finding workgroup(s)
 - ◆ `workgroup --menu`
 - ◆ `workgroup -q workgroups`
- Joining Workgroups
 - ◆ `workgroup -g [workgroup_name]`

```
[AVAILABLE WORKGROUPS]
1002  it  css
1102  unsponsored

COMMAND TO EXECUTE

[ ] Also change to group work directory
$WORKDIR = /lustre/it_css

Commands: [S]pawn workgroup shell [C]hange to work directory [Q]uit (also
use <TAB> to move between the menu and command areas
```

Joining Your Workgroup

How you join your workgroup?

- Finding workgroup(s)
 - ◆ `workgroup --menu`
 - ◆ `workgroup -q workgroups`
- Joining Workgroups
 - ◆ `workgroup -g`
[workgroup_name]

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

Joining Your Workgroup

How you join your workgroup?

- Finding workgroup(s)
 - ◆ `workgroup --menu`
 - ◆ `workgroup -q workgroups`
- Joining Workgroups
 - ◆ `workgroup -g`
[workgroup_name]

```
[mkyle@login00.darwin ~]$ workgroup --menu
[(it_css:mkyle)@login00.darwin ~]$ workgroup -q workgroups
1002  it_css
1102  unsponsored
[(it_css:mkyle)@login00.darwin ~]$
```

Joining Your Workgroup

How you join your workgroup?

→ Finding workgroup(s)

- ◆ `workgroup --menu`
- ◆ `workgroup -q workgroups`

→ Joining Workgroups

- ◆ `workgroup -g`
`[workgroup_name]`

```
[mkyle@login00.darwin ~]$ workgroup --menu
[(it_css:mkyle)@login00.darwin ~]$ workgroup -q workgroups
1002  it_css
1102  unsponsored
[mkyle@login00.darwin ~]$ workgroup -g <workgroup_name>
```

Joining Your Workgroup

How you join your workgroup?

- Finding workgroup(s)
 - ◆ `workgroup --menu`
 - ◆ `workgroup -q workgroups`
- Joining Workgroups
 - ◆ `workgroup -g [workgroup_name]`

```
[mkyle@login00.darwin ~]$ workgroup --menu
[(it_css:mkyle)@login00.darwin ~]$ workgroup -q workgroups
1002  it_css
1102  unsponsored
[mkyle@login00.darwin ~]$ workgroup -g it_css
```

Joining Your Workgroup

How you join your workgroup?

→ Finding workgroup(s)

- ◆ `workgroup --menu`
- ◆ `workgroup -q workgroups`

→ Joining Workgroups

- ◆ `workgroup -g`
`[workgroup_name]`

```
[mkyle@login00.darwin ~]$ workgroup --menu
[(it_css:mkyle)@login00.darwin ~]$ workgroup -q workgroups
1002  it_css
1102  unsponsored
[mkyle@login00.darwin ~]$ workgroup -g it_css
[(it_css:mkyle)@login00.darwin ~]$
```

File Systems & Storage



Home Storage

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



Lustre Storage

- 1.1PB
- Omni-path Infiniband

General Information

- > 13.5 TB Mem
- 100 GBit/s
- **No Back Ups!!**



Workgroup Storage

- ≥ 1 TB
- Shared within workgroup
- /lustre/<workgroup_name>/
- \$WORKDIR



Compute Node

- < 2TB
- Temporary job storage
- \$TMPDIR

File Systems & Storage

- Workgroup
 - ◆ `$WORKDIR`
 - ◆ `/lustre/<workgroup_name>`
- Home
 - ◆ Login Landing path
 - ◆ `$HOME`
 - ◆ `~/`
- Storage Usage Levels
 - ◆ `df -h $HOME`
 - ◆ `my_quotas`

```
[(it_css:mkyale)@login01.darwin ~]$ echo $WORKDIR
```


File Systems & Storage

- Workgroup
 - ◆ \$WORKDIR
 - ◆ /lustre/<workgroup_name>
- Home
 - ◆ Login Landing path
 - ◆ \$HOME
 - ◆ ~/
- Storage Usage Levels
 - ◆ df -h \$HOME
 - ◆ my_quotas

```
[(it_css:mkyale)@login01.darwin ~]$ echo $WORKDIR
/lustre/it_css
[(it_css:mkyale)@login01.darwin ~]$
```

File Systems & Storage

- Workgroup
 - ◆ \$WORKDIR
 - ◆ /lustre/<workgroup_name>
- Home
 - ◆ Login Landing path
 - ◆ \$HOME
 - ◆ ~/
- Storage Usage Levels
 - ◆ df -h \$HOME
 - ◆ my_quotas

```
[(it_css:mkyle)@login01.darwin ~]$ echo $WORKDIR
/lustre/it_css
[(it_css:mkyle)@login01.darwin ~]$ cd $WORKDIR
```

File Systems & Storage

- Workgroup
 - ◆ \$WORKDIR
 - ◆ /lustre/<workgroup_name>
- Home
 - ◆ Login Landing path
 - ◆ \$HOME
 - ◆ ~/
- Storage Usage Levels
 - ◆ df -h \$HOME
 - ◆ my_quotas

```
[(it_css:mkyle)@login01.darwin ~]$ echo $WORKDIR
/lustre/it_css
[(it_css:mkyle)@login01.darwin ~]$ cd $WORKDIR
[(it_css:mkyle)@login01.darwin it_css]$
```

File Systems & Storage

→ Workgroup

- ◆ \$WORKDIR
- ◆ /lustre/<workgroup_name>

→ Home

- ◆ Login Landing path
- ◆ **\$HOME**
- ◆ ~/

→ Storage Usage Levels

- ◆ df -h \$HOME
- ◆ my_quotas

```
[(it_css:mkyale)@login01.darwin ~]$ echo $WORKDIR
/lustre/it_css
[(it_css:mkyale)@login01.darwin ~]$ cd $WORKDIR
[(it_css:mkyale)@login01.darwin it_css]$ echo $HOME
```

File Systems & Storage

→ Workgroup

- ◆ \$WORKDIR
- ◆ /lustre/<workgroup_name>

→ Home

- ◆ Login Landing path
- ◆ **\$HOME**
- ◆ ~/

→ Storage Usage Levels

- ◆ df -h \$HOME
- ◆ my_quotas

```
[(it_css:mkyale)@login01.darwin ~]$ echo $WORKDIR
/lustre/it_css
[(it_css:mkyale)@login01.darwin ~]$ cd $WORKDIR
[(it_css:mkyale)@login01.darwin it_css]$ echo $HOME
/home/2179
[(it_css:mkyale)@login01.darwin it_css]$
```

File Systems & Storage

→ Workgroup

- ◆ \$WORKDIR
- ◆ /lustre/<workgroup_name>

→ Home

- ◆ Login Landing path
- ◆ **\$HOME**
- ◆ ~/

→ Storage Usage Levels

- ◆ df -h \$HOME
- ◆ my_quotas

```
[(it_css:mkyle)@login01.darwin ~]$ echo $WORKDIR
/lustre/it_css
[(it_css:mkyle)@login01.darwin ~]$ cd $WORKDIR
[(it_css:mkyle)@login01.darwin it_css]$ echo $HOME
/home/2179
[(it_css:mkyle)@login01.darwin it_css]$ cd $HOME
```

File Systems & Storage

→ Workgroup

- ◆ \$WORKDIR
- ◆ /lustre/<workgroup_name>

→ Home

- ◆ Login Landing path
- ◆ **\$HOME**
- ◆ ~/

→ Storage Usage Levels

- ◆ df -h \$HOME
- ◆ my_quotas

```
[(it_css:mkyale)@login01.darwin ~]$ echo $WORKDIR
/lustre/it_css
[(it_css:mkyale)@login01.darwin ~]$ cd $WORKDIR
[(it_css:mkyale)@login01.darwin it_css]$ echo $HOME
/home/2179
[(it_css:mkyale)@login01.darwin it_css]$ cd $HOME
[(it_css:mkyale)@login01.darwin ~]$
```

File Systems & Storage

→ Workgroup

- ◆ \$WORKDIR
- ◆ /lustre/<workgroup_name>

→ Home

- ◆ Login Landing path
- ◆ \$HOME
- ◆ ~/

→ Storage Usage Levels

- ◆ **df -h \$HOME**
- ◆ my_quotas

```
[(it_css:mkyale)@login01.darwin ~]$ echo $WORKDIR
/lustre/it_css
[(it_css:mkyale)@login01.darwin ~]$ cd $WORKDIR
[(it_css:mkyale)@login01.darwin it_css]$ echo $HOME
/home/2179
[(it_css:mkyale)@login01.darwin it_css]$ cd $HOME
[(it_css:mkyale)@login01.darwin ~]$ df -h $HOME
```


File Systems & Storage

→ Workgroup

- ◆ \$WORKDIR
- ◆ /lustre/<workgroup_name>

→ Home

- ◆ Login Landing path
- ◆ \$HOME
- ◆ ~/

→ Storage Usage Levels

- ◆ **df -h \$HOME**
- ◆ my_quotas

```
[(it_css:mkyale)@login01.darwin ~]$ echo $WORKDIR
/lustre/it_css
[(it_css:mkyale)@login01.darwin ~]$ cd $WORKDIR
[(it_css:mkyale)@login01.darwin it_css]$ echo $HOME
/home/2179
[(it_css:mkyale)@login01.darwin it_css]$ cd $HOME
[(it_css:mkyale)@login01.darwin ~]$ df -h $HOME
Filesystem                Size Used Avail Use% Mounted on
nfs0-ib:/beagle/home/2179  20G  6.7G  14G   34% /home/2179
[(it_css:mkyale)@login01.darwin ~]$
```

File Systems & Storage

→ Workgroup

- ◆ \$WORKDIR
- ◆ /lustre/<workgroup_name>

→ Home

- ◆ Login Landing path
- ◆ \$HOME
- ◆ ~/

→ Storage Usage Levels

- ◆ df -h \$HOME
- ◆ **my_quotas**

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

File Systems & Storage

→ Workgroup

- ◆ \$WORKDIR
- ◆ /lustre/<workgroup_name>

→ Home

- ◆ Login Landing path
- ◆ \$HOME
- ◆ ~/

→ Storage Usage Levels

- ◆ df -h \$HOME
- ◆ my_quotas

```
[(it_css:mkyle)@login01.darwin ~]$ echo $WORKDIR
/lustre/it_css
[(it_css:mkyle)@login01.darwin ~]$ cd $WORKDIR
[(it_css:mkyle)@login01.darwin it_css]$ echo $HOME
/home/2179
[(it_css:mkyle)@login01.darwin it_css]$ cd $HOME
[(it_css:mkyle)@login01.darwin ~]$ df -h $HOME
Filesystem                Size      Used Avail Use% Mounted on
nfs0-ib:/beagle/home/2179 20G   6.7G   14G   34% /home/2179
[(it_css:mkyle)@login01.darwin ~]$ my_quotas
Type Path                In-use / kiB Available / kiB Pct
----
user  /home/2179            6953984      20971520  33%
group /lustre/it_css        1625031711   3026866824  53%
[(it_css:mkyle)@login01.darwin ~]$
```

File Systems & Storage

Exercise #1

- Home
 - ◆ path?
 - ◆ storage amount?
- Workgroup
 - ◆ path?
 - ◆ storage amount?

```
mky1e@ITRC-G92D633:~$
```

File Systems & Storage

Exercise #1

- Home
 - ◆ path?
 - ◆ storage amount?
- Workgroup
 - ◆ path?
 - ◆ storage amount?

```
mky1e@ITRC-G92D633:~$
```

3:00

File Systems & Storage

Exercise #1

- **Home**
 - ◆ path?
 - ◆ storage amount?
- **Workgroup**
 - ◆ path?
 - ◆ storage amount?

```
[(it_css:mkyale)@login01.darwin ~]$ echo $HOME  
/home/2179  
[(it_css:mkyale)@login01.darwin ~]$
```

File Systems & Storage

Exercise #1

- **Home**
 - ◆ path?
 - ◆ storage amount?
- **Workgroup**
 - ◆ path?
 - ◆ storage amount?

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

File Systems & Storage

Exercise #1

- Home
 - ◆ path?
 - ◆ storage amount?
- Workgroup
 - ◆ path?
 - ◆ storage amount

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


File Systems & Storage

Exercise #1

- Home
 - ◆ path?
 - ◆ storage amount?
- Workgroup
 - ◆ path?
 - ◆ storage amount?

```
[(it_css:mkyale)@login01.darwin ~]$ echo $HOME
/home/2179
[(it_css:mkyale)@login01.darwin ~]$ df -h $HOME
Filesystem                Size      Used Avail Use% Mounted on
nfs0-ib:/beagle/home/2179 20G   6.7G   14G   34% /home/2179
[(it_css:mkyale)@login01.darwin ~]$ echo $WORKDIR
/lustre/it_css
[(it_css:mkyale)@login01.darwin ~]$ my_quotas
Type  Path                In-use / kiB  Available / kiB  Pct
-----
user  /home/2179          6953984      20971520         33%
group /lustre/it_css     1625031711   3026866824       53%
[(it_css:mkyale)@login01.darwin ~]$
```

Collecting Data On DARWIN

Data Transfer Methods

- Basic
 - ◆ **scp**
 - ◆ wget
 - ◆ curl
 - ◆ rsync
- Advanced
 - ◆ rclone
 - ◆ Globus

```
[(it_css:mkyale)@login01.darwin ~]$ cd $HOME
```

Collecting Data On DARWIN

Data Transfer Methods

- Basic
 - ◆ **scp**
 - ◆ wget
 - ◆ curl
 - ◆ rsync
- Advanced
 - ◆ rclone
 - ◆ Globus

```
[(it_css:mkyale)@login01.darwin ~]$ cd $HOME  
[(it_css:mkyale)@login01.darwin ~]$ mkdir workshop
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ **scp**
- ◆ wget
- ◆ curl
- ◆ rsync

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkye)@login01.darwin ~]$ mkdir workshop  
[(it_css:mkye)@login01.darwin ~]$ cd workshop  
[(it_css:mkye)@login01.darwin workshop]$
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ **scp**
- ◆ wget
- ◆ curl
- ◆ rsync

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkye)@login01.darwin ~]$ mkdir workshop  
[(it_css:mkye)@login01.darwin ~]$ cd workshop  
[(it_css:mkye)@login01.darwin workshop]$
```

Collecting Data On DARWIN

Data Transfer Methods

- Basic
 - ◆ **scp**
 - ◆ wget
 - ◆ curl
 - ◆ rsync
- Advanced
 - ◆ rclone
 - ◆ Globus

```
[(it_css:mkyale)@login01.darwin ~]$ mkdir workshop
[(it_css:mkyale)@login01.darwin ~]$ cd workshop
[(it_css:mkyale)@login01.darwin workshop]$ scp <file_to_copy> <destination>
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ **scp**
- ◆ wget
- ◆ curl
- ◆ rsync

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyale)@login01.darwin ~]$ mkdir workshop
[(it_css:mkyale)@login01.darwin ~]$ cd workshop
[(it_css:mkyale)@login01.darwin workshop]$ scp <file_to_copy> <destination>
```

Note:

SCP can work between different systems, IE your laptop to a cluster, or to/from cluster to cluster

```
$ scp mkyale@caviness.hpc.udel.edu:~/file1.txt ~/
$ scp fileX.txt mkyale@darwin.hpc.udel.edu:~/workshop
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ **scp**
- ◆ wget
- ◆ curl
- ◆ rsync

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkye)@login01.darwin ~]$ mkdir workshop
[(it_css:mkye)@login01.darwin ~]$ cd workshop
[(it_css:mkye)@login01.darwin workshop]$ scp
/opt/shared/help/DARWIN-Workshop/DARWIN_WrkShp.txt .
```


Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ **scp**
- ◆ wget
- ◆ curl
- ◆ rsync

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyale)@login01.darwin ~]$ mkdir workshop
[(it_css:mkyale)@login01.darwin ~]$ cd workshop
[(it_css:mkyale)@login01.darwin workshop]$ scp
/opt/shared/help/DARWIN-Workshop/DARWIN_WrkShp.txt .
[(it_css:mkyale)@login01.darwin workshop]$ ls
DARWIN_WrkShp.txt
[(it_css:mkyale)@login01.darwin workshop]$
```

Collecting Data On DARWIN

Exercise #1

- Open a second SSH session on DARWIN
- Use a text editor (nano, vim, etc) to open `DARWIN_WrkShp.txt`

```
mky1e@ITRC-G92D633:~$
```



Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ **wget**
- ◆ curl
- ◆ rsync

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyale)@login01.darwin ~]$ mkdir workshop
[(it_css:mkyale)@login01.darwin ~]$ cd workshop
[(it_css:mkyale)@login01.darwin workshop]$ scp
/opt/shared/help/DARWIN-Workshop/DARWIN_WrkShp.txt .
[(it_css:mkyale)@login01.darwin workshop]$ ls
DARWIN_WrkShp.txt
[(it_css:mkyale)@login01.darwin workshop]$
[(it_css:mkyale)@login01.darwin workshop]$ wget <URL_TO_FILE>
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ **wget**
- ◆ curl
- ◆ rsync

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyale)@login01.darwin ~]$ mkdir workshop
[(it_css:mkyale)@login01.darwin ~]$ cd workshop
[(it_css:mkyale)@login01.darwin workshop]$ scp
/opt/shared/help/DARWIN-Workshop/DARWIN_WrkShp.txt .
[(it_css:mkyale)@login01.darwin workshop]$ ls
DARWIN_WrkShp.txt
[(it_css:mkyale)@login01.darwin workshop]$
[(it_css:mkyale)@login01.darwin ~]$ wget
https://www.udel.edu/content/dam/udelImages/main/graphics/udelLogoImages/
logo-udel.png
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ curl
- ◆ rsync

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyale)@login01.darwin ~]$ mkdir workshop
[(it_css:mkyale)@login01.darwin ~]$ cd workshop
[(it_css:mkyale)@login01.darwin workshop]$ scp
/opt/shared/help/DARWIN-Workshop/DARWIN_WrkShp.txt .
[(it_css:mkyale)@login01.darwin workshop]$ ls
DARWIN_WrkShp.txt
[(it_css:mkyale)@login01.darwin workshop]$
[(it_css:mkyale)@login01.darwin ~]$ wget
https://www.udel.edu/content/dam/udelImages/main/graphics/udelLogoImages/
logo-udel.png
--2024-10-02 10:49:43--
https://www.udel.edu/content/dam/udelImages/main/graphics/udelLogoImages/
logo-udel.png
Resolving www.udel.edu (www.udel.edu)... 3.227.14.55, 34.207.18.216
Connecting to www.udel.edu (www.udel.edu)|3.227.14.55|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 31627 (31K) [image/png]
Saving to: 'logo-udel.png.1'

100% [=====]
=====>] 31,627      --.-K/s   in 0s

2024-10-02 10:49:43 (419 MB/s) - 'logo-udel.png' saved [31627/31627]
[(it_css:mkyale)@login01.darwin ~]$
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ **wget**
- ◆ curl
- ◆ rsync

→ Advanced

- ◆ rclone
- ◆ Globus

```
Resolving www.udel.edu (www.udel.edu)... 3.227.14.55, 34.207.18.216
Connecting to www.udel.edu (www.udel.edu)|3.227.14.55|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 31627 (31K) [image/png]
Saving to: 'logo-udel.png'
```

```
100%[=====] 31,627 --.-K/s in 0s
```

```
2024-10-02 10:49:43 (419 MB/s) - 'logo-udel.png' saved [31627/31627]
[(it_css:mkyale)@login01.darwin ~]$ ls
DARWIN_WrkShp.txt  logo-udel.png
[(it_css:mkyale)@login01.darwin workshop]$
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ **curl**
- ◆ rsync

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyale)@login01.darwin ~]$ ls
DARWIN_WrkShp.txt  logo-udel.png
[(it_css:mkyale)@login01.darwin workshop]$ mkdir sst_data
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ **curl**
- ◆ rsync

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyale)@login01.darwin ~]$ ls
DARWIN_WrkShp.txt  logo-udel.png
[(it_css:mkyale)@login01.darwin workshop]$ mkdir sst_data
[(it_css:mkyale)@login01.darwin workshop]$ cd sst_data
[(it_css:mkyale)@login01.darwin sst_data]$
```


Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ **curl**
- ◆ rsync

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyle)@login01.darwin ~]$ ls
DARWIN_WrkShp.txt  logo-udel.png
[(it_css:mkyle)@login01.darwin workshop]$ mkdir sst_data
[(it_css:mkyle)@login01.darwin workshop]$ cd sst_data
[(it_css:mkyle)@login01.darwin sst_data]$ curl
https://downloads.psl.noaa.gov/Datasets/noaa.oisst.v2.highres/sst.day.ano
m.[1981-1982].nc -O
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ **curl**
- ◆ rsync

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyale)@login01.darwin sst_data]$ curl  
https://downloads.psl.noaa.gov/Datasets/noaa.oisst.v2.highres/sst.day.anom.  
m.[1981-1982].nc -O
```

```
[1/2]:
```

```
https://downloads.psl.noaa.gov/Datasets/noaa.oisst.v2.highres/sst.day.anom.  
m.1981.nc --> sst.day.anom.1981.nc  
--_curl_--https://downloads.psl.noaa.gov/Datasets/noaa.oisst.v2.highres/s  
st.day.anom.1981.nc
```

% Total	% Received	% Xferd	Average Speed	Time	Time	Time
Current						

			Dload	Upload	Total	Spent	Left
Speed							

100	182M	100	182M	0	0	3790k	0	0:00:49	0:00:49
-----	------	-----	------	---	---	-------	---	---------	---------

```
--:--:-- 35.2M
```

```
[2/2]:
```

```
https://downloads.psl.noaa.gov/Datasets/noaa.oisst.v2.highres/sst.day.anom.  
m.1982.nc --> sst.day.anom.1982.nc  
--_curl_--https://downloads.psl.noaa.gov/Datasets/noaa.oisst.v2.highres/s  
st.day.anom.1982.nc
```

100	547M	100	547M	0	0	11.5M	0	0:00:47	0:00:47
-----	------	-----	------	---	---	-------	---	---------	---------

```
--:--:-- 51.1M
```

```
[(it_css:mkyale)@login01.darwin sst_data]$
```

Collecting Data On DARWIN

Data Transfer Methods

- Basic
 - ◆ scp
 - ◆ wget
 - ◆ curl
 - ◆ **rsync**
- Advanced
 - ◆ rclone
 - ◆ Globus

```
[(it_css:mkyale)@login01.darwin sst_data]$ cd $HOME  
[(it_css:mkyale)@login01.darwin ~]$
```

Collecting Data On DARWIN

Data Transfer Methods

- Basic
 - ◆ scp
 - ◆ wget
 - ◆ curl
 - ◆ **rsync**
- Advanced
 - ◆ rclone
 - ◆ Globus

```
[(it_css:mkye)@login01.darwin sst_data]$ cd $HOME
[(it_css:mkye)@login01.darwin ~]$ mkdir sample
[(it_css:mkye)@login01.darwin ~]$
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ curl
- ◆ **rsync**

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkye)@login01.darwin sst_data]$ cd $HOME
[(it_css:mkye)@login01.darwin ~]$ mkdir sample
[(it_css:mkye)@login01.darwin ~]$ cd sample
[(it_css:mkye)@login01.darwin sample]$
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ curl
- ◆ **rsync**

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkye)@login01.darwin sst_data]$ cd $HOME
[(it_css:mkye)@login01.darwin ~]$ mkdir sample
[(it_css:mkye)@login01.darwin ~]$ cd sample
[(it_css:mkye)@login01.darwin sample]$ cp
/opt/shared/help/DARWIN-Workshop/fileCreate.sh ./
[(it_css:mkye)@login01.darwin sample]$
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ curl
- ◆ **rsync**

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyale)@login01.darwin sst_data]$ cd $HOME
[(it_css:mkyale)@login01.darwin ~]$ mkdir sample
[(it_css:mkyale)@login01.darwin ~]$ cd sample
[(it_css:mkyale)@login01.darwin sample]$ cp
/opt/shared/help/DARWIN-Workshop/fileCreate.sh ./
[(it_css:mkyale)@login01.darwin sample]$ ls
fileCreate.sh
[(it_css:mkyale)@login01.darwin sample]$
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ curl
- ◆ **rsync**

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyale)@login01.darwin sst_data]$ cd $HOME
[(it_css:mkyale)@login01.darwin ~]$ mkdir sample
[(it_css:mkyale)@login01.darwin ~]$ cd sample
[(it_css:mkyale)@login01.darwin sample]$ cp
/opt/shared/help/DARWIN-Workshop/fileCreate.sh ./
[(it_css:mkyale)@login01.darwin sample]$ ls
fileCreate.sh
[(it_css:mkyale)@login01.darwin sample]$ chmod 744 fileCreate.sh
[(it_css:mkyale)@login01.darwin sample]$
```


Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ curl
- ◆ **rsync**

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyle)@login01.darwin sst_data]$ cd $HOME
[(it_css:mkyle)@login01.darwin ~]$ mkdir sample
[(it_css:mkyle)@login01.darwin ~]$ cd sample
[(it_css:mkyle)@login01.darwin sample]$ cp
/opt/shared/help/DARWIN-Workshop/fileCreate.sh ./
[(it_css:mkyle)@login01.darwin sample]$ ls
fileCreate.sh
[(it_css:mkyle)@login01.darwin sample]$ chmod 744 fileCreate.sh
[(it_css:mkyle)@login01.darwin sample]$ sh fileCreate.sh
[(it_css:mkyle)@login01.darwin sample]$
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ curl
- ◆ **rsync**

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyale)@login01.darwin sst_data]$ cd $HOME
[(it_css:mkyale)@login01.darwin ~]$ mkdir sample
[(it_css:mkyale)@login01.darwin ~]$ cd sample
[(it_css:mkyale)@login01.darwin sample]$ cp
/opt/shared/help/DARWIN-Workshop/fileCreate.sh ./
[(it_css:mkyale)@login01.darwin sample]$ ls
fileCreate.sh
[(it_css:mkyale)@login01.darwin sample]$ chmod 744 fileCreate.sh
[(it_css:mkyale)@login01.darwin sample]$ sh fileCreate.sh
[(it_css:mkyale)@login01.darwin sample]$ ls
file1.txt    file14.txt  file2.txt   file25.txt  file30.txt  file36.txt
file41.txt  file47.txt  file52.txt  file58.txt  file63.txt  file69.txt
file74.txt  ... MORE FILES REMOVED FOR SPACING ... fileCreate.sh
file13.txt  file19.txt  file24.txt  file3.txt   file35.txt  file40.txt
file46.txt  file51.txt  file57.txt  file62.txt  file68.txt  file73.txt
file79.txt  file84.txt  file9.txt   file95.txt
[(it_css:mkyale)@login01.darwin sample]$
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ curl
- ◆ **rsync**

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyle)@login01.darwin sst_data]$ cd $HOME
[(it_css:mkyle)@login01.darwin ~]$ mkdir sample
[(it_css:mkyle)@login01.darwin ~]$ cd sample
[(it_css:mkyle)@login01.darwin sample]$ cp
/opt/shared/help/DARWIN-Workshop/fileCreate.sh ./
[(it_css:mkyle)@login01.darwin sample]$ ls
fileCreate.sh
[(it_css:mkyle)@login01.darwin sample]$ chmod 744 fileCreate.sh
[(it_css:mkyle)@login01.darwin sample]$ sh fileCreate.sh
[(it_css:mkyle)@login01.darwin sample]$ ls
file1.txt    file14.txt  file2.txt   file25.txt  file30.txt  file36.txt
file41.txt  file47.txt  file52.txt  file58.txt  file63.txt  file69.txt
file74.txt  ... MORE FILES REMOVED FOR SPACING ... fileCreate.sh
file13.txt  file19.txt  file24.txt  file3.txt   file35.txt  file40.txt
file46.txt  file51.txt  file57.txt  file62.txt  file68.txt  file73.txt
file79.txt  file84.txt  file9.txt   file95.txt
[(it_css:mkyle)@login01.darwin sample]$ cd ../workshop/
[(it_css:mkyle)@login01.darwin workshop]$
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ curl
- ◆ **rsync**

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyle)@login01.darwin sst_data]$ cd $HOME
[(it_css:mkyle)@login01.darwin ~]$ mkdir sample
[(it_css:mkyle)@login01.darwin ~]$ cd sample
[(it_css:mkyle)@login01.darwin sample]$ cp
/opt/shared/help/DARWIN-Workshop/fileCreate.sh ./
[(it_css:mkyle)@login01.darwin sample]$ ls
fileCreate.sh
[(it_css:mkyle)@login01.darwin sample]$ chmod 744 fileCreate.sh
[(it_css:mkyle)@login01.darwin sample]$ sh fileCreate.sh
[(it_css:mkyle)@login01.darwin sample]$ ls
file1.txt    file14.txt  file2.txt   file25.txt  file30.txt  file36.txt
file41.txt  file47.txt  file52.txt  file58.txt  file63.txt  file69.txt
file74.txt  ... MORE FILES REMOVED FOR SPACING ... fileCreate.sh
file13.txt  file19.txt  file24.txt  file3.txt   file35.txt  file40.txt
file46.txt  file51.txt  file57.txt  file62.txt  file68.txt  file73.txt
file79.txt  file84.txt  file9.txt   file95.txt
[(it_css:mkyle)@login01.darwin sample]$ cd ../workshop/
[(it_css:mkyle)@login01.darwin workshop]$ mkdir dataFiles
[(it_css:mkyle)@login01.darwin workshop]$
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ curl
- ◆ **rsync**

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyle)@login01.darwin sst_data]$ cd $HOME
[(it_css:mkyle)@login01.darwin ~]$ mkdir sample
[(it_css:mkyle)@login01.darwin ~]$ cd sample
[(it_css:mkyle)@login01.darwin sample]$ cp
/opt/shared/help/DARWIN-Workshop/fileCreate.sh ./
[(it_css:mkyle)@login01.darwin sample]$ ls
fileCreate.sh
[(it_css:mkyle)@login01.darwin sample]$ chmod 744 fileCreate.sh
[(it_css:mkyle)@login01.darwin sample]$ sh fileCreate.sh
[(it_css:mkyle)@login01.darwin sample]$ ls
file1.txt    file14.txt  file2.txt   file25.txt  file30.txt  file36.txt
file41.txt  file47.txt  file52.txt  file58.txt  file63.txt  file69.txt
file74.txt  ... MORE FILES REMOVED FOR SPACING ... fileCreate.sh
file13.txt  file19.txt  file24.txt  file3.txt   file35.txt  file40.txt
file46.txt  file51.txt  file57.txt  file62.txt  file68.txt  file73.txt
file79.txt  file84.txt  file9.txt   file95.txt
[(it_css:mkyle)@login01.darwin sample]$ cd ../workshop/
[(it_css:mkyle)@login01.darwin workshop]$ mkdir dataFiles
[(it_css:mkyle)@login01.darwin workshop]$ cd dataFiles
[(it_css:mkyle)@login01.darwin dataFiles]$
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ curl
- ◆ **rsync**

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyle)@login01.darwin sample]$ ls
fileCreate.sh
[(it_css:mkyle)@login01.darwin sample]$ chmod 744 fileCreate.sh
[(it_css:mkyle)@login01.darwin sample]$ sh fileCreate.sh
[(it_css:mkyle)@login01.darwin sample]$ ls
file1.txt      file14.txt  file2.txt   file25.txt  file30.txt  file36.txt
file41.txt    file47.txt  file52.txt  file58.txt  file63.txt  file69.txt
file74.txt    ... MORE FILES REMOVED FOR SPACING ... fileCreate.sh
file13.txt    file19.txt  file24.txt  file3.txt   file35.txt  file40.txt
file46.txt    file51.txt  file57.txt  file62.txt  file68.txt  file73.txt
file79.txt    file84.txt  file9.txt   file95.txt
[(it_css:mkyle)@login01.darwin sample]$ cd ../workshop/
[(it_css:mkyle)@login01.darwin workshop]$ mkdir dataFiles
[(it_css:mkyle)@login01.darwin workshop]$ cd dataFiles
[(it_css:mkyle)@login01.darwin dataFiles]$ rsync -azP ~/sample/*.txt ./
file71.txt
0 100% 0.00kB/s      0:00:00 (xfr#70, to-chk=30/100)
file72.txt
0 100% 0.00kB/s      0:00:00 (xfr#71, to-chk=29/100)
file73.txt ... MORE FILES REMOVED FOR SPACING ...
0 100% 0.00kB/s      0:00:00 (xfr#99, to-chk=1/100)
file99.txt
0 100% 0.00kB/s      0:00:00 (xfr#100, to-chk=0/100)
[(it_css:mkyle)@login01.darwin dataFiles]$
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ curl
- ◆ **rsync**

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyale)@login01.darwin dataFiles]$ ls
file1.txt   file18.txt  file27.txt  file36.txt  file45.txt  file54.txt
file63.txt  file72.txt  file81.txt  file90.txt
file10.txt  file19.txt  file28.txt  file37.txt  file46.txt  file55.txt
file64.txt  file73.txt  file82.txt  file91.txt
file100.txt file2.txt   file29.txt
... MORE FILES REMOVED FOR SPACING ...
[(it_css:mkyale)@login01.darwin dataFiles]$ cd ~/sample/
[(it_css:mkyale)@login01.darwin sample]$
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ curl
- ◆ **rsync**

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyle)@login01.darwin dataFiles]$ ls
file1.txt   file18.txt  file27.txt  file36.txt  file45.txt  file54.txt
file63.txt  file72.txt  file81.txt  file90.txt
file10.txt  file19.txt  file28.txt  file37.txt  file46.txt  file55.txt
file64.txt  file73.txt  file82.txt  file91.txt
file100.txt file2.txt   file29.txt
... MORE FILES REMOVED FOR SPACING ...
[(it_css:mkyle)@login01.darwin dataFiles]$ cd ~/sample/
[(it_css:mkyle)@login01.darwin sample]$ nano file1.txt
[(it_css:mkyle)@login01.darwin sample]$ nano file2.txt
[(it_css:mkyle)@login01.darwin sample]$ nano file100.txt
[(it_css:mkyle)@login01.darwin sample]$
```


Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ curl
- ◆ **rsync**

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyle)@login01.darwin dataFiles]$ ls
file1.txt   file18.txt  file27.txt  file36.txt  file45.txt  file54.txt
file63.txt  file72.txt  file81.txt  file90.txt
file10.txt  file19.txt  file28.txt  file37.txt  file46.txt  file55.txt
file64.txt  file73.txt  file82.txt  file91.txt
file100.txt file2.txt   file29.txt
... MORE FILES REMOVED FOR SPACING ...
[(it_css:mkyle)@login01.darwin dataFiles]$ cd ~/sample/
[(it_css:mkyle)@login01.darwin sample]$ nano file1.txt
[(it_css:mkyle)@login01.darwin sample]$ nano file2.txt
[(it_css:mkyle)@login01.darwin sample]$ nano file100.txt
[(it_css:mkyle)@login01.darwin sample]$ cd ../workshop/dataFiles
[(it_css:mkyle)@login01.darwin dataFiles]$
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ curl
- ◆ **rsync**

→ Advanced

- ◆ rclone
- ◆ Globus

```
[(it_css:mkyle)@login01.darwin dataFiles]$ ls
file1.txt   file18.txt  file27.txt  file36.txt  file45.txt  file54.txt
file63.txt  file72.txt  file81.txt  file90.txt
file10.txt  file19.txt  file28.txt  file37.txt  file46.txt  file55.txt
file64.txt  file73.txt  file82.txt  file91.txt
file100.txt file2.txt   file29.txt
... MORE FILES REMOVED FOR SPACING ...
[(it_css:mkyle)@login01.darwin dataFiles]$ cd ~/sample/
[(it_css:mkyle)@login01.darwin sample]$ nano file1.txt
[(it_css:mkyle)@login01.darwin sample]$ nano file2.txt
[(it_css:mkyle)@login01.darwin sample]$ nano file100.txt
[(it_css:mkyle)@login01.darwin sample]$ cd ../workshop/dataFiles
[(it_css:mkyle)@login01.darwin dataFiles]$ rsync -azP ~/sample/*.txt .
sending incremental file list
file1.txt
              7 100% 0.00kB/s      0:00:00 (xfr#1, to-chk=99/100)
file100.txt
             12 100%  11.72kB/s 0:00:00 (xfr#2, to-chk=97/100)
file2.txt
             10 100%   9.77kB/s   0:00:00 (xfr#3, to-chk=87/100)
[(it_css:mkyle)@login01.darwin dataFiles]$
```

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ curl
- ◆ 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

Collecting Data On DARWIN

Data Transfer Methods

→ Basic

- ◆ scp
- ◆ wget
- ◆ curl
- ◆ 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

```
[(it_css:mkyle)@login01.darwin workshop]$ cd  
[(it_css:mkyle)@login01.darwin ~]$
```

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

```
[(it_css:mkyle)@login01.darwin workshop]$ cd
[(it_css:mkyle)@login01.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_history
- vpkg_rollback [all/#]
- vpkg_help

```
[(it_css:mkyale)@login01.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_history`
- `vpkg_rollback [all/#]`
- `vpkg_help`

```
[(it_css:mkyale)@login01.darwin ~]$ vpkg_versions intel-python
Available versions in package (* = default version):
```

```
[/opt/shared/valet/2.1/etc/intel-python.vpkg_yaml]
intel-python      Intel Distribution for Python
  2021             alias to intel-python/2021.4.0
  2021.1.1        2021 update 1 (Python 3.7)
  2021.1.1:unaltered 2021 update 1 (Python 3.7) ORIGINAL INSTALL
  2021.4.0        2021 update 4 (Python 3.7)
* 2022           alias to intel-python/2022.1.0
  2022.1.0       2022 update 1 (Python 3.9)
```


Software and VALET

VALET commands

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

```
[(it_css:mkyale)@login01.darwin ~]$ vpkg_info intel-python
[intel-python] {
  contexts: all
  flags: no-development-env,no-standard-paths
  actions: {
    executable(source, succeedThenFail, success=0) {
      sh : /opt/shared/valet/2.1/libexec/intel-python.sh
    } (contexts: all)
    SHADOW_PREFIX=${VALET_PATH_PREFIX} (contexts: all)
    INTEL_DASH_PYTHON_PREFIX=${VALET_PATH_PREFIX} (contexts:
development)
  }
  Intel Distribution for Python
  prefix: /opt/shared/intel-python
  source file: /opt/shared/valet/2.1/etc/intel-python.vpkg_yaml
  default version: intel-python/2022
  versions: {
    ...SKIPPING LINES...
    [intel-python/2022.1.0] {
      contexts: all
      2022 update 1 (Python 3.9)
      prefix: /opt/shared/intel-python/2022.1.0
    }
  }
}
```

Software and VALET

VALET commands

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

```
[(it_css:mkyale)@login01.darwin ~]$ python --version
Python 2.7.5
[(it_css:mkyale)@login01.darwin ~]$
```

Software and VALET

VALET commands

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

```
[(it_css:mkyale)@login01.darwin ~]$ python --version
Python 2.7.5
[(it_css:mkyale)@login01.darwin ~]$ vpkg_require intel-python/2022.1.0
Adding package `intel-python/2022.1.0` to your environment
[(it_css:mkyale)@login01.darwin ~]$
```

Software and VALET

VALET commands

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

```
[(it_css:mkyale)@login01.darwin ~]$ python --version
Python 2.7.5
[(it_css:mkyale)@login01.darwin ~]$ vpkg_require intel-python/2022.1.0
Adding package `intel-python/2022.1.0` to your environment
[(it_css:mkyale)@login01.darwin ~]$ vpkg_history
[standard]
  intel-python/2022.1.0
[(it_css:mkyale)@login01.darwin ~]$
```

Software and VALET

VALET commands

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

```
[(it_css:mkyale)@login01.darwin ~]$ python --version
Python 2.7.5
[(it_css:mkyale)@login01.darwin ~]$ vpkg_require intel-python/2022.1.0
Adding package `intel-python/2022.1.0` to your environment
[(it_css:mkyale)@login01.darwin ~]$ vpkg_history
[standard]
intel-python/2022.1.0
[(it_css:mkyale)@login01.darwin ~]$ vpkg_require gcc
Adding package `gcc/4.8.5` to your environment
[(it_css:mkyale)@login01.darwin ~]$
```

Software and VALET

VALET commands

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

```
[(it_css:mkyale)@login01.darwin ~]$ python --version
Python 2.7.5
[(it_css:mkyale)@login01.darwin ~]$ vpkg_require intel-python/2022.1.0
Adding package `intel-python/2022.1.0` to your environment
[(it_css:mkyale)@login01.darwin ~]$ vpkg_history
[standard]
intel-python/2022.1.0
[(it_css:mkyale)@login01.darwin ~]$ vpkg_require gcc
Adding package `gcc/4.8.5` to your environment
[(it_css:mkyale)@login01.darwin ~]$ vpkg_history
[standard]
intel-python/2022.1.0
[standard]
gcc/4.8.5
[(it_css:mkyale)@login01.darwin ~]$
```

Software and VALET

VALET commands

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

```
[(it_css:mkyale)@login01.darwin ~]$ python --version
Python 2.7.5
[(it_css:mkyale)@login01.darwin ~]$ vpkg_require intel-python/2022.1.0
Adding package `intel-python/2022.1.0` to your environment
[(it_css:mkyale)@login01.darwin ~]$ vpkg_history
[standard]
intel-python/2022.1.0
[(it_css:mkyale)@login01.darwin ~]$ vpkg_require gcc
Adding package `gcc/4.8.5` to your environment
[(it_css:mkyale)@login01.darwin ~]$ vpkg_history
[standard]
intel-python/2022.1.0
[standard]
gcc/4.8.5
[(it_css:mkyale)@login01.darwin ~]$ vpkg_rollback 1
[(it_css:mkyale)@login01.darwin ~]$
```

Software and VALET

VALET commands

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

```
[(it_css:mkyale)@login01.darwin ~]$ python --version
Python 2.7.5
[(it_css:mkyale)@login01.darwin ~]$ vpkg_require intel-python/2022.1.0
Adding package `intel-python/2022.1.0` to your environment
[(it_css:mkyale)@login01.darwin ~]$ vpkg_history
[standard]
intel-python/2022.1.0
[(it_css:mkyale)@login01.darwin ~]$ vpkg_require gcc
Adding package `gcc/4.8.5` to your environment
[(it_css:mkyale)@login01.darwin ~]$ vpkg_history
[standard]
intel-python/2022.1.0
[standard]
gcc/4.8.5
[(it_css:mkyale)@login01.darwin ~]$ vpkg_rollback 1
[(it_css:mkyale)@login01.darwin ~]$ vpkg_history
[standard]
intel-python/2022.1.0
[(it_css:mkyale)@login01.darwin ~]$
```


Software and VALET

VALET commands

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

```
[(it_css:mkyle)@login01.darwin ~]$ python --version
Python 2.7.5
[(it_css:mkyle)@login01.darwin ~]$ vpkg_require intel-python/2022.1.0
Adding package `intel-python/2022.1.0` to your environment
[(it_css:mkyle)@login01.darwin ~]$ vpkg_history
[standard]
intel-python/2022.1.0
[(it_css:mkyle)@login01.darwin ~]$ vpkg_require gcc
Adding package `gcc/4.8.5` to your environment
[(it_css:mkyle)@login01.darwin ~]$ vpkg_history
[standard]
intel-python/2022.1.0
gcc/4.8.5
[(it_css:mkyle)@login01.darwin ~]$ vpkg_rollback 1
[(it_css:mkyle)@login01.darwin ~]$ vpkg_history
[standard]
intel-python/2022.1.0
[(it_css:mkyle)@login01.darwin ~]$ vpkg_rollback all
[(it_css:mkyle)@login01.darwin ~]$
```

Software and VALET

VALET commands

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

```
[(it_css:mkyle)@login01.darwin ~]$ python --version
Python 2.7.5
[(it_css:mkyle)@login01.darwin ~]$ vpkg_require intel-python/2022.1.0
Adding package `intel-python/2022.1.0` to your environment
[(it_css:mkyle)@login01.darwin ~]$ vpkg_history
[standard]
intel-python/2022.1.0
[(it_css:mkyle)@login01.darwin ~]$ vpkg_require gcc
Adding package `gcc/4.8.5` to your environment
[(it_css:mkyle)@login01.darwin ~]$ vpkg_history
[standard]
intel-python/2022.1.0
gcc/4.8.5
[(it_css:mkyle)@login01.darwin ~]$ vpkg_rollback 1
[(it_css:mkyle)@login01.darwin ~]$ vpkg_history
[standard]
intel-python/2022.1.0
[(it_css:mkyle)@login01.darwin ~]$ vpkg_rollback all
[(it_css:mkyle)@login01.darwin ~]$ vpkg_help
```

Software and VALET

VALET commands

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

VALET(7)
Manual

VALET User
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 Automates Linux Environment Tasks -- is an alternative that strives to be as simple as possible to

Manual page valet(7) line 1 (press h for help or q to quit)

Setting Up Jobs

Job Templates

→ `/opt/shared/templates/slurm/`

- ◆ `applications/`
 - `R`
 - `comsol`
 - `abaqus`
 - `gaussian`
 - `lammmps`
 - `matlab`
 - `tensorflow`
- ◆ `generic/`
 - `serial.qs`
 - `thread.qs`
 - `mpi/`

```
[(it_css:mkyale)@login01.darwin ~]$ cd /opt/shared/templates/slurm/  
[(it_css:mkyale)@login01.darwin slurm]$
```

Setting Up Jobs

Job Templates

- /opt/shared/templates/slurm/
 - ◆ applications/
 - R
 - comsol
 - abaqus
 - gaussian
 - lammmps
 - matlab
 - tensorflow
 - ◆ generic/
 - serial.qs
 - thread.qs
 - mpi/

```
[(it_css:mkyale)@login01.darwin ~]$ cd /opt/shared/templates/slurm/  
[(it_css:mkyale)@login01.darwin slurm]$ ls  
README.md applications generic  
[(it_css:mkyale)@login01.darwin slurm]$
```

Setting Up Jobs

Job Templates

→ /opt/shared/templates/slurm/

- ◆ applications/
 - R
 - comsol
 - abaqus
 - gaussian
 - lammmps
 - matlab
 - tensorflow
- ◆ generic/
 - serial.qs
 - thread.qs
 - mpi/

```
[(it_css:mkyale)@login01.darwin ~]$ cd /opt/shared/templates/slurm/  
[(it_css:mkyale)@login01.darwin slurm]$ ls  
README.md  applications  generic  
[(it_css:mkyale)@login01.darwin slurm]$ cd applications  
[(it_css:mkyale)@login01.darwin applications]$
```

Setting Up Jobs

Job Templates

→ /opt/shared/templates/slurm/

- ◆ applications/
 - R
 - comsol
 - abaqus
 - gaussian
 - lammps
 - matlab
 - tensorflow
- ◆ generic/
 - serial.qs
 - thread.qs
 - mpi/

```
[(it_css:mkyale)@login01.darwin ~]$ cd /opt/shared/templates/slurm/  
[(it_css:mkyale)@login01.darwin slurm]$ ls  
README.md applications generic  
[(it_css:mkyale)@login01.darwin slurm]$ cd applications  
[(it_css:mkyale)@login01.darwin applications]$ ls  
RStudio-Server.qs abaqus.qs comsol.qs gaussian.qs gromacs.qs  
lammps.qs matlab-mcr.qs matlab.qs nwchem.qs tensorflow.qs  
[(it_css:mkyale)@login01.darwin applications]$
```

Setting Up Jobs

Job Templates

→ /opt/shared/templates/slurm/

- ◆ applications/
 - R
 - comsol
 - abaqus
 - gaussian
 - lammmps
 - matlab
 - tensorflow
- ◆ generic/
 - **serial.qs**
 - **thread.qs**
 - **mpi/**

```
[(it_css:mkyale)@login01.darwin ~]$ cd /opt/shared/templates/slurm/
[(it_css:mkyale)@login01.darwin slurm]$ ls
README.md  applications  generic
[(it_css:mkyale)@login01.darwin slurm]$ cd applications
[(it_css:mkyale)@login01.darwin applications]$ ls
RStudio-Server.qs  abaqus.qs  comsol.qs  gaussian.qs  gromacs.qs
lammmps.qs  matlab-mcr.qs  matlab.qs  nwchem.qs  tensorflow.qs
[(it_css:mkyale)@login01.darwin applications]$ cd ../generic/
[(it_css:mkyale)@login01.darwin generic]$
```


Setting Up Jobs

Job Templates

→ /opt/shared/templates/slurm/

- ◆ applications/
 - R
 - comsol
 - abaqus
 - gaussian
 - lammmps
 - matlab
 - tensorflow
- ◆ generic/
 - **serial.qs**
 - **thread.qs**
 - **mpi/**

```
[(it_css:mkyale)@login01.darwin ~]$ cd /opt/shared/templates/slurm/
[(it_css:mkyale)@login01.darwin slurm]$ ls
README.md  applications  generic
[(it_css:mkyale)@login01.darwin slurm]$ cd applications
[(it_css:mkyale)@login01.darwin applications]$ ls
RStudio-Server.qs  abaqus.qs  comsol.qs  gaussian.qs  gromacs.qs
lammmps.qs  matlab-mcr.qs  matlab.qs  nwchem.qs  tensorflow.qs
[(it_css:mkyale)@login01.darwin applications]$ cd ../generic/
[(it_css:mkyale)@login01.darwin generic]$ ls
mpi  serial.qs  threads.qs
[(it_css:mkyale)@login01.darwin generic]$
```

Setting Up Jobs

Job Templates

→ /opt/shared/templates/slurm/

- ◆ applications/
 - R
 - comsol
 - abaqus
 - gaussian
 - lammmps
 - matlab
 - tensorflow
- ◆ generic/
 - **serial.qs**
 - thread.qs
 - mpi/

```
[(it_css:mkyale)@login01.darwin ~]$ cd /opt/shared/templates/slurm/
[(it_css:mkyale)@login01.darwin slurm]$ ls
README.md  applications  generic
[(it_css:mkyale)@login01.darwin slurm]$ cd applications
[(it_css:mkyale)@login01.darwin applications]$ ls
RStudio-Server.qs  abaqus.qs  comsol.qs  gaussian.qs  gromacs.qs
lammmps.qs  matlab-mcr.qs  matlab.qs  nwchem.qs  tensorflow.qs
[(it_css:mkyale)@login01.darwin applications]$ cd ../generic/
[(it_css:mkyale)@login01.darwin generic]$ ls
mpi  serial.qs  threads.qs
[(it_css:mkyale)@login01.darwin generic]$ mkdir ~/workshop/job && cp
serial.qs ~/workshop/job/
[(it_css:mkyale)@login01.darwin generic]$
```

Setting Up Jobs

Job Templates

→ /opt/shared/templates/slurm/

- ◆ applications/
 - R
 - comsol
 - abaqus
 - gaussian
 - lammmps
 - matlab
 - tensorflow
- ◆ generic/
 - **serial.qs**
 - thread.qs
 - mpi/

```
[(it_css:mkyale)@login01.darwin ~]$ cd /opt/shared/templates/slurm/
[(it_css:mkyale)@login01.darwin slurm]$ ls
README.md  applications  generic
[(it_css:mkyale)@login01.darwin slurm]$ cd applications
[(it_css:mkyale)@login01.darwin applications]$ ls
RStudio-Server.qs  abaqus.qs  comsol.qs  gaussian.qs  gromacs.qs
lammmps.qs  matlab-mcr.qs  matlab.qs  nwchem.qs  tensorflow.qs
[(it_css:mkyale)@login01.darwin applications]$ cd ../generic/
[(it_css:mkyale)@login01.darwin generic]$ ls
mpi  serial.qs  threads.qs
[(it_css:mkyale)@login01.darwin generic]$ mkdir ~/workshop/job && cp
serial.qs ~/workshop/job/
[(it_css:mkyale)@login01.darwin generic]$ cd ~/workshop/job
[(it_css:mkyale)@login01.darwin job]$
```

Setting Up Jobs

Job Templates

→ /opt/shared/templates/slurm/

- ◆ applications/
 - R
 - comsol
 - abaqus
 - gaussian
 - lammmps
 - matlab
 - tensorflow
- ◆ generic/
 - **serial.qs**
 - thread.qs
 - mpi/

```
[(it_css:mkyle)@login01.darwin ~]$ cd /opt/shared/templates/slurm/
[(it_css:mkyle)@login01.darwin slurm]$ ls
README.md  applications  generic
[(it_css:mkyle)@login01.darwin slurm]$ cd applications
[(it_css:mkyle)@login01.darwin applications]$ ls
RStudio-Server.qs  abaqus.qs  comsol.qs  gaussian.qs  gromacs.qs
lammmps.qs  matlab-mcr.qs  matlab.qs  nwchem.qs  tensorflow.qs
[(it_css:mkyle)@login01.darwin applications]$ cd ../generic/
[(it_css:mkyle)@login01.darwin generic]$ ls
mpi  serial.qs  threads.qs
[(it_css:mkyle)@login01.darwin generic]$ mkdir ~/workshop/job && cp
serial.qs ~/workshop/job/
[(it_css:mkyle)@login01.darwin generic]$ cd ~/workshop/job
[(it_css:mkyle)@login01.darwin job]$ ls
serial.qs
[(it_css:mkyle)@login01.darwin job]$
```

Setting Up Jobs

Job Templates

→ /opt/shared/templates/slurm/

- ◆ applications/
 - R
 - comsol
 - abaqus
 - gaussian
 - lammmps
 - matlab
 - tensorflow
- ◆ generic/
 - **serial.qs**
 - thread.qs
 - mpi/

```
[(it_css:mkyle)@login01.darwin ~]$ cd /opt/shared/templates/slurm/
[(it_css:mkyle)@login01.darwin slurm]$ ls
README.md  applications  generic
[(it_css:mkyle)@login01.darwin slurm]$ cd applications
[(it_css:mkyle)@login01.darwin applications]$ ls
RStudio-Server.qs  abaqus.qs  comsol.qs  gaussian.qs  gromacs.qs
lammmps.qs  matlab-mcr.qs  matlab.qs  nwchem.qs  tensorflow.qs
[(it_css:mkyle)@login01.darwin applications]$ cd ../generic/
[(it_css:mkyle)@login01.darwin generic]$ ls
mpi  serial.qs  threads.qs
[(it_css:mkyle)@login01.darwin generic]$ mkdir ~/workshop/job && cp
serial.qs ~/workshop/job/
[(it_css:mkyle)@login01.darwin generic]$ cd ~/workshop/job
[(it_css:mkyle)@login01.darwin job]$ ls
serial.qs
[(it_css:mkyle)@login01.darwin job]$ less serial.qs
```

Setting Up Jobs

Job Templates

- /opt/shared/templates/slurm/
 - ◆ applications/
 - R
 - comsol
 - abaqus
 - gaussian
 - lammmps
 - matlab
 - tensorflow
 - ◆ generic/
 - **serial.qs**
 - thread.qs
 - mpi/

```
#!/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=1
#
# [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=1024M
```

Setting Up Jobs

- **Getting Jobs Files & Data**
- Creating a Copy
- Customizing Job Script

```
[(it_css:mkyale)@login01.darwin job]$ cp  
/opt/shared/help/DARWIN-Workshop/stockLRM.py .  
[(it_css:mkyale)@login01.darwin job]$
```

Setting Up Jobs

- Getting Jobs Files & Data
- Creating a Copy
- Customizing Job Script

```
[(it_css:mkyale)@login01.darwin job]$ cp
/opt/shared/help/DARWIN-Workshop/stockLRM.py .
[(it_css:mkyale)@login01.darwin job]$ tail -n1 ../DARWIN_WrkShp.txt
curl
'https://docs.google.com/spreadsheets/d/1upReABmkz8I-7X1RbqWzE6ILCiUUbbqs
L9LiQ6F8lQ8/gviz/tq?tqx=out:csv' -o sp500.csv
[(it_css:mkyale)@login01.darwin job]$
```


Setting Up Jobs

- Getting Jobs Files & Data
- Creating a Copy
- Customizing Job Script

```
[(it_css:mkyale)@login01.darwin job]$ cp
/opt/shared/help/DARWIN-Workshop/stockLRM.py .
[(it_css:mkyale)@login01.darwin job]$ tail -n1 ../DARWIN_WrkShp.txt
curl
'https://docs.google.com/spreadsheets/d/1upReABmkz8I-7X1RbqWzE6IlCiUUbbqs
L9LiQ6F8lQ8/gviz/tq?tqx=out:csv' -o sp500.csv
[(it_css:mkyale)@login01.darwin job]$ curl
'https://docs.google.com/spreadsheets/d/1upReABmkz8I-7X1RbqWzE6IlCiUUbbqs
L9LiQ6F8lQ8/gviz/tq?tqx=out:csv' -o sp500.csv
% Total          % Received % Xferd  Average Speed          Time          Time          Time
Current
                                Dload  Upload    Total   Spent    Left
Speed
100 152k    0 152k      0      0 48948      0 --:--:--   0:00:03
--:--:-- 48953
[(it_css:mkyale)@login01.darwin job]$
```

Setting Up Jobs

- Getting Jobs Files & Data
- Creating a Copy
- Customizing Job Script

```
[(it_css:mkyale)@login01.darwin job]$ cp
/opt/shared/help/DARWIN-Workshop/stockLRM.py .
[(it_css:mkyale)@login01.darwin job]$ tail -n1 ../DARWIN_WrkShp.txt
curl
'https://docs.google.com/spreadsheets/d/1upReABmkz8I-7X1RbqWzE6IlCiUUbbqs
L9LiQ6F8lQ8/gviz/tq?tqx=out:csv' -o sp500.csv
[(it_css:mkyale)@login01.darwin job]$ curl
'https://docs.google.com/spreadsheets/d/1upReABmkz8I-7X1RbqWzE6IlCiUUbbqs
L9LiQ6F8lQ8/gviz/tq?tqx=out:csv' -o sp500.csv
% Total      % Received % Xferd  Average Speed   Time    Time     Time
Current

                               Dload  Upload  Total  Spent  Left
Speed
100 152k    0 152k      0    0  48948    0  --:--:--  0:00:03
--:--:-- 48953
[(it_css:mkyale)@login01.darwin job]$ ls
serial.qs sp500.csv stockLRM.py
[(it_css:mkyale)@login01.darwin job]$
```

Setting Up Jobs

- Getting Jobs Files & Data
- Creating a Copy
- Customizing Job Script

```
[(it_css:mkyale)@login01.darwin job]$ ls
serial.qs  sp500.csv  stockLRM.py
[(it_css:mkyale)@login01.darwin job]$ cp serial.qs ${USER}_Stock.qs
[(it_css:mkyale)@login01.darwin job]$
```

Setting Up Jobs

- Getting Jobs Files & Data
- Creating a Copy
- **Customizing Job Script**

```
[(it_css:mkyle)@login01.darwin job]$ ls
serial.qs  sp500.csv  stockLRM.py
[(it_css:mkyle)@login01.darwin job]$ cp serial.qs ${USER}_Stock.qs
[(it_css:mkyle)@login01.darwin job]$ ls
mkyle_Stock.qs  serial.qs  sp500.csv  stockLRM.py
[(it_css:mkyle)@login01.darwin job]$ nano mkyle_Stock.qs
```

Setting Up Jobs

- Getting Jobs Files & Data
- Creating a Copy
- Customizing Job Script

```
#!/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=1
#
# [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=1024M
```

Setting Up Jobs

- Getting Jobs Files & Data
- Creating a Copy
- Customizing Job Script

```
#!/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=1
#
# [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=1G # template default 8G
# SBATCH --mem-per-cpu=1024M
```

Setting Up Jobs

- Getting Jobs Files & Data
- Creating a Copy
- Customizing Job Script

```
# [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
# --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=24G
#
# [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=stockLRM # Changed from serial_job
#
# [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!
#
```

Setting Up Jobs

- Getting Jobs Files & Data
- Creating a Copy
- Customizing Job Script

```
#SBATCH --partition=idle #standard
# [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
#
# 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=1
# SBATCH --gpus-per-task=1
# SBATCH --gpus-per-node=1
# SBATCH --gpus-per-socket=2
```


Setting Up Jobs

- Getting Jobs Files & Data
- Creating a Copy
- Customizing Job Script

```
# [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=0-00:05:00
#
#       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=0-00:02:00
#
# [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:
#
#SBATCH --output=%x-%j.out
#SBATCH --error=%x-%j.out
```

Setting Up Jobs

- Getting Jobs Files & Data
- Creating a Copy
- Customizing Job Script

```
# [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='my_address@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
```

Setting Up Jobs

- Getting Jobs Files & Data
- Creating a Copy
- **Customizing Job Script**

```
# [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.
#
#                 UD_EXEC mpirun vasp
#
#     If you do not use UD_EXEC, then the signals will not
#     get handled by the job shell!
#
#job_exit_handler() {
# # Copy all our output files back to the original job directory:
# cp * "$SLURM_SUBMIT_DIR"
#
# # Don't call again on EXIT signal, please:
# trap - EXIT
# exit 0
#}
#export UD_JOB_EXIT_FN=job_exit_handler
```

Setting Up Jobs

- Getting Jobs Files & Data
- Creating a Copy
- Customizing Job Script

```
#
# [EDIT] By default, the function defined above is registered
#       to respond to the SIGTERM signal that Slurm sends
#       when jobs reach their runtime limit or are
#       preempted.  You can override with your own list of
#       signals using this variable -- as in this example,
#       which registers for both SIGTERM and the EXIT
#       pseudo-signal that Bash sends when the script ends.
#       In effect, no matter whether the job is terminated
#       or completes, the UD_JOB_EXIT_FN will be called.
#
#export UD_JOB_EXIT_FN_SIGNALS="SIGTERM EXIT"

#
# [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
```

Setting Up Jobs

- Getting Jobs Files & Data
- Creating a Copy
- Customizing Job Script

```
#       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 intel-python/2022.1.0
#
# 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
python --version
python stockLRM.py
```

Setting Up Jobs

- Getting Jobs Files & Data
- Creating a Copy
- **Customizing Job Script**

```
[(it_css:mkyale)@login01.darwin job]$ ls
serial.qs  sp500.csv  stockLRM.py
[(it_css:mkyale)@login01.darwin job]$ cp serial.qs ${USER}_Stock.qs
[(it_css:mkyale)@login01.darwin job]$ ls
mkyle_Stock.qs  serial.qs  sp500.csv  stockLRM.py
[(it_css:mkyale)@login01.darwin job]$ nano mkyle_Stock.qs
[(it_css:mkyale)@login01.darwin job]$
```

Before we run the job...

Let's check our allocations SU's

SUs and Accounting

sproject

- **allocations**
 - ◆ --detail
- projects
- jobs
- failures

```
[(it_css:mkyale)@login00.darwin job]$ sproject allocations -w <workgroup>
```


SUs and Accounting

sproject

- allocations
 - ◆ --detail
- projects
- jobs
- failures

```
[(it_css:mkyale@login00.darwin job]$ sproject allocations -w it_css
Project id Alloc id Alloc descr Category RDR Start date          End date
-----
-----
          2          3 it_css::cpu startup  cpu 2021-07-12 00:00:00-04:00 2021-07-25
23:59:59-04:00
          2          4 it_css::gpu startup  gpu 2021-07-12 00:00:00-04:00 2021-07-25
23:59:59-04:00
          2         43 it_css::cpu startup  cpu 2021-07-26 00:00:00-04:00 2023-07-31
00:00:00-04:00
          2         44 it_css::gpu startup  gpu 2021-07-26 00:00:00-04:00 2023-07-31
00:00:00-04:00
          2        512 it_css::cpu startup  cpu 2023-08-01 00:00:00-04:00 2024-08-31
00:00:00-04:00
          2        514 it_css::gpu startup  gpu 2023-08-01 00:00:00-04:00 2025-08-31
00:00:00-04:00
[(it_css:mkyale@login00.darwin job]$
```

SUs and Accounting

sproject

→ **allocations**

◆ **--current-only**

◆ **--detail**

→ **projects**

→ **jobs**

→ **failures**

```
[(it_css:mkyale)@login00.darwin job]$ sproject allocations -w it_css
Project id Alloc id Alloc descr Category RDR Start date           End
date
-----
                2          3 it_css::cpu startup  cpu 2021-07-12 00:00:00-04:00 2021-07-25
23:59:59-04:00
                2          4 it_css::gpu startup  gpu 2021-07-12 00:00:00-04:00 2021-07-25
23:59:59-04:00
                2         43 it_css::cpu startup  cpu 2021-07-26 00:00:00-04:00 2023-07-31
00:00:00-04:00
                2         44 it_css::gpu startup  gpu 2021-07-26 00:00:00-04:00 2023-07-31
00:00:00-04:00
                2        512 it_css::cpu startup  cpu 2023-08-01 00:00:00-04:00 2024-08-31
00:00:00-04:00
                2        514 it_css::gpu startup  gpu 2023-08-01 00:00:00-04:00 2025-08-31
00:00:00-04:00
[(it_css:mkyale)@login00.darwin job]$ sproject allocations -w it_css
--current-only
Project id Alloc id Alloc descr Category RDR Start date           End date
-----
                2        680 it_css::cpu startup  cpu 2024-10-04 00:00:00-04:00 2025-09-30
00:00:00-04:00
                2        681 it_css::gpu startup  gpu 2024-10-04 00:00:00-04:00 2025-09-30
00:00:00-04:00
```

SUs and Accounting

sproject

→ allocations

◆ --current-only

◆ --detail

→ projects

→ jobs

→ failures

```
[(it_css:mkyale)@login00.darwin job]$ sproject allocations -w it_css --detail
Project id Alloc id Alloc descr Category RDR Credit Run+Cmplt Debit Balance
-----
2          3 it_css::cpu startup  cpu 108500          0 -1678  106822
2          4 it_css::gpu startup  gpu   417           0  -16     401
2         43 it_css::cpu startup  cpu 33333          0 -433  32900
2         44 it_css::gpu startup  gpu 33333          0  -22  33311
2        512 it_css::cpu startup  cpu 33333         -72 -6467  26794
2        514 it_css::gpu startup  gpu 33333          -3   -8  33322
2        680 it_css::cpu startup  cpu 33333           0    0  33333
2        681 it_css::gpu startup  gpu 33333           0    0  33333
[(it_css:mkyale)@login00.darwin job]$
```

SUs and Accounting

sproject

→ allocations

◆ --current-only

◆ --details

→ projects

→ jobs

→ failures

```
[(it_css:mkyale)@login00.darwin job]$ sproject allocations -w it_css --detail
Project id Alloc id Alloc descr Category RDR Credit Run+Cmplt Debit Balance
-----
2          3 it_css::cpu startup  cpu 108500          0 -1678  106822
2          4 it_css::gpu startup  gpu   417          0  -16    401
2          43 it_css::cpu startup  cpu 33333          0 -433  32900
2          44 it_css::gpu startup  gpu 33333          0  -22  33311
2          512 it_css::cpu startup  cpu 33333         -72 -6467  26794
2          514 it_css::gpu startup  gpu 33333          -3   -8  33322
2          680 it_css::cpu startup  cpu 33333          0    0  33333
2          681 it_css::gpu startup  gpu 33333          0    0  33333
[(it_css:mkyale)@login00.darwin job]$ sproject allocations -w it_css --detail
--current-only
Project id Alloc id Alloc descr Category RDR Credit Run+Cmplt Debit Balance
-----
2          680 it_css::cpu startup  cpu 33333          0    0  33333
2          681 it_css::gpu startup  gpu 33333          0    0  33333
[(it_css:mkyale)@login00.darwin job]$
```

SUs and Accounting

sproject

- allocations
- **projects**
 - ◆ --current-only
 - ◆ --detail
- jobs
- failures

```
[(it_css:mkyale)@login00.darwin job]$ sproject projects -w it_css
Project id Account Group id Group name  Creation date
-----
          2 it_css      1002 it_css  2021-07-12 14:51:57-04:00
[(it_css:mkyale)@login00.darwin job]$
```

SUs and Accounting

sproject

- allocations
- projects
 - ◆ --detail
- jobs
- failures

```
[(it_css:mkyale@login00.darwin job)]$ sproject projects -w it_css
Project id Account Group id Group name Creation date
-----
      2 it_css      1002 it_css      2021-07-12 14:51:57-04:00
[(it_css:mkyale@login00.darwin job)]$ sproject projects -w it_css --detail
Project id Account Group id Group name Allocation id Category RDR Start date
      End date Creation date
-----
      2 it_css      1002 it_css      3 startup cpu 2021-07-12
00:00:00-04:00 2021-07-25 23:59:59-04:00 2021-07-12 15:00:46-04:00
      4 startup gpu 2021-07-12
00:00:00-04:00 2021-07-25 23:59:59-04:00 2021-07-12 15:00:54-04:00
      43 startup cpu 2021-07-26
00:00:00-04:00 2023-07-31 00:00:00-04:00 2021-07-26 14:47:56-04:00
      44 startup gpu 2021-07-26
00:00:00-04:00 2023-07-31 00:00:00-04:00 2021-07-26 14:47:56-04:00
      512 startup cpu 2023-08-01
00:00:00-04:00 2024-08-31 00:00:00-04:00 2023-08-04 10:03:08-04:00
      514 startup gpu 2023-08-01
00:00:00-04:00 2024-08-31 00:00:00-04:00 2023-08-04 10:03:08-04:00
      680 startup cpu 2024-10-04
00:00:00-04:00 2025-09-30 00:00:00-04:00 2024-10-04 15:26:13-04:00
      681 startup gpu 2024-10-04
00:00:00-04:00 2025-09-30 00:00:00-04:00 2024-10-04 15:26:13-04:00
[(it_css:mkyale@login00.darwin job)]$
```

SUs and Accounting

sproject

- allocations
- projects
- jobs
- failures

```
[(it_css:mkyale)@login00.darwin job]$ sproject jobs -w it_css
sproject jobs -w it_css
Activity id Alloc id Alloc descr Job id Owner Status Amount Modification
date Creation date
-----
13:06:30-04:00 2024-06-25 12:31:28-04:00
4647038 514 it_css::gpu 5212043 thuachen completed -1 2024-06-25
13:34:30-04:00 2024-06-25 12:59:09-04:00
4647047 514 it_css::gpu 5212053 thuachen completed -1 2024-06-25
14:13:30-04:00 2024-06-25 13:38:30-04:00
4648461 514 it_css::gpu 5213473 thuachen completed 0 2024-06-25
15:28:44-04:00 2024-06-25 15:27:58-04:00
4648462 514 it_css::gpu 5213474 thuachen completed 0 2024-06-25
15:30:37-04:00 2024-06-25 15:30:06-04:00
4648464 514 it_css::gpu 5213476 thuachen completed 0 2024-06-25
15:53:32-04:00 2024-06-25 15:31:10-04:00
4697342 512 it_css::cpu 5264839 bkang completed -24 2024-07-23
10:43:14-04:00 2024-07-23 10:17:58-04:00
4697340 512 it_css::cpu 5264836 bkang completed 0 2024-07-23
10:16:50-04:00 2024-07-23 10:16:44-04:00
4697368 512 it_css::cpu 5264893 bkang completed -24 2024-07-23
13:11:17-04:00 2024-07-23 12:46:19-04:00
4697390 512 it_css::cpu 5264953 bkang completed -23 2024-07-23
14:19:05-04:00 2024-07-23 13:55:17-04:00
```

SUs and Accounting

sproject

- allocations
- projects
- jobs
- **failures**

--help Can be used with all
the above

```
[(it_css:mkyale@login00.darwin job] $ sproject failures -w it_css
Job id Error message
-----
5063695 Requested allocation has insufficient balance: 1612159 < 6451200
5063937 Requested allocation has insufficient balance: 1612159 < 6451200
[(it_css:mkyale@login00.darwin job] $
```


Running & Monitoring Jobs

Running Jobs

Required to join workgroup

→ `workgroup -g`
`<workgroup_name>`

Job types

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

```
[mkyle@login01.darwin ~]$ workgroup -g <workgroup>
```

Running & Monitoring Jobs

Running Jobs

Required to join workgroup

→ `workgroup -g`
`<workgroup_name>`

Job types

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

```
[mkyle@login01.darwin ~]$ workgroup -g it_css  
[(it_css:mkyle)@login01.darwin ~]$
```

Running & Monitoring Jobs

Running Jobs

Required to join workgroup

→ `workgroup -g`
`<workgroup_name>`

Job types

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

```
[mkyle@login01.darwin ~]$ workgroup -g it_css
[(it_css:mkyle)@login01.darwin ~]$ cd ~/workshop/job
[(it_css:mkyle)@login01.darwin ~]
```

Running & Monitoring Jobs

Running Jobs

Required to join workgroup

→ `workgroup -g`
`<workgroup_name>`

Job types

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

```
[mkyle@login01.darwin ~]$ workgroup -g it_css
[(it_css:mkyle)@login01.darwin ~]$ cd ~/workshop/job
[(it_css:mkyle)@login01.darwin job]$ ls
mkyle_Stock.qs  serial.qs  sp500.csv  stockLRM.py
[(it_css:mkyle)@login01.darwin job] $
```

Running & Monitoring Jobs

Running Jobs

Required to join workgroup

→ `workgroup -g`
`<workgroup_name>`

Job types

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

```
[mkyle@login01.darwin ~]$ workgroup -g it_css
[(it_css:mkyle)@login01.darwin ~]$ cd ~/workshop/job
[(it_css:mkyle)@login01.darwin job]$ ls
mkyle_Stock.qs  serial.qs  sp500.csv  stockLRM.py
[(it_css:mkyle)@login01.darwin job]$ sbatch ${USER}_Stock.qs
```

Running & Monitoring Jobs

Running Jobs

Required to join workgroup

→ `workgroup -g`
`<workgroup_name>`

Job types

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

```
[mkyle@login01.darwin ~]$ workgroup -g it_css
[(it_css:mkyle)@login01.darwin ~]$ cd ~/workshop/job
[(it_css:mkyle)@login01.darwin job]$ ls
mkyle_Stock.qs  serial.qs  sp500.csv  stockLRM.py
[(it_css:mkyle)@login01.darwin job]$ sbatch ${USER}_Stock.qs
Submitted batch job 5321840
[(it_css:mkyle)@login01.darwin job]$
```

Running & Monitoring Jobs

Running Jobs

Required to join workgroup

→ `workgroup -g`
`<workgroup_name>`

Job types

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

```
[mkyle@login01.darwin ~]$ workgroup -g it_css
[(it_css:mkyle)@login01.darwin ~]$ cd ~/workshop/job
[(it_css:mkyle)@login01.darwin job]$ ls
mkyle_Stock.qs  serial.qs  sp500.csv  stockLRM.py
[(it_css:mkyle)@login01.darwin job]$ sbatch ${USER}_Stock.qs
Submitted batch job 5321840
[(it_css:mkyle)@login01.darwin job]$ ls
meanSquareError.png  mkyle_Stock.qs  serial.qs  sp500.csv
stockLRM-5321840.out  stockLRM.py
[(it_css:mkyle)@login01.darwin job]$
```

Running & Monitoring Jobs

Running Jobs

Required to join workgroup

→ `workgroup -g`
`<workgroup_name>`

Job types

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

```
[mkyle@login01.darwin ~]$ workgroup -g it_css
[(it_css:mkyle)@login01.darwin ~]$ cd ~/workshop/job
[(it_css:mkyle)@login01.darwin job]$ ls
mkyle_Stock.qs  serial.qs  sp500.csv  stockLRM.py
[(it_css:mkyle)@login01.darwin job]$ sbatch ${USER}_Stock.qs
Submitted batch job 5321840
[(it_css:mkyle)@login01.darwin job]$ ls
meanSquareError.png  mkyle_Stock.qs  serial.qs  sp500.csv
stockLRM-5321840.out  stockLRM.py
[(it_css:mkyle)@login01.darwin job]$ display meanSquareError.png
```


Running & Monitoring Jobs

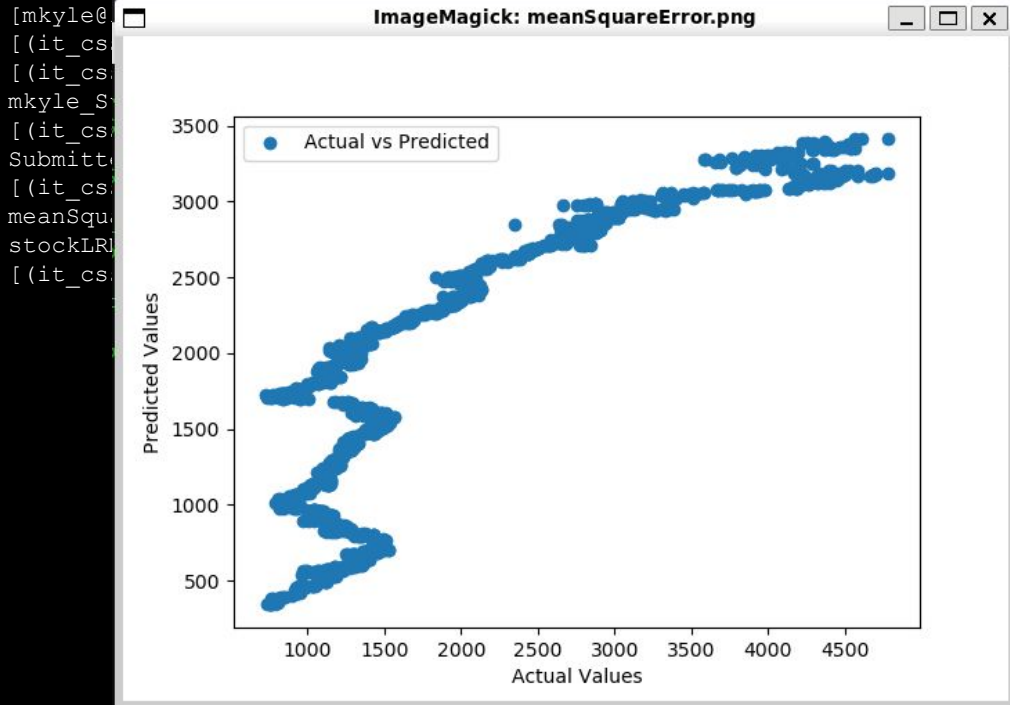
Running Jobs

Required to join workgroup

```
→ workgroup -g  
   <workgroup_name>
```

Job types

- Batch Jobs
 - **sbatch**
 <job_script>
- Interactive Jobs
 - salloc
 - slurm options



Running & Monitoring Jobs

Running Jobs

Required to join workgroup

→ `workgroup -g`
`<workgroup_name>`

Job types

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

```
[mkyle@login01.darwin ~]$ workgroup -g it_css
[(it_css:mkyle)@login01.darwin ~]$ cd ~/workshop/job
[(it_css:mkyle)@login01.darwin job]$ ls
mkyle_Stock.qs  serial.qs  sp500.csv  stockLRM.py
[(it_css:mkyle)@login01.darwin job]$ sbatch ${USER}_Stock.qs
Submitted batch job 5321840
[(it_css:mkyle)@login01.darwin job]$ ls
meanSquareError.png  mkyle_Stock.qs  serial.qs  sp500.csv
stockLRM-5321840.out  stockLRM.py
[(it_css:mkyle)@login01.darwin job]$ display meanSquareError.png
[(it_css:mkyle)@login01.darwin job]$ nano stockLRM-5321840.out
```

Running & Monitoring Jobs

Running Jobs Required to

→ workgr
<workg

Job types

- Batch J
 - s
 - ▲
- Interac
 - to
 - to

```
GNU nano 2.3.1 File: stockLRM-5321840.out
Adding package `intel-python/2022.1.0` to your environment
Python 3.9.10 :: Intel Corporation
Intel(R) Extension for Scikit-learn* enabled (https://github.com/intel/scikit-learn-intelex)
Mean Squared Error: 310439.05366691446
Predicted closing price for 2023-02-15: 3289.6151191285753

[ Read 5 lines ]
^G Get Help      ^O WriteOut     ^R Read File    ^Y Prev Page    ^K Cut Text     ^C Cur
Pos | ^X Exit        ^J Justify     ^W Where Is    ^V Next Page    ^U UnCut Text
^T To Spell
```

Running & Monitoring Jobs

Running Jobs

Required to join workgroup

→ `workgroup -g`
`<workgroup_name>`

Job types

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

```
[mkyle@login01.darwin ~]$ workgroup -g it_css
[(it_css:mkyle)@login01.darwin ~]$ cd ~/workshop/job
[(it_css:mkyle)@login01.darwin job]$ ls
mkyle_Stock.qs  serial.qs  sp500.csv  stockLRM.py
[(it_css:mkyle)@login01.darwin job]$ sbatch ${USER}_Stock.qs
Submitted batch job 5321840
[(it_css:mkyle)@login01.darwin job]$ ls
meanSquareError.png  mkyle_Stock.qs  serial.qs  sp500.csv
stockLRM-5321840.out  stockLRM.py
[(it_css:mkyle)@login01.darwin job]$ display meanSquareError.png
[(it_css:mkyle)@login01.darwin job]$ nano stockLRM-5321840.out
[(it_css:mkyle)@login01.darwin job]$ clear
```

Running & Monitoring Jobs

Running Jobs

Required to join workgroup

→ `workgroup -g`
`<workgroup_name>`

Job types

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

```
[(it_css:mkyale)@login01.darwin ~]$ salloc --partition=<partition>
```

Running & Monitoring Jobs

Running Jobs

Required to join workgroup

→ `workgroup -g`
`<workgroup_name>`

Job types

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

```
[(it_css:mkyale)@login01.darwin ~]$ salloc --partition=standard
```

Running & Monitoring Jobs

Running Jobs

Required to join workgroup

→ `workgroup -g`
`<workgroup_name>`

Job types

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

```
[(it_css:mkyale)@login01.darwin ~]$ salloc --partition=standard
salloc: Granted job allocation 5322677
salloc: Waiting for resource configuration
salloc: Nodes r1n02 are ready for job
```

Running & Monitoring Jobs

Running Jobs

Required to join workgroup

```
→ workgroup -g  
   <workgroup_name>
```

Job types

- Batch Jobs
 - sbatch
 <job_script>
- Interactive Jobs
 - **salloc**
 - slurm options

```
[(it_css:mkyale)@login01.darwin ~]$ salloc --partition=standard  
salloc: Granted job allocation 5322677  
salloc: Waiting for resource configuration  
salloc: Nodes r1n02 are ready for job  
[(it_css:mkyale)@r1n02 job]$
```


Running & Monitoring Jobs

Running Jobs

Required to join workgroup

→ `workgroup -g`
`<workgroup_name>`

Job types

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

```
[(it_css:mkyale)@login01.darwin job]$ salloc --partition=standard
salloc: Granted job allocation 5322677
salloc: Waiting for resource configuration
salloc: Nodes r1n02 are ready for job
[(it_css:mkyale)@r1n02 job]$ exit
logout
salloc: Relinquishing job allocation 5322677
salloc: Job allocation 5322677 has been revoked.
[(it_css:mkyale)@login01.darwin job]$
```

Running & Monitoring Jobs

Running Jobs

Required to join workgroup

→ `workgroup -g`
`<workgroup_name>`

Job types

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

```
# [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=0-00:05:00
#
#       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=0-00:02:00
#
# [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:
#
#SBATCH --output=%x-%j.out
#SBATCH --error=%x-%j.out
```

Running & Monitoring Jobs

Running Jobs

Exercise #3:

Change the stockLRM.py file's Random Seed value and rerun the program using slurm's sbatch or salloc method to run the job.

The goal is to get a lower Mean Squared Error: 310439.053

```
import pandas as pd
from sklearn.linear_model import LinearRegression
from sklearn.model_selection import train_test_split
from sklearn.metrics import mean_squared_error
import matplotlib.pyplot as plt

# Read the CSV file into a DataFrame
df = pd.read_csv('sp500.csv')

# Convert the 'Date' column to datetime and extract the year and month
df['Date'] = pd.to_datetime(df['Date'])
df['Year'] = df['Date'].dt.year
df['Month'] = df['Date'].dt.month

# Create a linear regression model
X = df[['Year', 'Month']] # Features
y = df['Close'] # Target variable

# Split the data into training and testing sets
#X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2,
random_state=42)
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.2,
random_state=3338)
```

Running & Monitoring Jobs

Running Jobs

Exercise:

Change the stockLRM.py file's Random Seed value and rerun the program using slurm's sbatch or salloc method to run the job.

The goal is to get a lower Mean Squared Error: 310439.053

```
[(it_css:mkyale)@login01.darwin job]$ salloc --partition=standard
salloc: Granted job allocation 5322677
salloc: Waiting for resource configuration
salloc: Nodes r1n02 are ready for job
[(it_css:mkyale)@r1n02 job]$ exit
logout
salloc: Relinquishing job allocation 5322677
salloc: Job all
[(it_css:mkyale)
```

3:00

Running & Monitoring Jobs

Monitoring Jobs

Active Jobs

→ **squeue**

→ **ssqueue (Pretty)**

Ended Jobs

→ **sacct**

→ **ssacct (Pretty)**

```
[(it_css:mkyale)@login01.darwin job]$ squeue
```

JOBID	PARTITION	NAME	USER	ST	TIME	NODES	NODELIST (REASON)
5322549	gpu-v100	MetSRGAN	xsedeu35	R	22:37:54	1	r2v01
5322589	gpu-v100	NNtrainL	xsedeu31	R	15:01:46	1	r2v00
5322606	idle	openmpi_	ruiqihu	PD	0:00	1	(QOSMaxCpuPerUserLimit)
5322641	idle	openmpi_	ruiqihu	PD	0:00	1	(QOSMaxCpuPerUserLimit)
5322686	idle	interact	mkyale	R	11:32	1	r2112
5322055	idle	mv	msnouri	R	2-03:38:57	1	r2x00
5322645	idle	interact	shard	R	3:58:01	1	r2112
5322665	idle	df	shard	R	2:37:09	1	r2121
5322626	idle	gromacs_	tianren	R	5:33:29	1	r2126

```
[(it_css:mkyale)@login01.darwin job]$
```

Running & Monitoring Jobs

Monitoring Jobs

Active Jobs

- `squeue`
- `ssqueue` (Pretty)

Ended Jobs

- `sacct`
- `ssacct` (Pretty)

```
[(it_css:mkyale)@login01.darwin job]$ squeue
JOBID PARTITION NAME USER ST TIME NODES NODELIST(REASON)
5322549 gpu-v100 MetSRGAN xsedeu35 R 22:37:54 1 r2v01
5322589 gpu-v100 NNtrainL xsedeu31 R 15:01:46 1 r2v00
5322606 idle openmpi_ ruiqihu PD 0:00 1 (QOSMaxCpuPerUserLimit)
5322641 idle openmpi_ ruiqihu PD 0:00 1 (QOSMaxCpuPerUserLimit)
5322686 idle interact mkyale R 11:32 1 r2112
5322055 idle mv msnouri R 2-03:38:57 1 r2x00
5322645 idle interact shard R 3:58:01 1 r2112
5322665 idle df shard R 2:37:09 1 r2121
5322626 idle gromacs_ tianren R 5:33:29 1 r2126
[(it_css:mkyale)@login01.darwin job]$
[(it_css:mkyale)@login01.darwin job]$ squeue -j 5322686
JOBID PARTITION NAME USER ST TIME NODES NODELIST(REASON)
5322686 idle interact mkyale R 20:13 1 r2112
[(it_css:mkyale)@login01.darwin job]$
```

Running & Monitoring Jobs

Monitoring Jobs

Active Jobs

→ **squeue**

→ **ssqueue (Pretty)**

Ended Jobs

→ **sacct**

→ **ssacct (Pretty)**

```
[(it_css:mkyle)@login01.darwin job]$ squeue
JOBID PARTITION NAME USER ST TIME NODES NODELIST(REASON)
5322549 gpu-v100 MetSRGAN xsedeu35 R 22:37:54 1 r2v01
5322589 gpu-v100 NNtrainL xsedeu31 R 15:01:46 1 r2v00
5322606 idle openmpi_ ruiqihu PD 0:00 1 (QOSMaxCpuPerUserLimit)
5322641 idle openmpi_ ruiqihu PD 0:00 1 (QOSMaxCpuPerUserLimit)
5322686 idle interact mkyle R 11:32 1 r2112
5322055 idle mv msnouri R 2-03:38:57 1 r2x00
5322645 idle interact shard R 3:58:01 1 r2112
5322665 idle df shard R 2:37:09 1 r2121
5322626 idle gromacs_ tianren R 5:33:29 1 r2126
[(it_css:mkyle)@login01.darwin job]$
[(it_css:mkyle)@login01.darwin job]$ squeue -j 5322686
JOBID PARTITION NAME USER ST TIME NODES NODELIST(REASON)
5322686 idle interact mkyle R 20:13 1 r2112
[(it_css:mkyle)@login01.darwin job]$
[(it_css:mkyle)@login01.darwin job]$ squeue -u mkyle
JOBID PARTITION NAME USER ST TIME NODES NODELIST(REASON)
5322686 idle interact mkyle R 22:34 1 r2112
[(it_css:mkyle)@login01.darwin job]$
```

Running & Monitoring Jobs

Monitoring Jobs

Active Jobs

→ `squeue`

→ `ssqueue` (Pretty)

Ended Jobs

→ `sacct`

→ `ssacct` (Pretty)

```
[(it_css:mkyale)@login01.darwin job]$ squeue --help
Usage: squeue [OPTIONS]
  -A, --account=account(s)    comma separated list of accounts
                              to view, default is all accounts
  -a, --all                    display jobs in hidden partitions
                              --array-unique    display one unique pending job array
                              element per line
  -D, --delimiter=DELIM      print DELIM between each output field
                              --federation      Report federated information if a member
                              of one
  -h, --noheader              no headers on output
                              --hide            do not display jobs in hidden partitions
  -i, --iterate=seconds       specify an interation period
  -j, --job=job(s)            comma separated list of jobs IDs
                              to view, default is all
                              --local           Report information only about jobs on the
                              local cluster. Overrides --federation.
  -l, --long                   long report
...More Information was skipped...
Help options:
  --help                       show this help message
  --usage                       display a brief summary of squeue options
[(it_css:mkyale)@login01.darwin job]$
```


Running & Monitoring Jobs

Monitoring Jobs

Active Jobs

- `queue`
- `ssqueue` (Pretty)

Ended Jobs

- `sacct`
- `ssacct` (Pretty)

```
[(it_css:mkye)@login01.darwin job]$ ssqueue
```

Running & Monitoring Jobs

Monitoring Jobs

Active Jobs

→ `squeue`

→ `ssqueue` (Pretty)

Ended Jobs

→ `sacct`

→ `ssacct` (Pretty)

JOBID	PARTITION	NAME	USER	ST	TIME	NODES	NODELIST
5322756	gpu-v100	NNtrainLSTM01	xsedeu3108	R	9:29:32	1	r2v00
5322774	idle	bulk	ruiqihu	PD	0:00	2	(QOSMaxCpuPerUserLimit)
5322055	idle	mv	msnouri	R	21:44:01	1	r2x00
5322744	idle	df_reduce	shard	R	12:06:30	1	r2100
5322645	idle	jupyter	shared	R	22:03:05	1	r2112
5322768	idle	gromacs_elp	tianren	R	2:02:15	1	r2104
5322769	idle	gromacs_elp	tianren	R	2:02:15	1	r2109
5322760	idle	gromacs_elp	tianren	R	2:02:16	1	r2100
5322761	idle	gromacs_elp	tianren	R	2:02:16	1	r2101
5322773	idle	openmpi_job	ruiqihu	R	51:27	1	r2120

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

Running & Monitoring Jobs

Monitoring Jobs

Active Jobs

- `squeue`
- `ssqueue` (Pretty)

Ended Jobs

- `sacct`
- `ssacct` (Pretty)

```
[(it_css:mkyale)@login01.darwin job]$ sacct -j <job_id>
```

Running & Monitoring Jobs

Monitoring Jobs

Active Jobs

- `squeue`
- `ssqueue (Pretty)`

Ended Jobs

- `sacct`
- `ssacct (Pretty)`

```
[(it_css:mkyale)@login01.darwin job]$ sacct -j 5322667
```

Running & Monitoring Jobs

Monitoring Jobs

Active Jobs

- `squeue`
- `ssqueue (Pretty)`

Ended Jobs

- `sacct`
- `ssacct (Pretty)`

```
[(it_css:mkyale)@login01.darwin job]$ sacct -j 5322667
-----
JobID   JobName  Partition  Account  AllocCPUS  State  ExitCode
-----
5322667  stockLRM  idle       it_css   1          COMPLETED  0:0
5322667.bat+  batch    it_css     1          COMPLETED  0:0
5322667.ext+  extern   it_css     1          COMPLETED  0:0
[(it_css:mkyale)@login01.darwin job]$
```

Running & Monitoring Jobs

Monitoring Jobs

Active Jobs

- `queue`
- `ssqueue` (Pretty)

Ended Jobs

- `sacct`
- `ssacct` (Pretty)

```
[(it_css:mkyle)@login01.darwin job]$ sacct --helpformat
Account                AdminComment           AllocCPUS              AllocNodes
AllocTRES              AssocID                AveCPU                 AveCPUFreq
AveDiskRead           AveDiskWrite           AvePages               AveRSS
AveVMSize              BlockID                Cluster                Comment
Constraints           ConsumedEnergy         ConsumedEnergyRaw     CPUTime
CPUTimeRAW            DBIndex                DerivedExitCode        Elapsed
ElapsedRaw            Eligible               End                    ExitCode
Flags                 GID                    Group                  JobID
JobIDRaw              JobName                Layout                 MaxDiskRead
MaxDiskReadNode       MaxDiskReadTask        MaxDiskWrite           MaxDiskWriteNode
MaxDiskWriteTask      MaxPages               MaxPagesNode           MaxPagesTask
MaxRSS                 MaxRSSNode             MaxRSSTask             MaxVMSize
MaxVMSizeNode         MaxVMSizeTask          McsLabel               MinCPU
MinCPUNode            MinCPUTask             NCPUS                  NNodes
NodeList              NTasks                Priority                Partition
QOS                   QOSRAW                 Reason                 ReqCPUFreq
ReqCPUFreqMin         ReqCPUFreqMax          ReqCPUFreqGov          ReqCPUS
ReqMem                 ReqNodes               ReqTRES                Reservation
ReservationId         Reserved               ResvCPU                ResvCPURAW
Suspended              State                  Submit                 Suspended
SystemCPU              SystemComment          Timelimit              TimelimitRaw
TotalCPU               TRESUsageInAve        TRESUsageInMax         TRESUsageInMaxNode
TRESUsageInMaxTask    TRESUsageInMin         TRESUsageInMinNode     TRESUsageInMinTask
TRESUsageInTot        TRESUsageOutAve       TRESUsageOutMax        TRESUsageOutMaxNode
TRESUsageOutMaxTask   TRESUsageOutMin       TRESUsageOutMinNode    TRESUsageOutMinTask
TRESUsageOutTot       UID                    User                   UserCPU
WCKey                 WCKeyID                WorkDir
```

Running & Monitoring Jobs

Monitoring Jobs

Active Jobs

- `squeue`
- `ssqueue` (Pretty)

Ended Jobs

- `sacct`
- `ssacct` (Pretty)

```
[(it_css:mkyale)@login01.darwin job]$ sacct --format="jobid, jobname,  
partition, account, user, state, alloccpus, totalcpu, elapsed" -j 5321840
```

Running & Monitoring Jobs

Monitoring Jobs

Active Jobs

- `squeue`
- `ssqueue` (Pretty)

Ended Jobs

- `sacct`
- `ssacct` (Pretty)

```
[(it_css:mkyale)@login01.darwin job]$ sacct --format="jobid, jobname, partition, account, user, state, alloccpus, totalcpu, elapsed" -j 5321840
```

JobID	JobName	Partition	Account	User	State	AllocCPUS	TotalCPU	Elapsed
5321840	stockLRM	idle	it_css	mkyale	COMPLETED	1	00:02.195	00:00:04
5321840.bat+	batch		it_css		COMPLETED	1	00:02.194	00:00:04
5321840.ext+	extern		it_css		COMPLETED	1	00:00.001	00:00:04

```
[(it_css:mkyale)@login01.darwin ~]$
```


Running & Monitoring Jobs

Monitoring Jobs

Active Jobs

- `squeue`
- `ssqueue` (Pretty)

Ended Jobs

- `sacct`
- `ssacct` (Pretty)

```
[(it_css:mkyale)@login01.darwin job]$ sacct --format="jobid, jobname,
partition, account, user, state, alloccpus, totalcpu, elapsed" -j 5321840
      JobID JobName Partition Account User State AllocCPUS
TotalCPU Elapsed
-----
5321840      stockLRM      idle   it_css  mkyale COMPLETED      1
00:02.195 00:00:04
5321840.bat+  batch                it_css                COMPLETED      1
00:02.194 00:00:04
5321840.ext+  extern                it_css                COMPLETED      1
00:00.001 00:00:04
[(it_css:mkyale)@login01.darwin ~]$ sacct --help
```

Running & Monitoring Jobs

Monitoring Jobs

Active Jobs

- `queue`
- `ssqueue` (Pretty)

Ended Jobs

- `sacct`
- `ssacct` (Pretty)

```
[(it_css:mkyle)@login01.darwin job]$ sacct --help
sacct [<OPTION>]
Valid <OPTION> values are:
-a, --allusers:
    Display jobs for all users. By default, only the
    current user's jobs are displayed. If ran by user root
    this is the default.
-A, --accounts:
    Use this comma separated list of accounts to select jobs
    to display. By default, all accounts are selected.
-b, --brief:
    Equivalent to '--format=jobstep,state,error'.
-c, --completion: Use job completion instead of accounting data.
-X, --allocations:
    Only show statistics relevant to the job allocation
    itself, not taking steps into consideration.
...More Lines that are Skipped...

Note, valid start/end time formats are...
HH:MM[:SS] [AM|PM]
MMDD[YY] or MM/DD[/YY] or MM.DD[.YY]
MM/DD[/YY]-HH:MM[:SS]
YYYY-MM-DD[THH:MM[:SS]]
now[{|+|-}count[seconds(default)|minutes|hours|days|weeks]]

[(it_css:mkyle)@login01.darwin job]$
```

Running & Monitoring Jobs

Monitoring Jobs

Active Jobs

- `queue`
- `ssqueue` (Pretty)

Ended Jobs

- `sacct`
- `ssacct` (Pretty)

```
[(it_css:mkyale)@login01.darwin job]$ ssacct -j 5321840
```


Monitoring Jobs

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

Monitoring Partitions

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)

Monitoring Jobs

Monitoring Jobs

→ `sinfo`

- ◆ Used to report status of all the nodes of cluster.

→ `scontrol`

```
[(it_css:mkyale)@login01.darwin job]$ sinfo

PARTITION  AVAIL  TIMELIMIT  NODES  STATE NODELIST
standard*  up 7-00:00:00      1 drain* r1n10
standard*  up 7-00:00:00      1 down* r1n19
standard*  up 7-00:00:00      9      mix
r1n[00,02,04-05,24,37-38,42,45]
standard*  up 7-00:00:00     36 alloc
r1n[01,03,06-09,11-16,18,20-23,25-36,39-41,43-44,46-47]
large-mem  up 7-00:00:00      1 down* r2l31
large-mem  up 7-00:00:00      3      mix r2l[00,09,12]
large-mem  up 7-00:00:00     24 alloc
r2l[01-08,10-11,13-18,20,22-25,27-29]
large-mem  up 7-00:00:00      4 idle r2l[19,21,26,30]
xlarge-mem up 7-00:00:00      1      mix r2x00
xlarge-mem up 7-00:00:00      3 alloc r2x[04,06-07]
xlarge-mem up 7-00:00:00      7 idle r2x[01-03,05,08-10]
extended-mem up 7-00:00:00      1 idle r2e00
gpu-t4     up 7-00:00:00      2 alloc r1t[00,05]
gpu-t4     up 7-00:00:00      7 idle r1t[01-04,06-07],r2t08
gpu-v100   up 7-00:00:00      1      mix r2v00
gpu-v100   up 7-00:00:00      2 idle r2v[01-02]
gpu-mi50   up 7-00:00:00      1 idle r2m00
gpu-mi100  up 7-00:00:00      1 idle r0m01
```

Monitoring Jobs

Monitoring Jobs

→ `sinfo`

→ `scontrol`

- ◆ Used for monitoring and modifying queued jobs, as well as holding and releasing jobs.

```
[(it_css:mkyle)@login01.darwin job]$ scontrol show job 5322641
  JobId=5322606 JobName=openmpi_job
  UserId=ruiqihu(2262) GroupId=mtg(1074) MCS_label=N/A
  Priority=90755271 Nice=0 Account=mtg QOS=normal
  JobState=RUNNING Reason=None Dependency=(null)
  Requeue=0 Restarts=0 BatchFlag=1 Reboot=0 ExitCode=0:0
  RunTime=12:28:21 TimeLimit=3-00:00:00 TimeMin=3-00:00:00
  SubmitTime=2024-10-07T07:34:43 EligibleTime=2024-10-07T07:34:43
  AccrueTime=2024-10-07T07:34:43
  StartTime=2024-10-08T01:29:08 EndTime=2024-10-11T01:29:08
  Deadline=N/A
  PreemptEligibleTime=2024-10-08T01:39:08 PreemptTime=None
  SuspendTime=None SecsPreSuspend=0 LastSchedEval=2024-10-08T01:29:08
  Partition=idle AllocNode:Sid=r0login1:113063
  ReqNodeList=(null) ExcNodeList=(null)
  NodeList=r2127
  BatchHost=r2127
  NumNodes=1 NumCPUs=64 NumTasks=64 CPUs/Task=1 ReqB:S:C:T=0:0:*:*
  TRES=cpu=64,mem=448G,node=1,billing=1
  Socks/Node=* NtasksPerN:B:S:C=64:0:*:* CoreSpec=*
  MinCPUsNode=64 MinMemoryCPU=7G MinTmpDiskNode=0
  Features=(null) DelayBoot=00:00:00
  OverSubscribe=OK Contiguous=0 Licenses=(null) Network=(null)
```


Office Hours

- Research Computing Office Hours
 - Time: Tuesdays and Wednesdays @ 11am – noon
 - Location: 002C Smith Hall, Conference Room or via Zoom
- Research Software Engineering (RSE) Support For The College Of Arts And Science (CAS)
 - Time: Tuesdays and Wednesdays @ 1pm – 2pm
 - Location: Room 110 Sharp Lab or via Zoom

Upcoming Workshops

- IT-RCI Workshop Series October 17-31
 - Getting Started with DARWIN: 10/17 @ 2-3 pm, 10/18 @ 10-11:30 am
 - Introduction to Slurm: 10/21 @ 1:30-3:30 pm
 - Shell Scripting: 10/23 @ 1:30-3:30 pm
 - HPC Software Development and Installation: 10/31 @ 10 am-noon



Helpful HPC Resources

Scan
4
url's

[hpc-ask](#)

[caviness-users](#)

[darwin-users](#)

[hpc-wiki](#)

[Help Desk](#)
askit@udel.edu



Thank you!

**For more information, contact:
askit@udel.edu**