HIGH PERFORMANCE RESEARCH COMPUTING

Introduction to HPRC Computing Resources

January 24, 2025 Josh Winchell



High Performance Research Computing DIVISION OF RESEARCH

High Performance Research Computing | hprc.tamu.edu | NSF Award #1925764

Ā M

Outline

Introduction

- Setup
- Usage Policies
- HPRC Cluster Overview

• Getting Started

Accessing HPRC ClustersHPRC Portal

• Software Infrastructure

- Module System
- Python Virtual Environments
- Cluster Computing with Slurm
- Need Help?

Introduction



Setup For This Course

To access the resources:

- HPRC Account: https://hprc.tamu.edu/user_services/#accounts
- TFA set up: <u>https://it.tamu.edu/duo/</u>

To do the exercises:

- Basic Linux/Unix skills
 - Slides from our LINUX short course are at: <u>hprc.tamu.edu/training/intro_linux.html</u>
 - Watch the relevant Introduction and Primer videos on our YouTube Channel: <u>https://www.youtube.com/texasamhprc</u>

Answers to frequently asked questions can be found on our KB: <u>https://hprc.tamu.edu/kb/FAQ/Accounts/</u>

HPRC Clusters

- We operate five clusters
- Today we'll focus on just one: Grace
 - It is our largest cluster
 - It is dedicated for TAMU usage



Others:

- Launch FASTER
- ViDaL ACES

See <u>hprc.tamu.edu/resources</u> for more info on all our clusters

Grace Overview





AM

Host Name:	grace.hprc.tamu.edu
Operating System:	Linux (CentOS 7)
Total Compute Cores/Nodes:	45,376 cores 940 nodes
Compute Nodes:	800 48-core compute nodes 100 48-core GPU nodes, each with two A100s 9 48-core GPU nodes, each with two RTX 6000s 8 48-core GPU nodes, each with 4 T4s 15 48-core GPU nodes, each with two A40s 8 80-core large memory nodes, each with 3TB RAM
Interconnect:	Mellanox HDR 100 InfiniBand
Peak Performance:	6.3 PFLOPS

More details at:

https://hprc.tamu.edu/kb/User-Guides/Grace/

Computing on HPRC Clusters

You interact with your own computer

Make login node send instructions to "compute nodes" to do the heavy-lifting

 Log into a "login" node, where you write instructions

(This image is not to-scale: we have way more machines than this!)

Clusters Are For You!

What kinds of problems are solved by cluster computing?

- Problems that are too big to fit in a laptop or workstation, due to limitation on memory, core count, gpu count, etc.
- Problems that scale well with more CPU cores or memory
- Single-threaded problems with many permutations
- Problems that require large high speed storage and/or interconnect

HPRC Knowledge Base

Home

See our Knowledge Base for announcements, more hardware details, and more about the subjects we cover today.

<u>hprc.tamu.edu/kb</u> <u>hprc.tamu.edu/kb/User-Guides/Grace</u>



Getting Started



Usage Policies (Be a good compute citizen)

- It is illegal to share computer passwords and accounts by state law and university regulation.
- It is prohibited to use HPRC clusters in any manner that violates the United States export control laws and regulations, EAR & ITAR.
- Abide by the expressed or implied restrictions in using commercial software .

hprc.tamu.edu/policies

Accessing the HPRC Portal

- (First, if you're off-campus, connect to TAMU VPN: <u>u.tamu.edu/VPnetwork</u>)
- HPRC webpage: <u>hprc.tamu.edu</u>: Portal dropdown menu



HPRC Portal Overview



Files: copy and edit files on the cluster's filesystems Jobs: submit and monitor cluster jobs Clusters: open a shell terminal (command line) on a login node Interactive Apps: start graphical software on a compute node and connect to it Dashboard: view file quotas and computing account allocations

Accessing Clusters via the HPRC Portal

In Portal: "Clusters" dropdown → "Shell Access"

TAMU HPRC OnDemand (Grace) Files - Jobs -	Clusters -	Interactive Apps -	Dashboard -	My Interactive Sessions
	>_grace Sh	rell Access	Access	
OnDemand provid	les an inte	grated, single ac	cess point f	or all of your HPC resou
Message of the	e Day			

(Works with most web browsers)



Accessing Clusters via SSH

- SSH (Secure Shell) allows users to establish a connection between their local machine and the TAMU HPRC clusters.
- SSH Programs:

Operating System	Windows	MacOS	Linux
Programs	<u>MobaXTerm</u> * <u>PowerShell</u> <u>Windows Command Prompt</u> <u>PuTTY SSH</u>	Terminal*	Terminal*
	* Recommende	ed	

- If off-campus, connect with VPN first: <u>u.tamu.edu/VPnetwork</u>
- If on-campus or connected to VPN, log in with:

ssh [NetID]@grace.hprc.tamu.edu



Logging In to Clusters

<pre>************************************</pre>	Type in your NetID password (it won't show anything as you type)
<pre>Inter a passedue of select one of the following options: 1. Duo Push to XXX-XXX-1234 2. Phone call to XXX-XXX-1234 3. SMS passedes to XXX-XXX-1234 Passede or option (1-3): 1 Success. Logging you in Last login: Tue Sep 3 16:54:20 2024 from 128.194.35.38</pre>	Then sign in with TFA

High Performance Research Computing | hprc.tamu.edu | NSF Award #1925764 16

A M

Logging In to Clusters



Hands-on Exercise

- 1. Log in to the HPRC portal: <u>hprc.tamu.edu</u>
- Open a terminal on Grace in the portal:
 Clusters → _Grace Shell Access

OR

- 1. Log in via Terminal or MobaXterm
- 2. Use your NetID password and your two-factor authentication method.

Logged-In

Remember, when you first log in, you are on dedicated *login nodes*.

Grace has 5 of them (check your shell prompt to see which one you are on). These are not for running big processes! There are rules:

- No processes longer than 1 hr
- Sessions idle for 1 hr will be killed.
- Do not use more than 8 cores.
- Do not use sudo.

File Systems and User Directories

Directory	Environment Variable	Space Limit	File Limit	Intended Use
/home/\$USER	\$HOME	10 GB	10,000	Small to modest amounts of processing. Backed-up.
/scratch/user/\$USER	\$SCRATCH	1 TB	250,000	Temporary storage of large files for on-going computations. Not intended to be a long-term storage area; NOT backed up.

\$SCRATCH is shared between FASTER (TAMU/ACCESS) and Grace (TAMU) clusters.

View file usage and quota limits on the command line with:

showquota

Do NOT share your home or scratch directories. Request a group directory for sharing files.

https://hprc.tamu.edu/kb/User-Guides/Grace/Filesystems_and_Files/

AM

HPRC Portal Quota Increase Request

Portal Homepage → Grace Dashboard



Request quota increases directly from the dashboard with a guided form.

	Disk		File		
Disk	Usage	Limit	Usage	Limit	
ome	927M	10.0G	8,148	10,000	
	(9.05 %)		(81.48 %)		
cratch	477.8G	1.0T	169,224	250,000	Request
	(46.66 %)		(67.69 %)		Quota

months?	chart o
OYes No	
Current Scratch Quota	
1 TB	
New Scratch Quota	
	ТВ
Current File Limit	
250000	
New File Limit	
Justification (Required)	
What data is stored with requested quota?	
What job requires this quota increase?	
What is the input/output size of the job?	
What is your long-term plan for this data?	
Comment (Optional)	
	11
□ I verify that I will remove any unnecessary	/ data and

Hands-on Exercise

Check your file Quota:

 In the Portal: Dashboard → Grace Dashboard



2. In the shell: (use showquota)

<u>Ă</u>M

3. Locate your /scratch disk usage stats.

Portal File Browser

The Portal includes a file browser in which you can do many of the things a terminal would do... or just launch a terminal.

Ă M



TAMU HPRC OnDemand (Grace) Files	s∓ Jobs∓	Clusters 🕶	Interactive Apps	 Dashboard - 	My Interactive Sessions		? Help 🔻	Logged in as jwinchell	🔂 Log Out
					>_ Open in Terminal • + New	File New Directo	ry 🗘 Upload 🛃	Download	Telete
Home Directory /scratch/user/jwinchell		/ home	/ jwinchell /	Change directory					Copy path
						□ Show Owner/I	Mode 🛛 Show [Dotfiles Filter: Showing 10 of 43 rows -	0 rows selected
		Туре	^_ Name	•	*	Size	Modifie	ed at	
			fortra	n_tutorial	••	-	4/8/202	2 10:06:50 AM	
			geant	4_workdir	•).ex	3/21/20	22 10:11:25 AM	

Hands-on Exercise

- 1. Enter your Scratch directory from the Portal: Files → /scratch/user/<NetID>
- 2. Create or upload a file.
- 3. Move the file to your home directory.



Software Infrastructure

Software

- HPRC provides both pre-installed software and installation assistance.
- See the Software pages for instructions and examples:
 - hprc.tamu.edu/kb/Software
 - <u>hprc.tamu.edu/software</u>
- License-restricted software
 - Check on command line with <a>license_status (example later)
 - Contact help@hprc.tamu.edu
- Contact us for software installation help/request (can use dashboard).
 - User can install software in their home/scratch directories.
 - Do NOT run the "sudo" command when installing software.

Computing Environment

- PATH: the location on disk where an executable or library may be found.
- Paths are saved as environment variables, so you can choose which libraries and executables will be used by modifying the variables.
- There is a lot of software, many versions, and many paths to manage.

..... How do you manage all these software versions?

https://hprc.tamu.edu/kb/Software/useful-tools/Modules/

Modules

Managing software versions using lmod

• Each version of a software, application, library, etc. is available as a <u>module</u>.

• Module names have the format:

software-name/version[-dependency-version (optional)] TensorFlow/2.11.0-CUDA-11.7.0

- Loading a module adds its location on disk to your Path environment variable.
 - Most of its dependencies will also be loaded automatically.

Modules

- Modules are organized *hierarchically*.
- The base modules that must be loaded before anything else are called "Toolchains":
 - The same software may be available from multiple toolchains.
 - Do not mix up software from different toolchains!
- Loading the toolchain makes other modules available.



Modules: Toolchain Details

- Toolchains are combinations of compilers, MPI libraries, and highly optimized math libraries.
- Toolchain components are primarily either Intel or Open Source.

Example toolchains for C++ development:

Components	Open Source	Intel Source	Mixed Source
Compiler only	GCCcore	iccifort	_
Compiler + MPI	gompi	iimpi	iompi
Compiler + MPI + MKL, BLAS, FFTW, LAPACK	foss	intel	iomkl
Compiler + all of the above + CUDA Compiler	fosscuda	intelcuda	iomklc

Example usage: module load foss/2022a

https://hprc.tamu.edu/kb/Software/useful-tools/Toolchains/

A M

Module Usage Basics

module avail

- Lists all available modules (may be slow).
- Navigation:
 -spacebar, arrows, or j and k
 -quit with q
- Case-sensitive search: /
- Use **mla** instead to save results to a file as well.

module load <module>

 Add <module> paths to the current environment variables.

module spider <word>

- Case-sensitive search for modules with "word" in name.
- Provide an exact name to see dependencies.

HPRC	Texas A&M H @TexasAMHPRC 821 subscribers	IPRC						· ~
HOME	VIDEOS	PLAYLISTS	COMMUNITY	CHANNELS	ABOUT	\bigcirc modules		
	A series of the	HPRC Intro: Texas A&M HPRC This video will cov system and availa CC	IPRC Intro: The Modules System exas A&M HPRC + 577 views + 2 years ago nis video will cover some of the HPRC modules system commands rstem and available toolchains at the following links: https://hprc.tr			our Yout hel for r	tube nore	
High Padromena Manager States and States	HPRC Tutorial Herarchical Module System 5:06	HPRC Intro: Texas A&M HPRC Learn how to load Homepage: https: CC	Hierarchical Modul • 350 views • 1 year ago modules on Grace with the F //hprc.tamu.edu/ - TAMU HPI	e System	iem. —Useful Links— - T, nu.edu/wiki/ - TAMU			

Module Commands Reference

- Commonly-used module commands
- Note that on there are some shorthands for you to use.

Command	Shorthand	Description
module avail	mla (saves a searchable file as well)	See what modules can be loaded now.
module spider		Search for modules and find dependencies.
module list	ml	See what modules have been loaded already.
module load foo bar	ml foo bar	Load specified modules.
module unload foo bar module load baz goo	ml -foo -bar baz goo	Load and unload specified modules.
module purge	ml purge	Unload all modules.

Module Usage Example

Search for the software "snakemake"; it saves a file of available modules.

mla snakemake

snakemake/5.26.1-Python-3.8.2

snakemake/6.1.0

snakemake/6.10.0

snakemake/7.22.0

snakemake/7.32.3

snakemake/8.4.2

Choose a specific version and use it to search for its dependencies.

module spider snakemake/8.4.2

snakemake: snakemake/8.4.2

Description: The Snakemake workflow management system is a tool to create reproducible and scalable data analyses.

You will need to load all module(s) on any one of the lines below before the "snakemake/8.4.2" module is available to load.

GCC/12.3.0 OpenMPI/4.1.5

module load GCC/12.3.0 OpenMPI/4.1.5 snakemake/8.4.2

Load the base dependency module(s) first then the full module name.



Hands-on Exercise: Module Loading



A M

Module Usage Practices

- Software installed as modules are available to all users.
 - (except for restricted modules)
- It's a good habit to unload unused modules before loading new modules: <u>module purge</u>
- It is recommended to load a specific software version instead of the defaults: <u>m1 GCC/13.2.0</u> instead of m1 GCC
- Avoid loading modules in your \$HOME/.bashrc
- Avoid mixing toolchains when loading multiple modules at the same time. This usually leads to one of them not working.

https://hprc.tamu.edu/kb/Software/useful-tools/Modules/

Python and Virtual Environments

Python is a language which supports many external libraries in the form of extensions–called Python Packages.

Some commonly used packages:

•SciPy •NumPy •Jupyter notebook •Scikit-learn

Once you have loaded the appropriate Python module, you can install these additional packages yourself using the *Virtual Environment* feature. Instructions are on our KB:

https://hprc.tamu.edu/kb/Software/Python/#hprc-venv-management-tools

Our HPRC student workers have developed a "create_venv" tool to help!
Hands-on Exercise: Python Software Install

ml purge

٦.

al GCCcore/13.2.0 Python/3.11.5

2. create_venv myEnv

- 3. source activate_venv myEnv
- 4. python -c "import pytime"
- 5. pip install python-time
- 6. python -c "import pytime; print(pytime)"
- 7. deactivate

A M

-Set up the underlying Python module

- -Use HPRC's *create_venv* tool; Create a virtual environment in \$SCRATCH/virtual_envs
- -Activate virtual environment.

-Check if python-time is installed (it is not).

-Install python-time.

-Where is python-time installed?

-Close virtual environment.



High Performance Research Computing | hprc.tamu.edu | NSF Award #1925764 38

A M

Batch Computing on HPRC Clusters: Overview



<u>Ă</u>M



Output Files

Consumable Computing Resources

Computing resources are shared! You will not be able to use an infinite amount!

- Resources external to jobs:
 - Service Unit (SU)
 - Software license/token
- Resources specified in a job file:
 - Processor cores
 - Memory
 - Wall time
 - GPU

Software Licenses

To check the availability of licensed software on the clusters, use the command **license_status**.

- List software the tool knows about:
- Help for this command:
- Find available license for "ansys":

Exact availability varies by license policy. Contact us if you have questions. license_status -l license status -h

license_statu	s -s ai	nsys	
License status for ANSYS:			
License Name	# Issued	# In Use #	Available
aa_mcad	50	0	50
aa_r	50	32	18
aim_mp1	50	0	50

https://hprc.tamu.edu/kb/Software/useful-tools/License_Checker/

Service Units

- Service Units (SUs) are part of our account management system (AMS).
- With a default (Basic) allocation, you start with 20,000 SUs.
- SUs are charged as your jobs spend time on the compute nodes (we will see how later). (Work on login nodes is not charged, but you cannot do big computing there!)
- You can check your SU balance on both:
 - The command line
 - The HPRC Portal

https://hprc.tamu.edu/kb/User-Guides/AMS/Service_Unit/ https://hprc.tamu.edu/kb/User-Guides/AMS/UI/

Check your (SU) Balance: Command Line

• List the SU Balance of your Account(s) with:

myproject

============ I	ist of	YourN	======= etID's P	============== roject Accou	nts	=======================================	============
Account	FY	De	fault .	Allocation	Used & Pending SUs	Balance	PI
1228000223136	2024		N	10000.00	0.00	10000.00	Doe, John
1428000243716	2024		Y	5000.00	-71.06	4928.94	Doe, Jane
1258000247058	2024		N	5000.00	-0.91	4999.09	Doe, Jane

Run myproject -d <Account#> to change default project account.

(Replace <Account#> with your number!)

Run myproject -h to see more options.

Check your (SU) Balance: Portal

SUs can also be checked in the Dashboard.

(The same place we checked file quotas previously)

Ā Ň

TAMU HPRC OnDemand (Grace) Files -Jobs 🕶 Clusters -Dashboard -My Interactive Sessions Interactive Apps -High Performance Create Help Ticket **Research Computing** TAMU DASHBOARD (GRACE) **Request Software** Accounts **Core Utilization** Node Utilization Account 1 Default 1 Allocation 1 Used Balance 1 132748756209 default 20000 5.1 19994.9 Account Management **Queue Availability Disk Quotas** Oueue

CPU Avail Nodes Avail Disk File Disk Usage Limit Usage Limit home 927M 10.0G 8.148 10.000 (9.05 %) (81.48 %) **Group Memberships** scratch 477.8G 1 OT 169,224 250,000 Request jwinchell staff hprc mitchcomp (46.66 %) (67.69 %) Quota Increase Manage Groups

Hands-on Exercise

Check your SUs in both the command line and the Portal.
 Check that you have a default account set.

5 DASHBOARD	(GRACE)	SUMMARY		Nequest 301
Accounts				
Account ↑↓	Default ^{↑↓}	Allocation $\uparrow\downarrow$	Used ^{↑↓}	Balance ^{↑↓}
132853914631	Set Default	200000	200000	0
132853918233	default	5000	-169.43	4830.57

We will return to SUs once we have talked about SLURM and the resources your jobs use.

myproject

Slurm Job Scripts

(batch manager)

Compute Nodes

 \gg \gg \gg

Output Files



Sample Job Script Structure

#!/bin/bash

```
##NECESSARY JOB SPECIFICATIONS
#SBATCH --job-name=hello_world
#SBATCH --time=00:15:00
#SBATCH --ntasks=2
#SBATCH --ntasks-per-node=2
#SBATCH --nodes=1
#SBATCH --mem=3G
#SBATCH --output=hello world log.%j
```

load required module(s)

module purge
module load GCCcore/13.2.0
module load Python/3.11.5
python hello world.py

```
# Job Environment variables
echo $SLURM_JOBID
echo $SLURM_SUBMIT_DIR
echo $TMPDIR
echo $SCRATCH
```

This is a single-line comment and not run as part of the script.

These parameters describe your job to the job scheduler. The lines starting with #SBATCH are NOT just comments! See the <u>Knowledge Base</u> for more info.

Whatever commands or scripts you want to run. Here, we set up the modules we need for our environment, run a python program, and print out some environment variables.

Submit a Job and Check Job Status

Submit job

sbatch example01.job

Submitted batch job 6853258 (from job_submit) your job is charged as below Project Account: 122792016265 Account Balance: 1687.066160 Requested SUs: 3

Check status

squeu	le -u	\$USER	or	squeu	le -	-me					
JOBID	NAME	USER	PARTITION	NODES	CPUS	STATE	TIME	TIME_LEFT	START_TIME	REASON	NODELIST
6853258	jobname	someuser	xlong	2	96	RUNNING	3-07:36:50	16:23:10	2024-01-23T17:27:3	None	c[180,202]
6853257	jobname	someuser	xlong	2	96	RUNNING	3-07:36:56	16:23:04	2024-01-23T17:27:2	None	c[523-524]



Hands-on Exercise

Submit a simple job file:

- Navigate to /scratch/training/Intro-to-Grace, either on the command line or in the Portal's file browser.
- Copy the files hello_world.py to your home directory.
- 3. Return to your home directory yourself.
- 4. Submit the job file with: **sbatch hello_world.slrm**
- 5. Check the job's status with: **squeue -u \$USER**
- 6. Check the output file (will be named like hello_world_log.#).

squeue --me

or

Job Environment Variables

Each job has access to several self-referential variables:

- **\$SLURM_JOBID** = job ID
- \$SLURM_SUBMIT_DIR = directory from which the job was submitted
- \$TMPDIR = /work/job.\$SLURM_JOBID

You can also still use non-Slurm variables like:

- \$SCRATCH = /scratch/user/NetID

https://hprc.tamu.edu/kb/Helpful-Pages/Batch-Translation/#environment-variables

Important Batch Job Parameters

Slurm	Comment
#SBATCHtime=HH:MM:SS	Specifies the upper time limit for the job
#SBATCHntasks=x	Total number of tasks (cores) for the job
#SBATCHntasks-per-node=xx	Specifies the maximum number of tasks (cores) to allocate per node
<pre>#SBATCHmem=xxxxM or #SBATCHmem=xG</pre>	Sets the maximum amount of memory per <i>node</i> Can use M for MB or G for GB
#SBATCHnodes=x	Specifies the number of nodes to use

These usually all go in your job script file

https://hprc.tamu.edu/kb/Helpful-Pages/Batch-Translation/#job-specifications



(Job Template)

```
#!/bin/bash
```

```
##NECESSARY JOB SPECIFICATIONS
#SBATCH --job-name=JobExample1 #Set the job name to "JobExample1"
#SBATCH --time=01:30:00 #Set the wall clock limit to 1hr and 30min
#SBATCH --ntasks=1 #Request 1 task
#SBATCH --ntasks-per-node=1 #Request 1 task/core per node
#SBATCH --mem=2560M #Request 2560MB (2.5GB) per node
#SBATCH --output=Example1Out.%j #Send stdout/err to "Example1Out.[jobID]"
#First Executable Line
```

```
-> module purge
```

I recommend resetting your environment first and loading the necessary modules in the job script. This helps with reproducibility.

https://hprc.tamu.edu/kb/Quick-Start/Grace/#running-your-program-preparing-a-job-file

52

Mapping Jobs to Cores per Node on Grace

Α.



48 cores on 1 compute node

#SBATCH --ntasks=48 #SBATCH --tasks-per-node=48

Preferred Mapping (if applicable)





48 cores on 2 compute nodes

#SBATCH --ntasks=48 #SBATCH --tasks-per-node=24

Ă M

Slurm Pop Quiz

- **#SBATCH** --job-name=stacks S2
- **#SBATCH** --ntasks=80
- **#SBATCH** --ntasks-per-node=20
- **#SBATCH** --mem=40G

Å M

- **#SBATCH** --time=48:00:00
- **#SBATCH** --output stdout.%J
- **#SBATCH** --error stderr.%J

How many nodes is this job requesting?

- A. 1600 C. 20
- B. 80 D. 4

Job Memory Requests on Grace

- Specify memory request based on memory per node: #SBATCH --mem=xxxxM # memory per node in MB
 - or #SBATCH --mem=xG # memory per node in GB
- On 384GB nodes, usable memory is at most 360 GB. The per-process memory limit should not exceed ~7500 MB for a 48-core job.
- On 3TB nodes, usable memory is at most 2900 GB. The per-process memory limit should not exceed 37120 MB for a 48-core job.

Slurm: Examples of SUs charged based on Job Cores, Time and Memory Requested

An SU is equivalent to one core hour *or* a proportional amount of memory on the node for one hour.

On **Grace**: a typical node has 48 cores and 360 GB usable for jobs. (7.5 GB per core)

Number of Cores	GB of memory per core	Total Memory (GB)	Hours	SUs charged	
]	7.5	7.5]]	
24	ask for More	24]	24	
]	360	360]	48	
48	7.5	360]	48	

https://hprc.tamu.edu/kb/User-Guides/AMS/Service_Unit/



Slurm: SU Charges for GPU jobs

GPUs will use more SUs per hour than CPUs.

Effective GPU	SU charge per one hour (wall_time)
A100	72
RTX 6000	48
T4	24

Job Submission and Tracking

Slurm commands	Description
sbatch jobfile1	Submit jobfile1 to batch system
squeue [-u user_name] [-j job_id]	List jobs
scancel job_id	Kill a job
sacct -X -j job_id	Show information for a job (can be when job is running or recently finished)
sacct -X -S YYYY-HH-MM	Show information for all of your jobs since YYYY-HH-MM
pestat -u \$USER	Show resource usage for a running job
seff job_id	Check CPU/memory efficiency for a job

https://hprc.tamu.edu/kb/Helpful-Pages/Batch-Translation/#job-specifications



Batch Queues

- Job submissions are auto-assigned to batch *queues* (also called *partitions*) based on the resources requested:
 - number of cores/nodes and walltime limit
 - specific resources requested
- Some jobs can be directly submitted to a queue:
 - E.g. if gpu nodes are needed, use the gpu partition/queue:
 #SBATCH --partition=gpu
- Batch queue policies are used to manage the workload and may be adjusted periodically.

https://hprc.tamu.edu/kb/User-Guides/Common/BatchProcessing/#batch-queues

Checking Resources for Jobs

sinfo pestat gpuavail maxconfig

There are several tools available to check what you can use in your jobs...



Checking Resources for Jobs: Queues

	S	sinfo per	stat gpu	lavail	maxconfi	g
username@login:	1:~\$ si:	nfo				
PARTITION	AVAIL	TIMELIMIT	JOB SIZE	NODES (A/I,	/О/Т) С	PUS(A/I/O/T)
short*	up	2:00:00	1-32	711/0/89/8	300 3	3689/439/4272/38400
medium	up	1-00:00:00	1-128	711/0/89/8	300 3	3689/439/4272/38400
long	up	7-00:00:00	1-64	711/0/89/8	300 3	3689/439/4272/38400
xlong	up	21-00:00:00	1-32	711/0/89/8	300 3	3689/439/4272/38400
vnc	up	12:00:00	1-32	100/5/12/1	L17 9	85/4055/576/5616
gpu	up	4-00:00:00	1-32	100/5/12/1	L17 9	85/4055/576/5616
bigmem	up	2-00:00:00	1-4	7/0/1/8	5	60/0/80/640
staff	up	infinite	1-infinite	817/14/101	L/932 3	4689/5199/4848/4473
special	up	7-00:00:00	1-infinite	811/5/101,	/917 3	4674/4494/4848/4401
gpu-a40	up	4-00:00:00	1-15	6/9/0/15	1	5/705/0/720

Shows job queue status

ĀŇ

For the NODES and CPUS columns: A = Active (in use by running jobs) I = Idle (available for jobs) O = Offline (unavailable for jobs) T = Total

Checking Resources for Jobs: Nodes

sinfo

pestat gpu

maxconfig

username@login1:~\$ pestat -p gpu -G Print only nodes in partition gpu GPU GRES (Generic Resource) is printed after each jobid												
Hostname	Partition	Node Nu	m CP	U	CPUload	Memsize	Freemem	GRES/node	Joblist			
		State U	se/T	ot	(15min)	(MB)	(MB)		JobID (Job	oArrayID)	User GRES/job	•
g001	gpu	mix	2	48	1.94	368640	271338	gpu:a100	:2 <jobid></jobid>	<usernam< td=""><td>ie> <gpus></gpus></td><td></td></usernam<>	ie> <gpus></gpus>	
g002	gpu	drain*	0	48	0.00*	368640) 0	gpu:a100:	2			
g003	gpu	drain*	0	48	0.01	368640	370427	gpu:a100	: 2			

Shows status of nodes

Notable options:

- -p <partition name> show only a specific partition
- -u \$USER

• -G

- show only specified user's jobs
 - show "generic resources" (e.g. the gpus used)

Checking Resources for Jobs: GPUs

sinfo

P

cat 🛛

gpuavail

maxconfig

See what resources in the GPU nodes are available

Ă M

username	eloginl	:~\$ gpuav	all			
CONF NODE	IGURATI	ON NODE				
TYPE		COUNT				
gpu:a1 gpu:a4 gpu:rt: gpu:t4	00:2 0:3 x:2 :4	100 15 9 8				
	AVA	AILABILIT	Y			
NODE	GPU	GPU	GPUs	CPUs	GB MEM	
NAME	TYPE 	COUNT	AVAIL	AVAIL	AVAIL	
g106	a40	3	1	40	260	
t001	t4	4	1	45	345	
t006	t4	4	2	20	310	

Checking Resources for Jobs: SUs

maxconfig If you ask for too many resources, your job will not run. You can check beforehand: username@login1:~\$ maxconfig Grace partitions: short medium long xlong vnc gpu bigmem special gpu-a40 Grace GPUs in qpu partition: a100:2 a40:3 rtx:2 t4:4 Showing max parameters (cores, mem, time) for partition long CPU-billing * hours * nodes = SUs Check specific partitions with: 48 * 168 * 1 = 8,064 #!/bin/bash maxconfig -p <partitionName> #SBATCH --job-name=my job #SBATCH --time=7-00:00:00 Estimate the SUs a job will require with: #SBATCH --nodes=1 #SBATCH --ntasks-per-node=1 maxconfig -f <jobScriptName>

#SBATCH --cpus-per-task=48

#SBATCH --output=stdout.%x.%j #SBATCH --error=stderr.%x.%j

#SBATCH --mem=360G

Job Summary: myjob

Command-line tool to show you details of a specific job:

myjob

Shows, e.g.:

- Submission details
- Resource usage of finished jobs
- Status of pending job

[u.ab12345@aces-login2 ~]\$ myjob 9814093

Job TD: 9814093 Cluster: grace User/Group: userid/userid Account: 123456789101 SUs charged: 283.82 State: COMPLETED (exit code 0) Partition: long Node Count: 1 NodeList: c121 Cores per node: 48 CPU Utilized: 01:07:56 CPU Efficiency: 0.42% of 11-08:28:00 core-walltime Submit time: 2024-02-24 17:32:41 Start time: 2024-02-24 17:32:50 End time: 2024-02-24 23:13:25 Job Wall-clock time: 05:40:35 Memory Utilized: 40.95 GB Memory Efficiency: 11.38% of 360.00 GB Job Name: parafold Job Submit Directory: /scratch/user/netid/myproject Submit Line: sbatch run parafold grace.sh

Portal Job Composers



Some job management can also be done from the Portal.

- "Active Jobs" is like squeue.
- "Job Composer" lets you:
 - Create job scripts
 - Save job templates
 - Monitor your own jobs

We will focus on the new "Drona" options:

- "Drona Joblisting" shows current and past jobs.
- "Drona Composer" helps you build Slurm files via a form.

Drona Composer

Clusters *

Int

Jobs

Ā M

Import additional environments



Drona Composer

	High Performance Research Computing DIVISION OF RESEARCH	DRONA COMPOSER (GRACE)
	Job Composer	
	Job Name	
Generic environment to	Location	ge /scratch/user/jwinchell/drona_composer/runs
create custom Slurm	Environments	ric +
batch job —	Upload files ⑦	an option v Add
	Add modules 💿	Default (foss/2023b) v
	Add	
Add modules for iob	Number of tasks ⑦	Advanced task options ⑦
	1	
	Use Accelerator 💿	
Fill out info that	Choose an option	
Slurm would need	TOTAL Memory ⑦ Expected	ed run time ⑦ Project Account ⑦ Additional Slurm parameters ⑦
	GB v Days	O Hour ○ Minu ○ Choose an option ✓
		Preview

High Performance Research Computing | hprc.tamu.edu | NSF Award #1925764 68

AM

Drona Composer



Other Type of Jobs

- Visualization:
 - <u>portal.hprc.tamu.edu</u>
 - Interactive Apps > choose application
- Large number of concurrent single-core jobs
 - Check out tamulauncher:
 - Useful for running many single core commands concurrently across multiple nodes within a job
 - Can be used with serial or multi-threaded programs
 - If a tamulauncher job gets killed, you can resubmit the same job to complete the unfinished commands in the input file.
 - <u>https://hprc.tamu.edu/kb/Software/tamulauncher/</u>

Jobs - Cluster	s Interactive Apps Dashboard	➡ ■ My Interact
OnDemand pro Message o	BIO BEO CRISPR-Local CRISPR-Local Gap5 BIO BIO RNAlysis Structure XtalOpt	aint for all of yo
 Unauthorized u Use of HPRC remembers are U Sharing HPRC Authorized use 	RC re GUI BRC : S ANSYS Workbench I use Abaqus/CAE AMATLAB IERE M ParaView W VNC	d and subject t tes export cont n is in violation s at: https://hp ISER HOME DIRI
	Imaging & ChimeraX Diffusion Toolkit & TrackVis Fiji VMD cisTEM cryoSPARC 3.3.1	
	Servers	

Interactive Apps

- Select an app from the dropdown menu.
- 2. Specify environment and resources.
- 3. Hit "Launch" at the bottom.
- 4. On the "My Interactive Sessions" screen:
 - a. wait for box to turn green
 - b. connect to app
 - c. quit from the app when you are done
 - d. box should turn gray

IPRC OnDemand (Grace) Files	s ▼ Jobs ▼ Clusters ▼ Inte	eractive Apps 🔻 Dashboard 🏲 🚽 💡 🝷 💄	6
Home /	My Interactive Sessions / Jup	pyter Notebook	
Interactiv	ve Apps Jupyte	er Notebook	
BIO	This app wil	ill launch a Jupyter Notebook server on the Grace	
📊 Beauti	i cluster.		
🖉 CRISP	R-Local Type of env	oad + Python virtualenv v	
∞ Gap5	Select the ty	ype of environment in which Jupyter is installed.	
∎= IGV	Module sele	lected	
😫 Mauve	e Python/3.	.8.6 ~	
O RNAW	Select a mod	dule to load. All modules listed will also load the	

J HPRC OnDemand (Grace) Files - Jobs	 Clusters Interactive Apps Dashboard	🤁 Help 👻 💄 Logged in as jwinchell 🛛 🖨 Log Ou
Home / My Interactive Session	S	
Interactive Apps	Jupyter Notebook (9018862)	1 node 1 core Running
BIO		
T Beauti		Delete
	Created at: 2023-09-14 11:59:50 CDT	
Children Contain	Time Remaining: 58 minutes	
∞ Gap5	Session ID: cc1adb21-9686-4703-b7f7-a26224f2e8b7	
ia≓ IGV		
😭 Mauve		
© RNAlysis	CRISPR-Local (9018853)	Completed
T Structure	C	
XtalOpt	Created at: 2023-09-14 11:52:44 CD1	The Delete
	Session ID: 089ab796-a8fd-4e9f-8558-828395af7d53	

Need Help?

First check the FAQ <u>https://hprc.tamu.edu/kb/FAQ/Accounts/</u>

- Grace User Guide <u>https://hprc.tamu.edu/kb/User-Guides/Grace/</u>
- Email your questions to help@hprc.tamu.edu

Help us help you -- we need more info:

- Which Cluster
- Username
- Job id(s) if any
- Location of your jobfile, input/output files
- Application used if any
- Module(s) loaded if any
- Error messages
- Steps you have taken, so we can reproduce the problem

Remember you can start help requests from the Dashboard!


High Performance Research Computing DIVISION OF RESEARCH Give us feedback on the class with this survey: <u>https://u.tamu.edu/hprc_shortcourse_survey</u>

Thank you

Questions?



High Performance Research Computing | hprc.tamu.edu | NSF Award #1925764

Batch Job Examples

High Performance Research Computing | hprc.tamu.edu | NSF Award #1925764 74

Slurm Job File (Serial Example)

#!/bin/bash

##NECESSARY JOB SPECIFICATIONS

#SBATCH --job-name=JobExample1

#SBATCH --time=01:30:00

#SBATCH --ntasks=1

#SBATCH --mem=6G #Re **#SBATCH --output=Example1Out.**%j

bExample1 #Set the job name to "JobExample1" 00 #Set the wall clock limit to 1hr and 30min #Request 1 task #Request 6GB per node ple1Out.%j #Send stdout/err to "Example1Out.[jobID]"

SUs = 1.5

##OPTIONAL JOB SPECIFICATIONS ##CHANGE ACCOUNT NUMBER AND EMAIL ADDRESS BEFORE USING ##SBATCH --account=123456 #Send email on all job events ##SBATCH --mail-type=ALL #Send email on all job events ##SBATCH --mail-user=email address #Send all emails to email address

this intel toolchain is just an example. recommended toolchain is TBD
module purge
module load intel/2022a

run your program ./helloworld.omp.C.exe

Slurm Job File (multi core, single node)

#!/bin/bash

##NECESSARY JOB SPECIFICATIONS

#SBATCH --job-name=JobExample2

#SBATCH --time=6:30:00
#SBATCH --nodes=1

#SBATCH --ntasks-per-node=8

#SBATCH --mem=8G

#SBATCH --output=Example2Out.%j

- # Set the job name to "JobExample2"
 # Set the wall clock limit to 6hr and 30min
 # Request 1 node
 # Request 8 tasks(cores) per node
 SUs = 52
- # Request 8GB per node
- # Send stdout/err to "Example2Out.[jobID]"

load required module(s)

module purge
module load intel/2022a

Ā Ň

run your program mpirun ./helloworld.mpi.C.exe

Slurm Job File (multi core, multi node)

#!/bin/bash

##NECESSARY JOB SPECIFICATIONS **#SBATCH** --job-name=JobExample3 # Set the job name to "JobExample3" #SBATCH --time=1-12:00:00 # Set the wall clock limit to 1 Day and 12hr **#SBATCH** --ntasks=8 # Request 8 tasks (cores) SUs = 288 **#SBATCH** --ntasks-per-node=2 # Request 2 tasks(cores) per node **#SBATCH** --mem=5G # Request 5GB per node **#SBATCH** --output=Example3Out.%j # Send stdout and stderr to "stdout.[jobID]" ##OPTIONAL JOB SPECIFICATIONS ##CHANGE ACCOUNT NUMBER AND EMAIL ADDRESS BEFORE USING ##SBATCH --account=123456 # Set billing account to 123456 #find your account with "myproject"

##SBATCH --mail-type=ALL # Send email on all job events
##SBATCH --mail-user=email address # Send all emails to email address

this intel toolchain is just an example. recommended toolchain is TBD
module purge
module load intel/2022a

run program with MPI
mpirun ./helloworld.mpi.C.exe

Slurm Job File (serial GPU)

#!/bin/bash







SUs = 72



load required module(s)

module purge
module load CUDA/11.8.0

run your program
my_cuda_enabled_program

Slurm Job File (multi GPU)

#!/bin/bash

##NECESSARY JOB SPECIFICATIONS

#SBATCH	job-name=JobExample5
#SBATCH	time=01:00:00
#SBATCH	ntasks=2
#SBATCH	nodes=1
#SBATCH	mem=250G
#SBATCH	output=Example4Out.%j
#SBATCH	gres=gpu:a100:2
#SBATCH	partition=gpu

```
# Set the job name to "JobExample4"
# Set the wall clock limit to 1hr
# Request 1 task (core)
                                               SUs = 179
# Request 250GB per node
# Send stdout and stderr to "Example4Out.[jobID]"
# Request 2 A100 GPUs
# Request the GPU partition/queue
```

##OPTIONAL JOB SPECIFICATIONS ##CHANGE ACCOUNT NUMBER AND EMAIL ADDRESS BEFORE USING **##SBATCH** --account=123456 **#** Set billing account to 123456 #find your account with "myproject" ##SBATCH --mail-type=ALL # Send email on all job events ##SBATCH --mail-user=email address # Send all emails to email address

load required module(s)

module purge module load CUDA/11.8.0

run your program my cuda enabled program

High Performance Research Computing | hprc.tamu.edu | NSF Award #1925764