HIGH PERFORMANCE RESEARCH COMPUTING

- HPRC Primer -Introduction to Grace: An HPRC Resource

February 7, 2025



High Performance Research Computing DIVISION OF RESEARCH



Grace Hardware

Grace is a 925-node Intel cluster from Dell with an InfiniBand HDR-100 interconnect, A100 GPUs, RTX 6000 GPUs and T4 GPUs. The 925 nodes are based on the Intel Cascade Lake processor.

48 cores/node

3TB Large Memory-80 cores/node Login Nodes: 10 GbE TAMU network connection

Resource	Count
Login Nodes	5
384GB memory general compute nodes	800
GPU - A100 nodes with 384GB memory	100
GPU - RTX 6000 nodes with 384GB memory	9
GPU - T4 nodes with 384GB memory	8
3TB Large Memory	8



For more information: <u>https://hprc.tamu.edu/kb/User-Guides/Grace/</u>



Computing on HPRC Clusters



Two sets of nodes on our clusters:

- Login nodes: log in, perform basic commands, write job scripts, and send job scripts to...
- **Compute nodes**: executes jobs, sends results back to you

Contains batch manager job parameters along with Unix and software commands.

Batch Jobs on HPRC Clusters

File Quotas and Resource Allocations

• Two things to keep track of when computing

Computing Resources Service Units (SUs)

Accessing Grace: Setup

• If off-campus:

Set up and start VPN (Virtual Private Network): <u>u.tamu.edu/VPnetwork</u>

- Two-Factor Authentication required
- Today we'll access Grace via the online Portal, but you can also use ssh.
- See https://hprc.tamu.edu/kb/User-Guides/Grace/Access/ for more details.

Accessing Grace via the Portal

Access the HPRC portals through most web browsers:

- 1. Go to <u>portal.hprc.tamu.edu</u> or use the <u>Portal dropdown menu</u> on the HPRC homepage: <u>https://hprc.tamu.edu/</u>
- 2. Choose Grace Portal

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

Accessing Grace via the Portal

*shell is also called terminal or command line

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

Hands-On Activity - 2 Minutes

Try to access a *shell** on Grace now, either through <u>portal.hprc.tamu.edu</u> or <u>hprc.tamu.edu</u>

*(also called *terminal* or *command line*)

What message do you see when you login?

Remember Grace has 5 login nodes. Which one does your command prompt say you got?

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 nightly.
/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.
				*Do NOT share your home or scratch directories.

Request a group directory for sharing files.

\$SCRATCH is shared between the FASTER and Grace clusters.

View file usage and quota limits in the shell using the command:

Your current disk quotas are: Disk Disk Usage Limit File Usage Limit /home/tjq 769M 10.0G 8146 10000 /scratch/user/tjq 1.6G 1.0T 28196 250000

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

High Performance Research Computing | hprc.tamu.edu

showquota

Portal: Grace Dashboard

- Easily view Cluster utilization, <u>Storage</u> <u>Quotas</u>, & <u>Allocation</u> <u>Balances</u>
- Ask for help and request software
- Also View current groups and job listings in the Grace Dashboard

Hands-On Activity - 2 Minutes

1. Please try to access dashboard now through the portal.

1. Check your quotas both on the command line and on the dashboard.

showquota

Checking Your Service Unit (SUs) Balance

- 1SU = 1 core/hr (GPUs are more expensive per-hour!)
- SUs are charged to default account when none is specified.

Checking Your SUs in the Shell

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

=======================================	List of	======================================	eeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeeee		
Account	 FY	Default	Allocation Used	l & Pending SUs	Balance PI
1228000223136	2023	N	10000.00	0.00	10000.00 Doe, John
1428000243716	2023	Y	5000.00	-71.06	4928.94 Doe, Jane
1258000247058	2023	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

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

myproject

Hands-On Activity - 2 Minutes

1. Use **myproject** to check the SU balance of your accounts.

2. Use the dashboard the check the same information.

Software

- Search for software modules on <u>https://hprc.tamu.edu/software/grace/</u>
- See the Software Knowledge Base page <u>https://hprc.tamu.edu/kb/Software/</u> for instructions and examples
- License-restricted software
 Contact <u>help@hprc.tamu.edu</u>

Last Updated: Mon Jun 2 The available software for	7 12:24:16 CDT	litional documentation if ava	
etc.			
Show 10 * entries	Search:	numpy	
Name	✓ Description		
awkward	'Awkward Array is a library for nested, variable-sized data, including arbitrary-length lists, records, mixed types, and missing da	ta, using NumPy-like idioms.	
bcolz	Tocolz provides columnar, chunked data containers that can be compressed either in-memory and on-disk. Column storage allows for efficiently querying tables, as well as for cheap column addition and removal. It is based on NumPy, and uses it as the standard data container to communicate with bcolz object but it also comes with support for import/export facilities tablem AIPS/Fighted stables and pondos dataframes."		
Bottleneck	'Fast NumPy array functions written in C'		
FablO	FablO is an I/O library for images produced by 2D X-ray detectors and written in Python. FablO support images detectors from a dozen of companies (includi Mar, Dectris, ADSC, Hamamatsu, Oxford,), for a total of 20 different file formats (like CBF, EDF, TIFF,) and offers an unified interface to their headers (as a python dictionary) and datasets (as a numpy ndarray of integers or floats). ¹		
jax	Composable transformations of Python+NumPy programs: differentiate, vectorize, JIT to GPU/TPU, and more		
mkl_fft	'NumPy-based Python interface to Intel(R) MKL FFT functionality'		
netcdf4-python	'Python/numpy interface to netCDF.'		
numba	Numba is an Open Source NumPy-aware optimizing compiler for Python sponsored by Continuum Analytics, Inc. It uses the re infrastructure to compile Python syntax to machine cade. ¹	markable LLVM compiler	
numexpr	The numexpr package evaluates multiple-operator array expressions many times faster than NumPy can. It accepts the expression as a string, analyzes it, rewrites it more efficiently, and compiles it on the fly into code for its internal virtual machine (VM). Due to its integrated just-in-time (IIT) compiler, it does n require a compiler at runtime.		
numpy	NumPy is the fundamental package for scientific computing with Python. It contains among other things: a powerful N-dimen (broadcasting) functions, tools for integrating C/C++ and Fortran code, useful linear algebra, Fourier transform, and random ni obvious scientific uses, NumPy can also be used as an efficient multi-dimensional container of generic data. Arbitrary data-typ NumPy to scentisely on also being the with a wide write with a dide write write with a dide write write write and write write a dide write write a dide write write a dide write write a dide write write and write write a dide write write and write write a dide write write and write write and write write and write write write and write write and write write and write write write and write write and write write and write write and write write write write and write write write and write write and write write and write write write and write write write and write write and write write write and write w	sional array object, sophistic umber capabilities. Besides i bes can be defined. This allow	

- Contact HPRC (can use the dashboard) for software installation help/request
 - User can install software in their home/scratch directory
 - Do NOT run the sudo command when installing software

Software: Application Modules

- Installed applications are made available with the module system
- Grace uses a *software hierarchy* inside the module system
- In this hierarchy, the user loads a compiler which then makes available Software built with the currently-loaded compiler

Software: Modules and Toolchains

- Toolchains are what we call groups of compilers & libraries
- There's a variety of toolchains available on the clusters:
 - o intel/2023b
 - o iomkl/2020a
 - o foss/2023b

- (more than just these versions)
- GCCcore/13.2.0
- Other module commands:

module spider module purge ← search for modules and their dependencies ← removes all loaded modules

Module System Youtube Video → <u>www.youtube.com/watch?v=drxpbrOCPFw</u>

Remember: module load

module purge

1. Please search for and load the following module:

OpenMPI/4.1.6

(Tip) Type this to show which compiler needs to be loaded:
 module spider OpenMPI/4.1.6
 (Tip) And check that it's been loaded with:
 module list

2. Next remove (unload) all your current modules.

Module System Youtube Video → <u>www.youtube.com/watch?v=drxpbrOCPFw</u>

The Drona Composer

• A simple app to assist you with composing jobs

The Drona Composer

	High Performance Research Computing DIVISION OF RESEARCH	DRONA COMPOSER (GRACE)	
Create your own jobs			
through the app:	Job Composer		
	Job Name		
(1) Pick a job name	Location	Change /scratch/user/jwinchell/drona_composer/runs	
(2) Chose location and	Environments	Generic	~ +
environments	Upload files 🕖	Select an option ~ Add	
(3) Add modules	Add modules 🕜		Default (foss/2023b) v
(4) Chose Resources	Number of tasks ⑦	Add Advanced task options (2)	
(5) Chose Time	1		
	Use Accelerator 📀		
	Choose an option		~
	TOTAL Memory 💿	Expected run time ⑦ Project Account ⑦	Additional Slurm parameters 📎
	GB V	Days C Hour Minu C Choose an option V	
		Preview	Hide History

Hands-On Activity - 3 Minutes

Using the Drona composer, create a job script with the following criteria:

- Generic Environment
- Modules Loaded: foss/2023b, Python/3.11.5
- 6 Cores
- 2 Nodes
- No GPU
- 30 GB of Memory
- 3 Minute Runtime

Preview the job script when you are done

Sample Job Script Structure

#!/bin/bash

```
##NECESSARY JOB SPECIFICATIONS
#SBATCH --export=NONE
#SBATCH --get-user-env=L
#SBATCH --job-name=JobExample1
#SBATCH --time=01:30:00
#SBATCH --time=01:30:00
#SBATCH --ntasks=1
#SBATCH --mem=2G
#SBATCH --output=stdout.%j
```

```
##OPTIONAL JOB SPECIFICATIONS
#SBATCH --account=123456
#SBATCH --mail-type=ALL
#SBATCH --mail-user=email_address
```

```
# load required module(s)
```

module purge
module load GCCcore/11.3.0 Python/3.10.4

```
# Run your program
python my_program.py
```

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

Account number to be charged

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

(We will practice with job files in a few slides)!

Submit a Job and Check Job Status

Submit job

Hands-On Activity

- 1. Navigate to /scratch/training/Intro-to-Grace
- Copy hello_world.slurm and hello_world.py to your home directory
- Return to your home directory and submit the job file using sbatch.
- 4. Check that the job is running in a Slurm queue with **squeue**.
- 5. When your job completes, check the contents of the output file.

Batch Queues

Job submissions are auto-assigned to batch queues based on the resources requested (e.g. number of cores/nodes and walltime limit)
Use sinfo to check their status:

[NetID@grace2 ~]\$ sinfo					
PARTITION	AVAIL	TIMELIMIT	JOB_SIZE	NODES $(A/I/O/T)$	CPUS(A/I/O/T)
short*	up	2:00:00	1-32	687/97/16/800	30786/6758/856/38400
medium	up	1-00:00:00	1-128	687/97/16/800	30786/6758/856/38400
long	up	7-00:00:00	1-64	687/97/16/800	30786/6758/856/38400
xlong	up	21-00:00:00	1-32	687/97/16/800	30786/6758/856/38400
vnc	up	12:00:00	1-32	104/12/1/117	895/4633/88/5616
gpu	up	4-00:00:00	1-32	104/12/1/117	895/4633/88/5616
bigmem	up	2-00:00:00	1-4	0/7/1/8	0/560/80/640
staff	up	infinite	1-infinite	791/109/17/917	31681/11391/944/4401
special	up	7-00:00:00	1-infinite	791/109/17/917	31681/11391/944/4401
gpu-a40	up	10-00:00:00	1-15	15/0/0/15	45/675/0/720

For the NODES and CPUS columns: A = Active (in use by running jobs) I = I

O= Offline (unavailable for jobs)

I = Idle (available for jobs) T = Total

https://hprc.tamu.edu/kb/User-Guides/Grace/Batch/#batch-queues

Job Submission and Tracking

Slurm queue command	Description
<mark>sbatch</mark> jobfile1	Submit jobfile1 to batch system
<mark>squeue</mark> [-u user_name] [-j job_id]	List jobs
<pre>scancel job_id</pre>	Kill a job
<pre>sacct -X -j job_id</pre>	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
lnu job_id	Show resource usage for a job
pestat -u \$USER	Show resource usage for a running job
<pre>seff job_id</pre>	Check CPU/memory efficiency for a job

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

Need Help?

First check the <u>FAQ</u>

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

Help us help you -- when you contact us, tell us:

- Which Cluster you're using
- Your 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

Continued Learning

Intro to HPRC Video Tutorial Series

HPRC's Knowledge Base

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>

Questions?

ĀМ