



The University of Texas at San Antonio™

## **Arc New User Training – with DGXs**

October 28, 2022

Brent League, Mando Rodriquez, Richard Zanni, Ritu Arora, Zhiwei Wang, Thomas Freeman

Email: [rcsg@utsa.edu](mailto:rcsg@utsa.edu)

# Overview

- High-Level Architecture of Arc
- Requesting accounts on Arc
- Connecting to Arc
- OnDemand Portal
- Arc user environment, Filesystems
- Job Submission
- Checkpointing and Distributed Computing
- Jupyter and Changing Defaults for File Writing
- DGX-related
- Trainings, Workshops, and Consulting
- Requesting support and Fair-Use Policies



# Introducing Arc – the Large-Scale Computing System at UTSA

- Arc is the primary large-scale computing system at The University of Texas at San Antonio (UTSA) that can be used for running data-intensive, memory-intensive, and compute-intensive jobs from a wide range of disciplines
- It is built using 111 new compute nodes and 56 old compute nodes from the now retired Shamu cluster
  - Out of the 56 old compute nodes from Shamu, 20 nodes are owned privately by a research group, 4 nodes are reserved for system administration and testing purposes
  - 111 new compute nodes including 2 big memory nodes, and 30 old nodes from Shamu are available for general use by the rest of the UTSA community
  - Not included in the 111 compute node count are the three DGX servers that are private to the Matrix AI Consortium
- Arc supports more than 239 customers from multiple departments at the time of this presentation
  - College of Business
  - College of Education and Human Development
  - College of Engineering and Integrated Design
  - College of Liberal and Fine Arts
  - College of Science
  - Strategic Research Initiatives
- Arc customers include UTSA undergraduate and graduate students, faculty, and staff and also some external collaborators of UTSA faculty

# Architecture of Arc

- Arc consists of 167 total compute/GPU nodes and 2 login nodes, majority of these have Intel Cascade Lake CPUs and some have AMD EPYC CPUs
  - 30 GPU nodes - each containing two CPUs with 20 cores each for a total of 40 cores, 384GB RAM, and each including one V100 Nvidia GPU accelerator
  - 5 GPU nodes - each containing two CPUs with 20 cores each for a total of 40 cores, 384GB RAM, and each including two V100 Nvidia GPU accelerators
  - 2 GPU nodes - each containing two CPUs and 4 V100 GPUs, and 384 GB RAM
  - 2 GPU nodes - each having two AMD EPYC CPUs and having one A100 80 GB GPU, and 1 TB RAM
  - 2 large-memory nodes, each containing four CPUs with 20 cores each for a total of 80 cores, and each including 1.5TB of RAM
  - 1 large-memory node, equipped with two AMD EPYC CPUs and 2 TB of RAM
  - 1 node equipped with two AMD EPYC CPUs and having 1 TB of RAM
  - 5 nodes - each equipped with two AMD EPYC CPUs and 1 NEC vector engine and 1 TB of RAM
- 100Gb/s Infiniband connectivity
- Two Lustre filesystems: /home and /work, where /home has 110 TBs capacity and /work has 1.1 PB of capacity
  - A secondary storage system named /vault is mounted on the login nodes as well
- Approximately 1.5 TB of /scratch space is available on most compute/GPU nodes)

# Thomas



# HOW CAN WE HELP?



ACCESS, INTERNET AND  
WEBSITE



BUSINESS DEPARTMENTS  
AND APPLICATIONS



COMMUNICATIONS (AUDIO  
VISUAL), COLLABORATION  
AND EVENTS



COMPUTERS AND DEVICES



DATA AND DOCUMENTS



HELP, ADVISORY AND  
CONSULTING



RESEARCH



PUBLIC SAFETY AND  
SECURITY



TEACHING (CLASSROOM),  
LEARNING AND STUDENTS  
SERVICES

HOME > ALL CATALOGS > RESEARCH

Search

**CATALOGS**

Research

**CATEGORIES**

- HPC Arc Services
- Training and Advising
- Additional Research Services
- Research IT

**Popular Items**

**New HPC Account**

Request a new HPC Account

[View Details](#)

**Arc HPC Application Support**

Request for Support or for new research applications related to HPC

[View Details](#)

**Other Support Request (Rese...**

Any other research related issues

[View Details](#)

**Mascot, Galaxy, RedCap**

To obtain support, access or to request a modification in Mascot, Galaxy, or RedCap

[View Details](#)

**Training / Consulting**

We provide training, advising, and consulting on the use of the HPC environment, as well as multiple applications

[View Details](#)

**Website Update / Application ...**

Request changes to the research website or other research applications

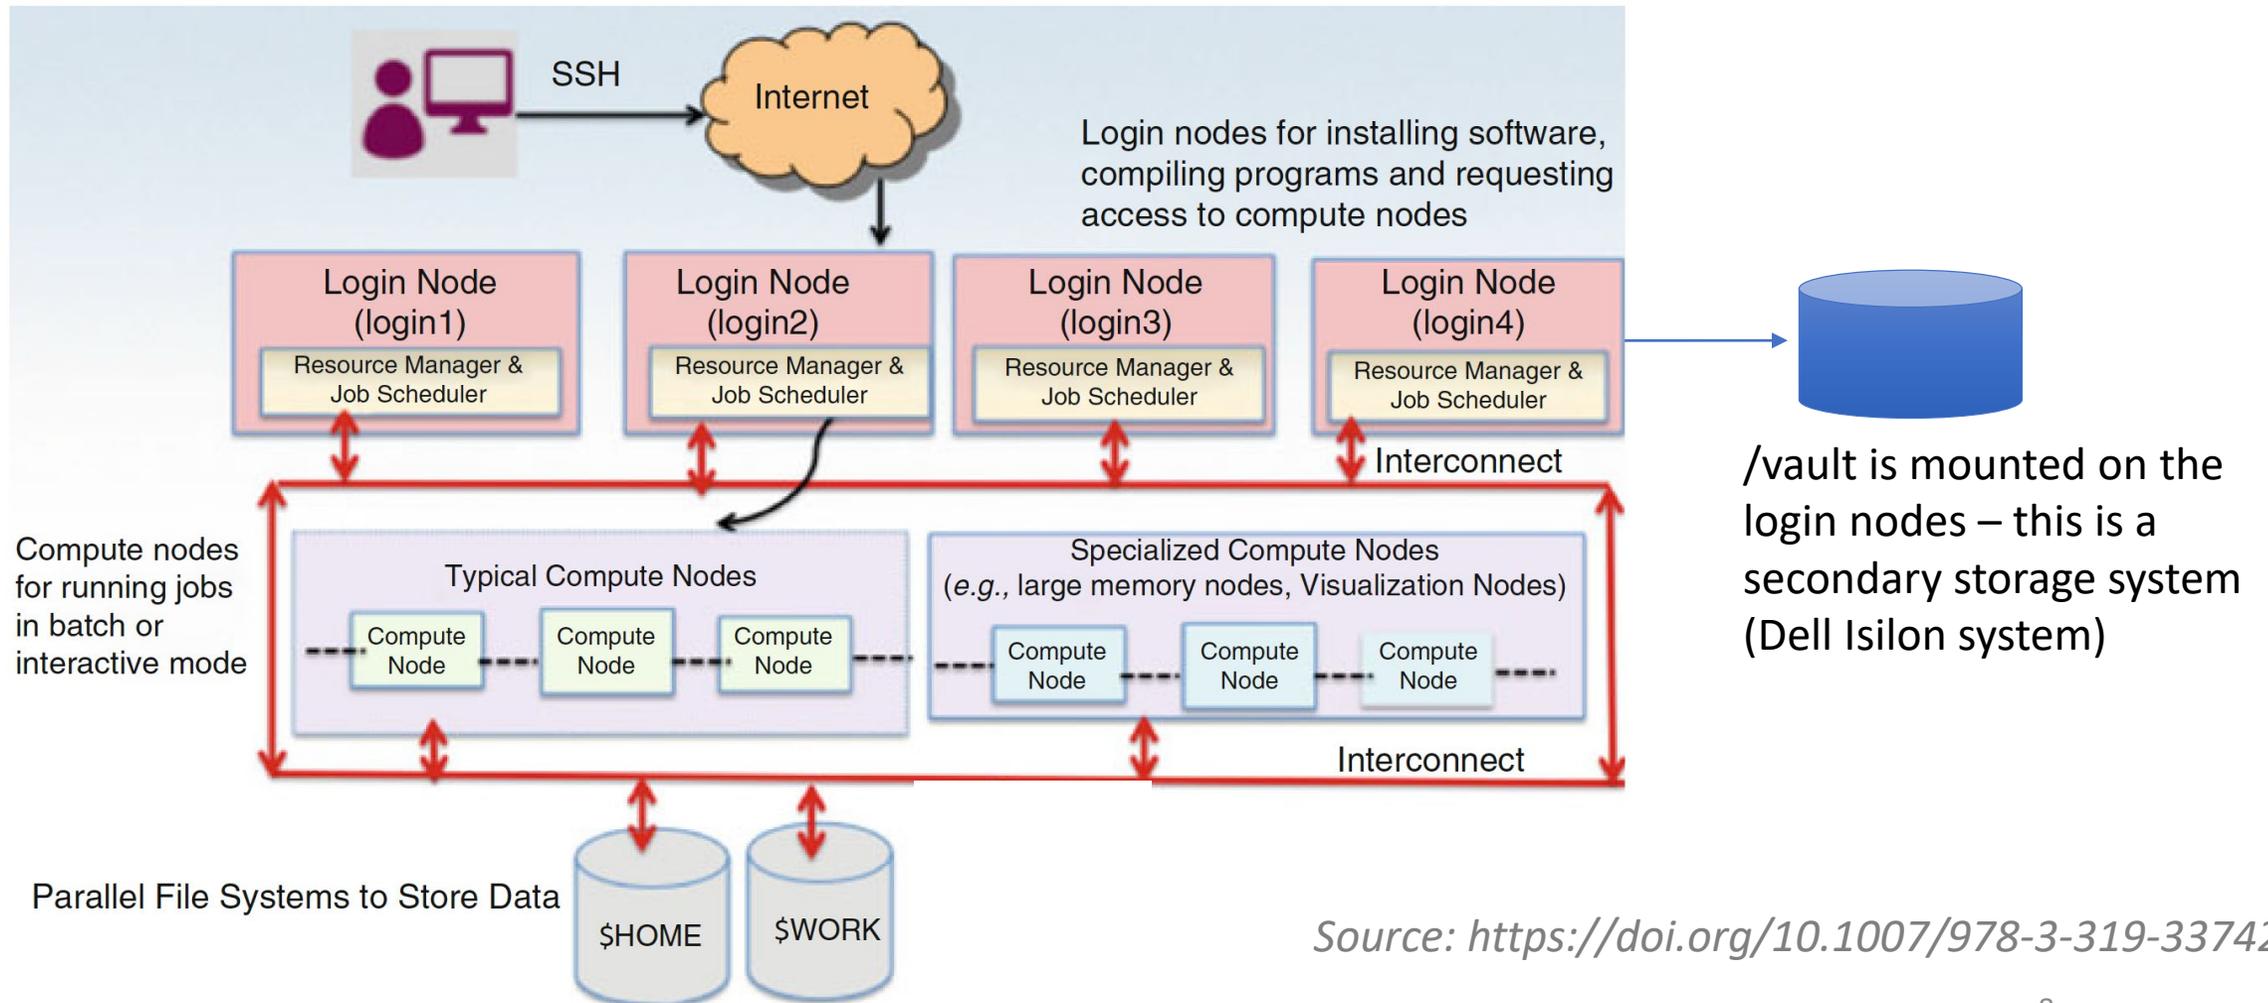
[View Details](#)

<https://tinyurl.com/HPCAccount>

Mando



# Overview of Working on Systems like Arc



Source: <https://doi.org/10.1007/978-3-319-33742-5>

# Connecting to Arc Using Two Factor Authentication with DUO

- To log into Arc you need to setup DUO on a device
- Go to <https://passphrase.utsa.edu/Home/Dashboard> and click on Manage/Set up DUO Authentication
- The host name to connect to Arc is **arc.utsa.edu**
- Use the ssh command with your abc123 UTSA id and the password is your campus passphrase

```
ssh abc123@arc.utsa.edu
```
- After entering in your password a prompt will appear so you can send a push to your device
- If you get an error, double check your password when entering it

# OnDemand Portal

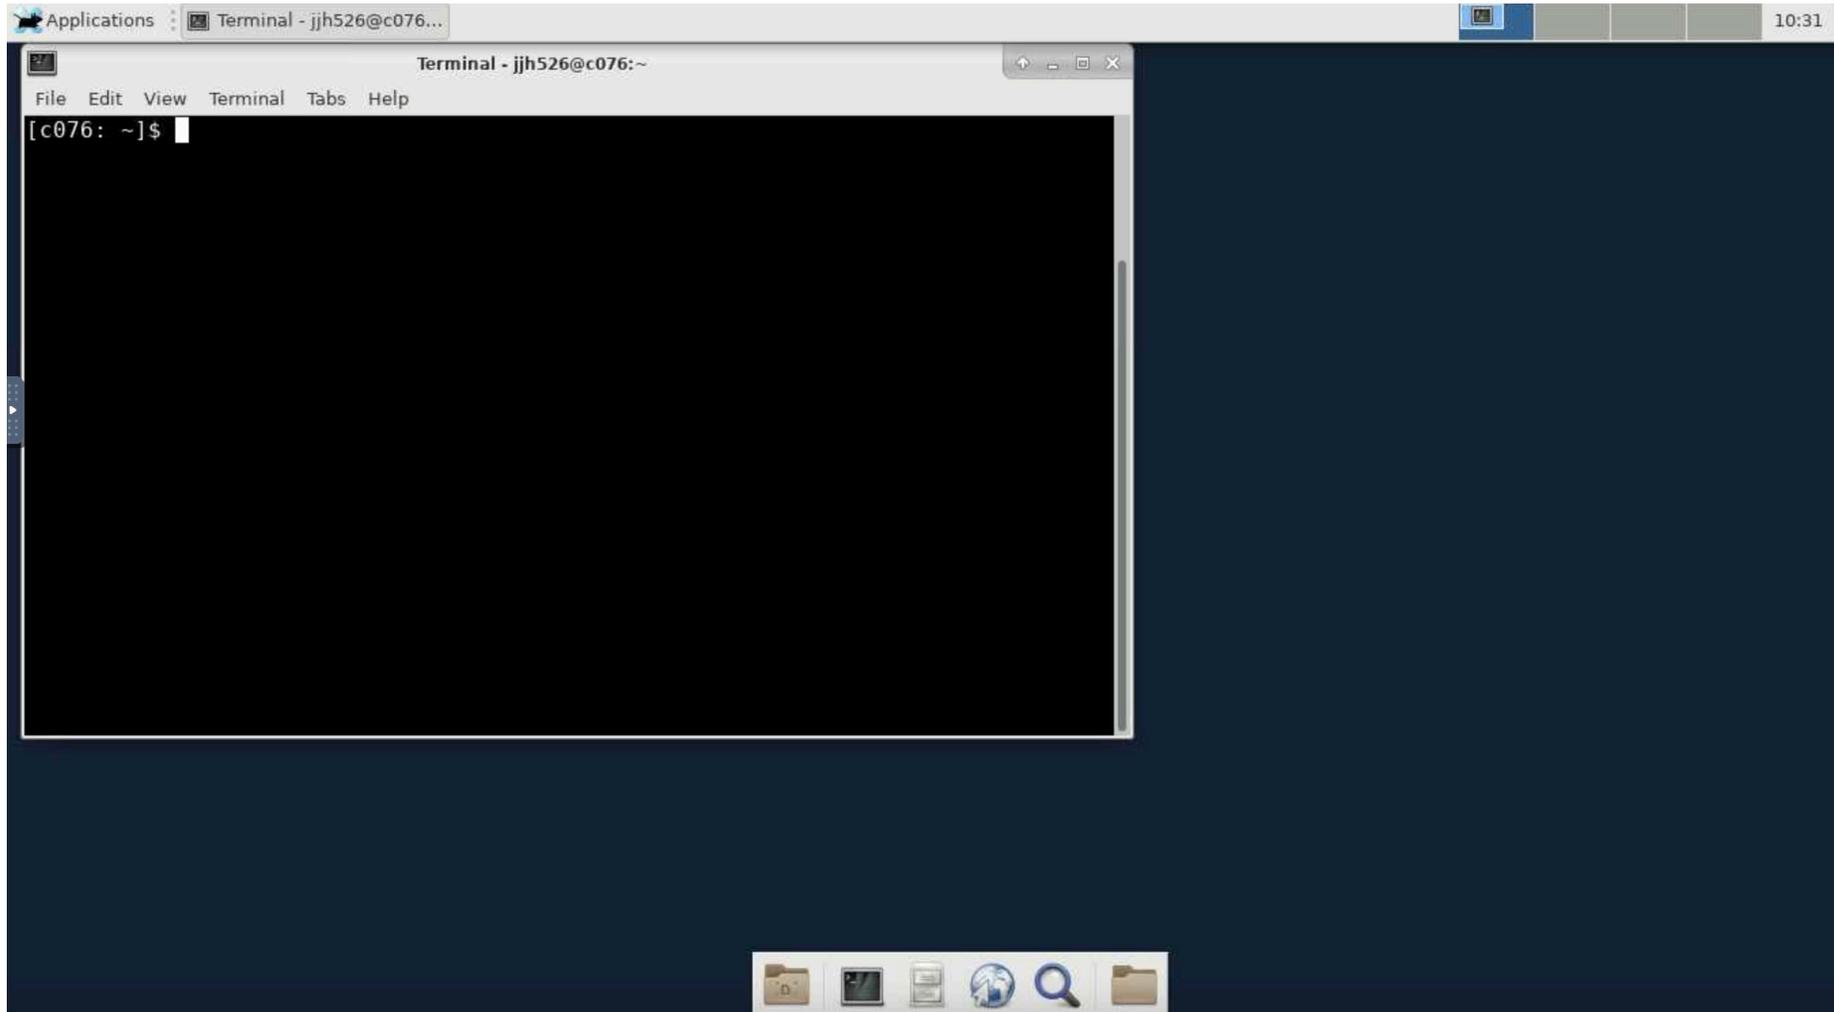
- Available at: <https://portal.arc.utsa.edu>
- Allows you to use a virtual desktop over your browser
- Allows you to connect to a shell over your browser
- Simple file transfers over the web
- Simple job management: monitor jobs, submit, cancel
- Also uses DUO over web login
- Documentation: <https://hpcsupport.utsa.edu/foswiki/bin/view/ARC/VirtualDesktop>

# OnDemand Portal Web Interface



Welcome to the ARC virtual desktop portal. If you have any issues please email [rcsg@utsa.edu](mailto:rcsg@utsa.edu) for assistance.

# OnDemand Portal Virtual Desktop



Ritu



# Partitions or Queues of Arc

- Arc is a shared resource and has the compute servers available in different partitions or queues according to server types and allowed usage policies
- **Some** of the different partitions (or queues) varying in characteristics and constraints are as follows
  - bigmem: 2 nodes
  - compute1: 65 nodes
  - compute2: 25 nodes
  - computedev: 5 nodes
  - gpu1v100: 28 nodes
  - gpu2v100: 5 nodes
  - gpudev: 2 nodes
  - two privately owned partitions consisting of 24 nodes
- Maximum time-period for which a node in the amdonly, amdbigmem, amdgpu, amdvector, compute1, compute2, gpu1v100, gpu2v100, gpu4v100, and bigmem partitions can be used is 72 hours
- Each user can have a maximum of 10 active jobs in a queue and can use maximum 20 nodes at a time through these active jobs

# Filesystems and Storage on Arc

- The two filesystems on Arc are the Lustre filesystems: /home and /work, default is /work
- Lustre enables parallel and distributed access to your files, which in turn improves the rate of IO of your parallel applications
- Lustre can be useful for your serial applications as well if they use large datasets that cannot fit on one individual storage target
- Some general guidelines for working on Lustre
  - Avoid opening and closing too many small files too quickly
  - Avoid putting several thousand files at the same level of directory hierarchy
  - Use parallel IO from your parallel applications to reap maximum benefit from the parallel and distributed nature of the filesystem
  - Stripe your files across multiple storage targets if they are very large and you can do parallel IO

# /vault Storage on Arc

Each user on Arc is provided 1TB of archival storage located in /vault.

Steps for accessing the Isilon from Arc:

1. Login to Arc using SSH with your abc123 account.

This can typically be done using MobaXTerm or Putty by creating a new ssh session for arc.utsa.edu.

If you're connecting from another linux client, the ssh command would be:

```
[linux client]$ ssh arc.utsa.edu
```

2. During the login process, you will be prompted for your password and then Duo factor authentication.

3. Once logged in, change directory to /vault/research

```
[login001: abc123]$ cd /vault/research
```

4. From here you can run the ls or ll command to see the folders and files that are available to you.

```
[login001: research]$ ll
```

```
total 0
```

```
drwxrwxr-x 2 abc123 abc123 0 Dec 16 14:55 ExampleFolder/
```

```
[login001: research]$
```

5. Change directory to your folder to access your files.

# Module System on Arc

- Multiple software packages are available on Arc so that you can start computing with them right away
- These packages can be loaded in your environment by selecting the right modules
  - A module system helps in managing access to the software packages and versions that you need
  - If there are multiple versions of a software available on a system, you can select the one that you need and enable it in your shell environment by loading the right module/s
  - By loading a module, certain environment variables like `$PATH` and `$LD_LIBRARY_PATH` get updated to include the required software and associated dependencies in your shell environment
- To see the modules that are currently loaded in your environment, you can use the following command: **`$ module list`**
- To see all modules that are compatible with the currently loaded modules in your shell, you can use the following command: **`$ module avail`**
- To see all installed modules on the system, you can use the following command: **`$ module spider`**
- You can use the "module load" command to load a module in your shell environment and use "module unload" to remove a module from your environment. For example, following are the commands to load and unload the R module: **`$ module load R/4.1.0`** , **`$ module unload R/4.1.0`**

# Modules Available on Arc

```

[c031: abc123]$ module available

----- /opt/ohpc/pub/modulefiles -----
EasyBuild/3.9.4   charliecloud/0.11   hwloc/2.1.0   papi/5.7.0   prun/1.3   valgrind/3.15.0
autotools        gnu8/8.3.0         llvm5/5.0.1   pmix/2.2.2   singularity/3.4.1

----- /apps/loginmodules -----
cmake/3.20.4     gmp/6.1.0          libctf/4.5.0   openmpi/3.1.4-gcc-10.3.0   sqlite/3.36.0   vscode/1623937300
gcc/4.8.5        gmp/6.2.1          mpb/1.11.1     openmpi/4.1.0-gcc-9.2.0   squashfs/current   xalt/2.10.34 (L)
gcc/7.3.0        harminv/1.4.1      mpc/1.0.3      openmpi/4.1.1-gcc-10.3.0   tcl/8.6.11       yaml-cpp/0.6.3
gcc/9.2.0        hdf5/1.12.0        mpfr/3.1.6     openmpi/4.1.1-gcc-11.2.0 (D)   tiff/4.3.0
gcc/9.3.0        hpcswttest/current (D)            petsc/3.15.2-gcc-10.3.0   tk/8.6.11
gcc/10.3.0       intel/oneapi/2021.2.0   mpich/4.0.a2   petsc/3.15.2-intel-2021.2.0 (D)   tmux/2.3a
gcc/11.2.0 (D)   isl/0.18            mvapich2/2.3.6-gcc_10.3.0   proj/4.9.3   valgrind/3.17.0 (D)
gdal/3.3.2       jpeglib/9d         mvapich2/2.3.6-gcc_11.2.0 (D)   proj/8.1.1 (D)   vim/8.1

----- /apps/gpumodules -----
cuda/cudnn/8.2.0.53   cuda/sdk/nvhpc-byo-compiler/21.5   cuda/sdk/nvhpc/21.5   cuda/toolkit/11.3
cuda/cudnn/8.2.1.32 (D)   cuda/sdk/nvhpc-nompi/21.5   cuda/toolkit/11.2   cuda/toolkit/11.4 (D)

----- /apps/computemodules -----
R/3.5.1           cellranger/2.0.0           geos/3.11.0           metashape-pro/1.7.3           orca/4.2.1
R/3.6.3           cellranger/2.1.1           glpk/4.54             miniconda/4.9.2             ospray/2.7.0
R/4.0.3           cellranger/3.0.1           gms/4.8.4             namd/2.14-gpu               parallel/20210622
R/4.0.5           cellranger/3.1.0           hisat2/2.2.1          namd/2.14-mpi               parallel/20210722 (D)
R/4.1.0           cellranger/3.1.0.1         homer/4.11            namd/2.14-multicore         paraview/5.9.1
R/4.1.1 (D)       cellranger/4.0.0           hoemd/cpu/2.9.6       ncbi/blast/2.11.0           perl/5.26.3
abaqus/2020       cellranger/5.0.0           hoemd/gpu/2.9.6       ncbi/sratoolkit/2.11.0     perl/5.34.0 (D)
anaconda3/current cellranger/6.0.1           imagej/1.8.0_172      ncbi/sratoolkit/2.11.3     rstudio/1.1.463
ansys/2019R3      cellranger/6.1.2 (D)      lammps/20210730-mpi/intel-oneapi/2021.2.0   necat/0.0.1                 rstudio/1.4.1717 (D)
autodock/4.2.6    columbus/6-falloc          lammps/20210730-mpi/openmpi/4.1.1/gcc/11.2.0   neuron/8.0.0                salmon/1.5.2
bcl2fastq/1.8.4   columbus/6                 lammps/20210730-serial/gcc/11.2.0             nwchem/6.8.1-intel          samtools/1.13
bcl2fastq2/2.17   columbus/7.0 (D)          lammps/20210730-serial/intel-oneapi/2021.2.0   openblas/0.3.17             sbt/1.5.3
bcl2fastq2/2.20 (D)   comsol56/5.6.0.401         ls-dyna/r10.1.0       openfoam/6.0.bak            scala3/3.0.0
cellranger-arc/2.0.0   cutadapt/1.18             lumerical/v212        openfoam/6.0-gcc-10.3.0-mpi-3.1.4   simple_test/1.0
cellranger-atac/1.0.1 fastqc/0.11.9             mathematica/12.3       openfoam/6.0 (D)            spaceranger/1.0.0
cellranger-atac/1.1.0 fds/6.7.6                 matlab/R2021a         opensees/3.3.0-mvapich2-2.3.6   subread/2.0.3
cellranger-atac/1.2.0 (D) fsl/6.0.5                 meep/1.17.1           opensees/3.3.0-openmpi-4.1.1 (D)   trimgalore/0.6.5

Where:
D: Default Module
L: Module is loaded

Use "module spider" to find all possible modules.
Use "module keyword key1 key2 ..." to search for all possible modules matching any of the "keys".

[c031: abc123]$

```

# More on Modules

<https://hpcsupport.utsa.edu/foswiki/bin/view/ARC/ModuleEnvironments>

Rich



# Interactive and Batch Computing Mode

[WebHome < ARC < UTSA Research Support Group](https://hpcsupport.utsa.edu/foswiki/bin/view/ARC/WebHome)

<https://hpcsupport.utsa.edu/foswiki/bin/view/ARC/WebHome>

[Running Jobs On Arc.pdf \(utsa.edu\)](https://hpcsupport.utsa.edu/foswiki/pub/ARC/WebHome/Running_Jobs_On_Arc.pdf)

[https://hpcsupport.utsa.edu/foswiki/pub/ARC/WebHome/Running\\_Jobs\\_On\\_Arc.pdf](https://hpcsupport.utsa.edu/foswiki/pub/ARC/WebHome/Running_Jobs_On_Arc.pdf)

- Login nodes should not be used to run jobs – they can be used for installing software and accessing files etc
- All jobs should be run on compute nodes either interactively or in batch computing mode
- Access to the compute nodes is obtained through the Slurm job scheduler
- What is interactive computing mode?
  - When you are actively developing, compiling and running jobs on the compute node, interactive access can be useful
  - Can suffer from disruptions due to internet issues
  - Let us see a demo of this
- What is batch computing mode?
  - Fire and forget
  - Results come back to you when your jobs have finished running
  - Let us see a demo of this

# Slurm Job Scheduler

- Submitting Batch Jobs using the SBATCH directives:

#SBATCH -J myjob # Defines the job name

#SBATCH -o myjob.o%j # Name of the output file

#SBATCH -e myjob.e%j # Defines the name of the file where any errors will be written to

- The following directives are mandatory for all jobs submitted:

#SBATCH -p compute1 # Queue (partition) name

#SBATCH -N 1 # Total # of nodes (must be 1 for serial)

#SBATCH -n 1 # Total # of mpi tasks (should be 1 for serial)

#SBATCH -t 01:30:00 # Run time (hh:mm:ss)

# Creating Job Dependencies via Slurm

<https://hpcsupport.utsa.edu/foswiki/bin/view/ARC/Checkpoint-and-RestartForR>

Ritu



# Checkpointing and Restart

- Checkpointing is the process of saving the execution state of an application such that this saved state can be used to continue the execution at a later time. Typically, the execution state is written to a file.
- Restart is the step that comes after checkpointing and helps in resuming the application from the saved state.
- Checkpointing not only **saves time by offering the capability to resume the execution of an application in case of a hardware failure**, but it also **helps in overcoming the time-limits associated with the different job queues/partitions**.
- Checkpointing and restart mechanism can be built into the applications that you are writing and some examples of achieving this are available at the following link: <https://tinyurl.com/4jr3zua8>

# Distributed Computing

- The form of computing in which the different components on an application run as a single system across multiple servers or nodes that are interconnected with each other
- Examples: Message Passing Interface (MPI), MPI+X computing paradigms
- Another example: Distributed Training in Deep Learning
  - Data Parallelism
  - Model Parallelism
  - <https://tinyurl.com/32y322fu>

Zhiwei



## Save Python virtual environments (VE ) in the Work directory

- Conda VEs are saved in the /home directory by default

- Use the following command to save a VE in the Work directory:

```
module load anaconda3
```

```
conda create --prefix $WORKDIR/test-env python=3.9.1
```

- Use the following command to activate the VE

```
conda activate $WORKDIR/test-env
```

# How to use Jupyter notebook on Arc

- Log in to the on-demand portal

<https://portal.arc.utsa.edu/>

- Launch an interactive session on a node, and launch Jupyter notebook using the following command:

```
jupyter notebook --notebook-dir /work/abc123/
```

Rich



# DGX-Related: <https://tinyurl.com/5x3525kp>

```
#!/bin/bash
#-----
# Sample DGX Slurm job script for Arc nodes

#SBATCH -J myjob          # Job Name
#SBATCH -o myjob.o%j      # Name of the stdout output file
#SBATCH -e myjob.e%j      # Name of stderr error file
#SBATCH -p dgxa100        # Queue (partition) name
#SBATCH -N 1              # Total # of nodes (must be 1 for serial)
#SBATCH -n 1              # Total # of mpi tasks (should be 1 for serial)
#SBATCH -t 01:30:00      # Run time (hh:mm:ss)

#Put other commands such as module load if needed after this line (All
#commands must follow #SBATCH directives)
# Launch code...

./mycode.exe
```

**Listing 3. Sample DGX Slurm Job-Script**

Brent



# Training Materials

Additional training materials can be found on our Wiki and include the following subjects:

- Express Linux Tutorial for New Users
- Bash Scripting
- Arc New User Training
- Submitting Jobs Using Slurm

<https://hpcsupport.utsa.edu/foswiki/bin/view/Main/Training>

Additional topics will be added in the future.

# Grant Proposal Preparations

We can assist with the following when you are creating a proposal:

- Technology recommendations
- Recommending technology vendors or assisting you when working with vendors
- Assist with writing data management plans
- Provide “facilities” documentation

# Level of Service

The RCSG team provides guidance and advice on software development and application optimization. Although we provide guidance, we are unable to perform the actual development, optimization and data analysis.

Thomas



# Technical Support and Service Now Ticket Creation

- To better serve our customers and make sure we track all service and support requests, we request that you begin entering any support or consulting questions into "MyPortal" (powered by Service Now).
- This will result in quicker response times, and allow you to follow up and communicate with us much easier: <https://www.utsa.edu/techsolutions/service-now/>
- The Research Computing Support Group is available between 8:00 AM to 5:00 PM on all business days to assist with the service requests.
- Our time-to-response on new incidents (break/fix) tickets is 4 business hours, and the time-to-resolution varies depending upon the complexity of the issue. For all other requests such as account creation, software package installations, data management, etc. our time-to-response is usually between 24 and 48 hours of receiving the ticket in Service Now.

# System Use Policies

- Because Arc is a shared computing resource, we request you to please familiarize yourself with the fair-use policies that are mentioned in the Arc user-guide. Not to mention the wealth of support information on how to best utilize the cluster:

<https://hpcsupport.utsa.edu/foswiki/bin/view/ARC/>

- The bigmem and gpu\* queues were reserved earlier and access was granted on a case-by-case basis. However, no special privileges are needed to access these queues now – Happy Computing!

# References

1. <https://hpcsupport.utsa.edu/foswiki/bin/view/ARC/WebHome>
2. <https://github.com/ritua2/documentation>
3. <https://hpcsupport.utsa.edu/foswiki/bin/view/ARC/VirtualDesktop>

*Any questions, comments, or concerns?*

Thank You!

Ritu Arora

Email: [ritu.arora@utsa.edu](mailto:ritu.arora@utsa.edu)



[utsa.edu](http://utsa.edu)