



The University of Texas at San Antonio™

Arc New User Training

February 4, 2022

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

Email: rcsg@utsa.edu

Overview

- Architecture of Arc
- Connecting to Arc
- Arc user environment
- Filesystem on Arc
- Interactive mode and batch computing mode - demo
- Slurm job scheduler
- Using Xdisplay/GUI and off-Arc storage
- OnDemand portal
- Policies around using Arc
- Requesting support



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 was built using 102 new compute nodes and 54 old compute nodes from the now retired Shamu cluster
- Out of the 54 old compute nodes from Shamu, 20 nodes are owned privately by a research group, 4 nodes are reserved for system administration and testing purposes
 - 102 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
- Arc supports more than 233 customers from the following 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 can deliver up to 387 TFLOPs of peak performance and is equipped with:

- 156 total compute servers or nodes (having 6032 total cores across all the nodes) and 2 login nodes - majority of these compute servers are Intel Cascade Lake CPUs
 - 30 nodes out of the 156 nodes are equipped with two CPUs with 20 cores each for a total of 40 cores, 384GB RAM, and each including one Nvidia V100 GPU accelerator
 - 5 GPU nodes are equipped with two CPUs with 20 cores each for a total of 40 cores, 384GB RAM, and each including two Nvidia V100 GPU accelerators
 - Two large-memory nodes, each containing four CPUs with 20 cores each for a total of 80 cores, and each including 1.5TB of RAM
- 100 Gb/s Infiniband connectivity
- Two parallel and distributed filesystems (Lustre): /home and /work, where /home has 110 TBs capacity and /work has 1.1 PB of capacity
- A cumulative total of 250TB of local scratch (approximately 1.5 TB of /scratch space on most compute nodes)

Partitions or Queues of Arc

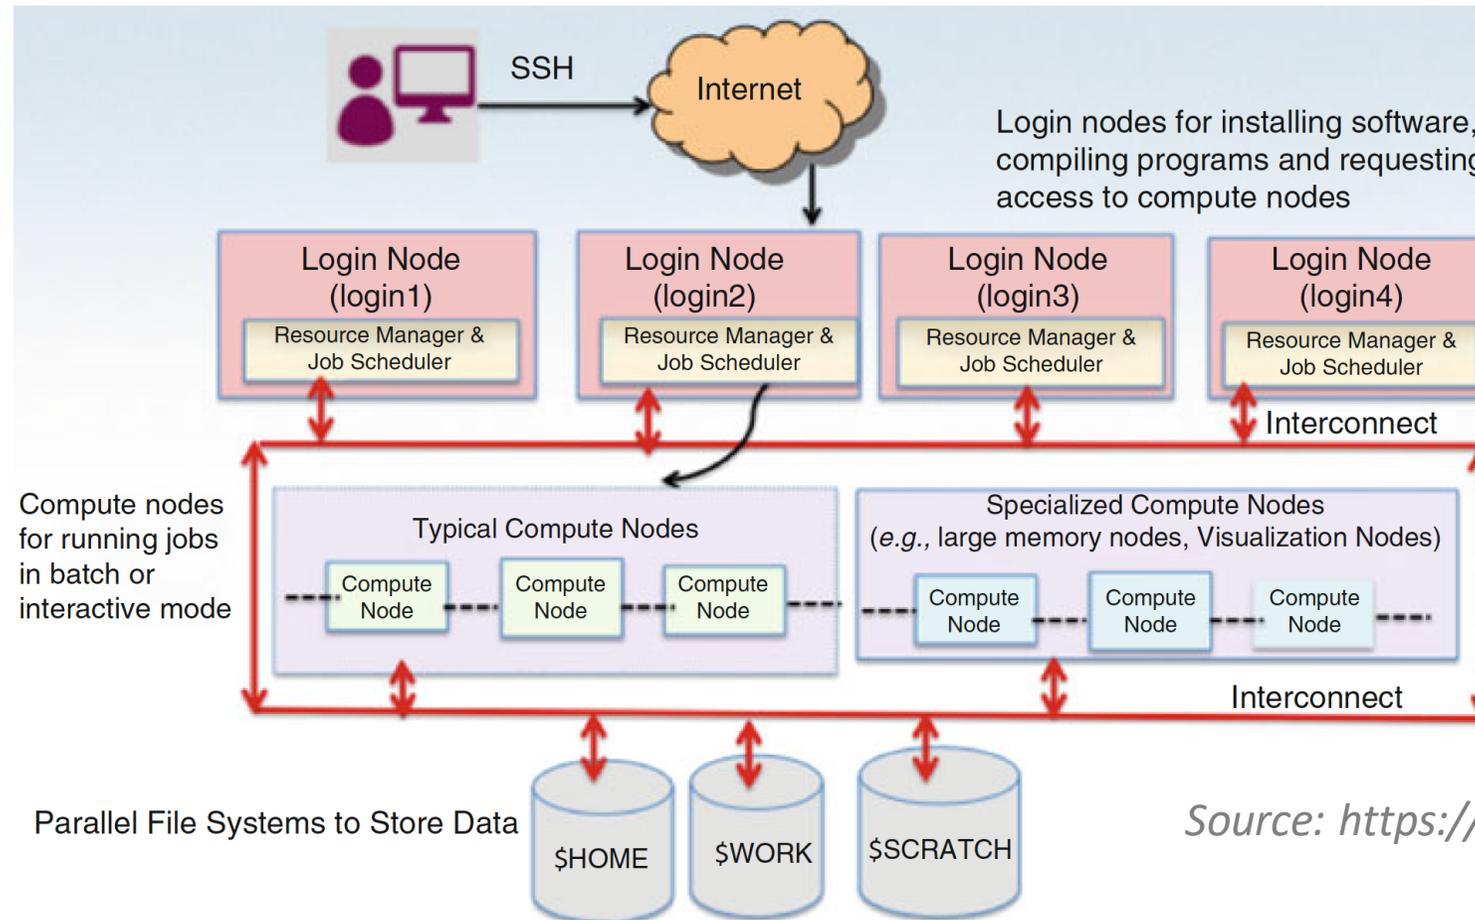
- Arc is a shared resource and has the compute servers available in different partitions or queues according to server types and allowed usage policies
- The different partitions (or queues) varying in characteristics and constraints are as follows
 - bigmem: 2 nodes
 - compute1: 65 nodes
 - compute2: 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 compute1, compute2, gpu1v100, gpu2v100, 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, and the 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

Connecting to Arc and General Overview of Working on Systems like Arc

- Arc is accessible over SSH using two-factor authentication with DUO
- Hostname for Arc is **arc.utsa.edu** and the SSH port number is 22
- In order to utilize DUO, you must register online at passphrase.utsa.edu



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

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

```

9. arc-Login001
[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        hpcswtest/current (D)            petsc/3.15.2-gcc-10.3.0   tk/8.6.11
gcc/10.3.0       intel/oneapi/2021.2.0   mpich/4.1.0   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   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   cellranger/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  cellranger/6              lammps/20210730-serial/gcc/11.2.0           nwchem/6.8.1-intel         samtools/1.13
bcl2fastq2/2.17  cellranger/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>

Modules and XALT

By default, a module named XALT is loaded into everyone's shell environment. XALT is a tool that allows the Arc HPC support staff to collect and understand job-level information about the libraries and executables that end-users access during their jobs. This assists us in tracking user executables and library usage on the cluster. If you experience an issue that may involve XALT, the module can be removed using the module unload command.

```
[c031: abc123]$ module list
Currently Loaded Modules:
  1) xalt/2.10.34

[c031: abc123]$ module load R/4.1.0
[c031: abc123]$ module list
Currently Loaded Modules:
  1) xalt/2.10.34  2) hdf5/1.12.0  3) gcc/4.8.5  4) R/4.1.0

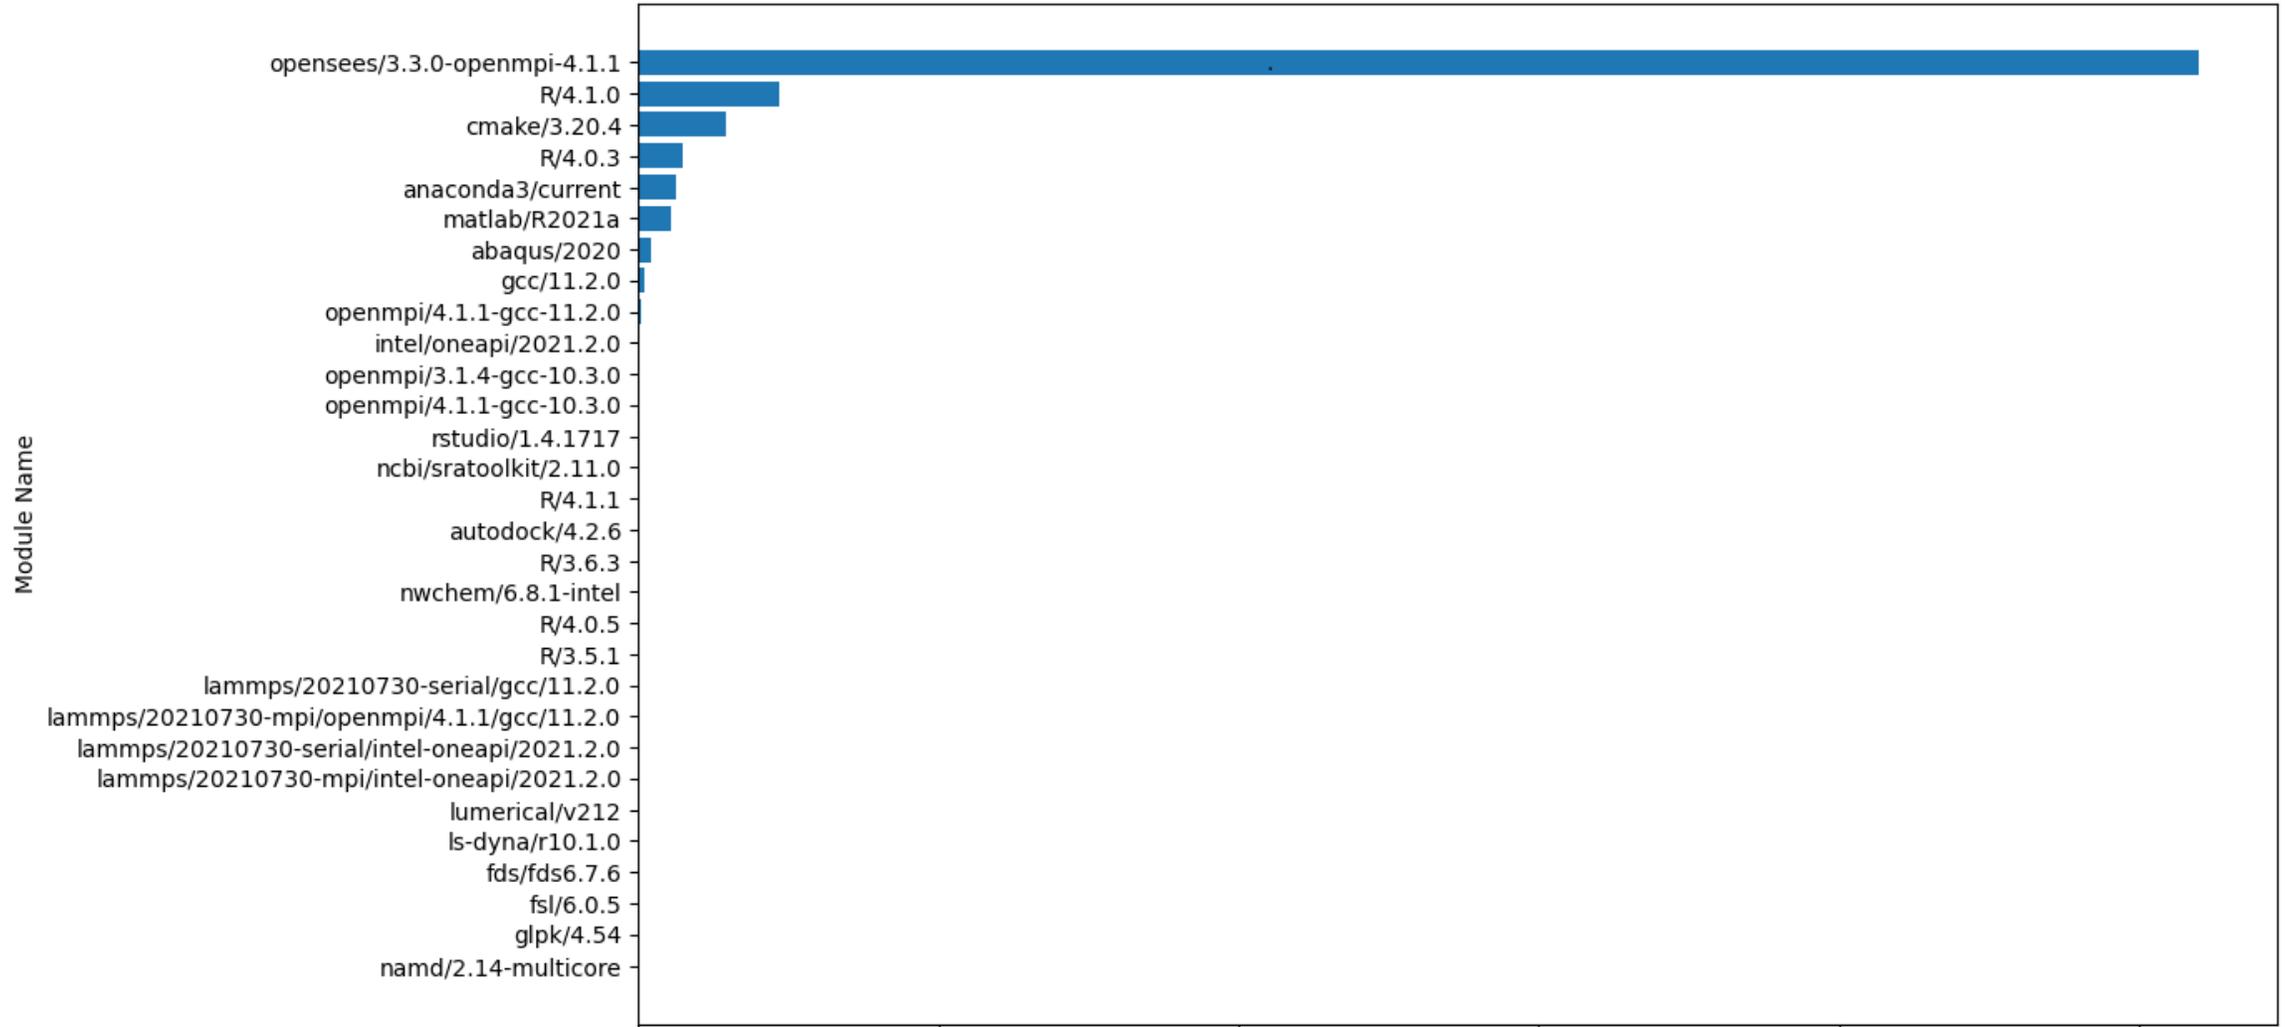
[c031: abc123]$ R --version
R version 4.1.0 (2021-05-18) -- "Camp Pontanezen"
Copyright (C) 2021 The R Foundation for Statistical Computing
Platform: x86_64-pc-linux-gnu (64-bit)

R is free software and comes with ABSOLUTELY NO WARRANTY.
You are welcome to redistribute it under the terms of the
GNU General Public License versions 2 or 3.
For more information about these matters see
https://www.gnu.org/licenses/.

[c031: abc123]$ █
```

XALT Monitoring

Module Usage Frequency: 2022-01-03 (Mon) 00:00:00 to 2022-01-08 (Sat) 23:59:59



Interactive and Batch Computing Mode

- 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 even if you disconnect from the system
 - 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>

Isilon Storage on Arc

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

Steps for accessing the Isilon archive 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 abc123@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 /isilon/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.

Viewing the Graphical User Interface

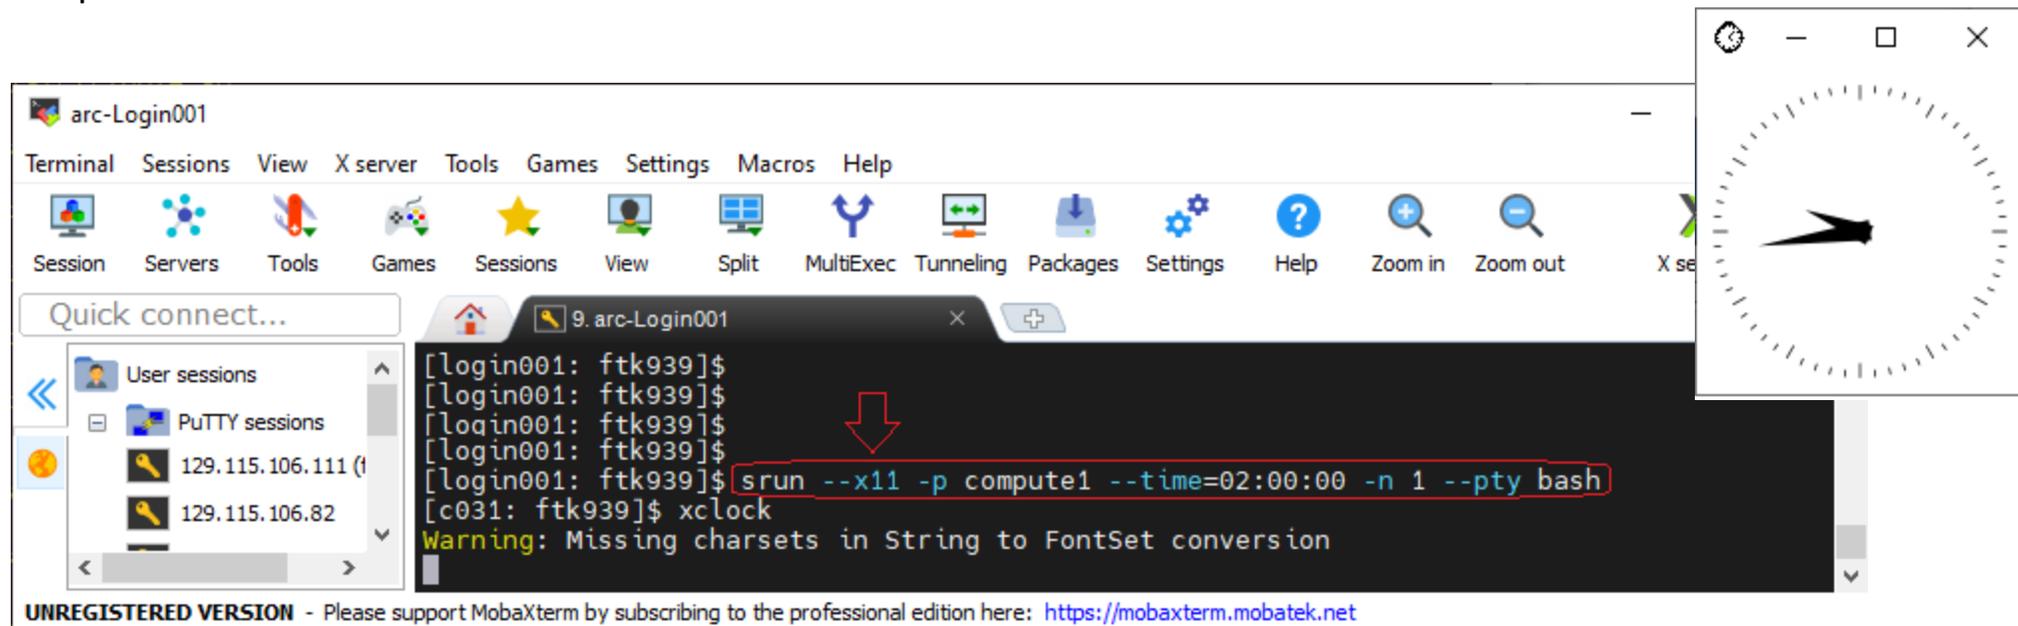
MobaXterm can be used for this purpose. It's an enhanced terminal for Windows with a built in X11 server.

[Installing MobaXterm](#)

MobaXterm can be downloaded here: <https://mobaxterm.mobatek.net/download.html>

Click the link for the Free version, MobaXterm Home Edition, to get the Installer Package for Windows.

MobaXterm comes with its own built in X server. X is a windowing system in Linux. To display GUI applications in Arc, just issue the command for the application in the MobaXterm terminal and the GUI window will be displayed on your Windows desktop.



OnDemand Portal

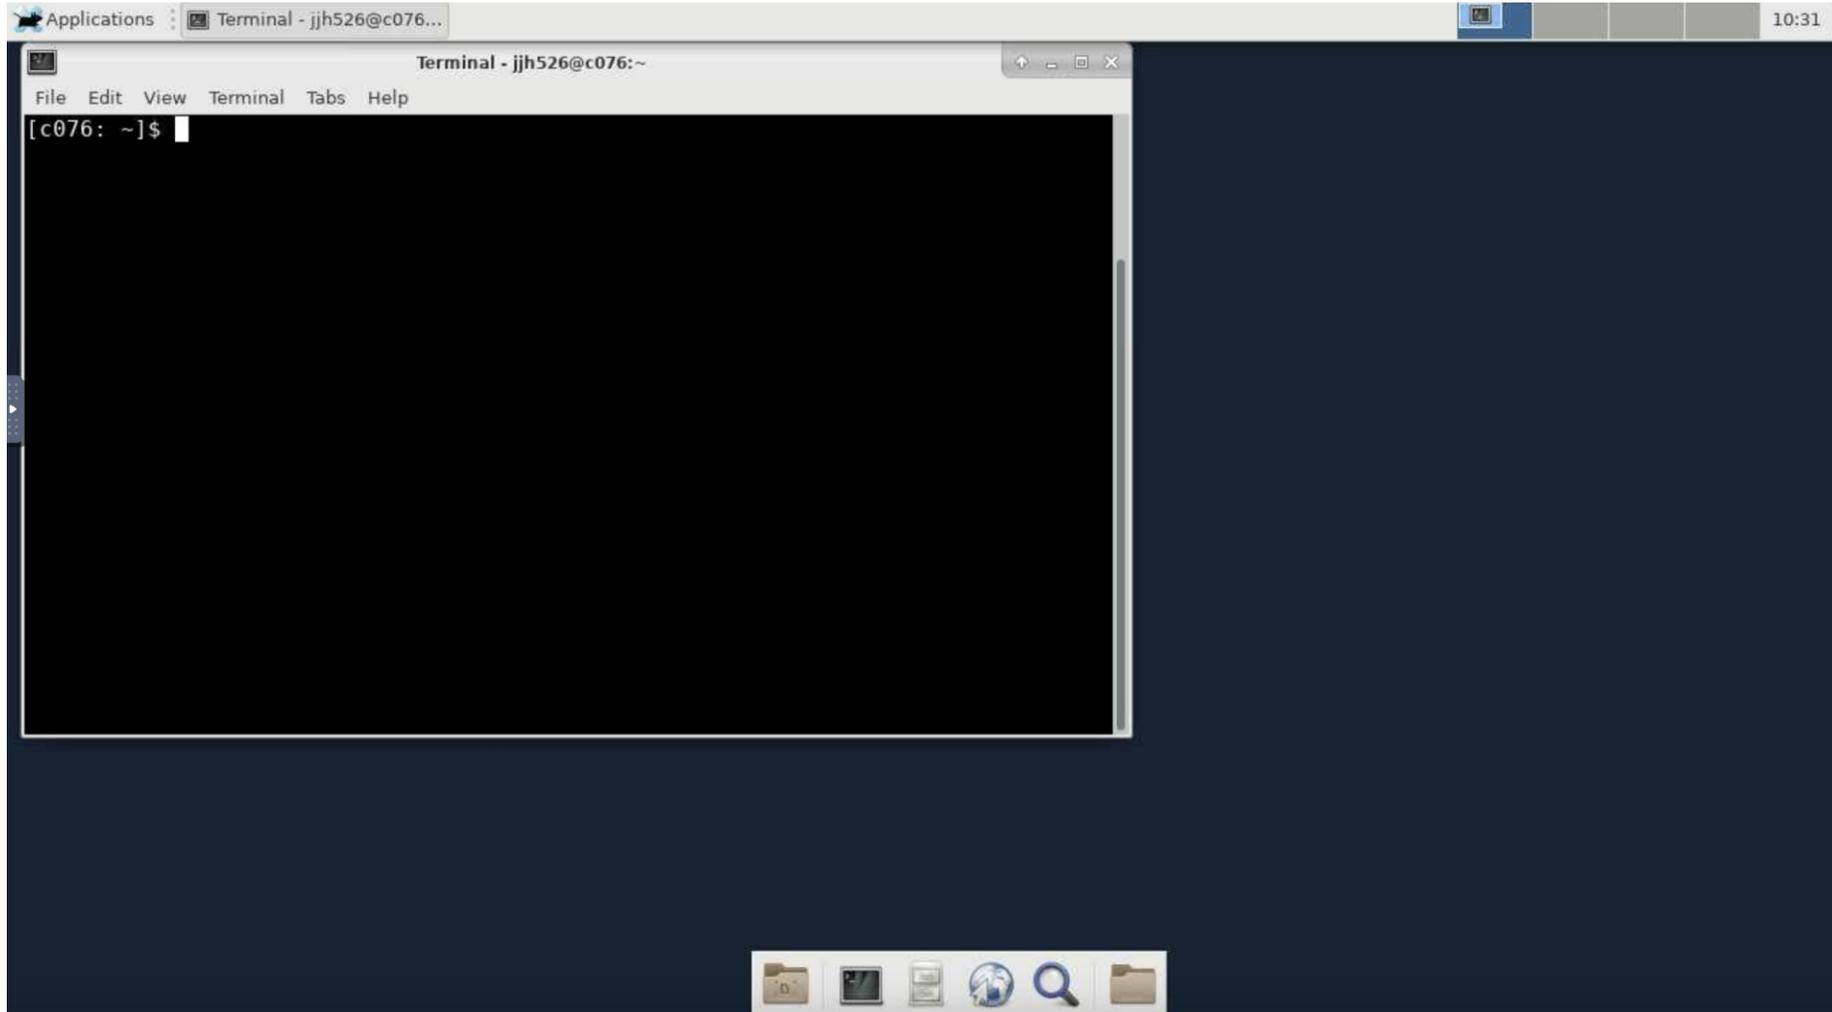
- <https://portal.arc.utsa.edu>
- Allows you to use a virtual desktop over your browser
- Allows you to connect to a shell over your browser
- Simple file transfers over the web
- Simple job management: monitor jobs, submit, cancel
- Documentation:
<https://hpcsupport.utsa.edu/foswiki/bin/view/ARC/VirtualDesktop>

OnDemand Portal Web Interface



Welcome to the ARC virtual desktop portal. If you have any issues please email rcsg@utsa.edu for assistance.

OnDemand Portal Virtual Desktop



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 are reserved on a case-by-case basis such as job requirements and needs of the end user. We ask that you open a Service Now ticket to request access to said queues with the type of job you are running and the technical reason you will require said resources.

(Very Near) Future Expansion of Arc

- DGX A100 80 GB appliances with AMD EPYC processors and in mini-pod architecture
 - Will be private to Matrix AI consortium
- AMD EPYC nodes with 1.5 TB of RAM
- AMD EPYC nodes with NEC Vector Engines
- AMD EPC nodes with A100 GPUs
- With the next round of expansions we are likely going to see over 0.5 PetaFLOPs of peak performance from Arc

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



utsa.edu