

Introduction to Cheyenne for New Users

December 11, 2020

Shiquan Su

Consulting Services Group



The Speaker

Shiquan Su, PhD in Physics,
helping users of NCAR's HPC resources since 2015

<https://staff.ucar.edu/users/shiquan>

Software Engineer, Consulting Services Group, High Performance Computing
Division

<https://staff.ucar.edu/browse/orgs/HPCD>

Topics to cover

- Available resources on Cheyenne
- Signing in and managing data
- Customizing your user environment
- Accessing and building software
- Managing jobs using Batch schedulers

Please type your questions in the chat. I will check the chat windows, and answer the questions at the end of my tutorial. It is perfectly ok that only part of the tutorial is relevant and useful to you.

This is only an introduction; for more usage details, see:
<https://www2.cisl.ucar.edu/resources/>

Cheyenne Supercomputer

**For Model Runs, Large Parallel Jobs
(Still rank at 60th up to November 2020)**

- 4032 batch compute nodes
- 2-socket 18-core Intel Broadwell CPUs, 2.3 GHz
- 3168 nodes with 64 GB memory
- 864 nodes with 128 GB memory
- EDR Infiniband Network
- SUSE Linux Enterprise Server 12 SP4
- PBS job scheduler to access Cheyenne compute nodes, Slurm job scheduler to access DAV nodes

SAM (Systems Accounting Manager)

- Web access: <https://sam.ucar.edu>
- Log in with Duo (YubiKey)
- Change some user settings (default project, shell, etc)
- Get information about available projects and remaining allocation balance
- See history of jobs and charges

SAM: user info

- SAM webpage -> login -> User tap -> Settings -> search username

NAME	Shiquan Su
USERNAME	shiquan
EMAIL	shiquan@ucar.edu
ORGANIZATION	CSG (Consulting Services Group)
INSTITUTION	

Shell

Username	Resource	Shell
shiquan	Cheyenne	tcsh
shiquan	GLADE	tcsh
shiquan	HPC_Futures_Lab	tcsh
shiquan	Laramie	tcsh

Edit

Resource Access

Username	Project Code	Resource
shiquan	SCSG0001	Anemone
shiquan	SCSG0001	Campaign_Store
shiquan	SCSG0001	Casper
shiquan	SCSG0001	Cheyenne
shiquan	SCSG0001	GLADE
shiquan	SCSG0001	HPC_Futures_Lab
shiquan	SCSG0001	HPSS
shiquan	SCSG0001	Laramie
shiquan	SSSG0001	Casper
shiquan	SSSG0001	Cheyenne
shiquan	SSSG0001	Laramie

Home Directory

Username	Resource	Home Directory
shiquan	Cheyenne	/glade/u/home/shiquan
shiquan	GLADE	/glade/u/home/shiquan
shiquan	HPC_Futures_Lab	/users/shiquan
shiquan	HPSS	/home/shiquan
shiquan	Laramie	/picnic/u/home/shiquan

Primary Group

Primary Unix group applies to all CISL resources.
Primary group changes do not currently propagate to HPSS.

Username	Primary GID	Primary Group Name
shiquan	1000	ncar

Edit

Groups

shiquan is a member of the following hpc-data tagged groups.

Group Name	GID
ncar	1000
csg	1564
csgteam	68122
sssg0001	68282

SAM: project info

- SAM webpage -> login -> Project tap -> Search -> search: Project Code/Project Lead/Username

Project Dashboard

PROJECT DETAILS Edit		ALLOCATIONS Edit			
Project Code	SCSG0001	Listed below are the active and future allocations associated with this project.			
Project Title	CSG systems project	Resource	Amount	Start Date	End Date
Project Lead	Siddhartha S Ghosh (sghosh)	Anemone	1	2020-10-01	2021-09-30
Project Administrator	Francis Coady (mickc)	Campaign_Store	10	2020-10-01	2021-09-30
Status	Active	Casper	20,000	2020-10-01	2021-09-30
Charge Type	EXEMPT	Cheyenne	5,000,000	2020-10-01	2021-09-30
Facility	CISL	GLADE	2	2020-10-01	2021-09-30
Panel	CISL USS	HPC_Futures_Lab	1	2020-10-01	2021-09-30
Allocation Type	Staff	Laramie	1	2020-10-01	2021-09-30
NCAR Organization	CSG	Stratus	10	2020-10-01	2021-09-30
CONTRACTS Edit		DISK DIRECTORY-PROJECT MAPPING Edit			
Listed below are the contracts associated with this project.		Listed below are all disk directory-project mappings for this project.			
Source	Contract Number	Start Date	End Date	Disk Resource	Directory Name
				Campaign_Store	/gpfs/csfs1/cisl/csg
				GLADE	/glade/p/CSG
				GLADE	/glade/p/cisl/CSG
				Stratus	/stratus/USS

USERS Edit			RATE-OF-USE THRESHOLDS Edit		
Listed below are all users associated with this project. Some users may not have access to some resources.			Listed below are the rate-of-use thresholds for this project.		
Name	Username	Status	Resource	30-Day Threshold	90-Day Threshold
Jared Baker	jbaker	Active	Anemone		
James Clark	jaclark	Active	Campaign_Store		
Francis Coady	mickc	Active	Casper		
Daide Del Vento	ddvento	Active	Cheyenne		
James Edwards	jedwards	Active	GLADE		
Thomas Engel	engel	Active	HPC_Futures_Lab		
David Flaherty	davidf	Active	HPSS		
Siddhartha S Ghosh	csgteam	Active	Laramie		
Siddhartha S Ghosh	sghosh	Active	Stratus		
David L. Hart	dhart	Active			
Timothy J Hoar	thoar	Active			
Ryan Hulguin	hulguin	Active			
SUB-PROJECTS Edit					
Listed below is the project tree for this project.					
SCSG0001					

SAM: account statement

- SAM webpage -> login -> Reports tap -> My Account Statements -> choose a project -> Detail

Account Statement: Project Activity

PROJECT CODE	SCSG0001
PROJECT TITLE	CSG systems project
PROJECT LEAD	Siddhartha S Ghosh
PROJECT ADMINISTRATOR	Mick Coady
DATES	2012-05-11 to 2021-09-30
<input type="button" value="More"/>	

Overall Usage

Resource	Status	Start	End	Allocation	Usage/Charges	Adjustments to Charges	Balance	
Cheyenne	Normal	2020-10-01	2021-09-30	5,000,000	180,886	0	4,819,114	Activity
HPSS	Expired				7	0		Activity
Casper	Normal	2020-10-01	2021-09-30	20,000	2,726	0	17,274	Activity
GLADE	Normal	2020-10-01	2021-09-30	2	2	0	0	Activity
Anemone	Normal	2020-10-01	2021-09-30	1	0	0	1	Activity
Campaign_Store	Normal	2020-10-01	2021-09-30	10	5	0	5	Activity
HPC_Futures_Lab	Normal	2020-10-01	2021-09-30	1	0	0	1	Activity
Laramie	Normal	2020-10-01	2021-09-30	1	0	0	1	Activity
Stratus	Normal	2020-10-01	2021-09-30	10	0	0	10	Activity

<input type="button" value="Cheyenne"/> <input type="button" value="HPSS"/> <input type="button" value="Casper"/> <input type="button" value="GLADE"/> <input type="button" value="Anemone"/> <input type="button" value="Campaign_Store"/> <input type="button" value="HPC_Futures_Lab"/> <input type="button" value="Laramie"/> <input type="button" value="Stratus"/>														
Project	Status	Allocation	Charges (with adjustments)	Balance	30-day Use	30-day Alloc	30-day % Use	30-day % Limit	90-day Use	90-day Alloc	90-day % Use	90-day % Limit	Users	
SCSG0001	Normal	5,000,000	180,886	4,819,114	88,211	412,088	21	-	180,886	1,236,264	15	-	25	

HPC System Access

Logging in to Cheyenne or Casper

- Use ssh along with your username to log in
ssh username@cheyenne.ucar.edu
ssh username@casper.ucar.edu
- X11 forwarding option:
 - -X for Linux/Windows
 - -Y for macOS
- You will be placed on one of 6 login nodes:
cheyenne[1-6] (check your host name by “echo \$HOST”)

User Authentication

Two factor Authentication:
CIT password +
Use Yubikey for authentication
Use Duo app for authentication



- app push 
 - “,” + passcodes 
- 

Be mindful of your usage on shared resources like the login nodes

- Your activities coexists with those of other users
- CPUs and memory are shared on the login nodes
- Limit your usage to:
 - Reading and writing text/code
 - Compiling smaller programs
 - Performing data transfers
 - Interacting with the job scheduler
- Programs that use excessive resources on the login nodes will be terminated

Data storage - GLADE

File spaces optimized for parallel IO, accessible from all HPC systems

File space	Quota	Backup	Uses
Home /glade/u/home/\$USER	25 GB	Yes	Settings, code, scripts
Work /glade/work/\$USER	1 TB	No	Compiled codes, models
Scratch /glade/scratch/\$USER	10 TB	Purged!	Run directories, temp output
Project /glade/p/entity/project_code	N/A	No	Project space allocations

*Keep track of usage with “**gladequota**”*

Data storage - Campaign Storage

- Resource for storing data on publication timescales
- Data purge after 5 years
- Multiple access methods
 - Globus (NCAR Campaign Storage)
 - Casper nodes (/glade/campaign/)
 - Data access nodes (/glade/campaign/)
- Allocated to and managed by NCAR labs and can be requested by University users

Data storage - Collections

- Curated data collections available on Cheyenne and Casper to facilitate easy access to research data sets
- RDA
 - Research Data Archive
 - /glade/collections/rda/
 - <https://www2.cisl.ucar.edu/data-portals/research-data-archive>
- CMIP6
 - Coupled Model Intercomparison Project
 - /glade/collections/cmip/CMIP6/
 - <https://www2.cisl.ucar.edu/resources/cmip-analysis-platform>

Moving data to and from GLADE

- For short transfers use **scp/sftp** to transfer files
- For large transfers use **Globus**
 - To use Globus, create a Globus ID if you need an account, and search for **NCAR GLADE** or **NCAR Campaign Storage** endpoints
 - CISL endpoints currently can be activated for up to 30-days
 - Globus has a web interface and a command-line interface
 - **Globus Connect Personal** can manage transfers from your local workstation as well.

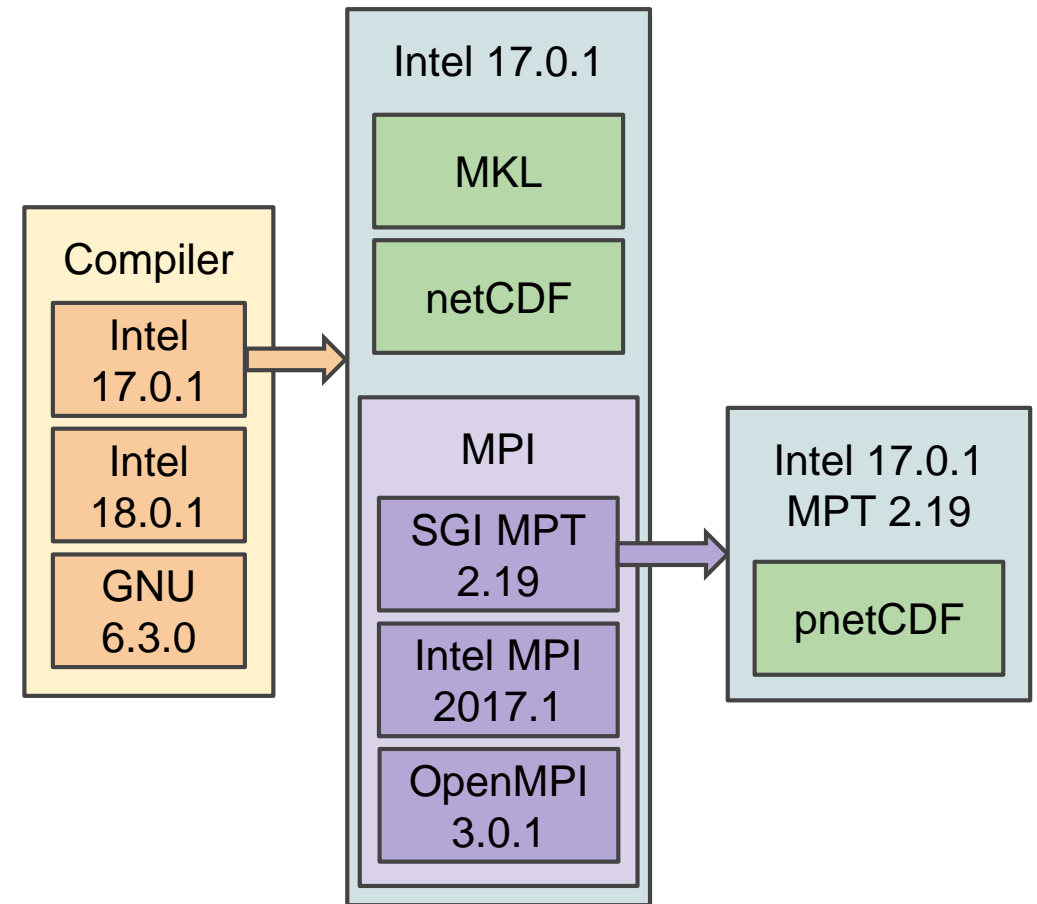


Environment Modules

- CISL installed software is provided as modules
- Modules provide access to runnable applications (compilers, debuggers, ...) as well as libraries (NetCDF, MPI, ...)
- Modules prevent loading incompatible software into your environment
- **Note that Cheyenne and DAV cluster each have independent collections of modules!**

Using modules

- **module load/unload <software>**
- **module avail** - show all currently-loadable modules
- **module list** - show loaded modules
- **module purge** - remove all loaded modules
- **module reset** - reset to default modules
- **module save/restore <name>**
create/load a saved set of software
- **module spider <software>**
search for a particular module



Changing your default modules

- If you commonly load certain modules, you may wish to have them load automatically when logging onto a cluster
- The right way to do so is with saved module sets:

```
module load ncl python nco mkl  
module save default
```

- Make multiple sets and load them using **module restore <set>**
- **Don't put module load commands in your shell startup files!**

Available Software

- Compilers (Intel, GNU, PGI)
- Debuggers / Performance Tools (DDT, MAP)
- MPI Libraries (MPT, Intel MPI, OpenMPI)
- IO Libraries (NetCDF, PNetCDF, HDF5)
- Analysis Languages (Python, Julia, R, IDL, Matlab)
- Convenience Tools (ncarccompilers, parallel, rclone)
- Many more ...

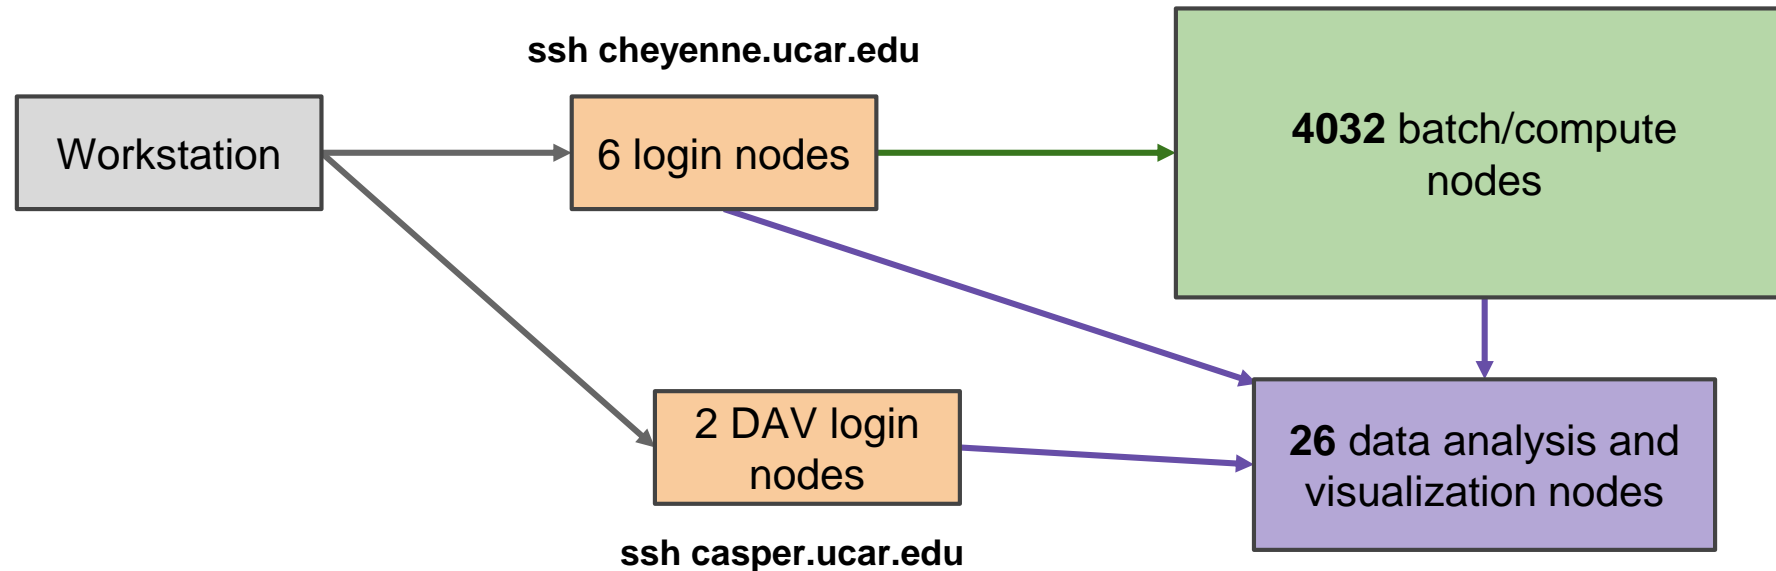
Considerations when compiling software

- Use **ncarc compilers** module along with library modules (e.g., netcdf) to simplify compiling and linking (*it adds include and link flags for you*)
- When using MPI, make sure you run with the same library with which you compiled your code
- **We strongly recommend you build code on the machine on which you will run**
 - Cheyenne and DAV cluster have different CPUs and operating systems

Run large tasks on compute nodes using batch jobs





- Most tasks require too many resources to run on a login node
- On a Cheyenne login node, the user can schedule the tasks to run on the Cheyenne compute nodes via **PBS** or on Casper nodes via **Slurm**



Supercomputer's nodes

- Each node is “like” your workstation, with:
 - 36-cores @2.3 GHz with either 64 or 128GB of RAM
- You can run a parallel “shared memory” job on a single node – up to 36/72 threads
- You can run a parallel “distributed memory” job on many nodes – whatever tasks you'd like

Shared vs Private Memory



Parallel Jobs

- Shared memory jobs usually use the OpenMP programming API
- Distributed memory usually use the MPI programming API
- Parallel jobs can use MPI, OpenMP or both (the latter are called hybrid)

Run large tasks on compute nodes using batch jobs



- Jobs request a given number of compute tasks for an estimated wall-time on specified hardware
- Jobs use core-hours, which are charged against your selected project/account. Remaining resources are viewable in SAM
- Temporary files are often written by programs - set TMPDIR variable to scratch space to avoid job failures
- You can run interactive jobs on batch nodes

Example PBS batch job script

```
$ cat basic_mpi.pbs
#!/bin/tcsh
#PBS -N hello_pbs
#PBS -A <project_code>
#PBS -j oe
#PBS -o pbsjob.log
#PBS -q regular
#PBS -l walltime=00:05:00
#PBS -l select=2:ncpus=36:mpiprocs=36

### Set temp to scratch
setenv TMPDIR /glade/scratch/${USER}/temp
mkdir -p $TMPDIR

module load mpt/2.19

### Run MPT MPI Program
mpiexec_mpt ./hello_world
```

Interacting with the job schedulers

PBS on Cheyenne

qsub *<script>* - submit batch job

qstat *<jobid>* - query job status

qstat -u *<your_username>* - query your job status

qstat -q - query queue status

qdel *<jobid>* - delete/kill a job

qinteractive -A *<project>*

Run an interactive job

qcmd -A *<project>* -- *cmd.exe*

Run a command on a single compute node

Using threads/OpenMP parallelism on Cheyenne with MPT

Only OpenMP

```
#!/bin/tcsh
#PBS -l select=1:ncpus=10:ompthreads=10

# Run program with 10 threads
./executable_name
```

Hybrid MPI/OpenMP

```
#!/bin/tcsh
#PBS -l select=2:ncpus=36:mpiprocs=12:ompthreads=3

module load mpt/2.19

# Run program with one MPI task and 36 OpenMP
# threads per node (two nodes)
mpiexec_mpt omplace ./executable_name
```

Running serial code on multiple data files using command file jobs

```
./cmd1.exe < input1 > output1  
./cmd2.exe < input2 > output2  
./cmd3.exe < input3 > output3  
./cmd4.exe < input4 > output4
```

cmdfile contents

```
#!/bin/tcsh  
#PBS -l select=1:ncpus=4:mpiprocs=4  
  
module load mpt/2.19  
  
# This setting is required to use command files  
setenv MPI_SHEPHERD true  
  
mpiexec_mpt launch_cf.sh cmdfile
```

PBS Job script

Optimal if commands have similar runtimes

PBS job queues on Cheyenne

PBS Queue	Priority	Wall clock	Details
premium	1	12 h	Jobs are charged at 150% of regular rate
regular	2	12 h	Most production compute jobs go here
economy	3	12 h	Jobs are charged at 70% of regular rate
share	N/A	6 h	Memory is shared among all users on a node Jobs are limited to 18 cores or less

Job charges depend on the queue:

Exclusive: wall-clock hours \times nodes \times 36 cores/node \times queue factor

Shared: core-seconds / 3600

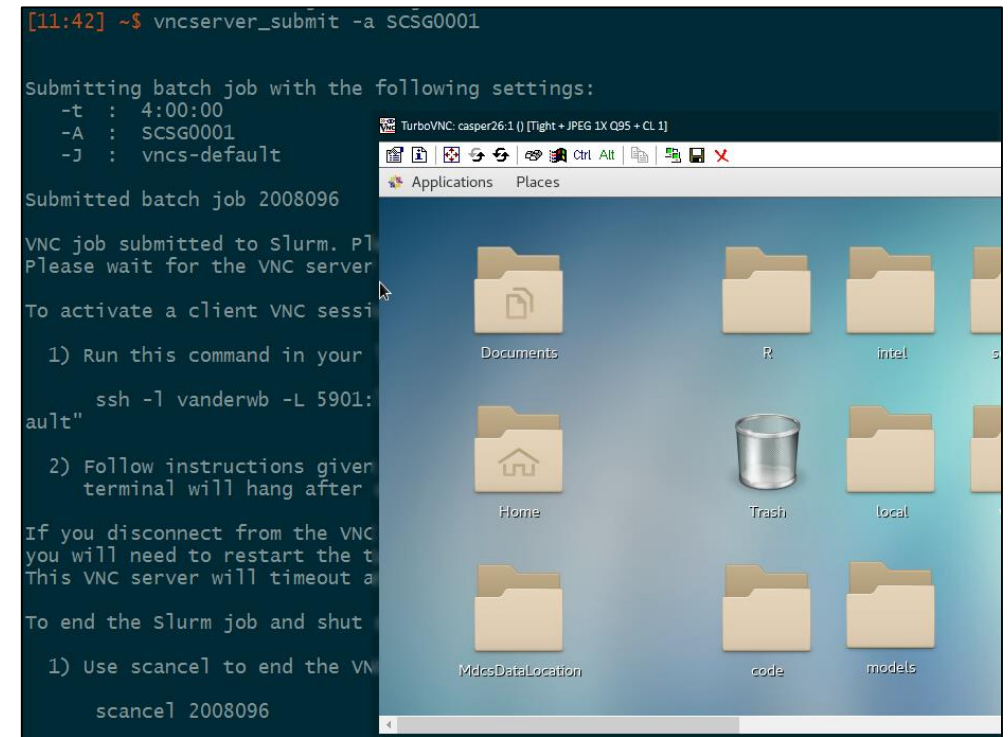
Run GUI Programs with TigerVNC

VNC can be used to run a remote GNOME/KDE desktop

Need to install a VNC client first - We recommend TigerVNC, but other VNC clients such as TurboVNC will also work

Usage:

```
vncserver_submit -a <project>
```

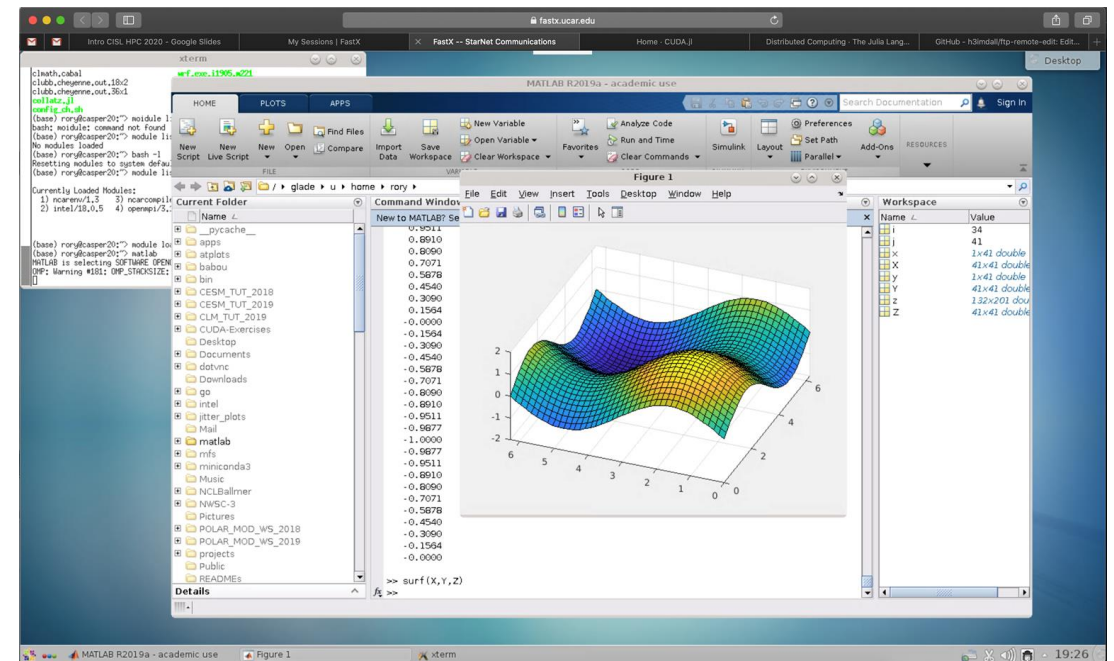


Run GUI Programs with FastX

FastX can also be used to run a remote desktop or terminal session

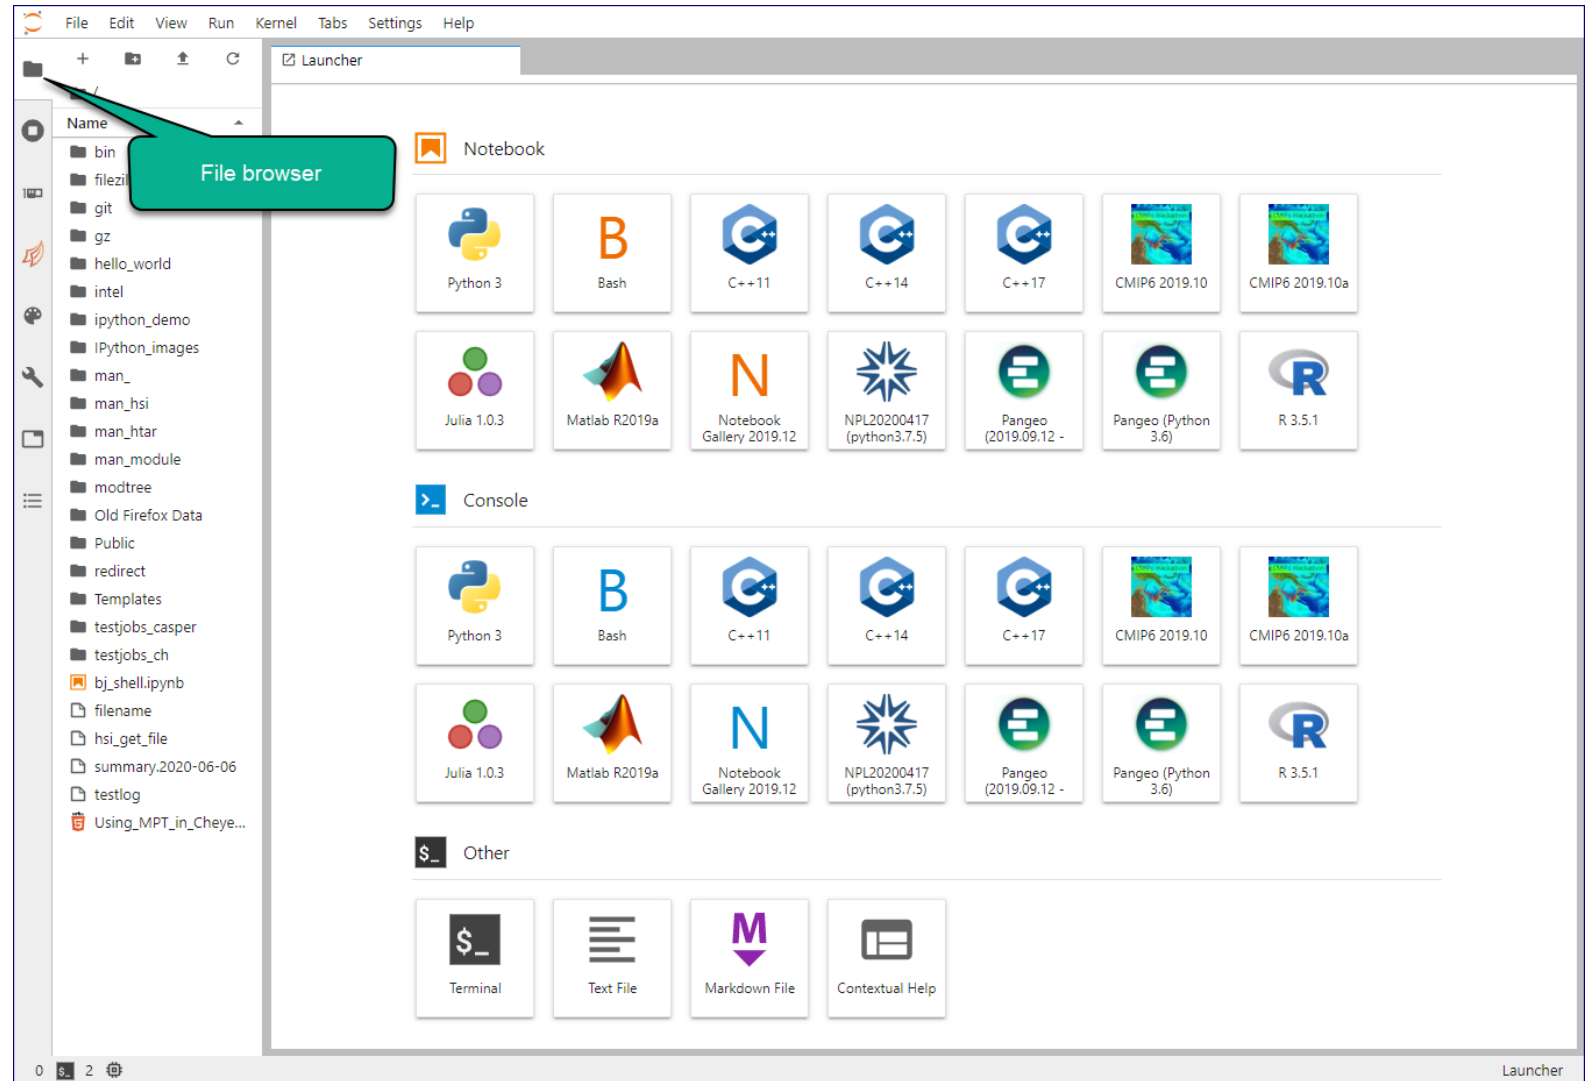
Can be run in a browser without a client

Connect to the NCAR VPN, then go to <https://fastx.ucar.edu:3300>



JupyterHub

- Supports:
 - Python
 - C++
 - Julia
 - Matlab
 - R
 - Shell (bash)
- Familiar environment
- Casper / Cheyenne



DATA ANALYSIS AND VISUALIZATION PACKAGES

These are among the more frequently used data analysis and visualization packages available to Cheyenne users:

GrADS / IDL / MATLAB (Octave) / NCL / ParaView / PyNGL and PyNIO / VAPOR

Additional applications and tools are available through the CISL Research Data Archive. Such as:

Mathematica / Vis5d / VTK.

License Use Guidelines

The CISL user community shares a limited number of licenses on some software and compilers. Avoid monopolizing these licenses. CISL reserves the right to kill jobs/tasks of users who monopolize these licenses. To see how many licenses are being used, at your command line, run command:

licstats

What's next?

- Request an allocation:
<https://www2.cisl.ucar.edu/user-support/training>
(a grad student at a US university can easily request a small allocation)
- Course library: <https://www2.cisl.ucar.edu/user-support/training>

CISL Help Desk / Consulting

Support Page:

<https://www2.cisl.ucar.edu/user-support/getting-help>

- File support tickets
- Call Help Desk: 303-497-2400
- Walk-in Service Closed indefinitely due to the pandemic

Thanks for the attention!

Any more questions?