ACCRE
Advanced Computing Center for Research & Education

INTRODUCTION TO THE CLUSTER

What is a cluster?

A **cluster** is a bunch of computers that are networked together to perform intensive computing jobs.

If it takes a long time to run on a laptop…

…it can run on ACCRE faster and with more memory, doing more in less time

WHERE IS THE ACCRE CLUSTER?

The cluster is downstairs!

You don’t need to be there to operate it, in fact we restrict physical access to the cluster and we only go there for maintenance

10 REASONS TO USE THE ACCRE CLUSTER

ACCRE
Advanced Computing Center for Research & Education

Accessing larger processor or memory resources	It’s available anywhere
Getting more done in less time	You can share files and results with others
It’s reliable	It’s reasonably priced
All the software is installed	It’s backed up
We’re here to help	We can handle sensitive data

STARTING ACCRE FOR THE FIRST TIME

In most cases you use a **SSH** client (**s**ecure **sh**ell client)

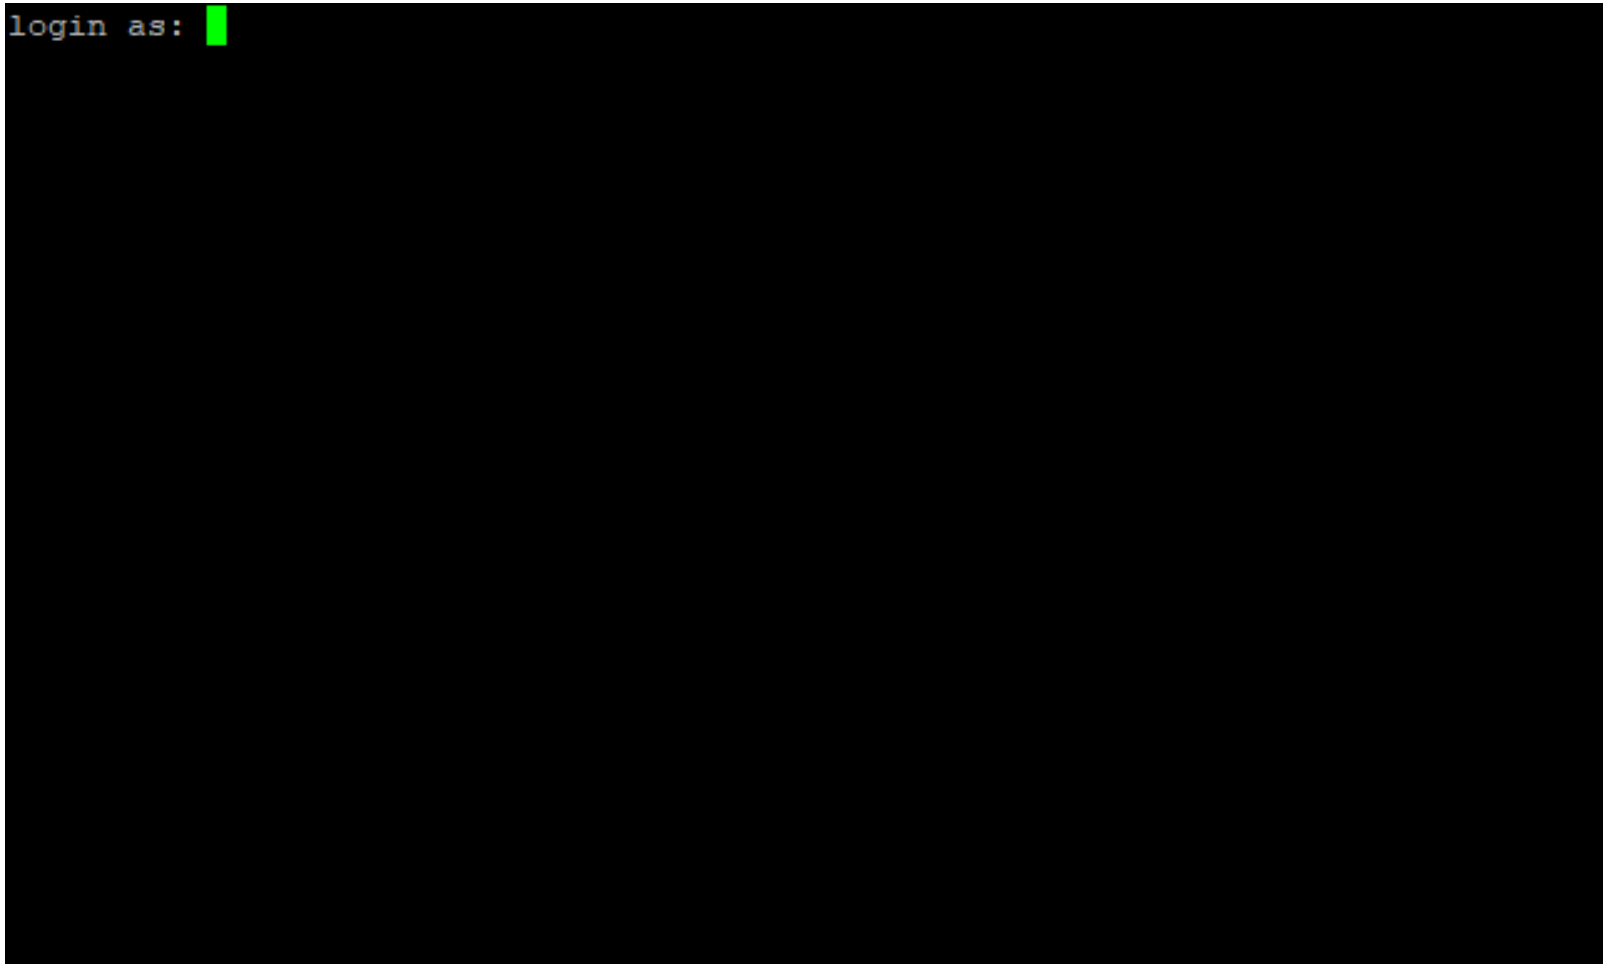

STARTING ACCRE FOR THE FIRST TIME

ACCRE Advanced Computing Center for Research & Education

Users can log into the cluster with a Secure Shell (**ssh**) client.

From a terminal: **`ssh` *`vunetid`*`@login.accre.vanderbilt.edu`**

PuTTY

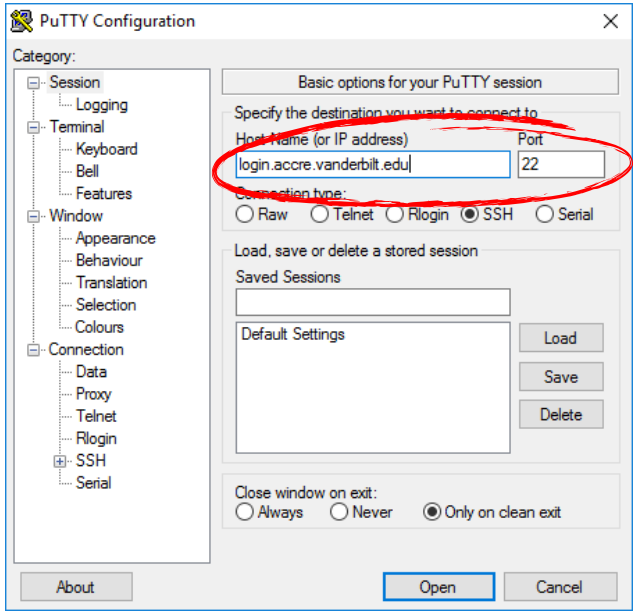

Bash on Windows

Full Ubuntu-based Bash shell

To install:

https://goo.gl/tAsj8U

Windows 10 only

Account credentials

You will need to set up a password if you haven't already.

Password criteria is very strict!

- At least 14 characters long (but less than 4096)
- Perfect score on the ZXCVBN test
- You may use non-ASCII characters (emoji, Chinese, Japanese, Korean, etc.)
- This may cause issues with some terminal clients
- All characters must be printable and in UTF-8

If you need a password, type "accre_password generate"

Account credentials

How can I change my password?

Log into the cluster

Get a suggestion: `accre_password generate`

Change your password: `accre_password change`

When prompted, follow instructions

Your new password will be immediately active

I forgot my password
or
My password is expired

Open a helpdesk ticket:

www.vanderbilt.edu/accre/help

Do not use your VU E-password!

Account credentials

Use common sense!

- Choose a strong password
- Use a unique password for ACCRE, and do not use it for anything else
- If you write down your password, keep the note hidden and secure
- Consider using a password manager
- Never share your password with anyone for any reason
- If anyone asks for your password, report the incident to ACCRE and to VUIT security
- If you suspect your password is compromised, alert ACCRE and change it

The ACCRE Visualization Portal

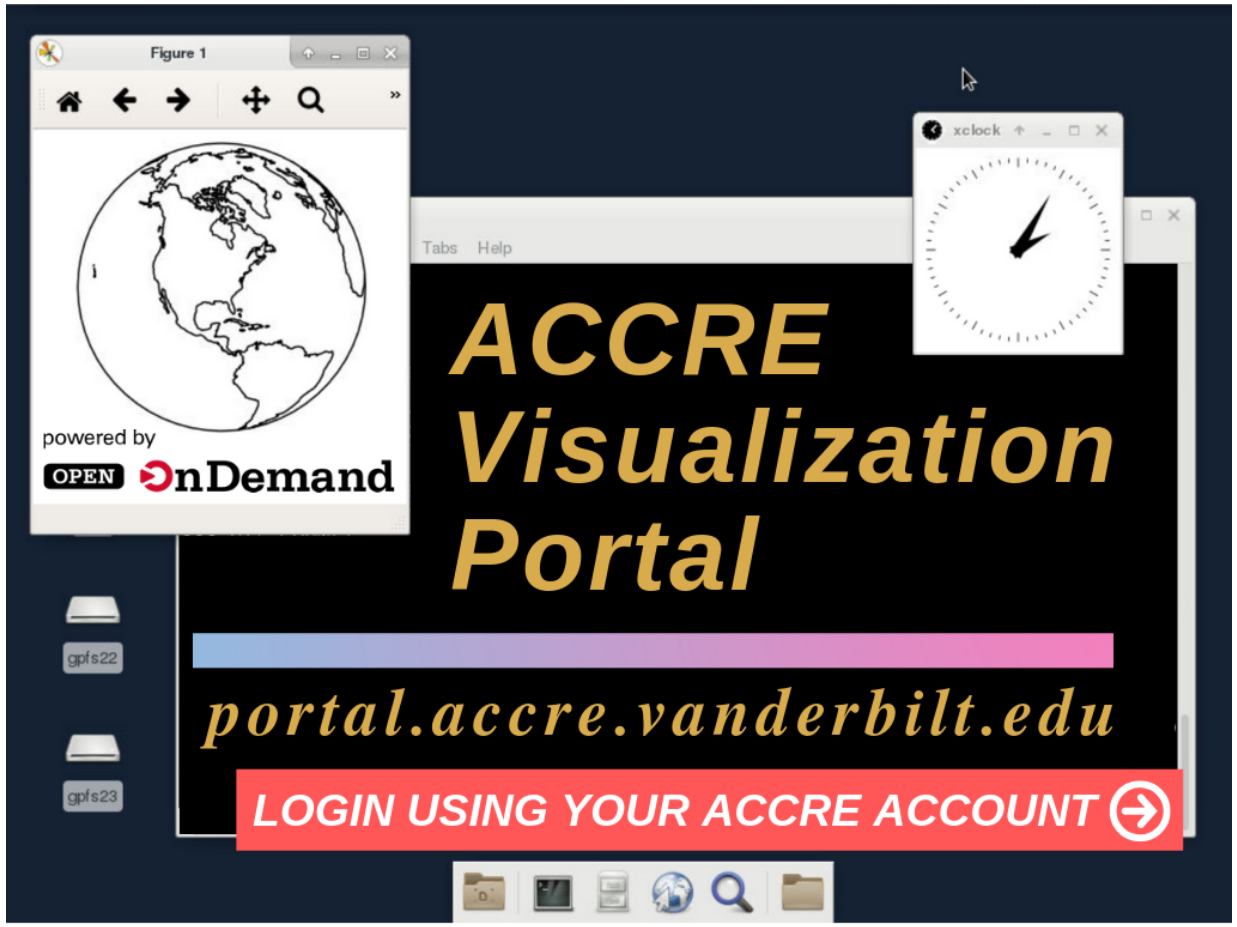

The Visualization Portal allows for cluster access using any modern web browser:

- File management, upload, download, editing
- Shell access to gateways in browser
- View running jobs
- Spawn interactive desktop jobs on compute nodes
- Spawn Jupyter notebooks on compute nodes

WATCH VIDEO

STARTING THE VISUALIZATION PORTAL

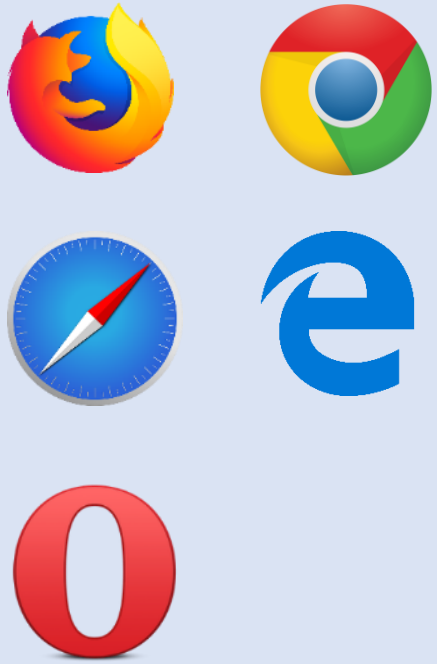

Log into the Visualization Portal using an up to date web browser:

https://portal.accre.vanderbilt.edu

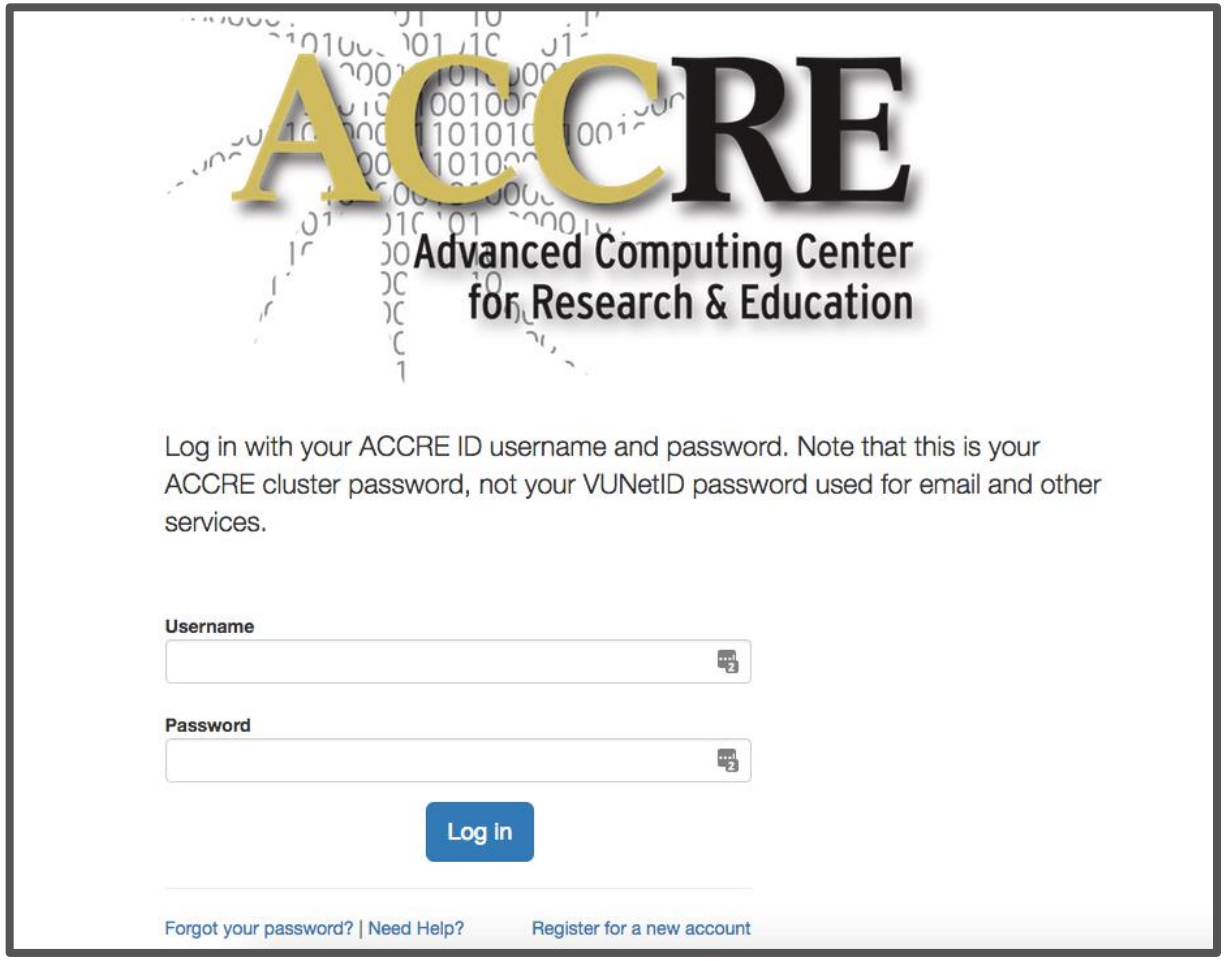

Check for browser updates and a valid site certificate before entering credentials!

Data Transfer - Portal

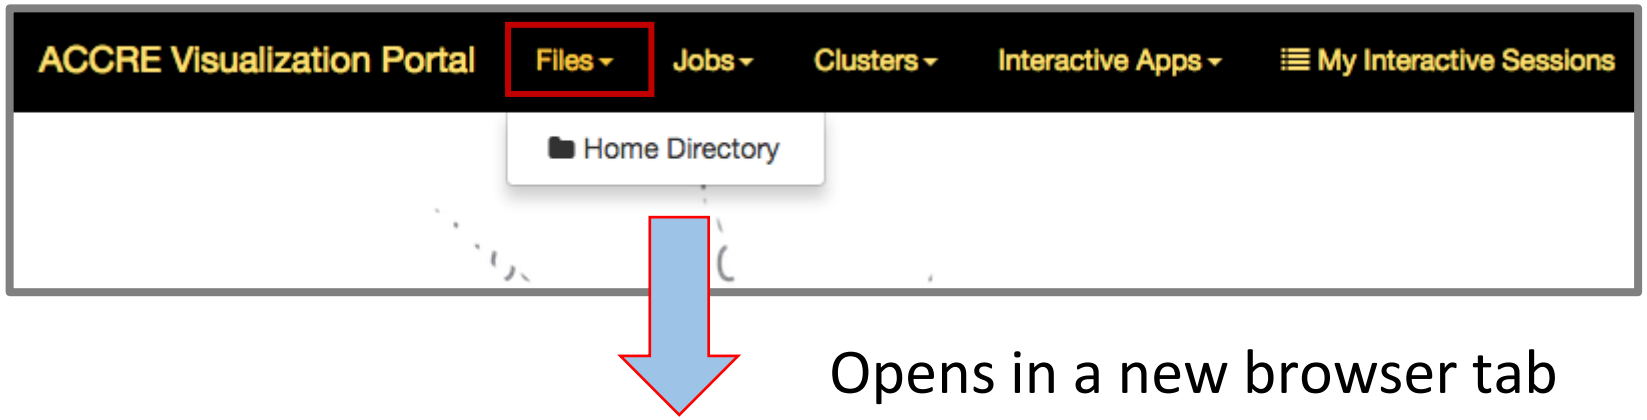

Opens in a new browser tab

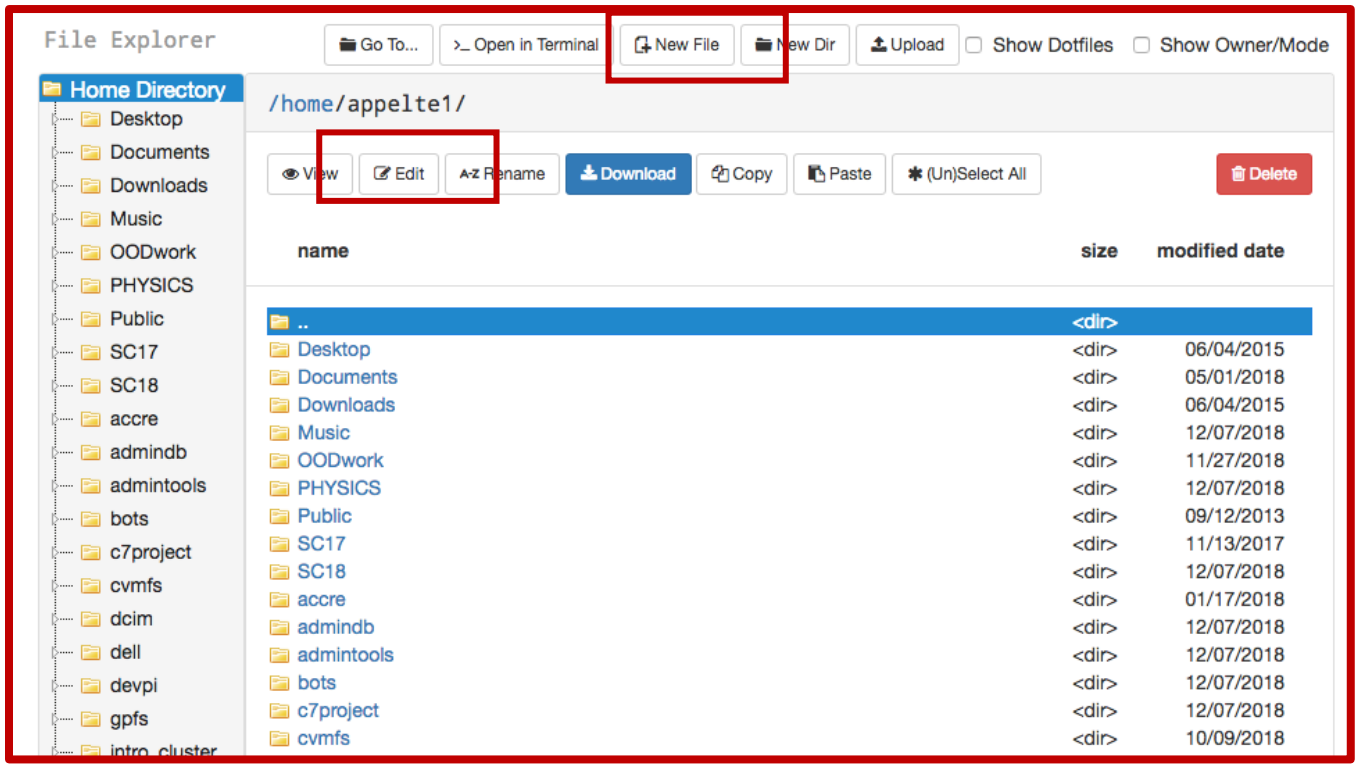

Shell Access to the Cluster – Portal

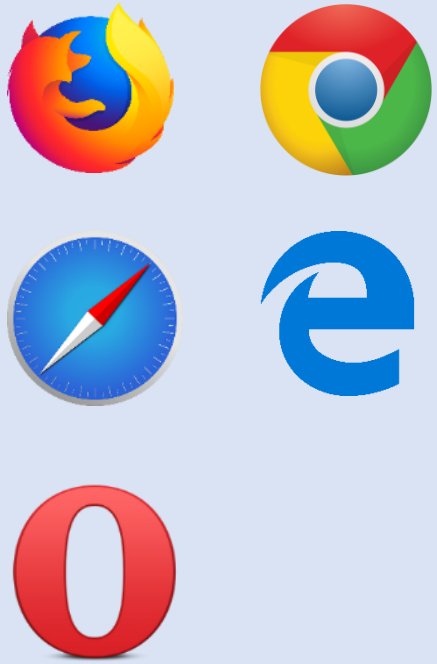

Opens in a new browser tab

```
Last login: Fri Dec  7 08:52:49 2018 from 10.0.20.109
Vanderbilt University - Advanced Computing Center for Research and Education

 _   ___ ____ ____ _____    ____ _           _
 / \  / __/ ___|  _ \| ____|  / ___| |_   _ ___| |_ ___ _ __
 / _ \| |  | |   | |_) | _|   | |   | | | / __| __/ _ \ '__|
 / ___ \ |__| |___|  _ <| |___  | |___| | |_| \__ \ ||  __/ |
/_/   \_\___\____|_| \_\_____|  \____|_|\__,_|___/\__\___|_|

================================================================

Go forth and compute!

This is a shared gateway node designed for interactive use and small test jobs.
Please restrict your total system memory usage to less than 31 GB,
and do not run individual processes exceeding 20 minutes of CPU-time.

To list useful cluster commands type:      accre_help
To view your current storage type:         accre_storage
To list basic Linux commands type:         commands101
[appelte1@gw341 ~]$
```

Remote Interactive Desktop

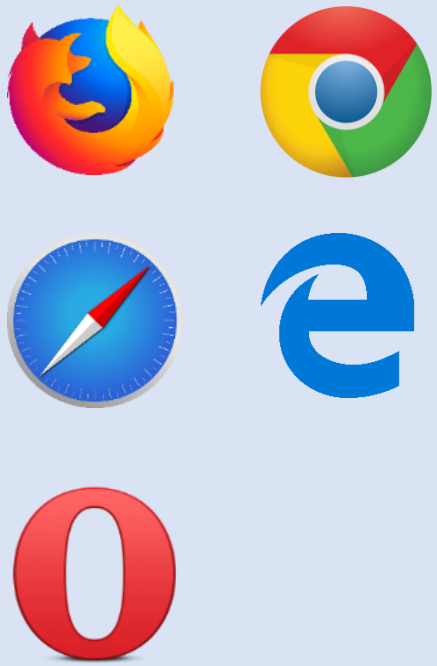

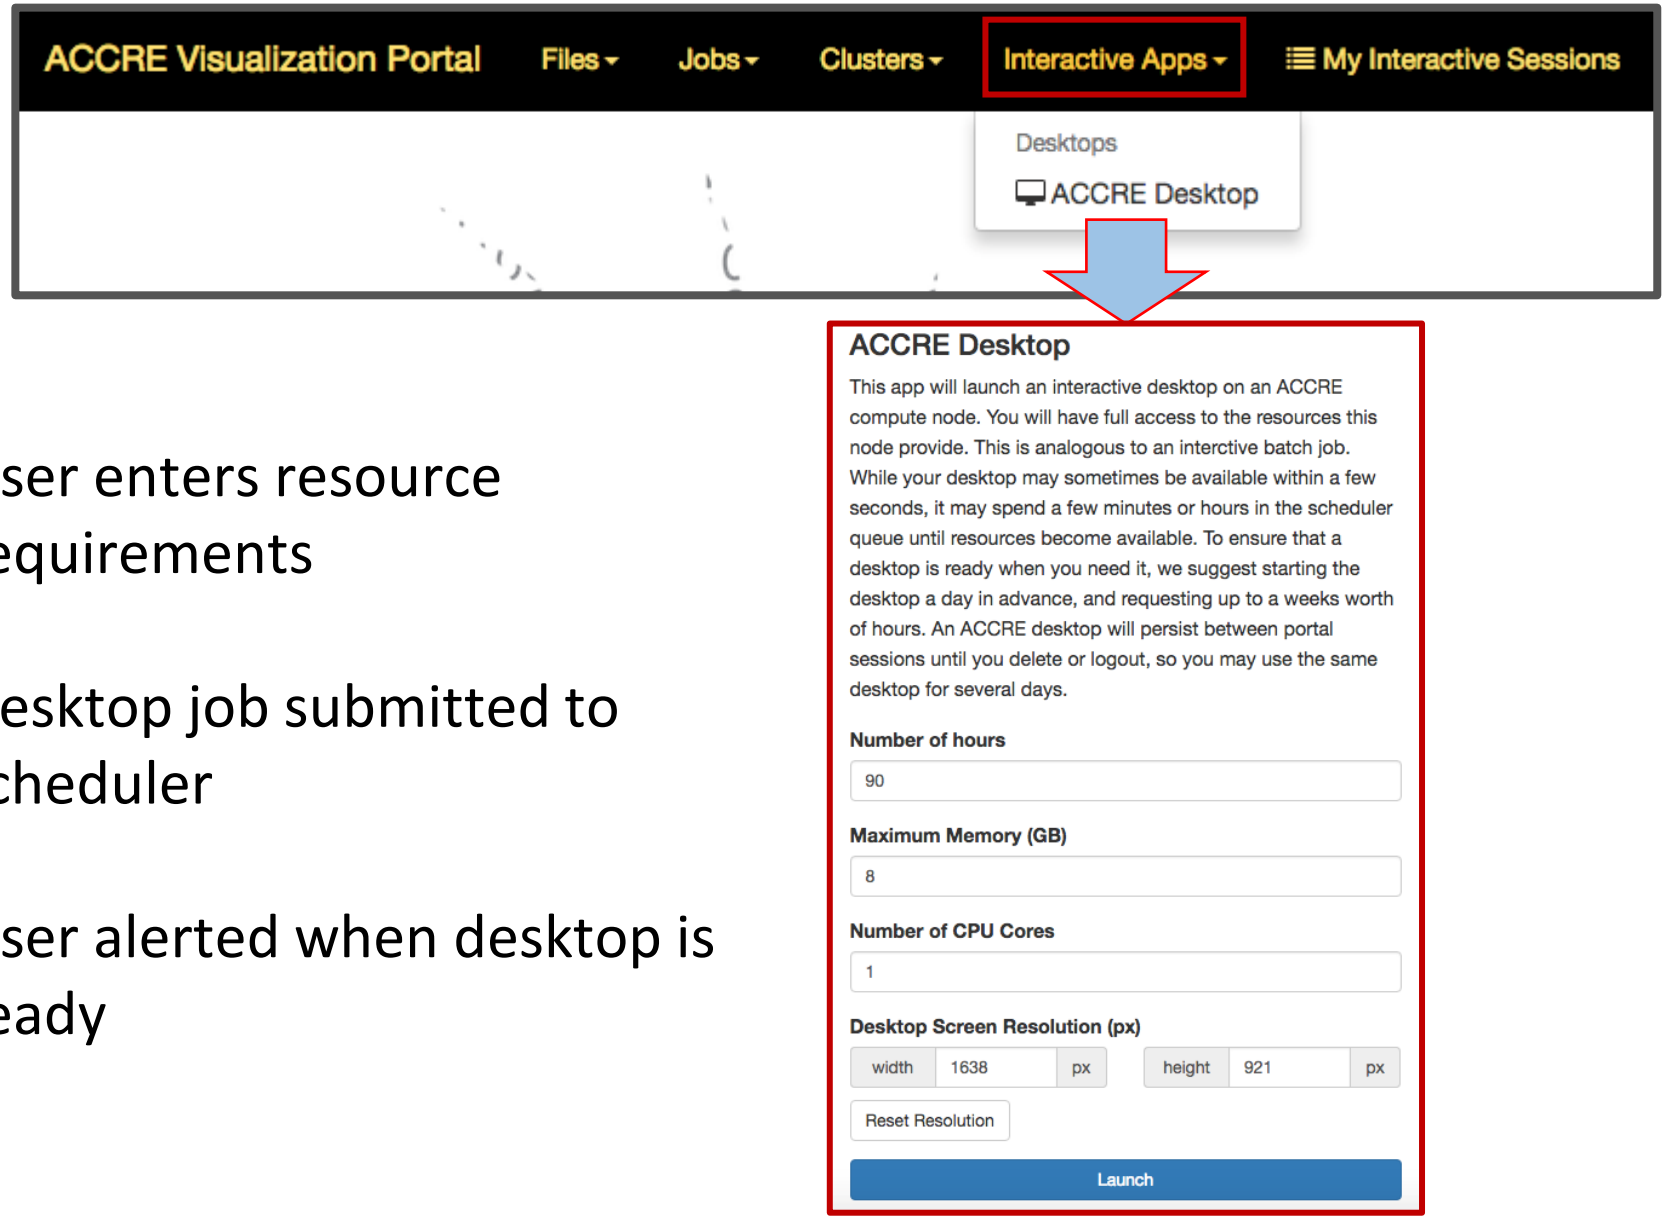

- User enters resource requirements
- Desktop job submitted to scheduler
- User alerted when desktop is ready

Remote Interactive Desktop

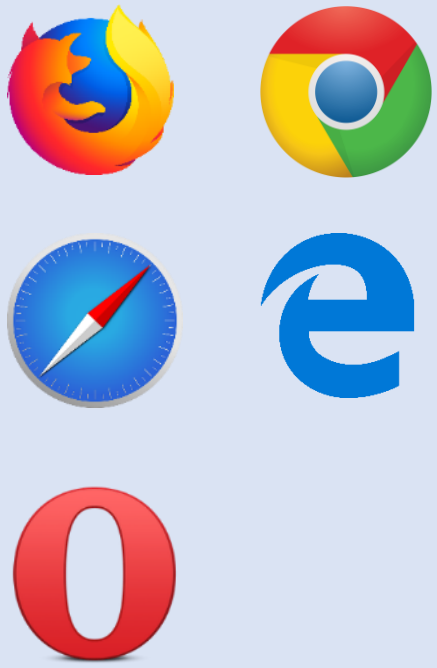

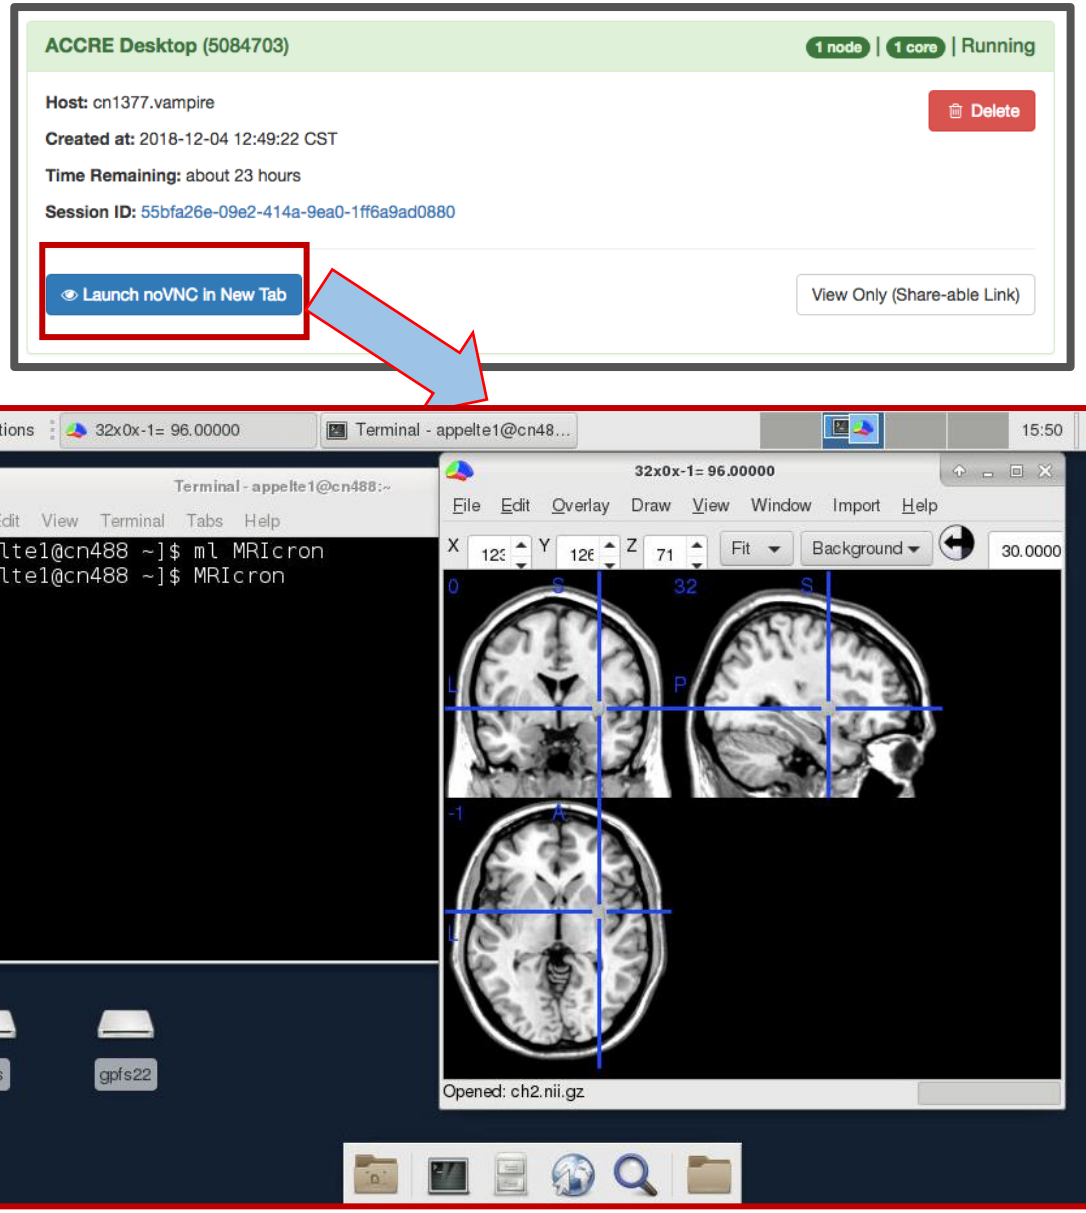

Desktop opens in a new browser tab.

The desktop will persist between browser sessions.

Data transfer

Best way – use the file explorer on the ACCRE Visualization Portal!

If you must use a terminal:

`scp source destination`

- Copy data from *source* to *destination*.
- Remote source or destination must be preceded by "*vunetid*@login.accre.vanderbilt.edu:"

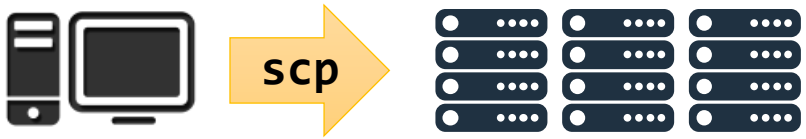


```
scp local_path vunetid@login.accre.vanderbilt.edu:remote_path
```

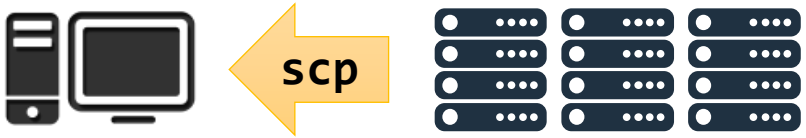


```
scp vunetid@login.accre.vanderbilt.edu:remote_path local_path
```

Data transfer

WinSCP

https://winscp.net

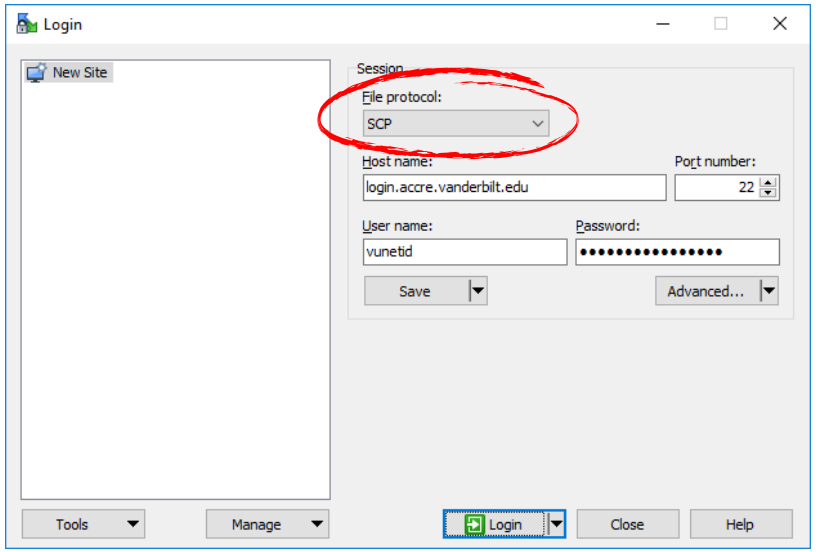

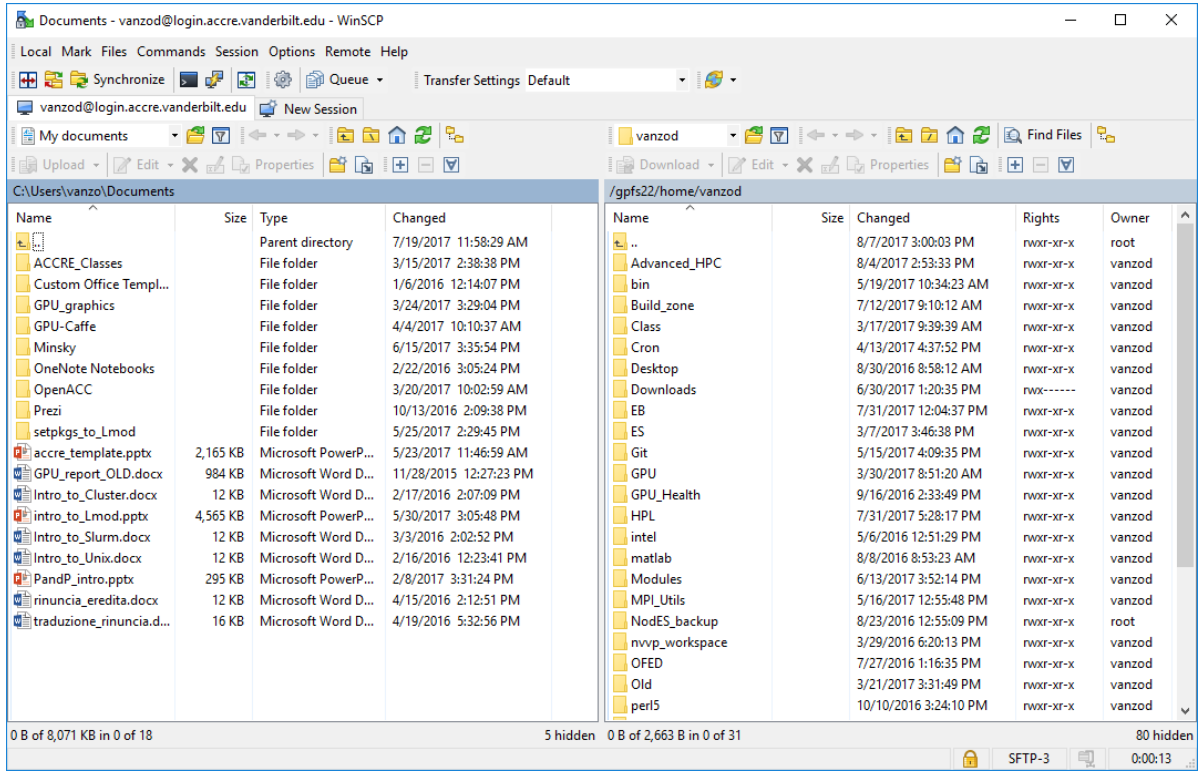

CONNECT WITH REMOTE DISPLAY SUPPORT

Best way – create a virtual desktop on the ACCRE Visualization Portal to access graphical programs.

Remote displays on the terminal are slow and use up resources for other users.

If you must use a terminal:

From a terminal:

```
ssh -X vunetid@login.accre.vanderbilt.edu
```

MacOS

Install **XQuartz** and connect as for Linux.

www.xquartz.org

1. Install and launch **Xming**

https://sourceforge.net/projects/xming

2. Configure PuTTY

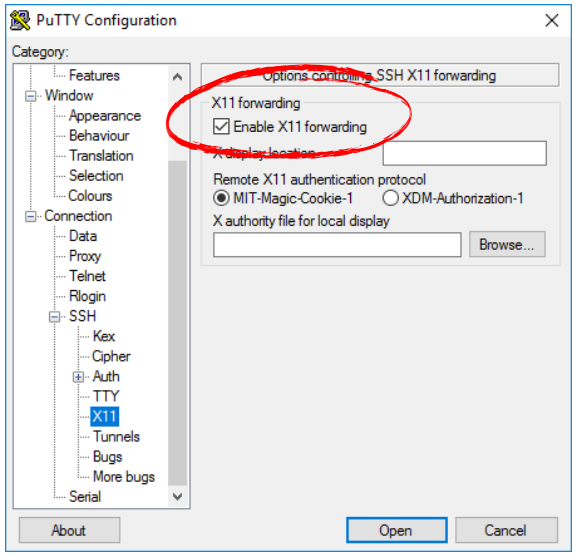

What is Lmod?

Lua-based module system

Allows you to run software packages and build environments based on them

Developed at Texas Advanced Computing Center

SOFTWARE HIERARCHY

Software is organized in a tree structure and displayed accordingly to the loaded dependencies.

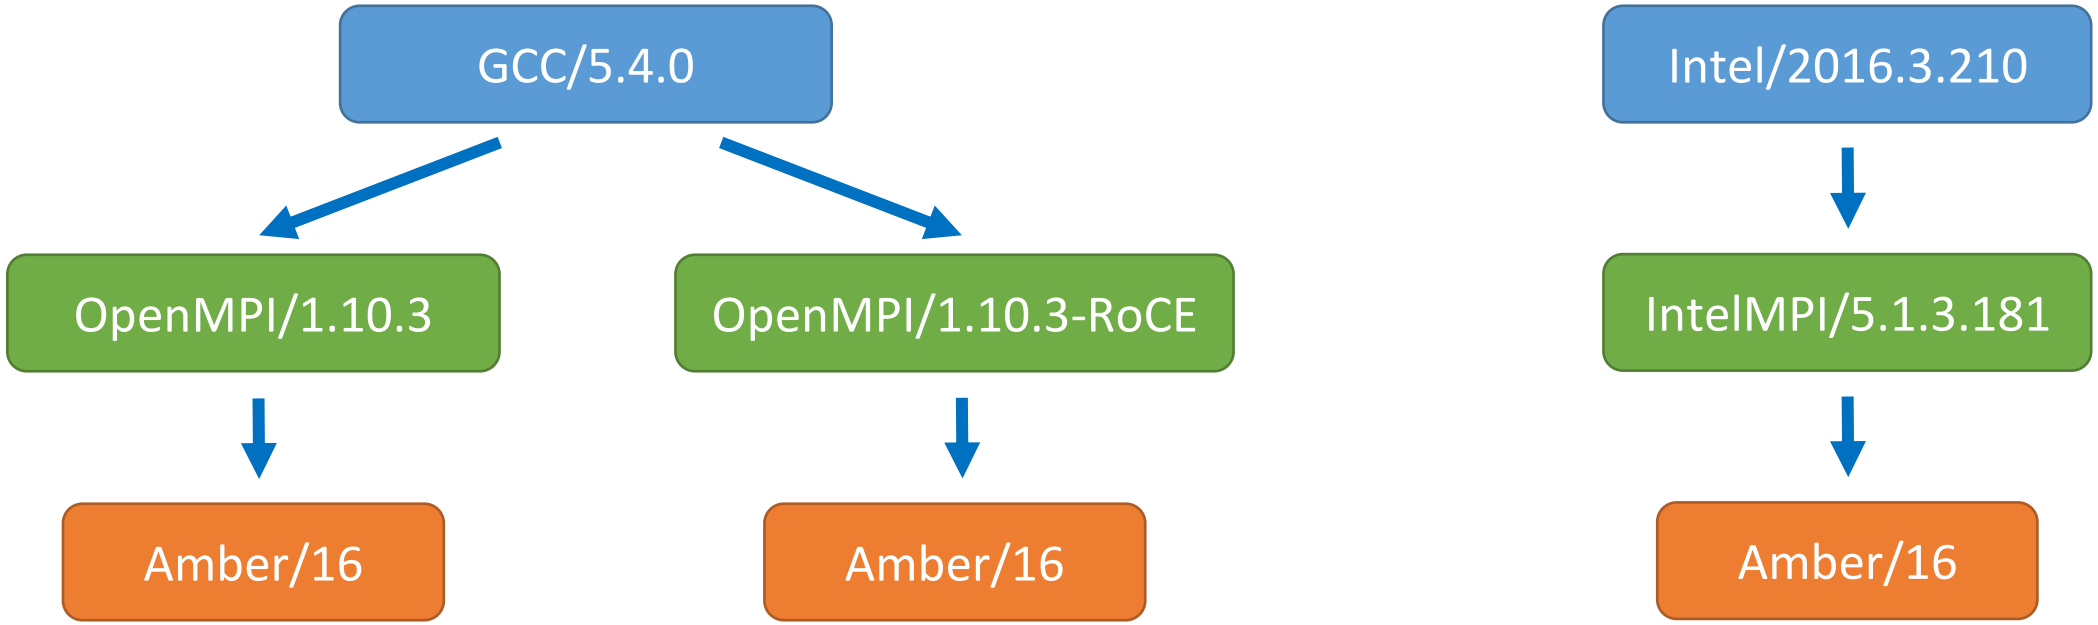

Different ways to load Python

~~Module load python~~

Module load Anaconda3

Module load GCC python

Module load Intel python

While helpful for beginners, Anaconda Python isn't optimized for our hardware and can quickly use up the filesystem quota

You will need to load the GCC or Intel module prior to loading the Python module. For starting out, GCC and Intel are interchangeable

LMOD - THE ESSENTIALS

ACCRE Advanced Computing Center for Research & Education

`module avail <mod>`

- If no module is passed, print a list of all modules that are available to be loaded.
- If a module is specified, show all available modules with that name.

`module load mod1 mod2 …`

- Load the specified modules.

`module unload mod1 mod2 …`

- Unload the specified modules.

`module list`

- Show all modules loaded in the current environment.

`module purge`

- Remove all loaded modules from the environment.

LMOD - A FEW SHORTCUTS

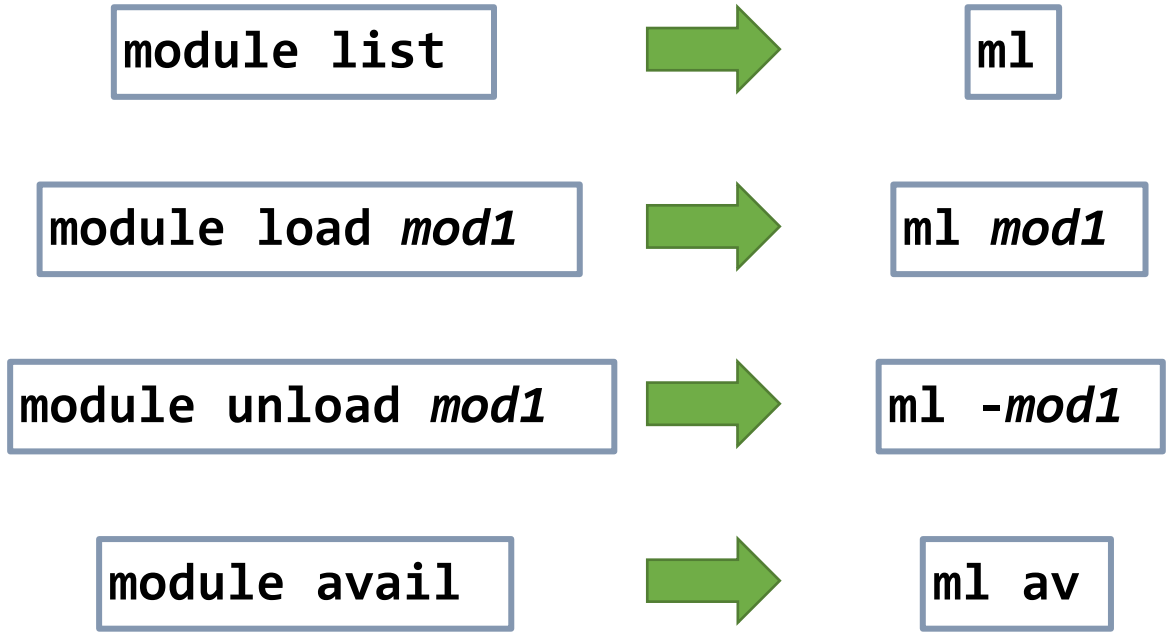

"Traditional science is all about finding shortcuts"

Rudy Rucker

LMOD - RULES OF THE GAME

ACCRE
Advanced Computing Center for Research & Education

1

Only one version of a given software can be loaded at any time.

```
module load GCC Python=3.7
module load GCC Python=3.8 # unloads Python 3.7
```

2

Incompatible programs are prevented to be loaded at the same time.

```
module load Anaconda3
module load GCC Python # doesn't work
```

3

If the version is not specified, the most recent version is assumed.

```
module load GCC Python # loads the latest GCC-compiled version of Python
```

4

When unloading a module, Lmod does not automatically unload its dependencies.

```
module load GCC Python
module unload Python # GCC is still loaded
```

LMOD - SEARCHING FOR MODULES

How can I search among the visible modules?

`module avail` *`pattern`*

- Show only the visible modules that contain the desired pattern.

How can I search through all the modules, even the non visible ones?

`module spider` *`pattern`*

- Search all the modules that contain the desired pattern.

LMOD - SAVE LOADED MODULES

I always need the same set of modules.
How can I have them loaded automatically?

OPTION 1:

Add `module load` statements in your `~/.bashrc` file.

This is the primary cause of software errors for our cluster users!

OPTION 2:

Save loaded modules in **named collections**.

```
module save collection_name
```

- Save the list of current loaded modules in ~/.lmod.d/collection_name.

```
module restore collection_name
```

- Restore the desired named collection in the current environment.

LMOD - HIDDEN MODULES

All the dependencies are built from source with the available compilers.

The whole software stack will be (mostly) independent from OS libraries.

All non-essential dependencies are hidden for user clarity.

```
module --show-hidden avail
```

- Show all the visible modules, including hidden ones.

```
module --show-hidden spider pattern
```

- Search across all modules, including hidden ones.

To load a hidden module, the version must be specified.

SOFTWARE POLICIES

New compilers/MPI with relative software stacks are available every 12 months.

Software stacks older than 3 years will be removed.

What if the software I need is not available via Lmod?

ACCRE uses **EasyBuild** to build the software stack.

Open a ticket to request the installation.

If not available via EasyBuild, we will discuss the alternatives.

SLURM: INTERFACE TO THE COMPUTE NODES

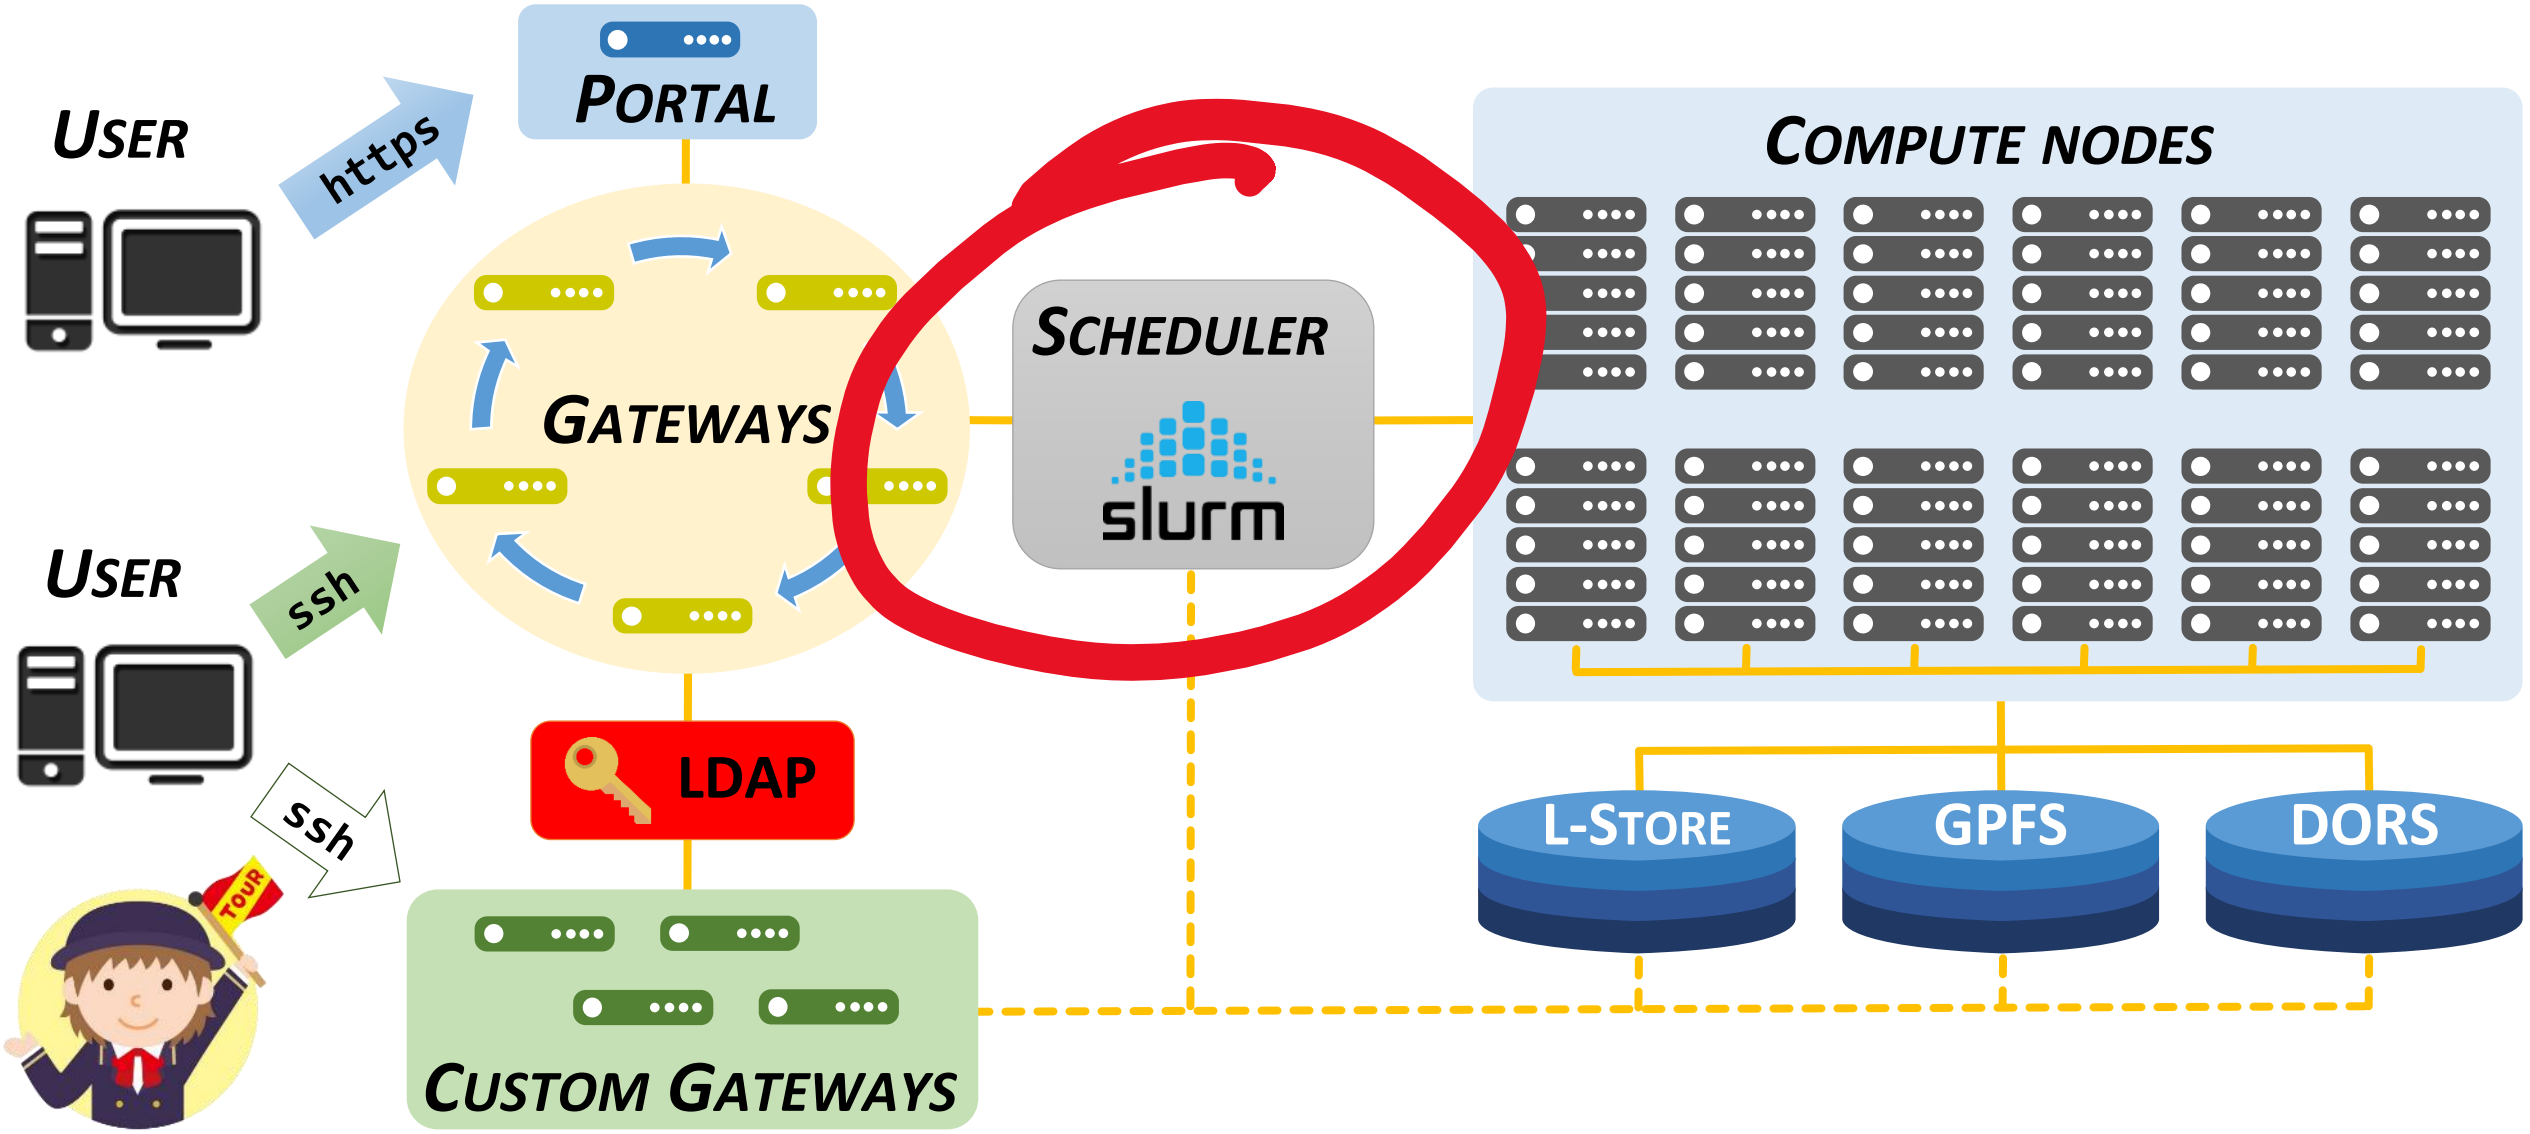

THE GATEWAYS

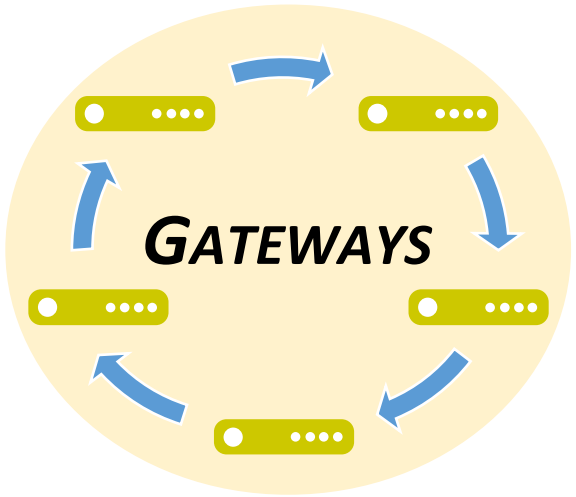

The **cluster gateways** are the main entry point to the cluster.
Multiple gateways in round-robin rotation to guarantee access.

- Manage and/or edit files
- Code development
- Jobs submission
- Lightweight debugging

- **Run resource intensive processes**

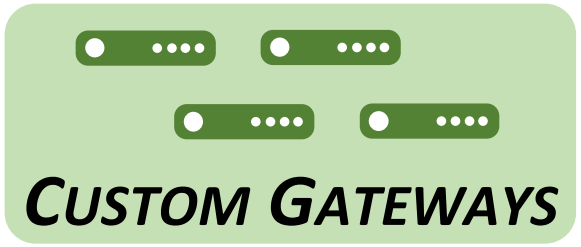

Custom gateways are gateways paid by and dedicated to specific research groups. They differ from cluster gateways for:

- Restricted group access
- Users can run resources intensive processes

THE SCHEDULER

1. Execute user's workloads in the right priority order
2. Provide requested resources on compute nodes
3. Optimize cluster utilization

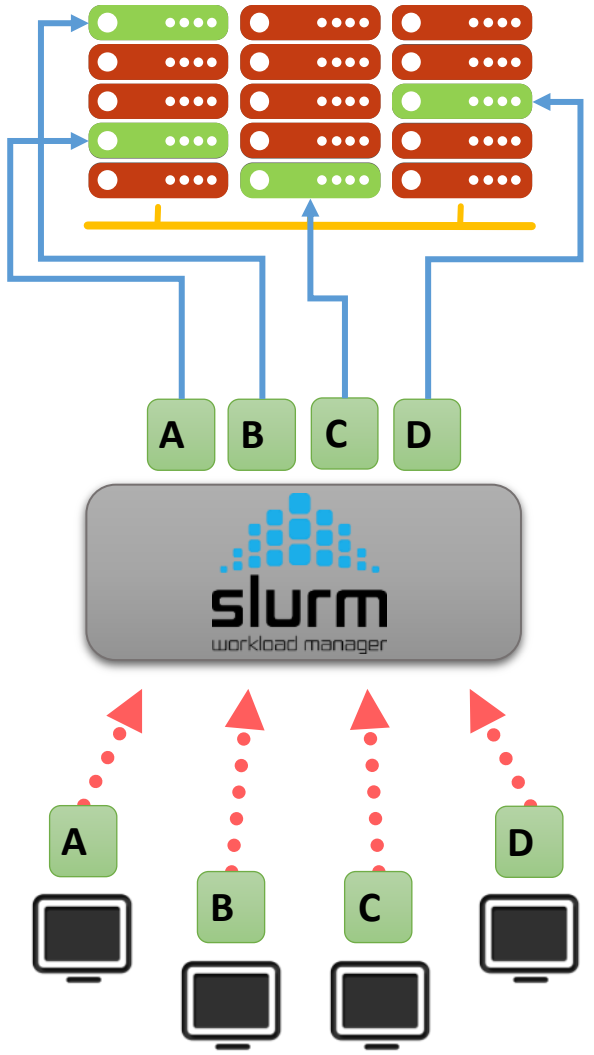

Users do not access compute nodes directly (unless requested via the portal)!

THE COMPUTE NODES

Regular nodes

Dual multicore CPUs

Random Access Memory

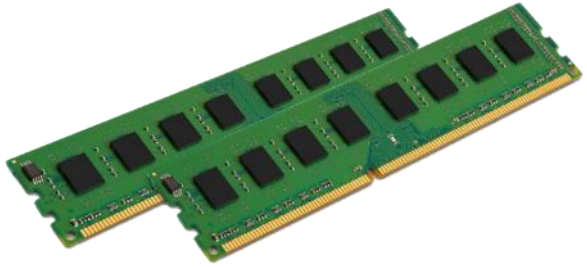

Newer

Older

Family	*No. of cores*	*RAM* / GB	*No. of nodes*
Skylake	**16**	*256*	41
	24	*128*	52
Haswell	**12**	*128*	41
	16	*128*	120
		256	50
Sandy Bridge	**12**	*64*	31
		96	2
		128	193
		256	4
	16	*128*	3
Westmere	**8**	*128*	22
	12	*48*	16
Total	**8,292**	**82,432**	**575**

THE COMPUTE NODES

Accelerated nodes

Dual multicore CPUs

Random Access Memory

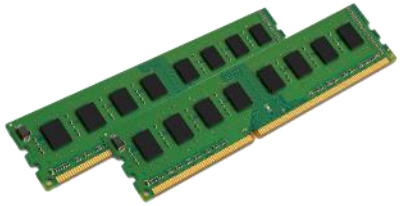

4 x Nvidia GPU

40 Gbit/s RoCE Network

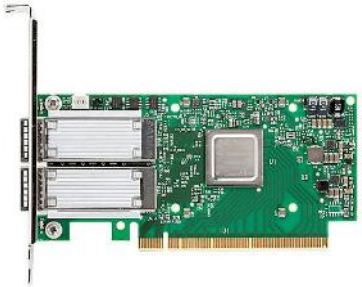

Newer

Older

Family	*No. of cores*	*RAM* / GB	*No. of nodes (GPUs)*
Nvidia Pascal ***Intel Broadwell***	**8**	*256*	24 (96)
Nvidia Maxwell ***Intel Haswell***	**12**	*128*	10 (40)
Total	**312**	**7,424**	**34 (136)**

Data Storage

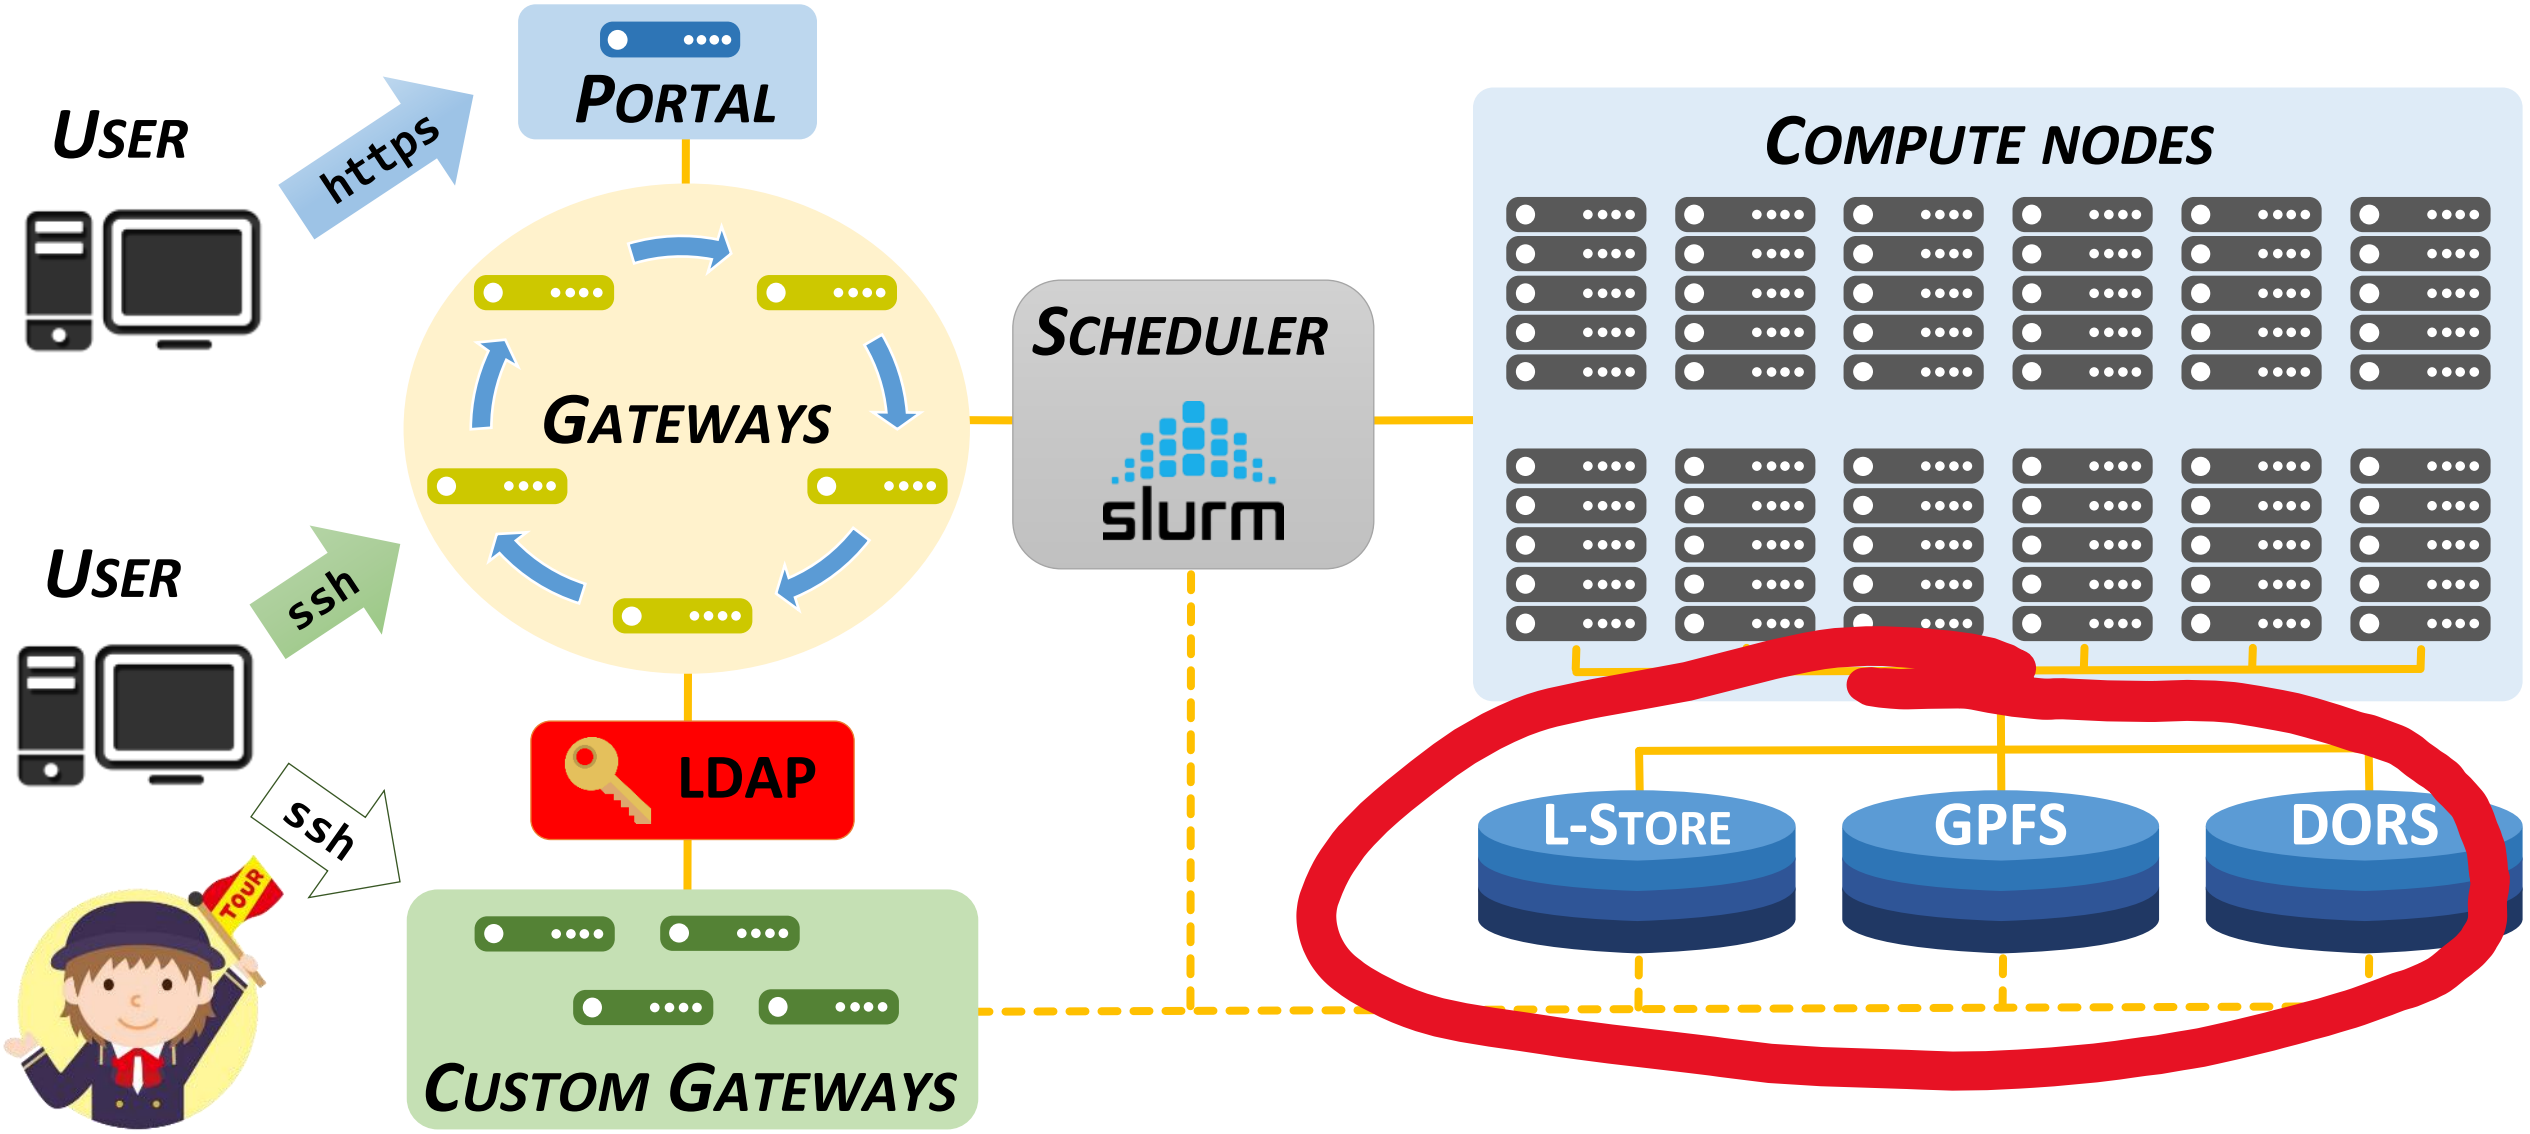

CLUSTER STORAGE

GPFS and **DORS** are distributed parallel filesystems that allow users to get access to the same set of directories on all nodes and all gateways on the cluster.

	Nightly backup	*Included with account*	*For purchase*
`/home`	✓	✓	✗
`/scratch`	✗	✓	✓
`/data`	✓	✗	✓

`/dors`

Managed by ***Center for Structural Biology***, supported by ACCRE.

Provides easy access to data from both desktops and cluster.

UNDERSTANDING QUOTAS AND LIMITS

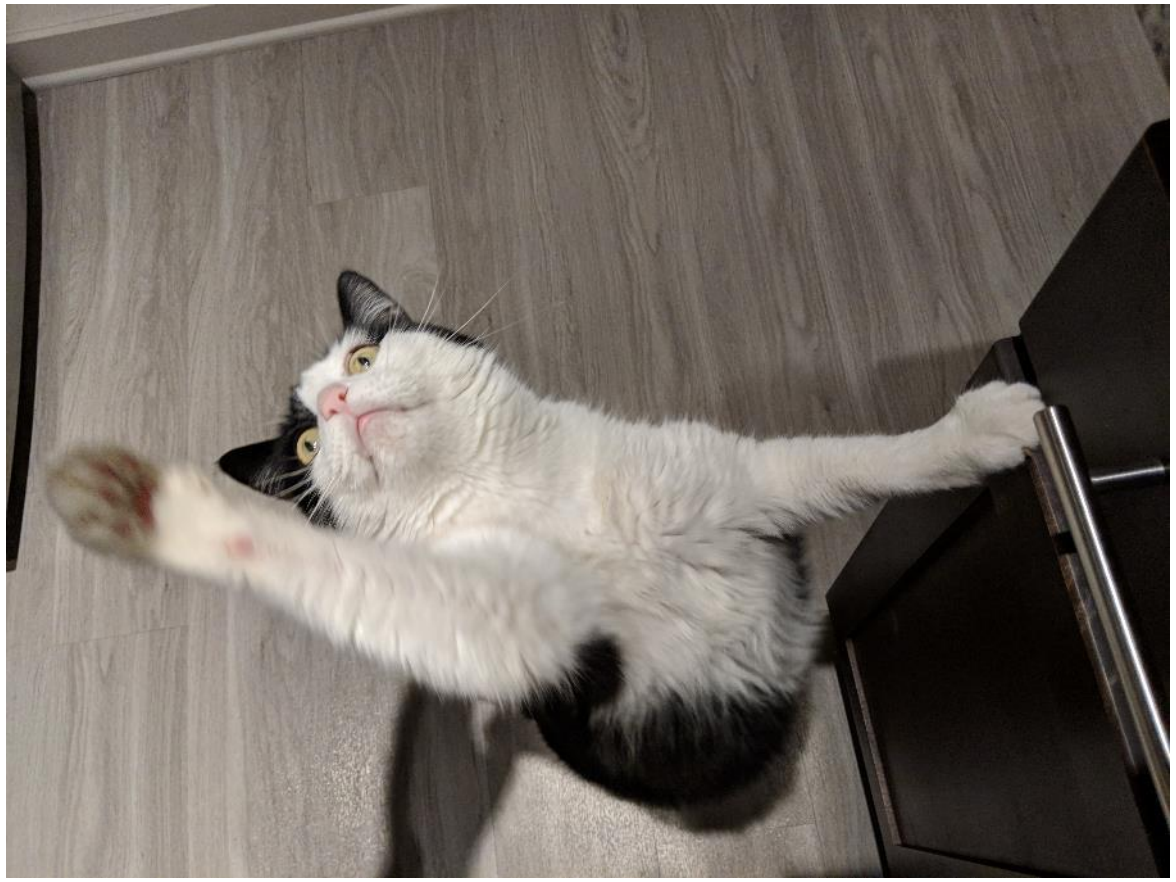

Rabo exceeded his quota on /scratch:

/scratch quota: **50 GB**
/scratch limit: **200 GB**
/scratch grace period: **14 days**

He has **14 days** to get /scratch below **50 GB**, and he can't go above **200 GB**, or he won't be able to add more data.

GPFS STORAGE QUOTA

QUOTA: When exceeded the user receives a warning message.
Usage has to return below the quota within the **GRACE PERIOD**.

LIMIT: Cannot be exceeded.
Automatically set to the actual quota usage when grace period expires.

	Data size		*Number of files*		GRACE PERIOD
	QUOTA	LIMIT	QUOTA	LIMIT	
/home	15 GB	20 GB	200,000	300,000	7 days
/scratch	50 GB	200 GB	200,000	1,000,000	14 days
/data	Can be purchased at 1 TB increments.				

Note that groups may also purchase additional scratch quota in 1 TB increments.

GPFS STORAGE QUOTA

QUOTA: When exceeded the user receives a warning message.
Usage has to return below the quota within the **GRACE PERIOD**.

LIMIT: Cannot be exceeded.
Automatically set to the actual quota usage when grace period expires.

Quota usage on GPFS is accounted in different ways.

/home	→	**USER:**	User ownership
/scratch	→	**USER/GROUP:**	User or group ownership
/data	→	**FILESET:**	Group's data directory content

GPFS STORAGE QUOTA

How can I check my current quota usage?

`accre_storage`

- Shows the current usage for all quotas associated with the user.

```
                          Space                                  Files
                 Usage    Quota    Limit          Usage    Quota    Limit
Home (user):     12.41G   15G      20G            120304   200000   300000
Scratch (user):  36.23G   50G      200G           180276   200000   1000000
```

GPFS STORAGE QUOTA

How can I check my current quota usage?

`accre_storage`

- Shows the current usage for all quotas associated with the user.

```
                        Space                       Files
              Usage     Quota    Limit       Usage     Quota    Limit

Home (user):      20.01G    15G      20G         120304    200000   300000
Scratch (user):   93.45G    50G      200G        180276    200000   1000000

Scratch (group):
             accre    7.85T     9T       10T         287562    0        0

Data (fileset):
             accre    6.11G     2T       3T          4538712 0        0
```

FAIRSHARE/BURSTING LIMITS

showLimits -g *group*

- Show the cluster resources limits for a specific *group*.

```
[vanzod@vmps09 ~]$ showLimits -g capra_lab
ACCOUNT          GROUP   FAIRSHARE    MAXCPUS   MAXMEM(GB)  MAXCPUTIME(HRS)
-------------------------------------------------------------------------------
capra_lab_account            16         272       2720           26112
                capra_lab     1           -          -               -
-------------------------------------------------------------------------------
```

Users in the same group share the same amount of resources.

Cluster Utilization

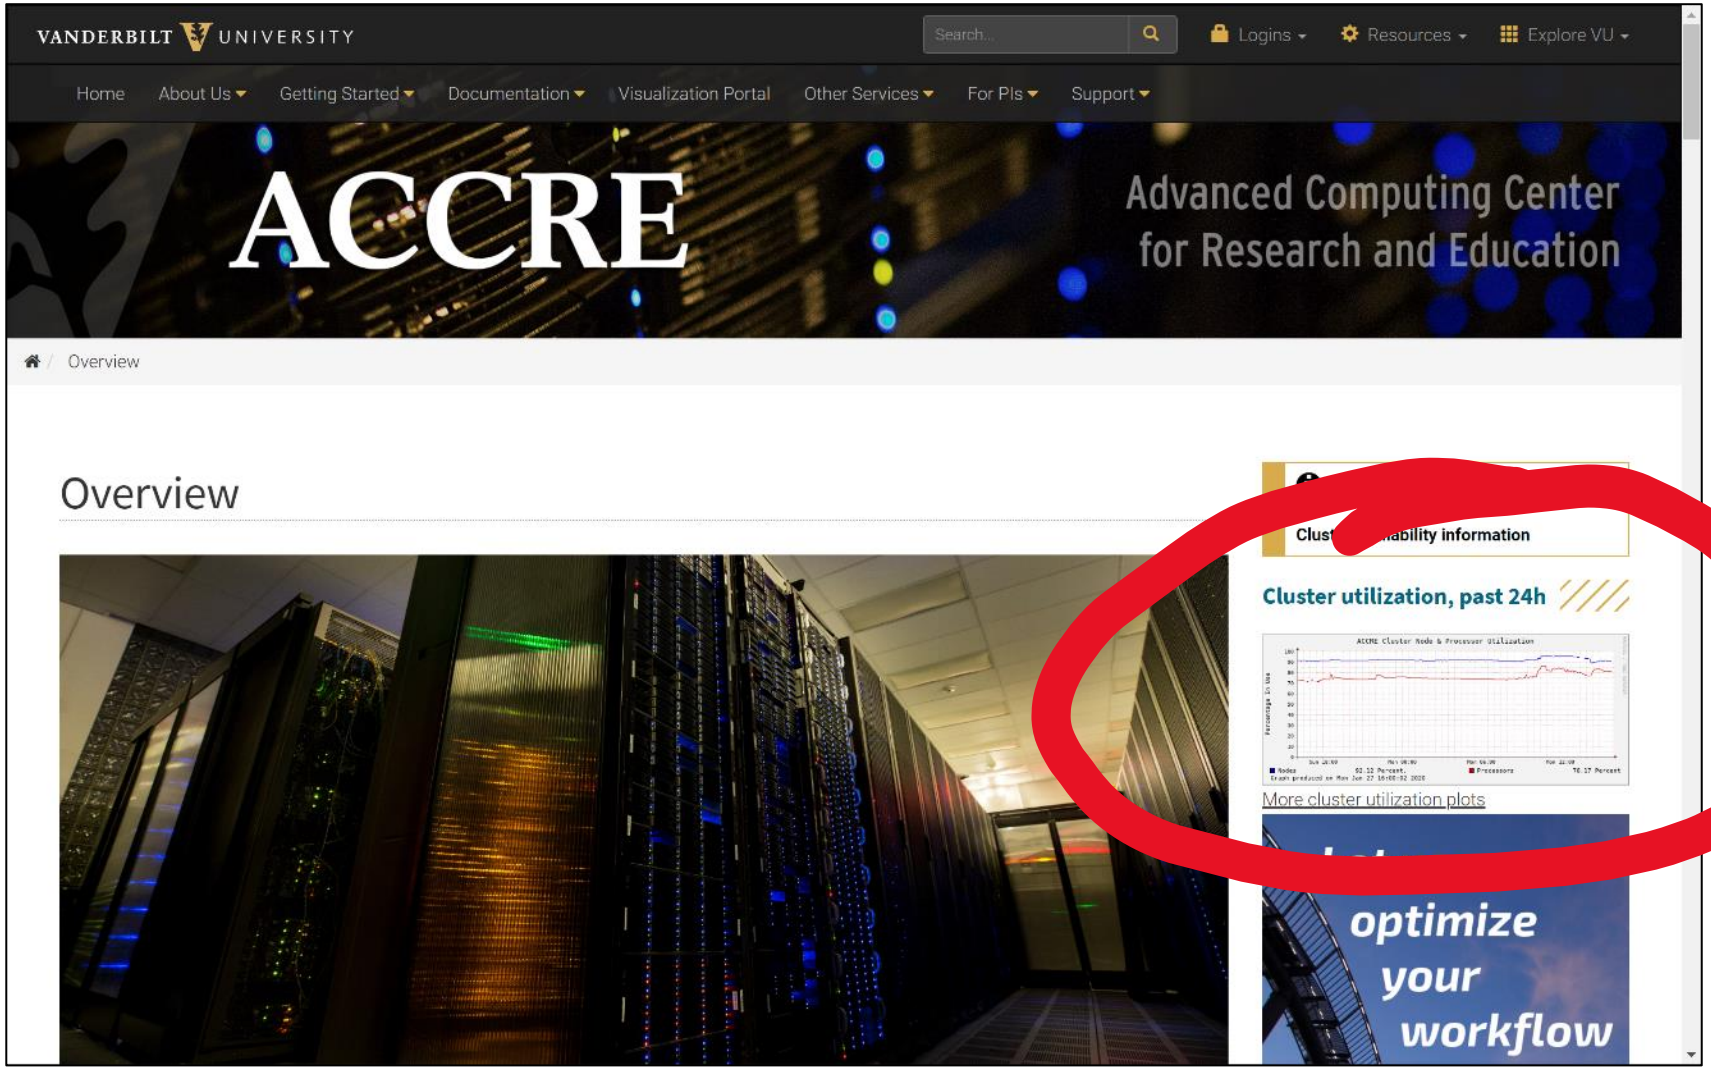

CLUSTER UTILIZATION

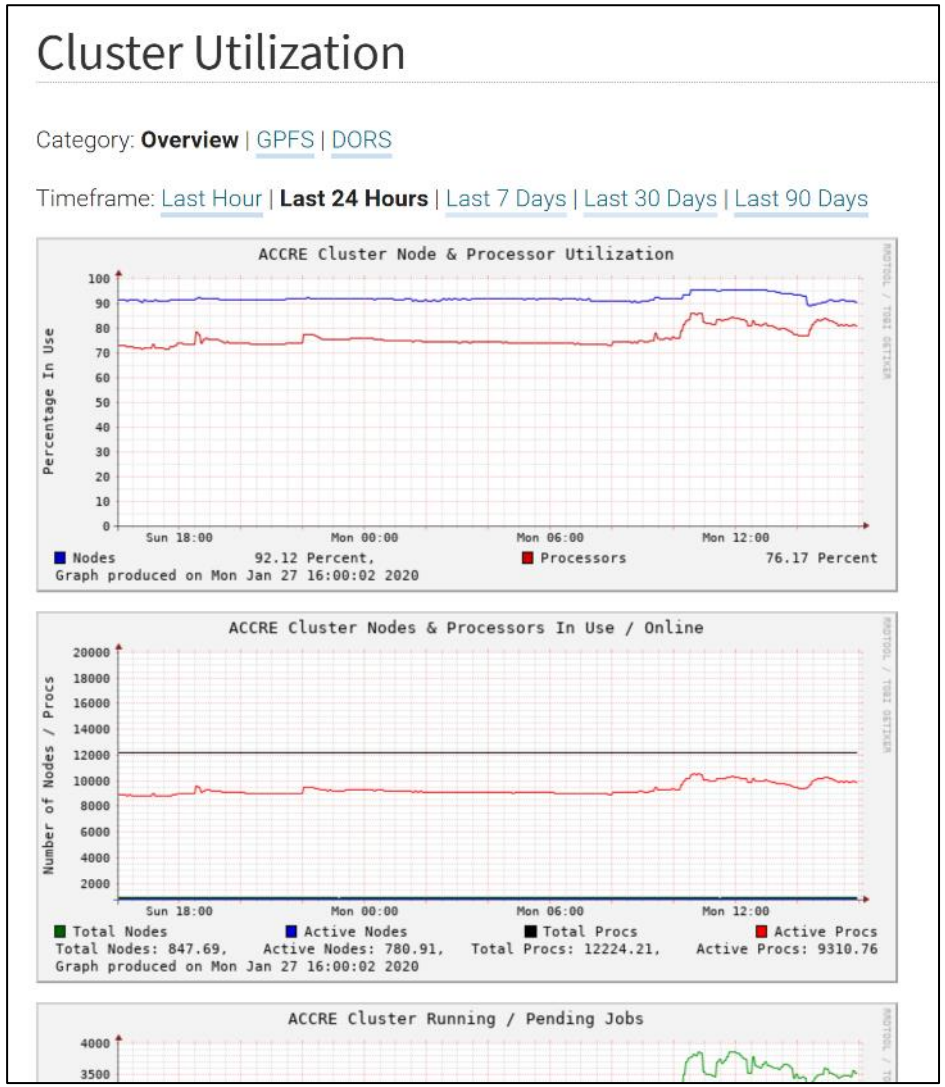

Provides live status on cluster usage and file system usage

Need more help?

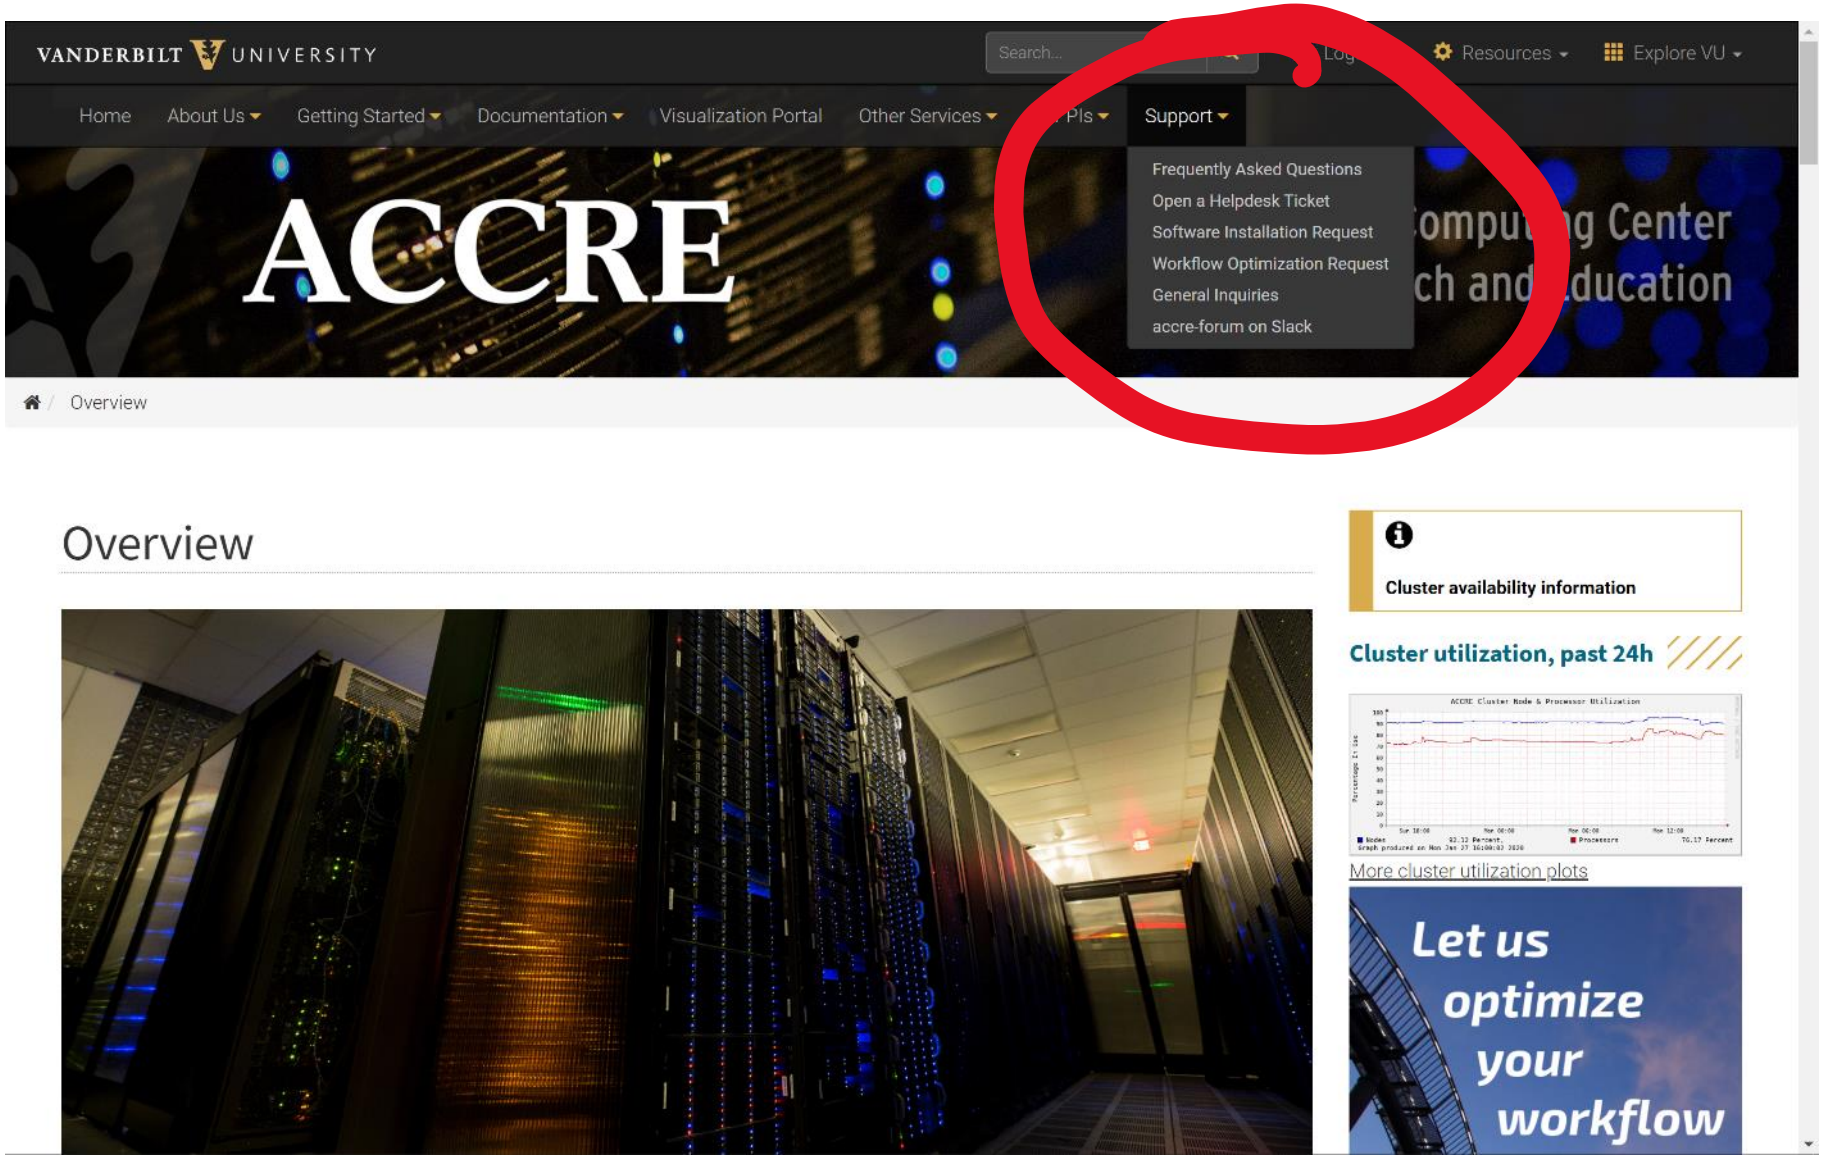

Need more help?

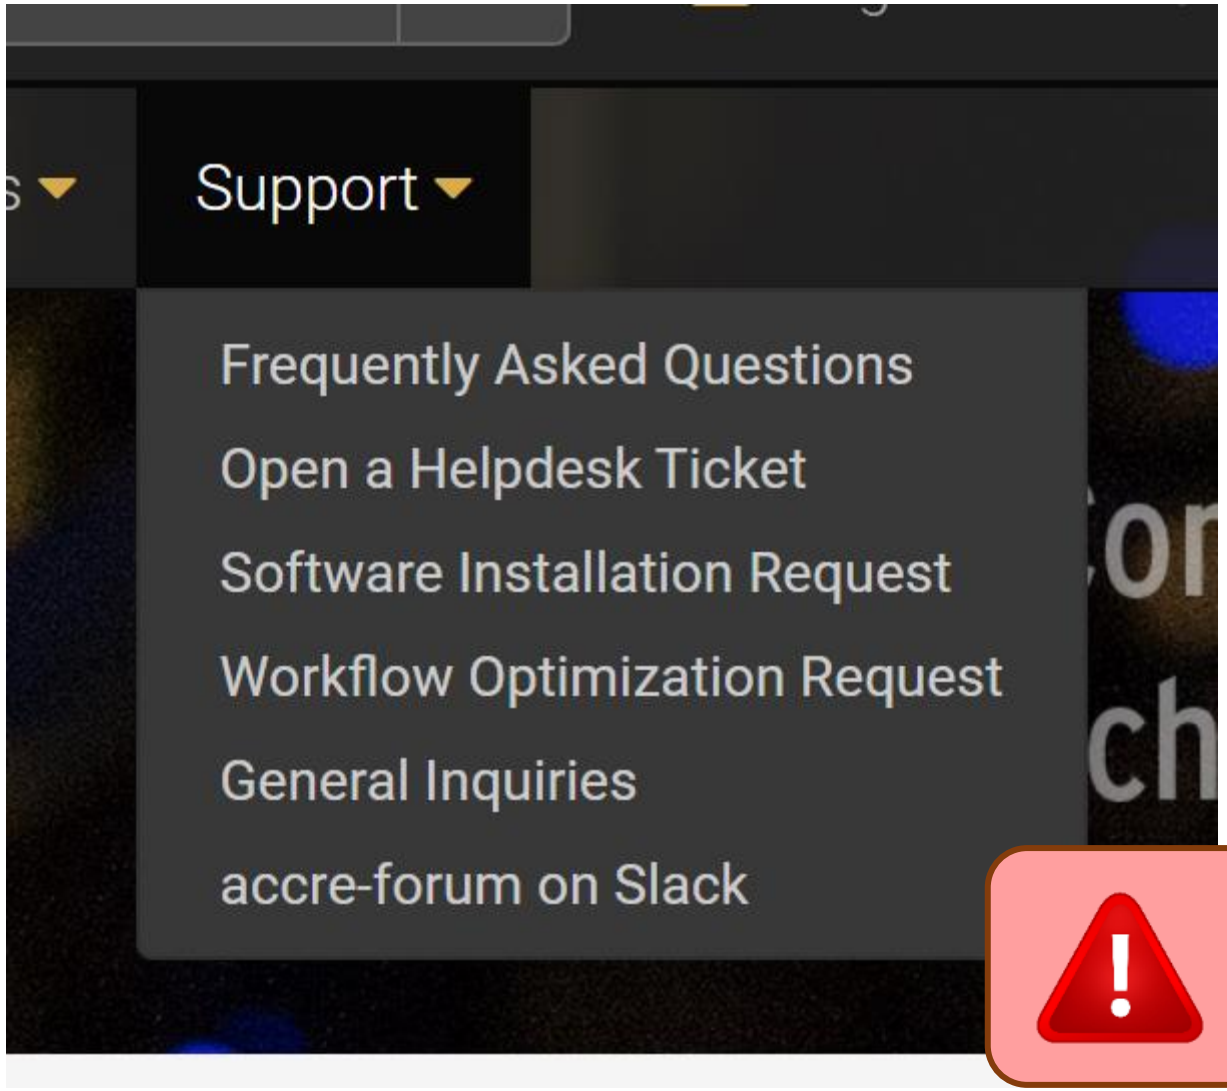

Check the FAQs before submitting a ticket!

If you need software installed, use the **Software Installation Request** form

Rush tickets wake up the ACCRE on-call staff member. Only open a rush ticket if it affects all ACCRE users!

For the next class...

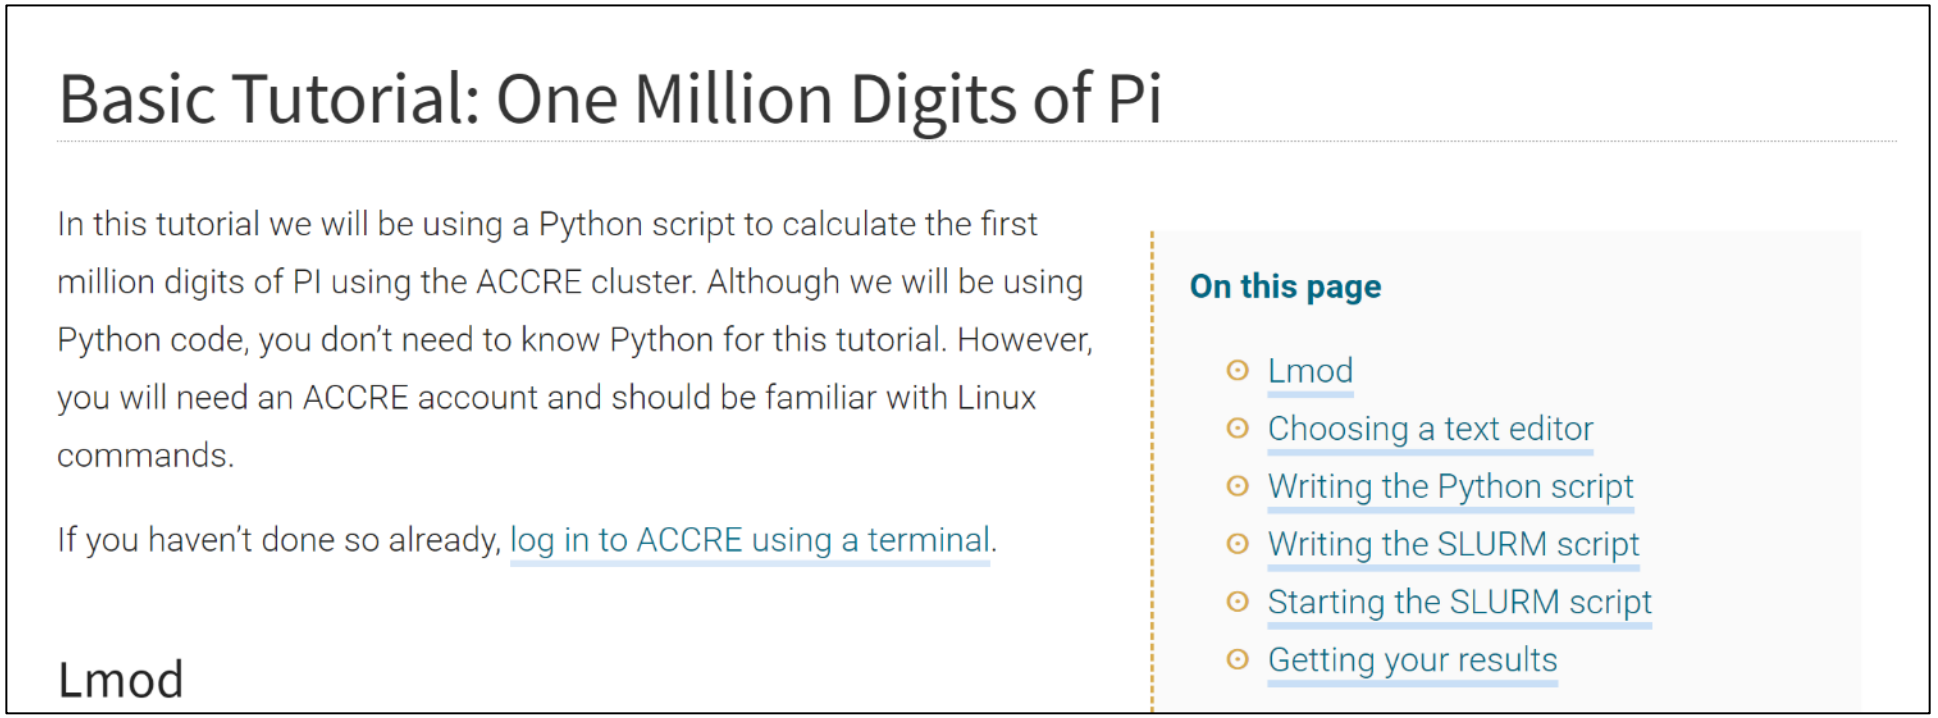

Basic Tutorial: One Million Digits of Pi

In this tutorial we will be using a Python script to calculate the first million digits of PI using the ACCRE cluster. Although we will be using Python code, you don't need to know Python for this tutorial. However, you will need an ACCRE account and should be familiar with Linux commands.

If you haven't done so already, log in to ACCRE using a terminal.

Lmod

On this page

- Lmod
- Choosing a text editor
- Writing the Python script
- Writing the SLURM script
- Starting the SLURM script
- Getting your results

Complete the One Million Digits of Pi tutorial
vanderbilt.edu/accre/getting-started/tutorial

This isn't required but it will give you a sense of how Lmod and SLURM work!

It's cluster tour time!

ACCRE
Advanced Computing Center for Research & Education