XSEDE

Extreme Science and Engineering Discovery Environment

New User Tutorial

Tom Maiden

PSC

tmaiden@psc.edu

XSEDE New User Tutorial

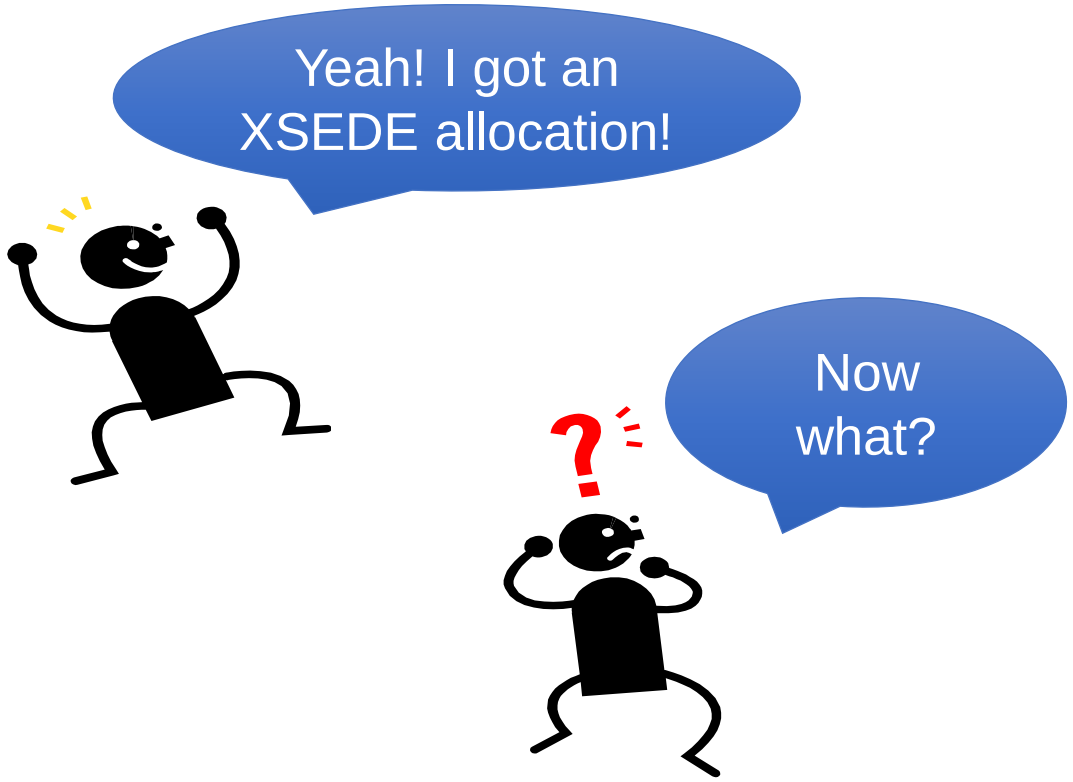

Today's session is an overview of XSEDE for new users. We will assume you have an allocation and will show you how to start using it.

The aims are to help you:

- Get started using XSEDE resources, and
- Ask for help if you run into difficulties.

This lecture is on-line at: portal.xsede.org/online-training

Code of Conduct

XSEDE has an external code of conduct which represents our commitment to providing an inclusive and harassment-free environment in all interactions regardless of race, age, ethnicity, national origin, language, gender, gender identity, sexual orientation, disability, physical appearance, political views, military service, health status, or religion. The code of conduct extends to all XSEDE-sponsored events, services, and interactions.

Code of Conduct: https://www.xsede.org/codeofconduct

Contact:

- Event organizer: Tom Maiden, Pittsburgh Supercomputing Center (tmaiden@psc.edu)
- XSEDE ombudspersons:
  - Linda Akli, Southeastern Universities Research Association (akli@sura.org)
  - Lizanne Destefano, Georgia Tech (lizanne.destefano@ceismc.gatech.edu)
  - Ken Hackworth, Pittsburgh Supercomputing Center (hackworth@psc.edu)
  - Bryan Snead, Texas Advanced Computing Center (jbsnead@tacc.utexas.edu)
- Anonymous reporting form available at https://www.xsede.org/codeofconduct.

Terminology Statement

In line with XSEDE's Code of Conduct, XSEDE is committed to providing training events that foster inclusion and show respect for all. This commitment applies not only to how we interact during the event; it also applies to the training materials and presentation. It is not XSEDE's position to use, condone, or promote offensive terminology.

XSEDE instructors strive to keep inclusive language at the forefront. In the event that we have included inappropriate materials, verbal or written, please let us know at terminology@xsede.org

While XSEDE has no control over external third-party documentation, we are taking steps to effect change by contacting the relevant organizations; we hope this will be addressed by all third parties soon.

If you see any terminology concerns in the following presentation or slides, we want to know!
Please contact the Terminology Task Force: terminology@xsede.org

Outline

1. What is:
   - XSEDE
   - User Portal. Getting Help. Training
   - XSEDE Resources
2. Using the Computational Resources:
   - Accessing and Running
   - Managing Your Software Environment (Modules)
   - File Systems and Transfers
3. What to do when your allocation is over

What is XSEDE?

- XSEDE 2.0 is a 5-year program funded by the NSF Office of Advanced Infrastructure (OAC). Although it is the follow-on to the original XSEDE award, we generally just call it XSEDE rather than XSEDE 2.0.
- **XSEDE funds people** (not computers) **who coordinate NSF's high-end computing, visualization, storage and networking resources**.
  - Community Engagement and Enrichment (CEE)
  - Extended Collaborative Support Services (ECSS)
  - XSEDE Cyberinfrastructure Integration (XCI)
  - Operations (OPS)
  - Resource Allocations Services (RAS)
  - Program Office (PO)
- The compute resources (Bridges, Bridges-2, Comet, Expanse, Stampede2, Jetstream, etc.) are all funded by separate NSF grants

How can XSEDE help advance your project?

Some of the resources coordinated by XSEDE include:

- Massively parallel clusters
- GPUs, large memory, shared nodes
- Cloud Computing
- Parallel file systems
- Fast networking
- Very large databases
- Efficient data movement tools
- Consulting services provided by experts (expertise on optimization, parallelization, science gateway, scientific database, visualization, and more!)

XSEDE Home Page: xsede.org

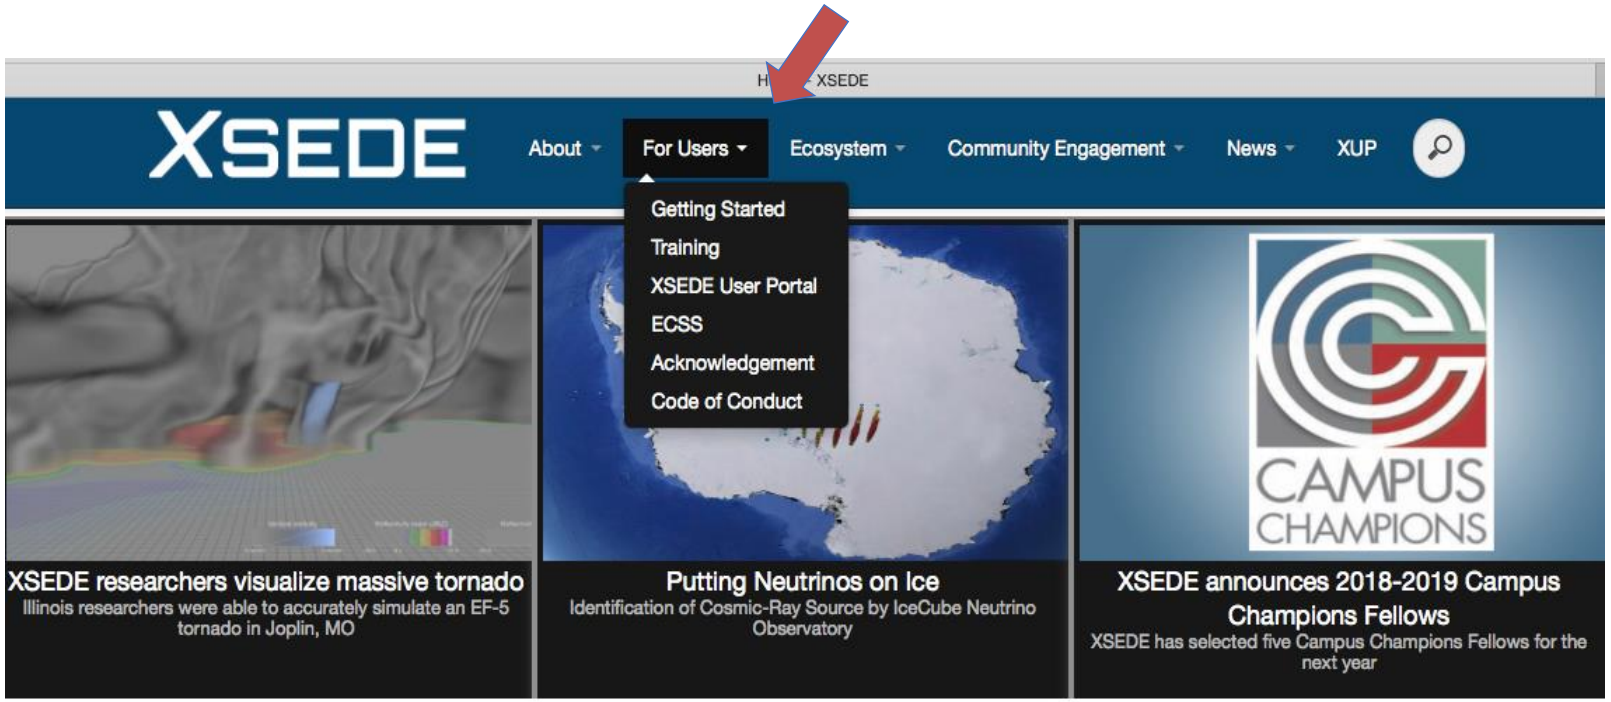

XSEDE is a single virtual system that scientists can use to interactively share computing resources, data and expertise. People around the world use these resources and services — things like supercomputers, collections of data and new tools — to improve our planet.

- Getting Started Guide
- Science Success Stories
- Diversity and Inclusion Programs
- XSEDE's Project Documents

For Users tab:

- Getting Started
- Acknowledgement (to include in publications)

XSEDE User Portal portal.xsede.org

Web-based interface that:

- Provides information about your accounts and allocations. From the Portal you can:
  - Monitor your allocation
  - Submit a renewal request for your allocation
- Access:
  - Science Gateways
  - Documentation, User Guides, Training
  - Help Desk
- We will discuss in detail during the first part of this lecture

More info: Using the XSEDE User Portal (CI-Tutor)

Mobile User Portal mobile.xsede.org

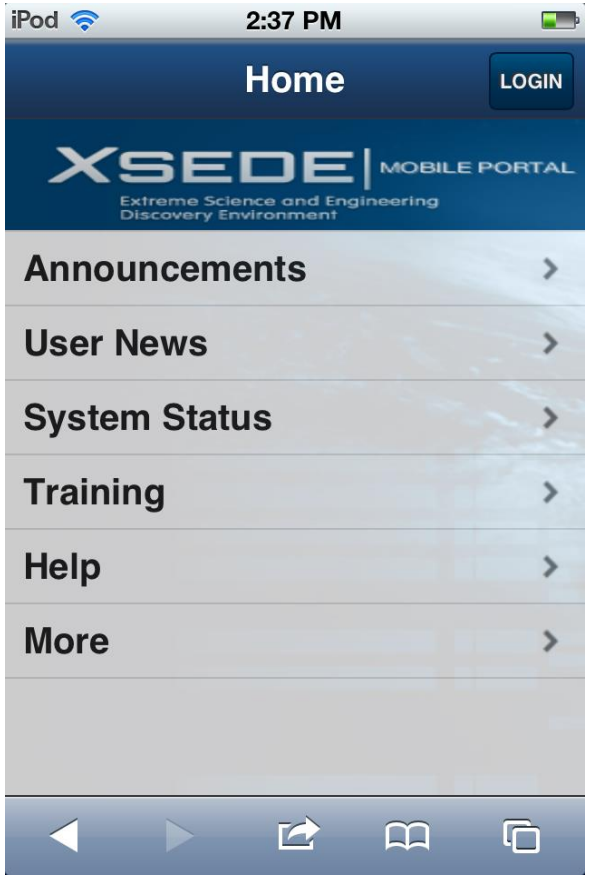

Secure Password

- Easy for you to remember but difficult for others to guess:
  - Example of complex password easy to remember:
    Use the first letter in a phrase:
  - Mswbof7702 = “My son was born on Friday July 7th 2002”
- Avoid words than can be easily guessed:
  - Found in dictionary, names of your children, etc.
- Don’t share your password with others.
- Use different passwords for each account.

More info: XSEDE Cybersecurity (CI-Tutor)

First portal login each allocation term:

You will you will be asked to accept the usage policy on the User Responsibilities form:

- Explains acceptable use to protect shared resources and intellectual property. You are responsible for your account and for protecting your passwords.
- Responsible use includes:
  - Choose a strong password.
    Not sharing passwords, Not writing passwords down where they can be easily found. Not using tools which expose passwords on the network (such as telnet). Closing SSH terminals and logging out of the User Portal when done, etc.
  - Do not share accounts.
    One account per user. Can add members of your group to your grant through the Portal.
  - Report Suspicious Activity:
    If you have any suspicion that your account or personal computer has been compromised send email to help@xsede.org or call 24/7: 1-866-907-2383 immediately.
  - Examples of Suspicious Activity:
    Finding files that you did not create. Jobs have run from your allocation without your knowledge.

MY XSEDE tab on portal.xsede.org

- ***My XSEDE*** tab is about you:
  - your allocation,
  - accounts,
  - publications,
  - usernames,
  - jobs running

- **Profile**
  - Update your email address
  - Subscribe to mailing lists

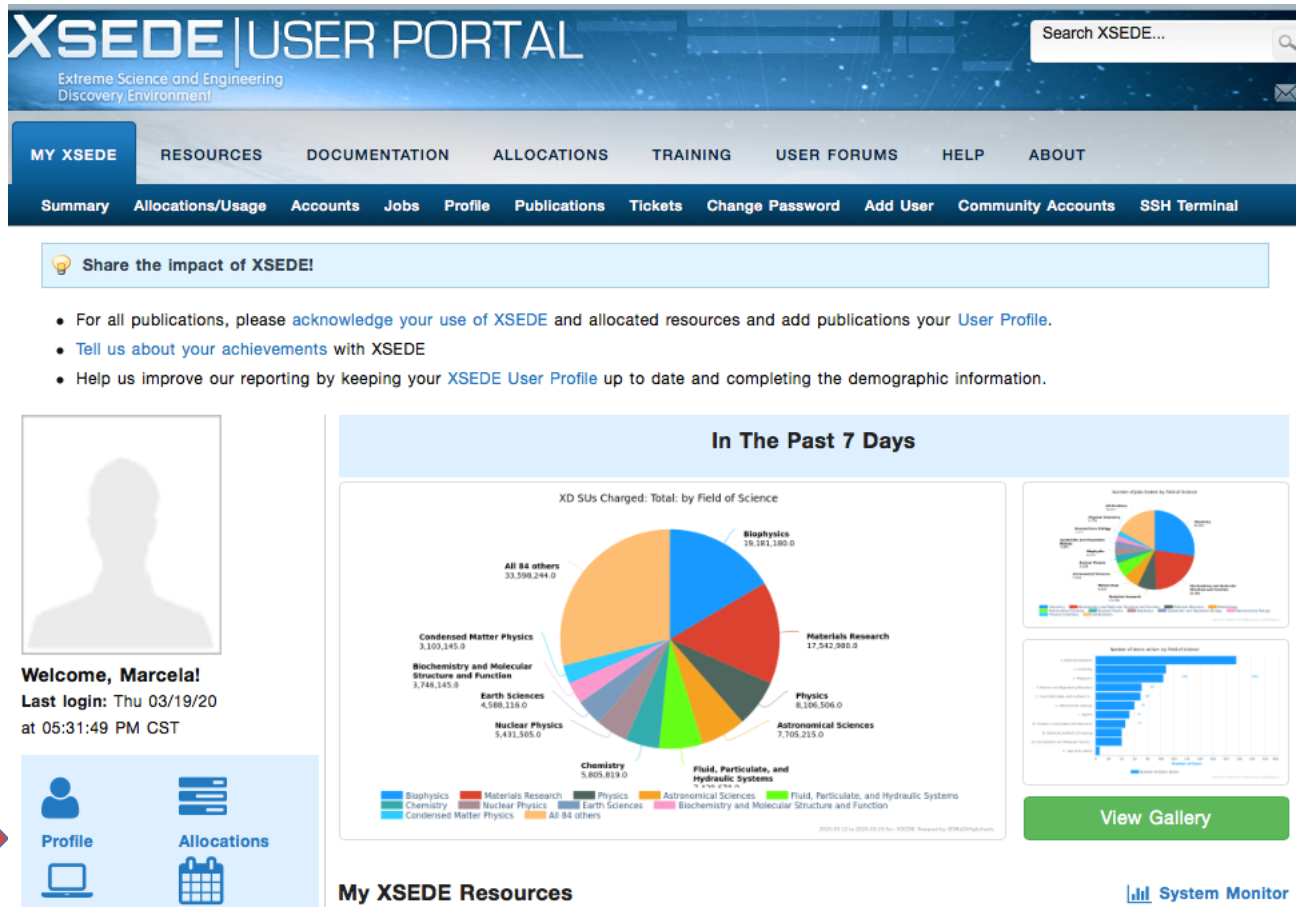

My XSEDE

(scroll down the page)

Please add here:

- Your Scientific progress. It does not need to be finished or published. Bullet or Paragraph.
- Your publications. Anyone who is on an XSEDE project can add a publication and associate it with their project.

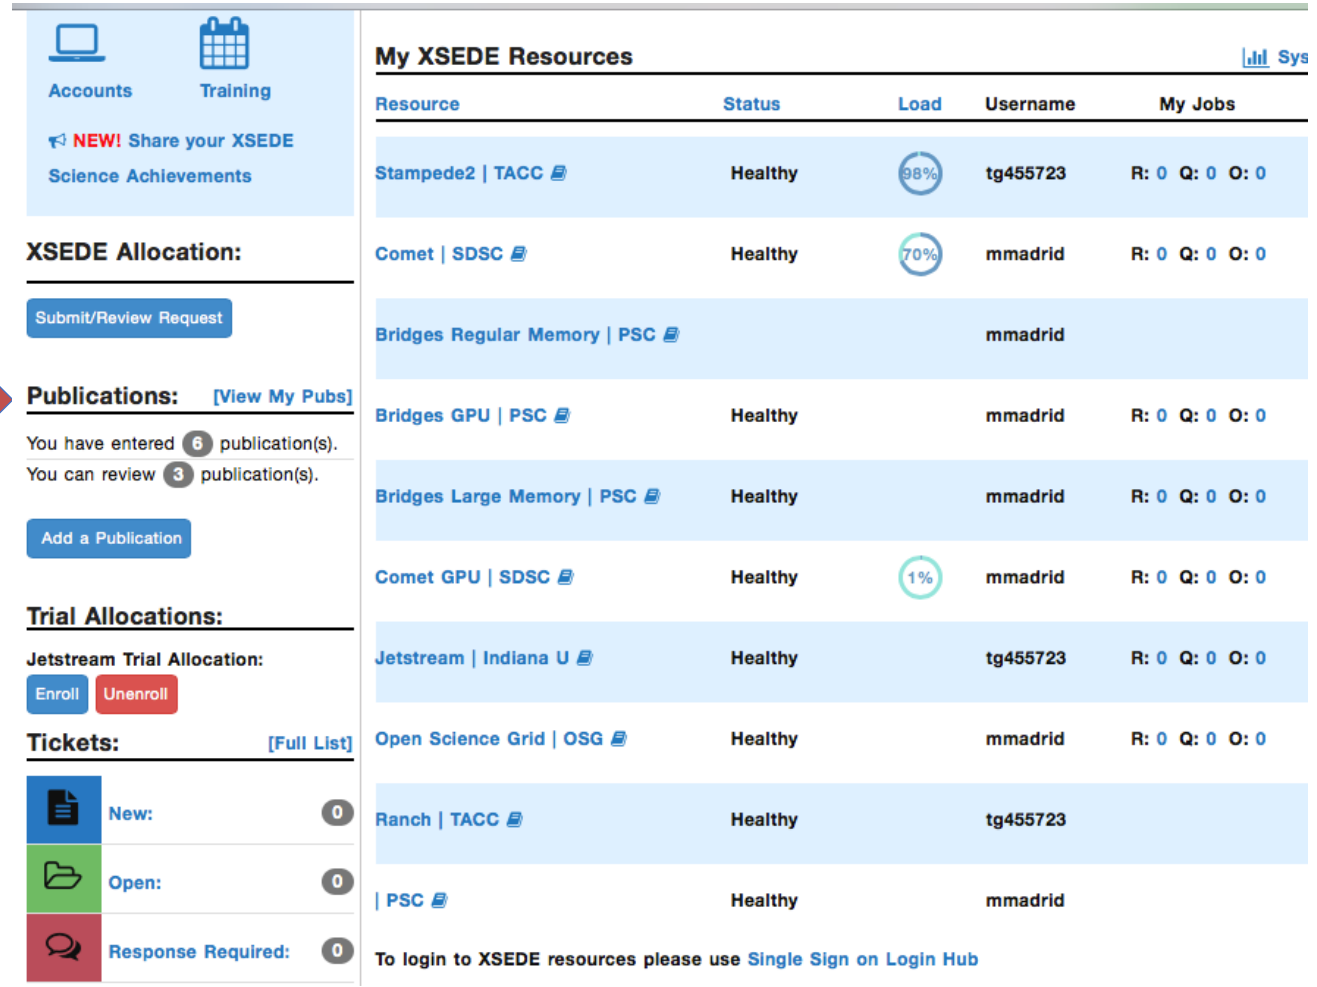

Publications or Presentations to Conferences!

- Acknowledge XSEDE:
  - Cite: Computing in Science & Engineering, vol.16, no. 5, pp. 62-74, Sept.-Oct. 2014, doi:10.1109/MCSE.2014.80, and:
  - Include the acknowledgement: This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1548562
  - and a sentence including the resources used along with the your allocation ID.
- If you have received Extended Collaborative Support Services (ECSS):
  - Cite: High Performance Computer Applications 6th International Conference, ISUM 2015, Mexico, March 9-13, 2015, Revised Selected Papers Gitler, Isidoro, Klapp, Jaime (Eds.) Springer International Publishing. ISBN 978-3-319-32243-8, 3-13, 2016. 10.1007/978-3-319-32243-8.
  - Include the acknowledgement: "We thank consultant name(s) for their assistance with (describe tasks such as porting, optimization, visualization, etc.), which was made possible through the XSEDE Extended Collaborative Support Service (ECSS) program."
- Add your publications to the Portal
- See xsede.org → users tab for complete proper citations

Reporting and Tracking Issues

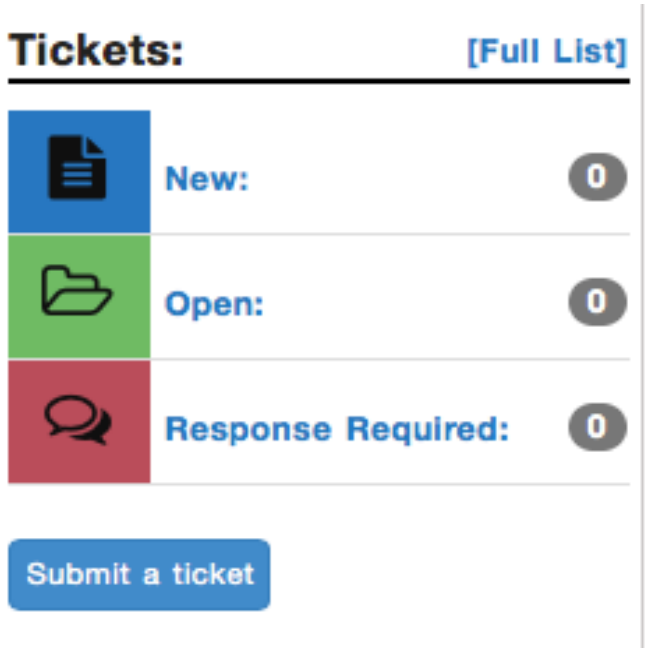

We are here to help you…

- Preferred:
  - portal.xsede.org → My XSEDE → Tickets →Submit ticket
  - Helps ensure all relevant details are provided for quick routing
- Can also:
  - email help@xsede.org or
  - Call 1-866-907-2383, at any hour (24/7)

User Portal: User Forums

Discussing your problems…

The User Forums are a great place to ask questions, get help, or discuss ideas about XSEDE.

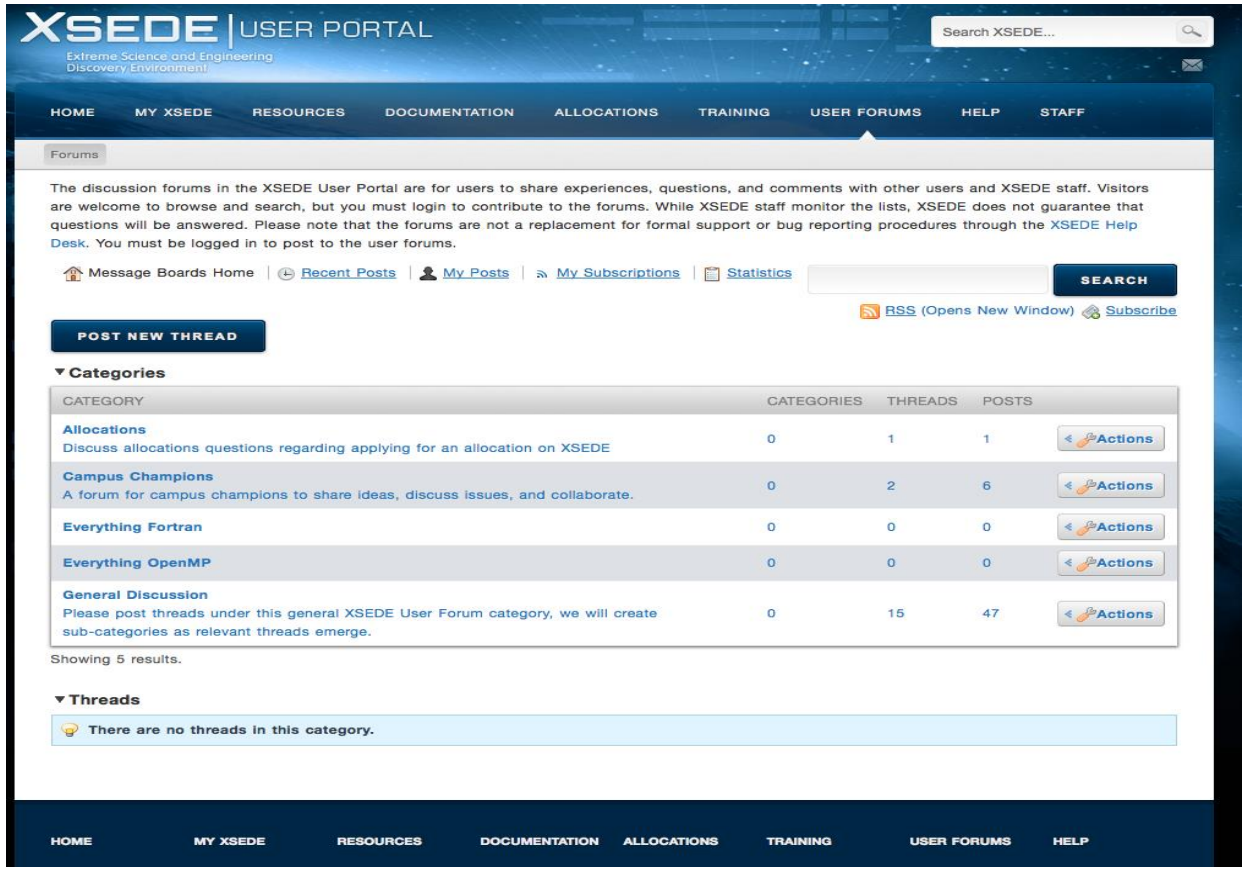

My XSEDE

You will see the tabs:

- Allocations/Usage
- Accounts

An *allocation* is associated with a Project.

When the allocation is active, it will show in the Allocations/Usage tab.

An *account* is associated with a person.

Multiple accounts can charge to a project (usually members of the research group).

An account can charge to multiple projects (if you have more than one research project).

Adding a user to your allocation

- The PI or the co-PI can designate an Allocation Manager

- The Allocation Manager can:
  - Add other users (add accounts to the allocation).
  - Submit renewal

Training

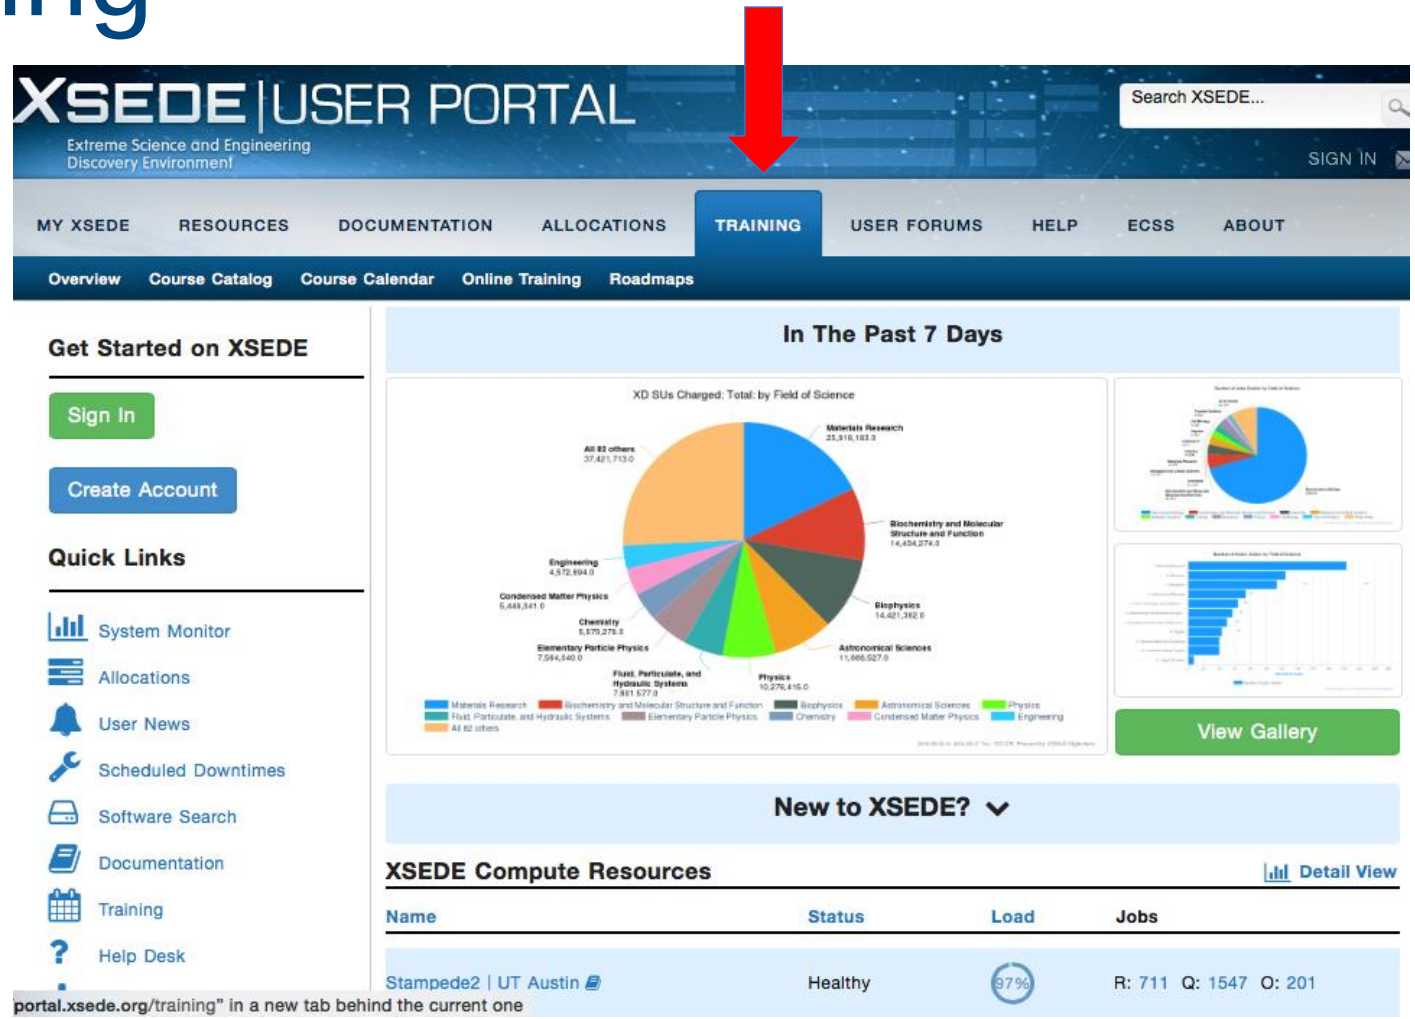
XSEDE | USER PORTAL
Extreme Science and Engineering Discovery Environment
Search XSEDE...
SIGN IN
MY XSEDE
RESOURCES
DOCUMENTATION
ALLOCATIONS
TRAINING
USER FORUMS
HELP
ECSS
ABOUT
Overview
Course Catalog
Course Calendar
Online Training
Roadmaps
Get Started on XSEDE
Sign In
Create Account
Quick Links
System Monitor
Allocations
User News
Scheduled Downtimes
Software Search
Documentation
Training
Help Desk
In The Past 7 Days
XD SUs Charged: Total: by Field of Science
View Gallery
New to XSEDE?
XSEDE Compute Resources
Detail View
Name
Status
Load
Jobs
Stampede2 | UT Austin
Healthy
97%
R: 711 Q: 1547 O: 201
portal.xsede.org/training" in a new tab behind the current one

XSEDE Training

Training is focused on systems and software supported by XSEDE, covering programming principles and techniques for using resources and services.

Available on a variety of topics:

- Using XSEDE resources
- high performance computing
- visualization
- data management
- distributed and grid computing
- science gateways
- and more

Available in a variety of formats:

- Multicast
- Webinars
- Online training
- In person workshops

XSEDE Training: Events Calendar

Synchronous events (multicast, webcast, in-person) are listed here:

https://portal.xsede.org/course-calendar

Use your XSEDE username to register

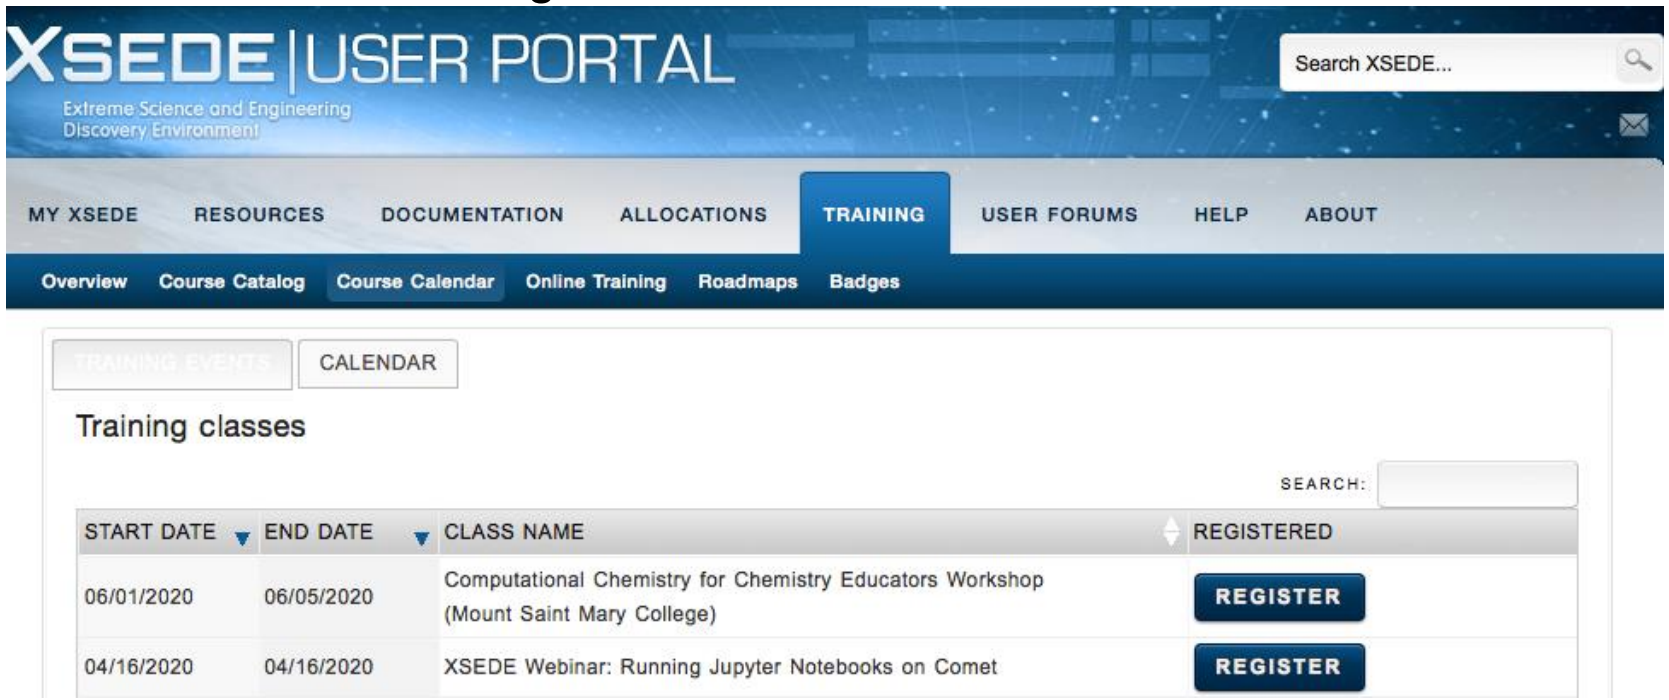

XSEDE Training: Online

~70 standalone, online, asynchronous tutorials, e.g.

- XSEDE: Customizing Your Computing Environment (CI-Tutor)
- Python for High Performance (Cornell Virtual Workshop)
- Vectorization (Cornell Virtual Workshop)
- Introduction to Multi-core Performance (CI-Tutor)

https://portal.xsede.org/online-training

XSEDE Training: Monthly Workshop Series

- Rotating (Nuts and Bolts) Topics
  - MPI, OpenMP, OpenACC, Big Data, Summer Boot Camp
- Up to 25 satellite sites per session (zoom only during covid-19)
  - Sites are spread geographically and include MSIs and National Labs
- Register VIA XSEDE Portal:
  - https://portal.xsede.org/course-calendar
- To Become a Satellite Site or Questions Contact:
  - Tom Maiden – tmaiden@psc.edu

XSEDE Broadening Participation Program

Expand awareness

- Campus Visits
- Conference Exhibiting
- Papers
- News

Identify programs and researchers who can benefit from advanced computing services

- Conference Exhibiting
- Campus Visits
- Training Events
- Consulting

Enable institutions and faculty to use advanced computing to increase their research productivity

- Build and Maintain a Thriving Peer Support Community
- **Deliver training mapped to needs**
- Connect researchers with XSEDE services and expertise

Create scalable and sustainable models and best practices

- Enhance curriculum
- Foster productive campus champions
- Create connections to the CI Ecosystem

To learn more about the program or request a visit email Linda Akli
akli@sura.org

XSEDE

XSEDE Education Program

- Outreach to faculty to **help add computational science and HPC materials to existing courses**

- Support for changing curriculum – add new courses or programs
  - Campus visits to promote department engagement and administration buy-in
  - Database of computational science programs
  - Resources on Computational Science Education
  - Competencies for Computational Science, Data Science, and Computational Physics programs
  - See these resources and more at HPCUniversity.org

- Collaborative Online Courses – Offer a foundational computational science course at your own institution
  - Contact Kate Cahill kcahill@osc.edu for more information

- Summer workshops for faculty – teaching computational science & curriculum development
  - check XSEDE training calendar for new events

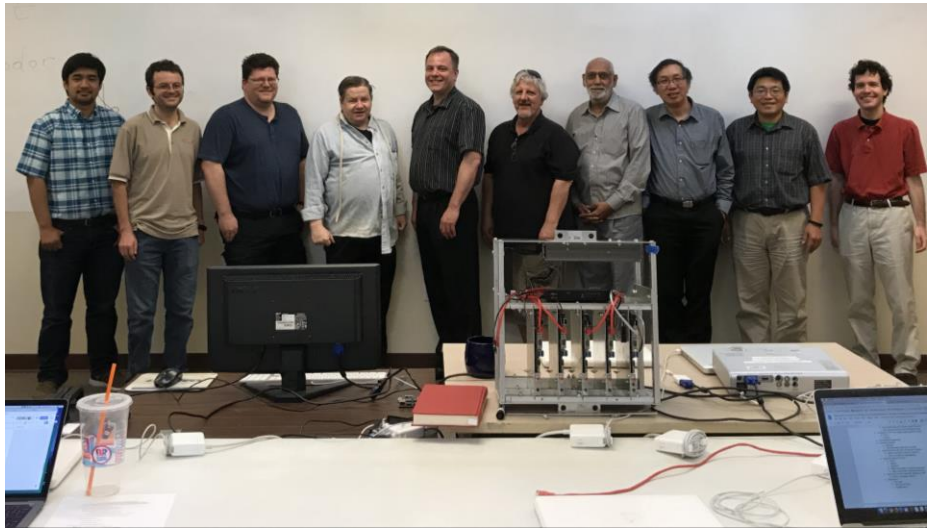

Outline

1. What is:
   - XSEDE
   - User Portal. Getting Help. Training
   - XSEDE Resources
     - Human (Staff support in ECSS, NIP, etc)
     - Compute, Storage and High Throughput (HTC)
     - Software
     - Science Gateways
2. Using the Computational Resources:
   - Accessing and Running
   - Managing Your Software Environment (Modules)
   - File Systems and Transfers
3. What to do when your allocation is over

Resources: Human Component (ECSS)

Extended Collaborative Support Services (ECSS)

- Collaborations between XSEDE researchers and staff
- Expertise is available in a wide range of areas:
  - Performance analysis and optimization
  - Software parallelization and scalability improvements
  - Gateway and web portal development
  - Specialized scientific software
  - Visualization
  - Workflows
- Can solicit Support at any time through the Allocations tab at the XSEDE User Portal

More info: Extended Collaborative Support Services (ECSS)

Who qualifies for ECSS?

- **All** XSEDE PIs can qualify for ECSS

- It is NOT limited to only the most advanced users or projects with the largest allocations

- Users who are new to XSEDE, in the early stages of their computational work or making the transition from small-scale to large-scale computing often benefit the most from ECSS.

ECSS areas

ECSS consists of five areas. As an end user, you won't need to be aware of the areas and our staff will figure out where you fit best:

- Extended Support for Research Teams (ESRT): Traditional projects to improve software and make best use of XSEDE resources
- Extended Support for Community Codes (ESCC): Emphasis on codes available for public use
- Extended Support for Science Gateways (ESSGW): Development of web interfaces to XSEDE resources
- Novel and Innovative Projects (NIP): Assistance to users from domains that are relatively new to XSEDE and high performance computing
- Extended Support for Training, Education and Outreach (ESTEO): Technical support for use of advanced cyberinfrastructure

Requesting ECSS

Can solicit Extended Collaborative Support Services at any time,

through the Allocations tab at the XSEDE User Portal

You will be asked to answer five questions:

1. What do you want to accomplish with the help of expert staff? Have you already done any work on this aspect of your software?
2. How would the success of this collaboration benefit your project?
3. Which member(s) of your team would collaborate with ECSS staff?
4. Have you had significant interaction on previous projects related to your current proposal or discussed your extended support needs with any XSEDE staff? If so, please indicate with whom.
5. Have you received XSEDE advanced support in the past? If so, please indicate the time period, and how the support you received then relates to the support you request now.

ECSS projects

Research Teams

Project Name	PI	PI Institution	ECSS Consultant(s)	Allocation End Date
An implicit, Chimera-based discontinuous Galerkin solver: development and application	Paul David Orkwis	University of Cincinnati	Davide Del Vento, Shiquan Su	9/30/18
Statistical Analysis for Partially-Observed Markov Processes with Marked Point Process Obs, Y4	Yong Zeng	University of Missouri, Kansas City	Mitchel DeWayne Horton	6/30/18
Genome-Wide microRNAs and Single Gamete Based Genetic Profiling of Sweet Sorghum Varieties for Biofuel Production	Ahmad Naseer Aziz	Tennessee State University	Alex Ropelewski	5/16/18
DISSCO, a Digital Instrument for Sound Synthesis and Composition, Y2	Sever Tipei	University of Illinois at Urbana-Champaign	Paul Rodriguez	6/13/18
Six Degrees of Francis Bacon, Y2	Christopher Norton Warren	Carnegie Mellon University	David Walling	6/30/18
Allocation Request on Bridges for Joint Analysis of Metagenomics and Metabolomics Data, Y4	Ping Ma	University of Georgia	Paul Rodriguez, Philip Blood	9/30/18
Analysis of human cortical electrophysiological data: database design for rapid analysis	Max Novelli	University of Pittsburgh	Roberto Gomez	7/16/18
Assessment of Competition in the US Markets Based on Retail Scanner Data	Philip Garland Gayle	Kansas State University	Kwai Wong, Od Odbadrakh	9/4/18
Simulation for 2D Semiconductor with Parallel Uniform and Adaptive Multigrid Method for Multi-component Phase Field Crystal Models, Y2	Zhen Guan	University of California, Irvine	David Bock, Dmitry Pekurovsky, Sudhakar Pamidighantam	9/30/18
Turbulent Mixing in a Magnetic Field and Flow strcuture under Successive Axisymmetric Straining, Y2	Pui-kuen Yeung	Georgia Institute of Technology	Lars Koesterke	9/30/18
Modeling Heliospheric Phenomena with MS-FLUKSS and Observational Boundary Conditions	Nikolai Pogorelov	University of Alabama, Huntsville	Laura Carrington	6/30/18
The "Morelli Machine": A Proposal Testing a Critical, Algorithmic Approach to Art History	Christopher James Nygren	University of Pittsburgh	Alan Craig, Paul Rodriguez	10/8/18
New Frontiers of Direct Laser Acceleration in Megatesla Magnetic Fields	Alex Arefiev	University of California, San Diego	Amit Chourasia, Shiquan Su	2/6/19
Method Development and Application of Electronic Structure Calculations for Complex Nanostructures	Kaushik Dayal	Carnegie Mellon University	Yang Wang	3/31/19

List of ECSS projects (current and completed) can be found at:

https://www.xsede.org/web/site/for-users/ecss/ecss-projects

Novel and Innovative Projects

NIP reaches out and supports areas or communities that have not traditionally used high performance computing:

- Economics: finance trading data
- Census data
- Text mining
- Social sciences
- Humanities
- Art
- And more

Campus Champions

Campus Champions are campus representatives, affiliated with your college or university, that can provide face-to-face guidance. Their role is to help researchers at their institutions with research computing.

One of the easiest ways to get started is to contact the local campus champion at your institution:

https://www.xsede.org/community-engagement/campus-champions

She/he will:

- Provide information about high-performance computing resources, including resources available from XSEDE
- Assist you to quickly get access to allocations of high-performance computing resources
- Facilitate awareness and training workshops about the use of high-performance computing resources and services
- Provide contacts within the high-performance computing community for quick problem resolution

More info: Campus Champions

How do I become a Campus Champion?

- Write to champion-info@xsede.org and ask to get involved

- Follow us on Twitter at @XSEDECC

More info: Campus Champions

Resources

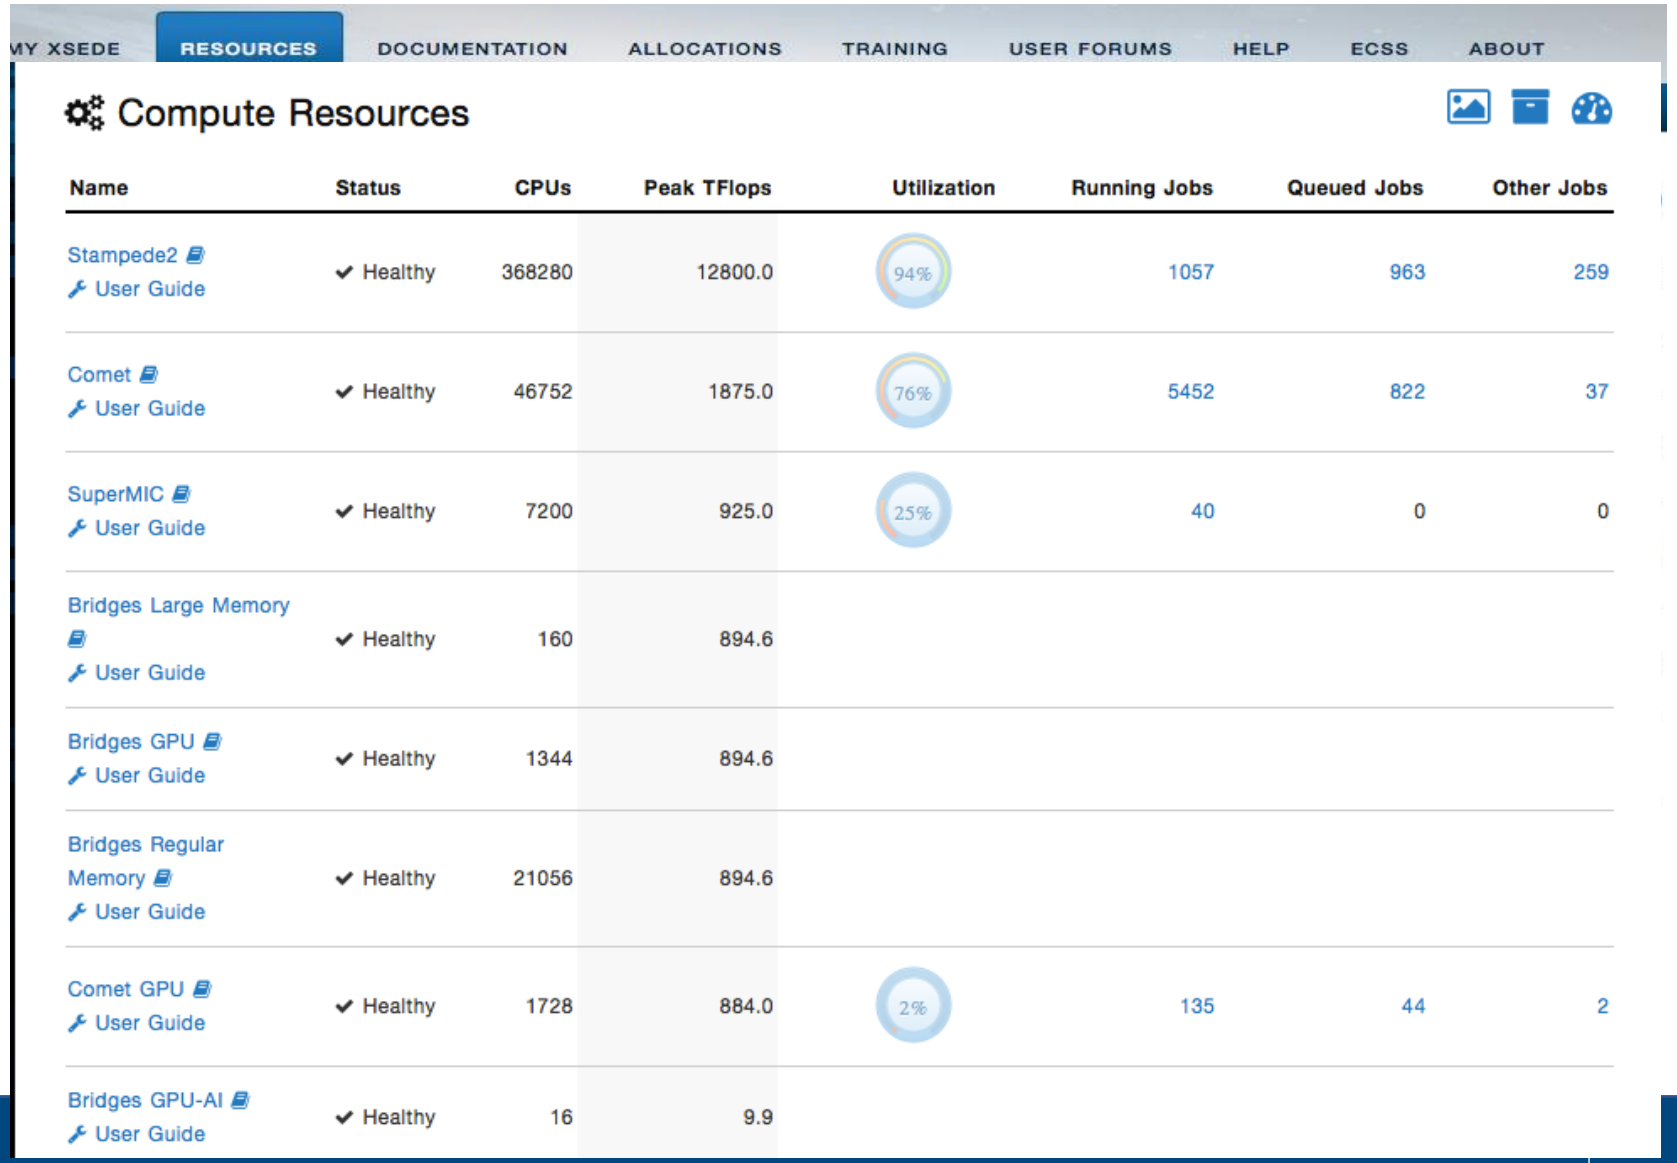

Name	Status	CPUs	Peak TFlops	Utilization	Running Jobs	Queued Jobs	Other Jobs
Stampede2 User Guide	Healthy	368280	12800.0	94%	1057	963	259
Comet User Guide	Healthy	46752	1875.0	76%	5452	822	37
SuperMIC User Guide	Healthy	7200	925.0	25%	40	0	0
Bridges Large Memory User Guide	Healthy	160	894.6				
Bridges GPU User Guide	Healthy	1344	894.6				
Bridges Regular Memory User Guide	Healthy	21056	894.6				
Comet GPU User Guide	Healthy	1728	884.0	2%	135	44	2
Bridges GPU-AI User Guide	Healthy	16	9.9				

XSEDE

Scrolling down:

Storage Resources

Name	File Space (TB)
TACC Long-term tape Archival Storage (Ranch) User Guide	61440.0
TACC Long-term Storage (Wrangler Storage) User Guide	10000.0
PSC Storage (Bridges Pylon) User Guide	10000.0
SDSC Medium-term disk storage (Data Oasis) User Guide	4000.0
IU/TACC Storage (Jetstream Storage)	960.0

HTC and Cloud Resources

Name	Status	CPUs	Peak TFlops
Jetstream User Guide	✔ Healthy	15360	516.1
Open Science Grid User Guide	✔ Healthy	60000	50.0

Go to top ↑

Resources: Available Software

- Comprehensive software search feature at the portal: resources→software takes you to the link: https://portal.xsede.org/software/

- Software classified by Resource or by Science Domain

- Questions/requests:
  - Submit a ticket through the Portal or
  - help@xsede.org

Resources: Science Gateways

- Advanced web interfaces that accelerate and democratize research
- Gateways simplify access to computing resources by hiding any infrastructure complexities through the use of a graphical interface
- Many gateways available for use right now in many different fields

Gateways

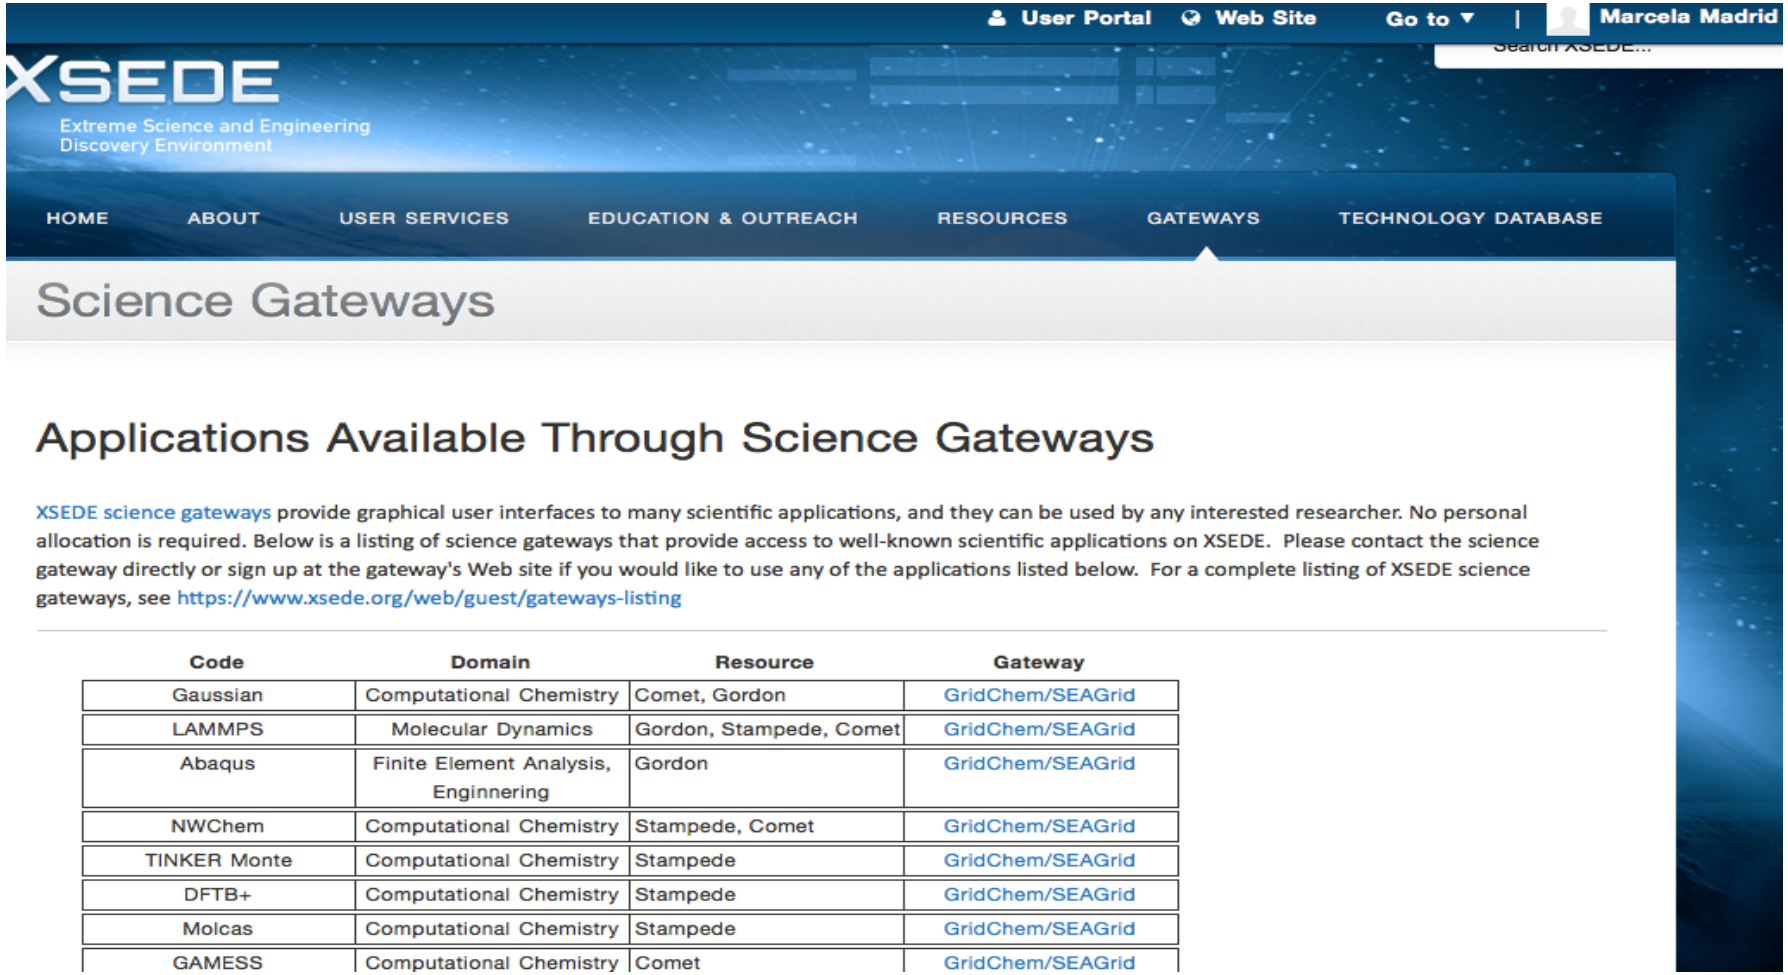

Code	Domain	Resource	Gateway
Gaussian	Computational Chemistry	Comet, Gordon	GridChem/SEAGrid
LAMMPS	Molecular Dynamics	Gordon, Stampede, Comet	GridChem/SEAGrid
Abaqus	Finite Element Analysis, Enginnering	Gordon	GridChem/SEAGrid
NWChem	Computational Chemistry	Stampede, Comet	GridChem/SEAGrid
TINKER Monte	Computational Chemistry	Stampede	GridChem/SEAGrid
DFTB+	Computational Chemistry	Stampede	GridChem/SEAGrid
Molcas	Computational Chemistry	Stampede	GridChem/SEAGrid
GAMESS	Computational Chemistry	Comet	GridChem/SEAGrid

Have an existing gateway or want to build one?

- Help with your Gateway is available through the XSEDE allocation process
- If you have an existing gateway that you would like to integrate into XSEDE: request ECSS* support.
- If you'd like to build a new gateway: request SGCI** support:
- Don't worry if you don't know which to request, we'll sort it out!
- Both provide 25% of an FTE devoted to your project for 6-12 months

*Extended Collaborative Support Service

**Science Gateways Community Institute

More info: How to Create a Science Gateway Application (Cornell Virtual Workshop)

Outline

1. What is:
   - XSEDE
   - User Portal. Getting Help. Training
   - XSEDE Resources

2. Using the Computational Resources:
   - Accessing and Running
     - Gateway or Command line
     - Running: file systems, batch jobs
       - Why has my job not run?
       - Batch schedulers. Improving your job turnaround
   - Managing Your Software Environment (Modules)
   - File Systems and Transfers

3. What to do when your allocation is over

Computing on XSEDE is easy: choose a ramp:

- As we saw, one of the easiest way to compute on XSEDE is through a Gateway:
- Scientific Gateways (Certain specific scientific domains)
- OnDemand (Bridges). A graphical interface that makes it easier to run without knowledge of Unix. Not limited to a particular scientific application

- We will next learn how to login and submit jobs through the command line.
- Login to a Compute Resource:
  - Interact command
  - Batch mode

Login to XSEDE Computational Resources

- Open up a terminal on your laptop or workstation. Windows users may download one of many ssh clients, free ones include "putty".
- Connect to the Hub:

  ssh login.xsede.org

  Enter your **Portal username and password**

- Multi-Factor Authentication (MFA): (you will be sent a push to Duo on your phone). Only the very first time: Need to install Duo on your phone (if not installed already) and enroll your XSEDE account.
  Step-by-step instructions:
  Portal → Documentation → MFA tab https://portal.xsede.org/mfa

- After this, you can connect to those machines on which you have an allocation, using the command:

  gsissh machine-name

- No need to enter another passwd. Single Sign On.

Accessing Resources: Direct Access

- The Single Sign On method just described is the default. However, if you choose direct access:

- Please request the direct access password for the machine that you need (by submitting a ticket).

  - Then use a secure shell (ssh) client:
  - From Linux or Mac terminal window:
    - ssh –l <username> <machinename>
  - From Windows: putty

Working on a Remote Computational Resource (simplified)

Login nodes for installing software, compiling programs and requesting access to compute nodes

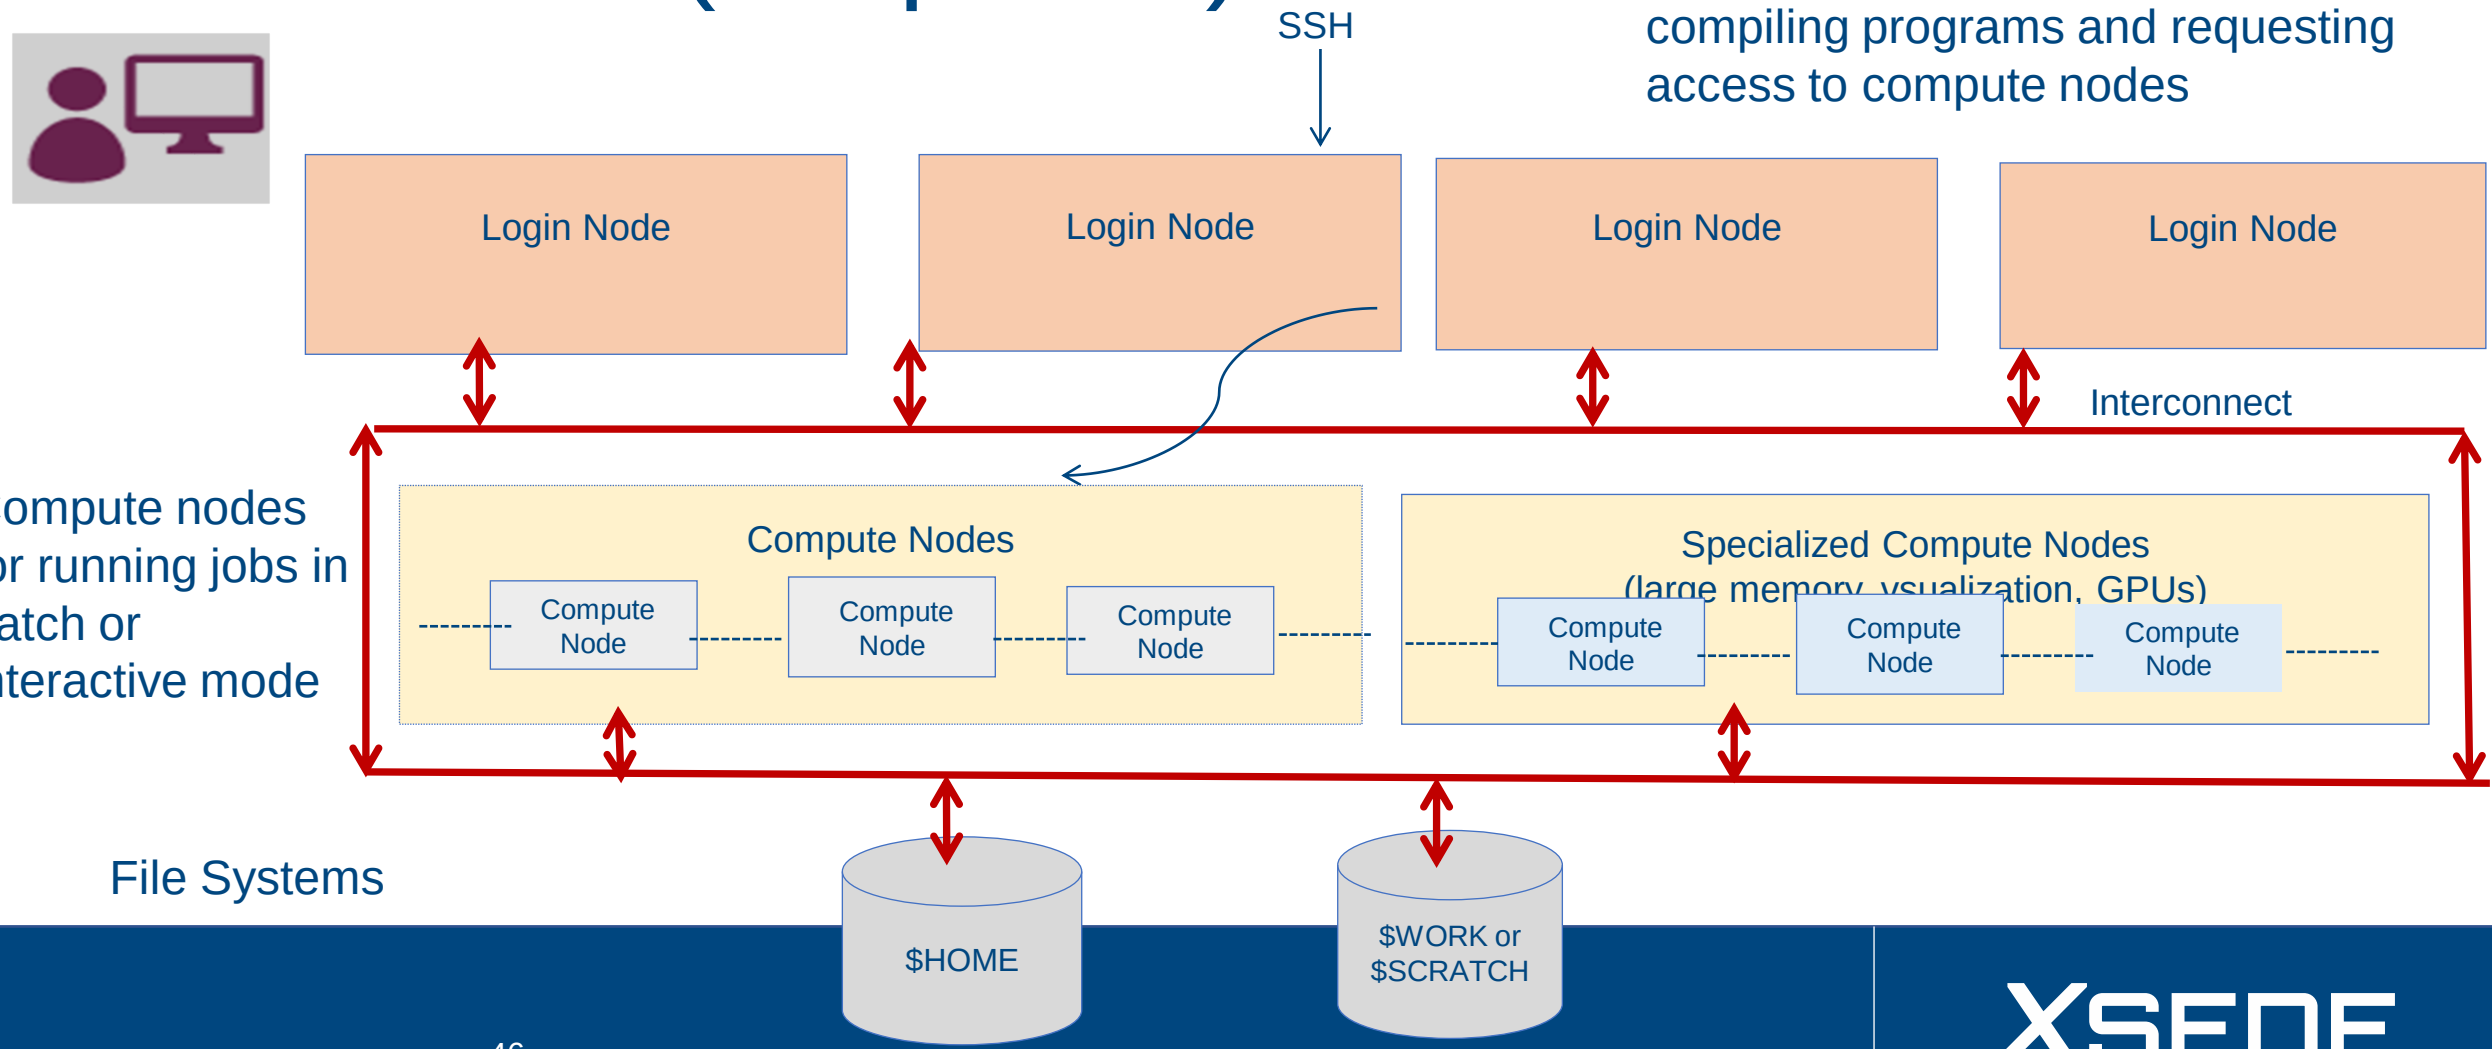

XSEDE

Appropriate use of login nodes

- When you connect to a compute resource, you will be on a login node. A small number of nodes shared by all users.

- Login nodes are intended for non-demanding tasks
  - Compiling code
  - File editing
  - Simple data analysis
  - Small data transfers
  - Submitting and managing jobs

- Do not run computationally intensive jobs on the login nodes! Always use the compute nodes instead. Two ways: Interact or Batch

Running: interact

- You type commands and receive output back to your screen as the commands complete
- Run quickly on a few dedicated nodes. Debug.
- After you login type:
  interact or
  interact –options
  Options: time, memory, partition requested.
- A command prompt will appear when your session begins
- "Ctrl+d" or "exit” to end your session
- Check user guides for availability/policies. Bridges has prioritized the interact command.

Running: batch jobs

- XSEDE compute resources use a batch scheduler, called SLURM, to submit, monitor & cancel jobs
- Configuration details vary from site to site, but basic functionality is consistent:
    - sbatch to submit jobs
    - squeue to view information about jobs
    - scancel to cancel jobs
    - sinfo to view information about nodes and partitions
- See resource User Guide for details/options: Portal→Documentation→User Guides
- General: slurm.schedmd.com

sbatch – basic job script

```
#!/bin/bash
#SBATCH -J “hellompi”
#SBATCH -N 8
#SBATCH --ntasks-per-node=24
#SBATCH -t 1:00:00


ibrun ./hellompi.exe
```

Job specification

Executable statements

Key parameters include:

- Number of nodes (N) and of tasks/node or total number of tasks
- Partition (queue)
- Job duration (how long to run for).
- Job name
- Account

sbatch – selecting an account

- As a new user, you will probably have access to a single account (allocation)
- If you are on multiple allocations (research projects), be sure to explicitly specify the account that you want to charge to – the default won't necessarily be what you expect

```
#SBATCH --account=<account>
 -- or --
#SBATCH -A <account>
```

sbatch – submit a batch script to SLURM

- Submit your jobscript:

```
$ sbatch myjobscript
```

- See slurm.schedmd.com/sbatch.html for more details

squeue – monitor jobs

- With squeue, you can monitor the state of jobs that have been submitted to the queues.
- Without any arguments, squeue returns information on the job status for all users. In most cases you will probably want to restrict to yourself using the -u option
- Many options for formatting output and restricting to particular job states, partitions or job ids
- See slurm.schedmd.com/squeue.html for more details

squeue – monitor jobs

```
$ squeue
 JOBID PARTITION     NAME     USER   ST       TIME  NODES NODELIST(REASON)
18912381 gpu-share     bash    rynlm PD       0:00      1 (Resources)
18941470 gpu-share      efe   rbnjko PD       0:00      1 (Priority)
18937286 gpu-share     aout      xyzj PD       0:00      1 (Dependency)
18915882   compute     dask  willc97 PD       0:00      8 (Dependency)
18911406   compute NGBW-JOB   cipres  R 3-17:30:45      2 comet-26-[01-02]
18918197    shared NGBW-JOB   cipres  R 1-06:30:41      1 comet-08-16
```

- For running jobs (state R), squeue lists the nodes being used
- For pending jobs (state P), squeue states why job is not running
- Other job states include Completing (CG), Failed (F) and Cancelled (CA). See slurm.schedmd.com/squeue.html for full list

scancel – cancel jobs

- Slurm allows you to cancel jobs that are running or queued
- Use squeue to find jobid
- Comes in handy if you realize job is not progressing as expected, wrong input files were used, etc.

```
$ scancel jobid [jobid] [jobid]
$ scancel -u username # cancel all my jobs
```

- See slurm.schedmd.com/scancel.html for more details

sinfo – information about partitions

```
$ sinfo -s
PARTITION    AVAIL  TIMELIMIT   NODES(A/I/O/T)  NODELIST
compute*        up 2-00:00:00  639/1207/0/1846  comet-01-[01-54]
debug           up      30:00          1/7/0/8  comet-14-[01-08]
shared          up 2-00:00:00    401/229/0/630  comet-01-[01-54]
gpu             up 2-00:00:00       13/56/3/72  comet-30-[03-20]
gpu-shared      up 2-00:00:00       12/20/2/34  comet-30-[03-11]
large-shared    up 2-00:00:00          2/2/0/4  comet-30-[01-02]
```

* Indicates default queue

-s option used to limit to summary information

See slurm.schedmd.com/sinfo.html for more details

Site specific batch scheduling details

We touched briefly on batch scheduling and the SLURM commands, but many of the details are site specific:

- Partitions: CPU, GPU, shared, regular memory, large memory nodes
- Maximum job sizes and wall times
- Scheduler tuning (e.g. optimized for throughput or for large jobs)
- Ability to run shared jobs (e.g. use less than all cores on node)

See the Stampede2, Comet, Bridges or other user guides on the Portal for more information

XSEDE User Guides

Scheduler tuning: job priority

- Job priority in the batch queues is based on the number of cores and wall clock time requested. Differs by site. Examples:

- FIFO system. (Mechanisms in place to prevent a single user from dominating the batch queue and to prevent idle time on the machine).

- Some sites may be optimized for throughput or for large number of cores. May give priority to jobs that request large number of cores. Smaller core counts can run on other systems, or, they can run effectively as *backfill*.

Queue Structure: improving turnaround time

- Try to be as accurate as possible in estimating the wall time for your jobs.
- Note that the default wall time is normally set to the maximum wall time.
- But don’t underestimate the time since your job will be killed if you exceed the time limit. Any results that have not been checkpointed will be lost. Checkpoint your run (write restart files)
- SLURM uses a strategy called backfill to improve throughput:
- If there is a pending job, the scheduler can “loan out” nodes that are being collected for this job, as long as the duration of the loan is less than the time remaining until all nodes are available
- Pack your jobs if you have many small jobs:
  - Running many small jobs places a great burden on the scheduler.
  - Pack many executions into a single job, which you then submit with a single squeue command.
- Check user guides for examples or submit a ticket for help.

My job was not successful

- Job not accepted by the queue: The job script is requesting incorrect number of cores or memory for the resource.
- Job did not complete:
- Check the log files created: job.e.89890 job.o.89890
- One common problem: job run out of CPU time. If you have checkpointed or written a restart file→ restart
- Check the job script: time and memory requested, directory where you are writing files to.
- Do the input files exist in the directory where you specified?
- Do you have permission to use software?
- Waiting a very long time in the queue…
- If all fails… Ask for help by submitting a ticket

Managing Your Environment: **Modules**

- Allows you to manipulate your software environment.
- It gives you access to executables that have been built by XSEDE staff to achieve optimal performance.
- Uses best choice of compilers, compiler options, libraries, etc.
- Can be used to choose which version of software to run. Example:

module list

Currently Loaded Modulefiles:

1) intel/2013_sp1.2.144 2) mvapich2_ib/2.1 3) gnutools/2.69

module avail

abyss/2.0.2(default) diamond/0.7.12(default) lammps/20170331(default) R/3.4.0(default)

amber/16(default) drive-data/6.2(default) llvm/3.6.2(default)

amber/18

Modules (continued)

module show amber

```
module-whatis     amber
module-whatis     Version: 16
setenv            AMBERHOME /opt/amber/16
prepend-path      PATH /opt/amber//16/bin
prepend-path      LD_LIBRARY_PATH /opt/amber/16/lib
prepend-path      PYTHONPATH /opt/amber/16/lib/python2.7/site-packages
module            load python scipy gnu mvapich2_ib fftw netcdf
```

module load amber

module list

```
1) intel/2013_sp1.2.144   2) mvapich2_ib/2.1      3) gnutools/2.69   4) amber16
```

module unload amber

module load amber/18

http://modules.sourceforge.net

More info: Customizing Your Computing Environment (CI-Tutor)

Outline

1. What is:
   - XSEDE
   - User Portal. Getting Help. Training
   - XSEDE Resources
2. Using the Computational Resources:
   - Accessing and Running
   - Managing Your Software Environment (Modules)
   - File Systems and Transfers
3. What to do when your allocation is over

File Systems on XSEDE Resources

Where your data resides on XSEDE and the appropriate storage is your responsibility. In general, all resources provide:

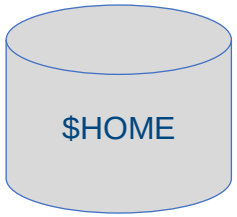

$HOME: Permanent space, but small. A good choice for building software and working file collections of small to medium sized files, where a medium sized file is less than 50 MB.

$WORK: More space, but TEMPORARY. Recommended for running jobs. **Backup your files that are here! They get deleted (purged).**

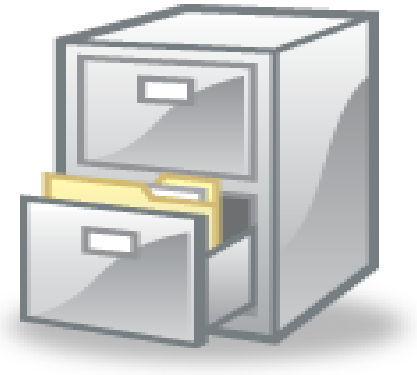

- **Archival (mass) storage:** Must request through the Allocation Process.
- Storage of large amounts of data, is accessible from all sites, slower access time.
- Tar files before storing.

Managing your XSEDE files

How to move files:

- Command line using
  rsync, uberftp, scp, or sftp
- **Globus** web app, command line

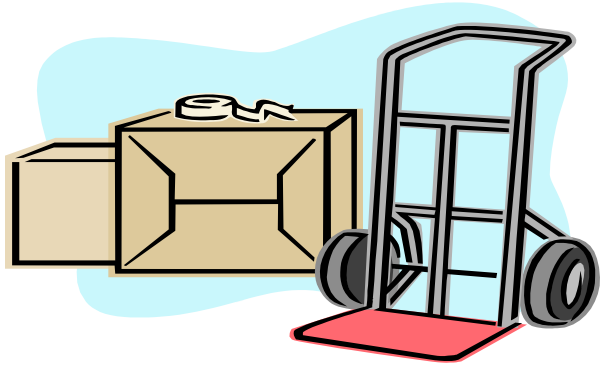

File Transfer

File transfers between XSEDE resources

Or

Between your laptop and an XSEDE resource

USAGE MODE	TRANSFER METHOD	PROS	CONS
GRAPHICAL USER INTERFACE	Globus	easy to use web interface, can use XUP login (SSO), desktop download available	none
COMMAND LINE INTERFACE	Globus Command Line Interface (CLI)	managed, reliable and auto tuned transfer; advanced syntax for scripting; can use XSEDE single sign on	need to setup SSH key in Globus profile; advanced knowledge required for authentication and scripting capabilities
	`globus-url-copy & uberftp`	high performance transfer with tuning options; command line interface	advanced knowledge required for authentication and performance tuning, and increased reliability
	`scp & sftp`	easy command line interface	must user local (resource-specific) username and password

- Portal: Documentation→Manage Data

File Transfers: Small (< 2 GB) Files

- To transfer ***small files*** between XSEDE Resources and/or your own workstation you can:
- *scp or sftp* from the Unix command line. Easy to use. May provide poor performance for very large files.
- *rsync -rltpDv -e 'ssh -l joeuser' source_directory data.bridges.psc.edu:target_directory* Provides fast validation mechanism.
  - From Linux or Mac, you can run these commands directly from the terminal.
  - From Windows, use your ssh client to do this (putty has free downloads for these tools, too! – just Google “putty sftp”).
- These are easy to use and secure, but provide poor performance for large files. May need direct passwd to the resource.
- Check user guide for fast speed interconnectors where available.

Globus

- Easy to use Graphical Interface.
- High performance -- move terabytes of data in many (thousands) or few (1) file(s)
- Designed for researchers -- easy "fire and forget" transfers
- Ability to move files to any machine (even your laptop) with ease
- No client software installation -- new features automatically available
- Beyond transfer -- also offers easy file sharing, publication, search, cloud storage support, and more

To access Globus: Use your existing XSEDE (or other institutional) credentials

For help: Contact support@globus.org

When your allocation is over…

- **If run out of Service Units:** Supplement request: A request for additional resources during an existing allocation's one-year time frame. Supplement awards are highly dependent upon availability of resources.
- Renewal request. Four deadlines per year:

SUBMISSION PERIOD	ALLOCATION BEGIN DATE
Dec 15 thru Jan 15	April 1
Mar 15 thru Apr 15	July 1
Jun 15 thru Jul 15	Oct 1
Sep 15 thru Oct 15	Jan 1

- PI or grant administrator submits request through the Portal.
- Examples of successful requests are on-line at the Allocations tab of the Portal. Need Progress Report. Benchmarks, code performance. Collect these!
- Enter your publications in the Portal.

Summary: How to Get Started on XSEDE?

- Sign in to the Portal→ https://portal.xsede.org/
- https://portal.xsede.org/documentation-overview
- Check the User Guides for Expanse, Stampede2, Bridges-2, etc.
- Get help:
  - Submit a ticket through the Portal
  - Send email help@xsede.org
  - Face-to-face help from your local Campus Champion.
  - Can request Extended Collaborative Support Services (ECSS)
- portal.xsede.org → Training
  - Course Calendar https://portal.xsede.org/course-calendar
  - On-line training https://portal.xsede.org/online-training

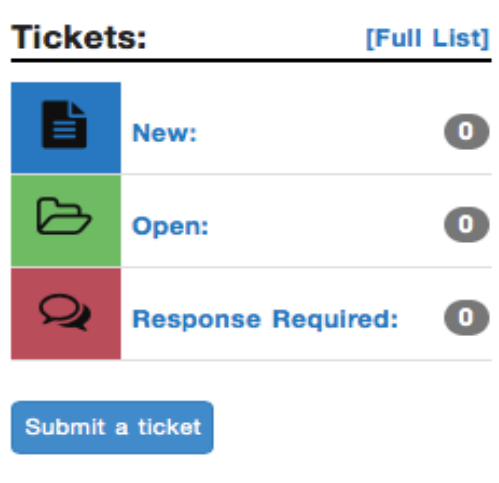

More info: Getting Started (Roadmap)

Thanks for listening and welcome to

XSEDE

Extreme Science and Engineering Discovery Environment