August 9, 2012

# Introduction to the National Science Foundation's High Performance Computing Program

Richard Moore
rlm@sdsc.edu

XSEDE
Extreme Science and Engineering Discovery Environment

NSF

# National Science Foundation

*Our sponsor*

- Created in 1950 "to promote the progress of science; to advance the national health, prosperity, and welfare; to secure the national defense…"
- $7B annual budget
- Funds ~20% of basic research at US colleges and universities

# Diverse areas of fundamental science covered by 7 directorates + 7 ‘offices’

- Biological Sciences
- Computer and Information Science and Engineering
- Education and Human Resources
- Engineering
- Geosciences
- Mathematical and Physical Sciences
- Social, Behavioral and Economic Sciences

- **Office of Cyberinfrastructure**
- Office of the General Counsel
- Office of Integrative Activities
- Office of International Science and Engineering
- Office of Diversity and Inclusion
- Office of Legislative and Public Affairs
- Office of Polar Programs

# Office of Cyberinfrastructure (OCI)

- Funds all supercomputers for non mission-specific academic research
  - Note that many large users of OCI resources also use DOE Office of Science systems under the INCITE program
- ~$200M annual budget
  - Hardware (a lot of the money)
  - People (some of the money)
  - Software (little of the money)
- Not enough to do all the office would like to do
- http://www.nsf.gov/awards/award_visualization.jsp?org=OCI#showAwardDollars=true

# So what is XSEDE?

- A 5-year program that funds **people** who **coordinate** the NSF's high-end computing, visualization, storage, and network resources
  - Coordination & Management Service (CMS) $12M/year
  - Extended Collaborative Support Service (ECSS) $8M/year
  - Training, Education and Outreach Service (TEOS) $3M/year
- The **resources** coordinated by the XSEDE program are funded separately by OCI to Resource Providers
- 'XSEDE Federation' =
  - XSEDE program per se + Resource Providers + XDMoD Audit Program + Additional Partners (e.g. PRACE)

# Huge variety of resources to request

- Leading-edge distributed memory systems
- Very large shared memory systems
- High throughput systems, including now OSG
- Visualization engines
- Accelerators like GPUs
- Long-term and medium-term storage
- Advanced Support – people to help you (ECSS)!

# High-end Computing Resources within the XSEDE Federation

- Kraken @ NICS
  - 1.2 PF Cray XT5
- Ranger @ TACC
  - 580 TF Sun Cluster
- Gordon @ SDSC
  - 341 TF Appro Cluster
- Lonestar 4 @ TACC
  - 302 TF Dell Cluster
- Trestles @ SDSC
  - 100TF Appro Cluster
- Steele @ Purdue
  - 67 TF Dell Cluster
- Blacklight @ PSC
  - 36 TF SGI UV (SMP)
- Keeneland @ GaTech/NICS
  - GPU Cluster

*See Resources tab at www.xsede.org*

And early 2013

- Stampede @ TACC
  - Dell/Intel cluster
  - 2PF CPU +
    8 PF Intel Many-Integrated Core +
    16 1TB + 2 GPU nodes +
    128 GPUs (visualization)
  - http://www.tacc.utexas.edu/stampede

# Special Purpose Compute Resources

- Condor Pool @ Purdue
  - 150 TF, 27k cores
- Quarry @ IU
  - Virtual machines
- FutureGrid (IU + many) (www.futuregrid.org)
  - Experimental/development distributed grid environment
- Open Science Grid (www.opensciencegrid.org)
  - For example, in 24 hours (Mar 14 2012):
    - CPU Hours / day: 2,075,000
    - TB Transferred: 667
  - In the last year:
    - Total CPU Hours: 603,150,000
    - TB Transferred: 269,263

*See Resources tab at www.xsede.org*

# Visualization and Data Resources within the XSEDE Federation

- Visualization & Data Analysis
  - Longhorn @ TACC
    - 20.7 TF Dell/NVIDIA cluster
    - 18.7 TB disk
  - Nautilus @ NICS
    - 8.2 TF SGI/NVIDIA SMP
    - 960 TB disk
  - Spur @ TACC
    - 1.1 TF Sun cluster
    - 1.7 PB disk

- Storage
  - Albedo
    - 1 PB Lustre distributed WAN filesystem
  - Data Supercell @ PSC
    - 4 PB disk archive
  - HPSS @ NICS
    - 6.2 PB tape
  - MSS @ NCSA
    - 10 PB tape
  - Project Storage @ SDSC
    - 400 TB Lustre disk
  - Ranch @ TACC
    - 70 PB tape

*See Resources tab at [www.xsede.org](http://www.xsede.org)*

# How do researchers use XSEDE?

It's **easy** to get started as an XSEDE user:

1. Go to the main web site: www.xsede.org
2. Select 'How Do I Use XSEDE?' under the "User Services" menu

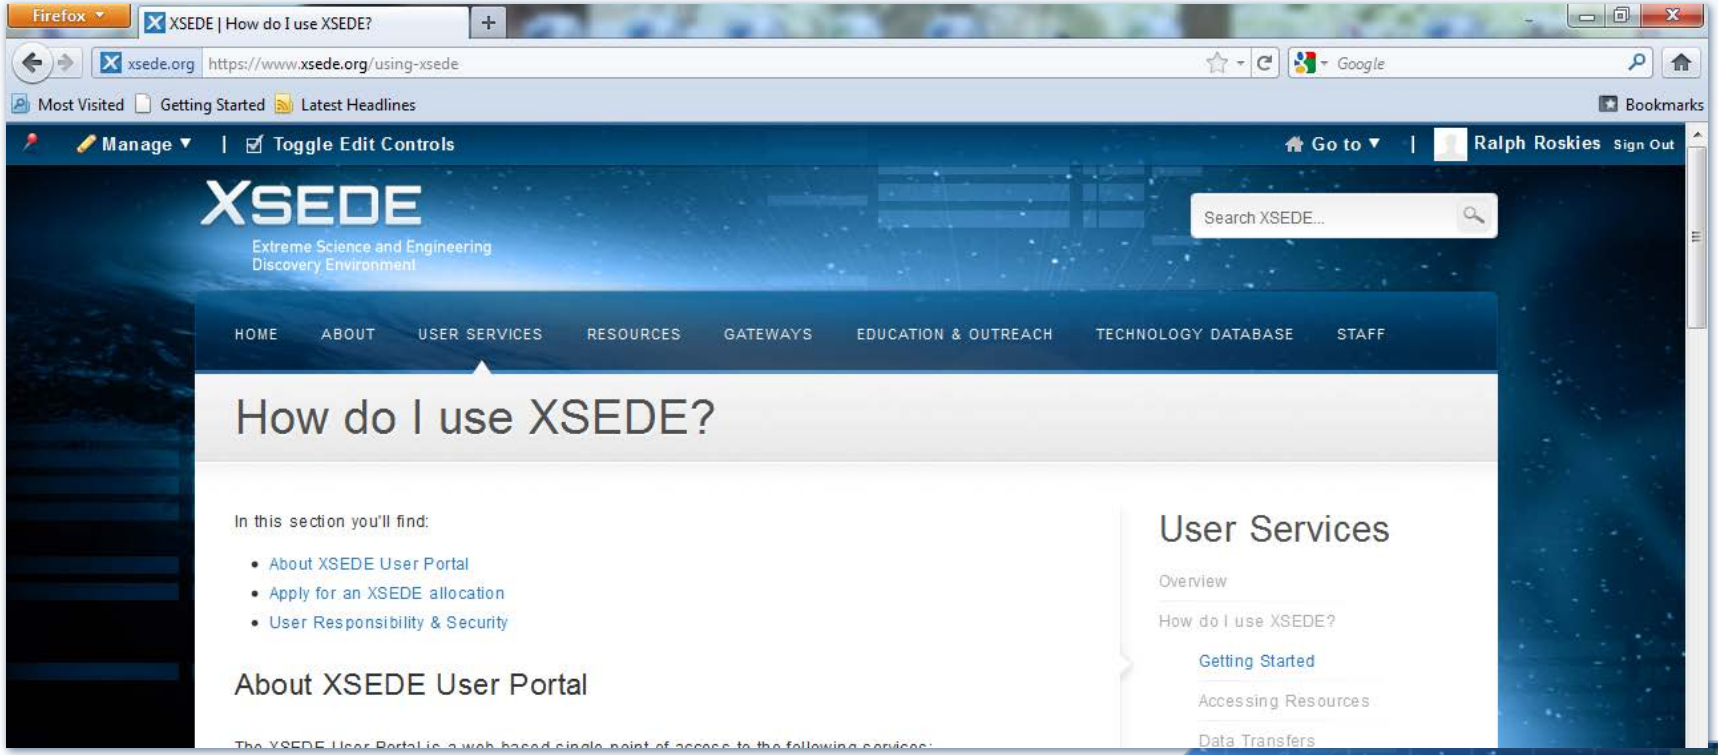

# Simple Enough

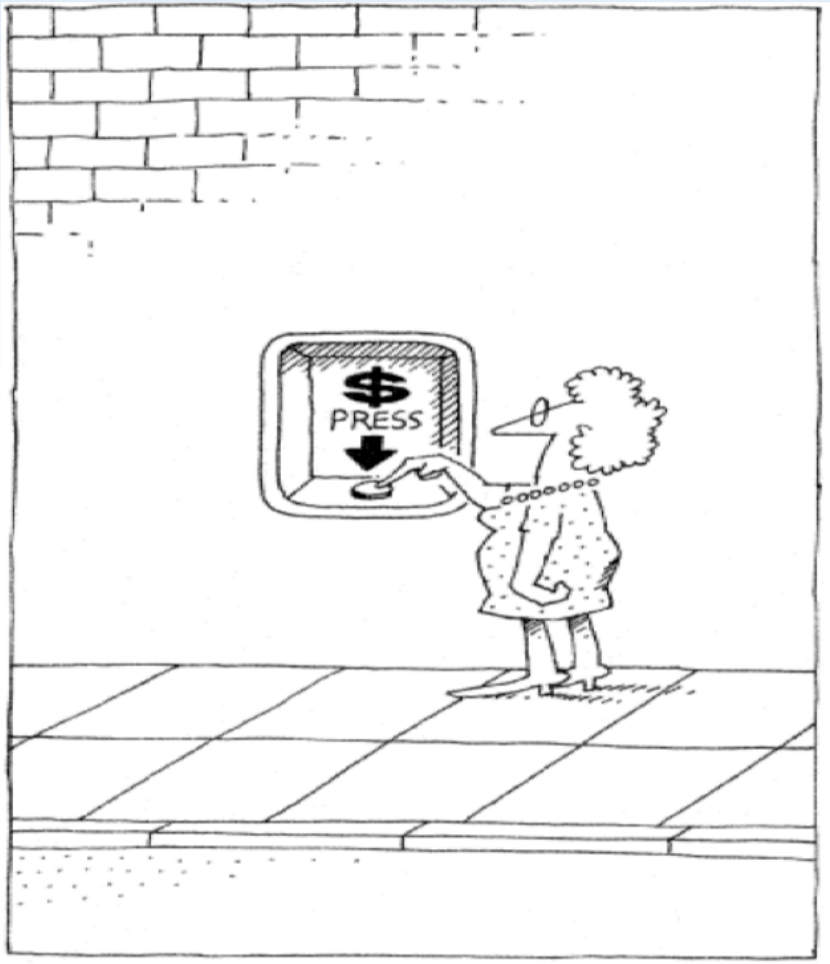

# XSEDE Allocations (www.xsede.org/allocations)

- 'Start-up' requests (up to 200k CPU hours, depends on system) granted continuously with a straightforward request
  - Gordon startup is 100K SUs and Trestles is 50K
- 'Research' allocations are generally larger, require proposals, and are peer-reviewed 4X/year
  - Proposal process (covered later)
  - Very high awards to justified proposals (tens of millions of CPU hours)
- 'Education' requests also available for classes/training – similar process/limits to start-ups

# XSEDE Training (www.xsede.org/training)

- XSEDE provides extensive training
  - Covering every major resource
  - From beginner to advanced classes
  - At locations across the country
  - Online via
    - asynchronous technologies
    - webcasts
- Signing up is **simple--**in the XSEDE User Portal!

# XSEDE offers more in-depth support via Extended Collaborative Support Service (www.xsede.org/ecss)

- Support people who understand the discipline as well as the systems (perhaps more than one support person working with a project).
- 37 FTEs, spread over >70 people at more than half a dozen sites.
- Distributed support
  - Easier to find the right expert for the project
  - Allows us to cover many more disciplines than if every site had to staff the common applications.
  - Support does not have to move with platform change

# What kind of expertise does ECSS have?

- Optimization
  - Profiling
    - TAU, CrayPAT
  - Scaling to higher core count
    - MPI, OpenMP
  - Improving I/O
    - MPI-I/O
  - Finding better solvers (what's better often depends on the degree of parallelization)
    - ScaLAPACK, FFTs, PetSC
- GPU programming
  - CUDA
- Visualization
  - Visit, Paraview
- Grid Computing
  - Globus, Unicore, Condor
- Workflows
  - Kepler, Taverna, Pegasus
- Portals/Gateways
  - OGCE, HUBzero
- Data management
  - SQL

You can request ECSS support as part of your allocation proposal (FTE-months, rather than core-hours)

# XD Technology Audit Service - XDMoD

- A separate award to U Buffalo for XD Metrics on Demand (XDMoD) tool
- xdmod.ccr.buffalo.edu
- Lots of interesting metrics about resource usage and job characteristics

# Questions?