XSEDE: An Advanced and Integrated Set of Digital Resources for Science and Engineering

Scott Lathrop, lathrop@illinois.edu

XSEDE

Extreme Science and Engineering Discovery Environment

Presentations for Today

- Posted for your use

  http://www.hpcuniversity.org/trainingMaterials/18

XSEDE Vision/Mission/Goals

- **Tag line:**
  - *XSEDE – accelerating scientific discovery*
- **Vision:**
  - XSEDE aspires to be ***the*** place to go to access digital research services.
- **Mission:**
  - Accelerate scientific discovery by enhancing the productivity of researchers, engineers, and scholars by deepening and extending the use of XSEDE's ecosystem of advanced digital, services and by advancing and sustaining the XSEDE advanced digital infrastructure.
- **Goals:**
  - *Deepen* and *extend* the use of the XSEDE ecosystem.
  - *Advance* the XSEDE infrastructure.
  - *Sustain* the XSEDE infrastructure.

Science requires diverse digital capabilities

- XSEDE is a comprehensive, expertly managed and evolving set of advanced heterogeneous high-end digital services, integrated into a general-purpose infrastructure.
- XSEDE is about increased user productivity
  - increased productivity leads to more science
  - increased productivity is sometimes the difference between a feasible project and an impractical one

Range of Advanced Digital Capabilities

- Often use the terms “resources” and “services”
  - these should be interpreted very broadly
  - most are likely not operated by XSEDE
- Examples of resources
  - compute engines: HPC, HTC (high throughput computing), campus, departmental, research group, project, …
  - data: simulation output, input files, instrument data, repositories, public databases, private databases, …
  - instruments: telescopes, beam lines, sensor nets, shake tables, microscopes, …
  - infrastructure: local networks, wide-area networks, …
- Examples of services
  - collaboration: wikis, forums, telepresence, …
  - data: data transport, data management, sharing, curation, provenance, …
  - access/used: authentication, authorization, accounting, …
  - coordination: meta-queuing, …
  - support: helpdesk, consulting, ECSS, training, …
  - And many more: education, outreach, community building, …

XSEDE supports a breadth of research

From direct contact with user community as part of requirements collections

- Earthquake Science and Civil Engineering
- Molecular Dynamics
- Nanotechnology
- Plant Science
- Storm modeling
- Epidemiology
- Particle Physics
- Economic analysis of phone network patterns
- Brain science
- Analysis of large cosmological simulations
- DNA sequencing
- Computational Molecular Sciences
- Neutron Science
- International Collaboration in Cosmology and Plasma Physics
- Social Sciences
- Humanities

XSEDE supports thousands of such projects – there are sample domains.

XSEDE offers a variety of resources

- Leading-edge distributed memory systems
- Very large shared memory systems
- High throughput systems, e.g. OSG
- Visualization servers
- Accelerators and co-processors including NVIDIA GPUs and XEON Phi (MICs)

Many scientific problems have components that call for use of more than one architecture.

Current XSEDE Compute Resources

- Stampede @ TACC
  - 6 PFLOPS (PF) Dell Cluster w/ GPUs and Xeon PHIs
- Gordon @ SDSC
  - 341 TF Appro Distributed SMP cluster
- Lonestar (4) @ TACC
  - 302 TF Dell Cluster
- Trestles @ SDSC
  - 100TF Appro Cluster
- Steele @ Purdue
  - 60 TF Dell Cluster
- Blacklight @ PSC
  - 37 TF SGI UV (2 x 16TB shared memory SMP)
- Mason
  - 3.8 TF HP Cluster with large memory nodes (2TB/node)

https://www.xsede.org/web/xup/resource-monitor

Current XSEDE Visualization and Data Resources

- Visualization
  - Nautilus @ UTK
    - 8.2 TF SGI/NVIDIA SMP
    - 960 TB disk
  - Longhorn @ TACC
    - 20.7 TF Dell/NVIDIA cluster
    - 18.7 TB disk

https://www.xsede.org/web/xup/resource-monitor#advanced_vis_systems

- Storage
  - HPSS @ NICS
    - 6.2 PB tape
  - Data Supercell @ PSC
    - 4 PB tape
  - Ranch @ TACC
    - 40 PB tape
  - Data Oasis @ SDSC
    - 4 PB tape

https://www.xsede.org/web/xup/resource-monitor#storage_systems

NCSA Blue Waters System

- Funded by the NSF to support very large scale computational science and engineering
- Cray systems
  - 22,640 Cray XE6 nodes - 64 GB of memory per node
  - 3,072 Cray XK7 nodes include NVIDIA processors with 32 GB of memory
  - 26 petabytes of online storage
  - 380 petabytes of tape storage
- Allocations are made via:
  - Applications to the NSF PRAC proposal process
  - Applications to Blue Waters education allocations

XSEDE User Services

XSEDE User Services are grouped into four main areas:

- Technical information
  - Always available via web site and XSEDE user portal
- Allocations
  - Request access to XSEDE's systems
- Training
  - Sign up for classes to learn to use XSEDE resources
- User Engagement
  - Includes 'consulting support' to answer questions
  - Also includes user interviews, focus groups, and surveys

Getting Started with 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

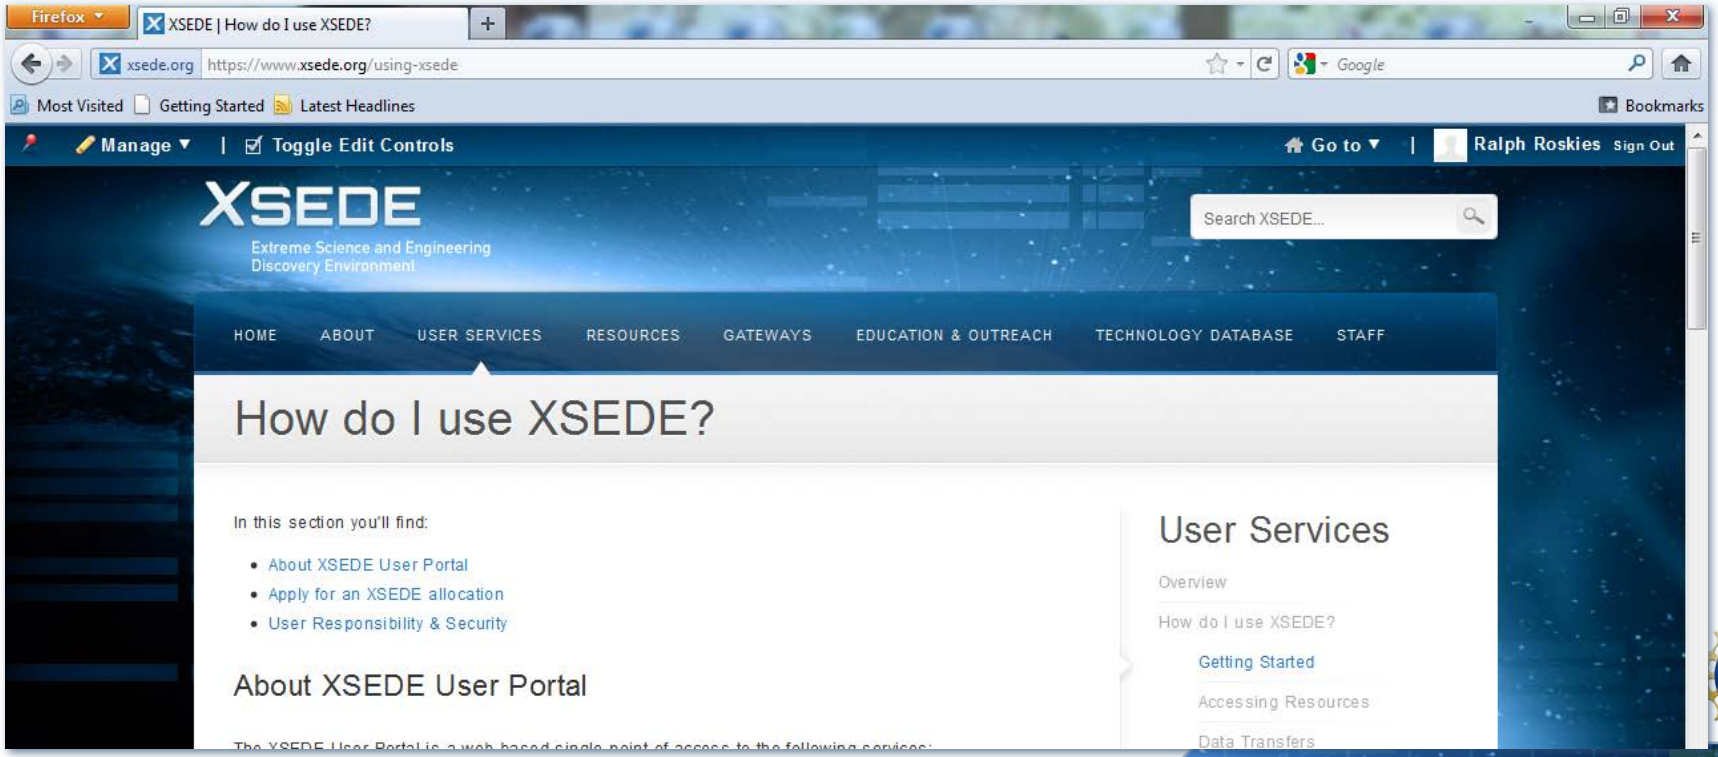

XSEDE User Portal: THE User Site
portal.xsede.org

- XSEDE User Portal (XUP) is designed to be the *only* site a **user** needs to use XSEDE
- XUP presents *only* info relevant to users
  - nothing else, so user info is easier to find
  - XUP also provides dynamic data about XSEDE systems
  - capabilities to manage usage, files, data
- After creating an account, a user can
  - request an allocation, and manage allocations
  - sign up for training
  - request help
  - manage file and data
  - and much more!
- Portal provides single sign-on to all XSEDE resources

XSEDE User Guides and News

- XSEDE provides intro user guides for every XSEDE-allocated system—no matter where it is actually hosted
  - Consistently structured and formatted
  - All available from website and XUP
  - Prepared using expertise of host sites
- XSEDE also provides up-to-date User News about every system, and XSEDE-wide services, available via:
  - Web/portal
  - Email
  - RSS feeds
  - Calendar feeds (for downtimes, training events, etc.)

XSEDE Allocations

- XSEDE allocates access/time on powerful, valuable systems providing different capabilities at NO COST TO YOU
  - HPC
  - High throughput computing
  - Remote visualization
  - Data storage
  - Etc.
- Users may request XSEDE staff support to assist with optimization of research codes, visualization, workflows, novel projects, and science gateways
- Single Sign-On allows you to use just one username and password (your User Portal one). You will be recognized by all XSEDE services on which you have an account, without having to enter your login information again for each resource.

XSEDE Allocations (2)

- Request allocations through the XSEDE User Portal
- It's **easy** to get a 'Startup allocation'—best way to get started
- Education allocations for classroom use via XSEDE and Blue Waters
- Larger year-long 'research' allocations can be requested 4 times/year, are peer reviewed, and have a longer lead-time
- Quarterly webinars on writing allocations

XSEDE Training

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

Getting Help

- Getting help is **easy—again**, via the XUP
  - Can also call the helpdesk **1-866-907-2383** 24x7 to request assistance

Science Gateways

- Researchers using tools where inputs could be standardized
  - Same executables (no need to recompile)
    - GridChem, CHARMM
- Creating standardized workflows
- Input is streamed data to the web

e.g. LEAD takes radar data, and determines whether the pattern suggests possible formation of a tornado cell, in which case more fine-grained simulation is needed

Today, there are approximately 35 gateways using XSEDE

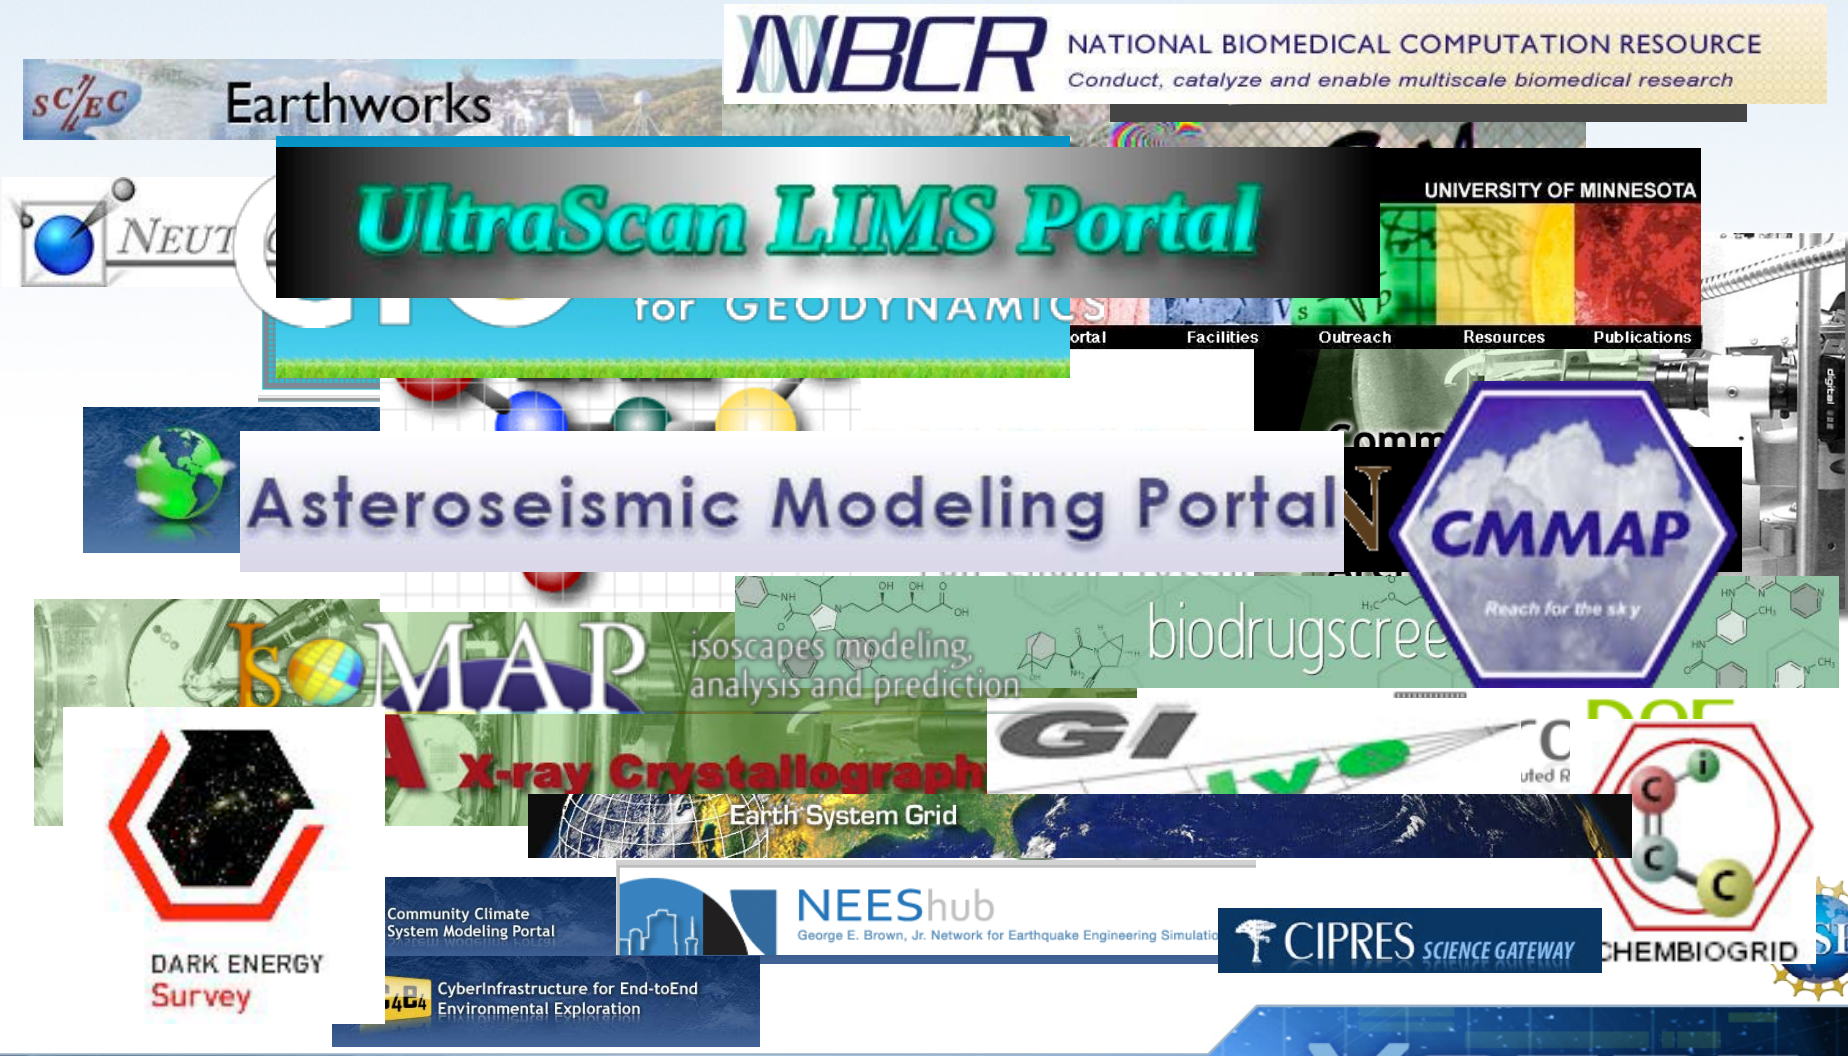

XSEDE

Community Engagement Activities

- Student Programs
- Under-represented Community Engagement
- Champions Program
- Campus Bridging
- Campus Visits
- HPC University portal
- Annual XSEDE Conference

Student Programs

- XSEDE Scholars
  - engaging undergraduates and graduates in year-long series of webinars attend annual XSEDE Conference
- XSEDE Summer Research Experience
  - Summer internship with XSEDE staff or user
- Blue Waters Internship
  - 2 week training institute for undergrads and grads
  - year-long computational science problem solving
- Blue Waters Graduate Fellowship
  - similar to NSF Graduate Fellowships
  - year-long engagement
- XSEDE Annual Conference
  - travel support for students to attend the annual Conference
- HPC University
  - Lists other student engagement opportunities

Under-represented Community Engagement

- Outreach to faculty and students at Minority Serving institutions
- Assist faculty with conducting their research using XSEDE resources
- Assist faculty with incorporating computational tools, resources and methods into the curriculum
- Minority Research Committee – faculty assisting one another
- Engaging students various programs

Campus Champions Program

- Over 160 campuses are members – no cost to join
- Champions receive monthly training and updates
- Champions provided with start-up accounts
- Champions are members of User Services team
- Forum for sharing and interactions
- Access to information on usage by local users
- Registrations for annual XSEDE14 Conference waived
- Community building across campuses

Campus Champions Role

- Raise awareness locally
- Provide training
- Get users started with access quickly
- Represent needs of local community
- Provide feedback to improve services
- Attend annual XSEDE14 conference
- Share their training and education materials
- Build community among all Champions

Champions Program

- Campus Champions
  - Representatives to spread information about XSEDE to local faculty, students and staff
- Student Champions
  - Students assist the Campus Champions
- Regional Champions
  - Representatives to spread information about XSEDE to other campuses in the area
- Domain Champions
  - Disciplinary people able to assist others with domain specific HPC questions

University of Minnesota Champion

Pedro da Silveira

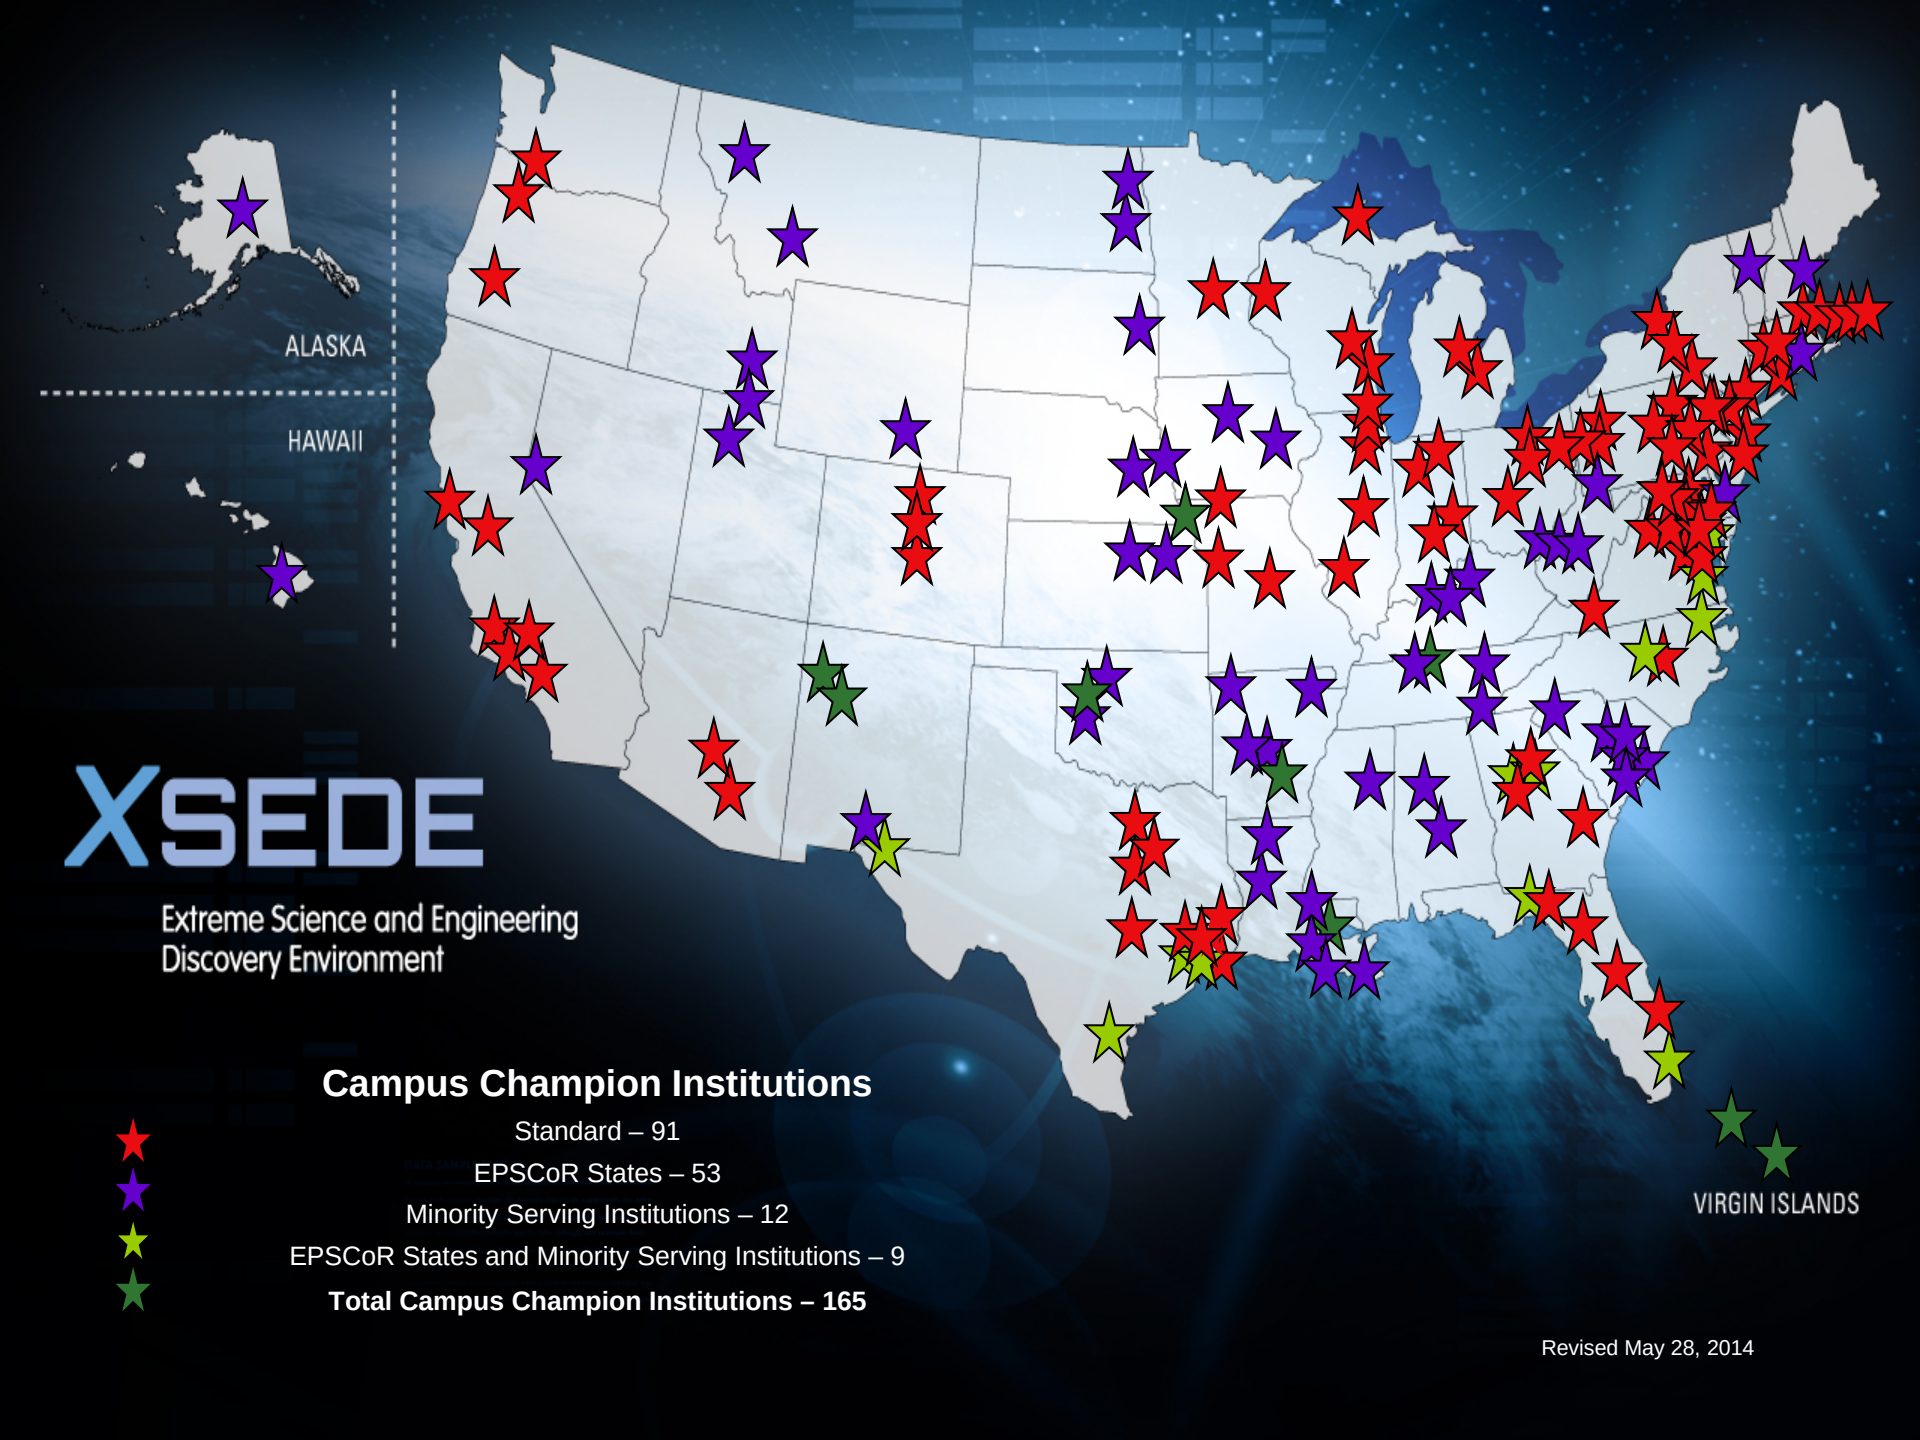
ALASKA
HAWAII
XSEDE
Extreme Science and Engineering Discovery Environment
Campus Champion Institutions
Standard – 91
EPSCoR States – 53
Minority Serving Institutions – 12
EPSCoR States and Minority Serving Institutions – 9
Total Campus Champion Institutions – 165
VIRGIN ISLANDS
Revised May 28, 2014

What is “Campus Bridging”?

The goal of campus bridging in general is to create a sense of “virtual proximity.” Any resource should feel as if it’s just a peripheral to their laptop or workstation.

The goal is to make it convenient and intuitive to simultaneously use your personal computing systems, departmental and campus systems (at your campus and others), and national resources liked XSEDE . . . all (almost) transparently and easily.

Campus Visits

- XSEDE visits campuses to
  - raise awareness
  - conduct professional development and curriculum development sessions,
  - assist with incorporating campus bridging tools and resources
  - meet with administrators, faculty, staff and students to effect institutional change
- Let us know how we can assist your campus

HPC University portal

- Training and education resources
- HPC training roadmap
- Events worldwide
- Internship and fellowship opportunities
- Career opportunities
- Computational science and education blog

www.hpcuniversity.org

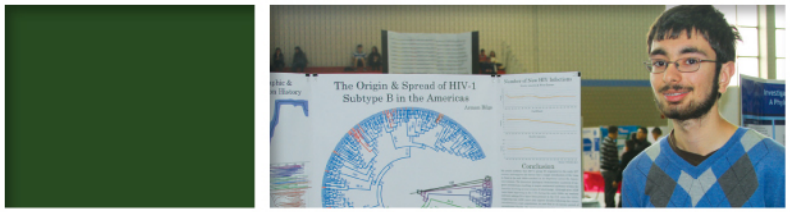

XSEDE[14]

July 13-18, 2014
Atlanta Marriott Marquis

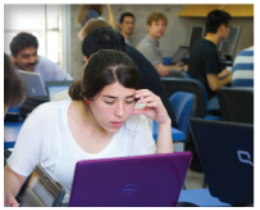

Registration before June 16 for lower rates!

www.xsede.org/xsede14

XSEDE14 Conference

- Theme is Engaging Communities
- Topics span accelerating discovery, advanced technologies, software, science gateways and portals, and education, outreach and training
- Tutorials, papers, BOFs, posters, sponsors
- Over 700 people from academia, industry, government, and other organizations
- Interact with experts in the field and people from all XSEDE partner organizations

Stay Connected

- XSEDE's public web site is www.xsede.org
- Create an XSEDE User Portal signon and receive news and notices
- Training events are announced via the public web site; and registrations via the XSEDE User Portal
- For access to additional training and educational resources www.hpcuniversity.org


Our reach will forever
exceed our grasp, but,
in stretching our horizon,
we forever improve our world.
XSEDE
Extreme Science and Engineering
Discovery Environment