October 20, 2011

XSEDE: the eXtreme Science and Engineering Discovery Environment

XSEDE

Extreme Science and Engineering Discovery Environment

October 20, 2011

Introduction to XSEDE

John Towns
PI and Project Director, XSEDE
Director, Persistent Infrastructure, NCSA
jtowns@ncsa.illinois.edu

XSEDE

Extreme Science and Engineering Discovery Environment

NSF

XD Solicitation/XD Program

- eXtreme Digital Resources for Science and Engineering (NSF 08-571) -- Extremely Complicated
  - High-Performance Computing and Storage Services
    - aka Track 2 awardees and others
  - High-Performance Remote Visualization and Data Analysis Services
    - 2 awards; 5 years; $3M/year
    - proposals due November 4, 2008
  - Integrating Services (5 years, $26M/year)
    - Coordination and Management Service (CMS)
      - 5 years; $12M/year
    - Technology Audit and Insertion Service (TAIS)
      - 5 years; $3M/year
    - Advanced User Support Service (AUSS)
      - 5 years; $8M/year
    - Training, Education and Outreach Service (TEOS)
      - 5 years, $3M/year
  - two phase proposal process for IS
    - pre-proposals November 4, 2008
    - final proposals due June 15, 2009

Science requires diverse digital capabilities

- XSEDE will be a comprehensive, expertly managed 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

XSEDE Vision

The eXtreme Science and Engineering Discovery Environment (XSEDE) will:

- enhance the productivity of scientists and engineers by providing them with new and innovative capabilities

and thus

- facilitate scientific discovery while enabling transformational science/engineering and innovative educational programs

XSEDE will support 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

Sampling of much larger set. Many examples are new to TeraGrid/HPC. Range from petascale to disjoint HTC, many are data driven. XSEDE will support thousands of projects.

XSEDE's Distinguishing Characteristics

- Foundation for a national CI ecosystem
  - comprehensive suite of advanced digital services will federate with other high-end facilities and campus-based resources
- Unprecedented integration of diverse digital resources
  - innovative, open architecture making possible the continuous addition of new technology capabilities and services

Infrastructure Designed for Innovation & Evolution

- An environment in which all resources, data and services relevant to a researcher can be embedded and shared
  - campus bridging creating a single virtual system with interactive data transfer and resource sharing capabilities
    - "make my data accessible everywhere I want to be"
  - coordinated archival approach to ensure persistence of important datasets beyond the lifetime of particular service providers
- An underlying infrastructure to support this
  - open architecture with judicious use of standards designed to evolve in a non-disruptive way
  - interoperability of XSEDE with other CIs

XSEDE's Distinguishing Characteristics - Governance

- World-class leadership from CI centers with deep experience: partnership led by NCSA, NICS, PSC, TACC and SDSC
  - PI: John Towns, NCSA/Univ of Illinois
  - Co-PIs: Jay Boisseau, TACC/Univ of Texas Austin
    TBD, NICS/Univ of Tenn-Knoxville
    Ralph Roskies, PSC/CMU
    Nancy Wilkins-Diehr, SDSC/UC-San Diego
- Partners who strongly complement these CI centers with expertise in science, engineering, technology and education
  - Univ of Virginia
    SURA
    Indiana Univ
    Univ of Chicago
    Berkeley
    Shodor
    Ohio Supercomputer Center
    Cornell
    Purdue
    Rice
    NCAR
    Jülich Supercomputing Centre

How we propose to engage stakeholders

- Collection of stakeholder needs:
  - surveys, ticket mining, …
  - focus groups, usability panels, …
  - interviews, shoulder surfing, …
- Prioritization of identified need and derived requirements
  - User Requirements Evaluation and Prioritization (UREP) Working Group
    - broad participation across architecture, deployment, operations, users, and service providers
- Assessing plans and deployments
  - through a variety of stakeholder-focused, facilitated workshops
    - e.g., interactive ATAM sessions focused on identifying, quantifying, discussing tradeoffs
- Representation in the management of XSEDE
  - XSEDE Advisory Board
  - User Advisory Committee
  - Service Providers Forum

XSEDE Distributed Systems Architecture

- Architecture defines the XSEDE system's components and how they interact
  - each component is motivated by one or more requirements
  - each component is defined in terms of required capabilities: interfaces and qualities of service
- Equally important is the process by which we revise the architecture over time
  - key point: driven by new or revised requirements

Initial XSEDE architecture: High-order bits

- Don’t disrupt the user community! Maintain existing TeraGrid services
- Focus on user-facing access layer
  - for power users: “first, do no harm”
  - for other users: expand use via new hosted XSEDE User Access Services (XUAS) and Global Federated File System (GFFS)
- Promote standards and best practices to enhance interoperability, portability, and implementation choice

XSEDE's Distinguishing Characteristics - Architecture

- XSEDE is *designed* for innovation & evolution
  - there ***is*** an architecture defined
    - based on set of design principles
    - rooted in the judicious use of standards and best practices
    - clearly defined transition plan from TeraGrid to XSEDE
- Professional systems engineering approach
  - responds to evolving needs of existing, emerging, and new communities
    - incremental development/deployment model
  - new requirements gathering processes
    - ticket mining, focus groups, usability panels, shoulder surfing
  - ensure robustness and security while incorporating new and improved technologies and services
  - process control, quality assurance, baseline management, stakeholder involvement

Systems Architecture: Basic Components

- Functional components
  - think operating systems
  - processes, inter-process communication, security, file systems, memory management
- Non-functional – "ilities"
  - reliability, availability, extensibility, usability, "performability," etc.
  - note trade-offs

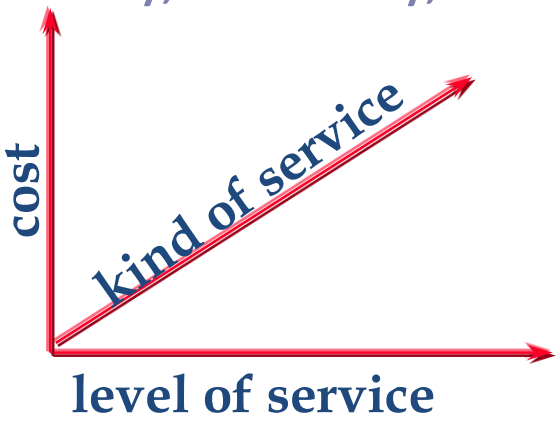

"Give me simple abstractions and make them work reliably."
--Kent Blackburn

How we describe the XSEDE architecture

- A set of "views" describing the elephant from the perspectives of different stakeholders
  - Not (only) immensely detailed documentation!
- Different stakeholders require different views, e.g.,
  - Service provider
  - System administrator
  - Power user
  - Occasional user
  - Gateway developer
  - Security officer
  - NSF program manager
  - Campus CIO
  - Trainer
  - …
- Tell us what views **you** think are important

Initial Structural Views

- Capabilities view
  - primarily for management stakeholders
  - a set of capabilities and the definition of each
  - mapping from capabilities to requirements
  - mapping from capabilities to the pieces of the architecture
- Component-and-connector view
  - intended for a very technical audience
    - detailed information about how the system works when running
  - run-time entities that execute and cooperate to perform the work of the system
- Module decomposition view
  - all of the pieces of the system that have to be developed, maintained, integrated, and tested
- Deployment view
  - where the running software executes. More formally, the view is a mapping from components (such as a process or service) onto the processing node that hosts it

Initial Quality Views

- Security view
  - shows how architecture achieves security requirements
  - re-packages “security aspects” of design into special form for the security stakeholders
- Operations (Monitoring) view
  - how site-specific and XSEDE-wide monitoring will be carried out
    - e.g., how system reports/keeps track of user jobs
  - how the architecture makes information available to people monitoring the system’s operation
- Availability view
  - how does system react to a failure?
    - how the architecture records/reports/recovers from faults and failures
  - for those many stakeholders to whom availability is of high concern and wish to see how availability is provided by the architecture
- Performance view
  - answers questions about the various kinds of performance
    - e.g., job throughput, transfer latency, bandwidth, compute capability…
  - for stakeholders concerned with performance
    - first-order performance analysis can be carried out using C&C and Deployment views
- Install view
  - maps/associates) components of C&C view to file management system production environment

Current XSEDE Compute Resources

- Kraken @ NICS
  - 1.2 PF Cray XT5
- Ranger @ TACC
  - 580 TF Sun Cluster
- Lonestar (4) @ TACC
  - 302 TF Dell Cluster
- Forge @ NCSA
  - 150 TF Dell/NVIDIA GPU Cluster
- Trestles @ SDSC
  - 100TF Appro Cluster
- Steele @ Purdue
  - 67 TF Dell Cluster
- Blacklight @ PSC
  - 36 TF SGI UV (SMP)
- Dash @ SDSC
  - 5 TF Appro Distributes SMP cluster

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
  - Spur @ TACC
    - 1.1 TF Sun cluster
    - 1.7 PB disk

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

- Storage
  - Albedo
    - 1 PB Lustre distributed WAN filesystem
  - Data Capacitor @ Indiana
    - 535 TB Lustre WAN filesystem
  - Data Replication Service
    - 1PB iRODS distributed storage
  - HPSS @ NICS
    - 6.2 PB tape
  - MSS @ NCSA
    - 10 PB tape
  - Ranch @ TACC
    - 70 PB tape
  - HPSS @ SDSC
    - 25 PB tape

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

Current XSEDE Special Purpose Resources

- Condor Pool @ Purdue
  - 150 TF, 27k cores
- Keeneland @ GaTech/NICS
  - developmental GPU cluster platform
- FutureGrid
  - Experimental/development distributed grid environment

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

OSG and XSEDE

- Realizing a national CI ecosystems requires confederation of CI providers
  - OSG and XSEDE working together represents a substantial step in this direction
  - still much to address though!
- OSG is a significant CI in the US
  - ties to CI (eScience infrastructure) providers internationally
- OSG represents a direction of expanding the scope of XSEDE
  - OSG has a focus on HTC opportunistic resources
  - TeraGrid was mostly about large parallel applications of HPC capability and capacity resources
  - these compliment each other well and present real opportunities for leverage and integration that will benefit the research community

OSG Relationship

- OSG will be a Service Provider in XSEDE
  - anticipated to be a Tier 1 SP
- OSG resources will be made available via XSEDE allocations processes
  - primarily HTC resources
  - OSG very interested in leveraging the XSEDE review process
  - opportunistic nature of OSG resource will present a new twist to allocations processes and review
- OSG has two other interaction points with XSEDE
  - participation in outreach/campus bridging/campus champions activities
    - assure incorporation of the OSG cyberinfrastructure resources and services into campus research and education endeavors
  - effort in ECSS specifically to work with applications making use of both OSG and XSEDE resources

Additional Activities between XSEDE and OSG

- EXTenCI
  - joint proposal between OSG and TeraGrid
  - pursuing technologies of mutual benefit for communities served by both OSG and TeraGrid
- Involving participants in both projects in each other's planning processes
- Developing additional joint proposal ideas

Campus bridging role

- To be conscientiously targeted at Data, HPC, and HTC – probably in that order
- Working closely with architecture and security teams to help disseminate XSEDE's plans
  - XSEDE architecture plans => out to campus champions and community
  - Funnel community response => back to XSEDE Architecture Team through DOORS, requirements gathering, and evaluations
- Work with TEOS teams to promote adoption of approaches that create a better integrated (and in aggregate larger) suite of resources for use by the national engineering and research community

XSEDE campus bridging tactics year 1

- There is an important value proposition that does not involve cash
- Support Installers created by Architecture group
  - Call for participation coming soon – pursue a small number of campuses as part of a pilot program to affect uptake by working within the them with diligence, and reap economies of scale (if things go right) or clear learning experiences (otherwise)
  - Planning to visit campuses to raise awareness and work with campus personnel
- Documentation & training
  - Science Gateways (document by Surresh Marru & Marlon Pierce)
  - Promote use of systems template created by TACC
- Serve as 'connectors' in discussions that should or could play a role in campus bridging activities (sit in in Arch and User Services calls)
- Work closely with OSG (OSG is plenty good at what they do!)

http://pti.iu.edu/campusbridging/

NSF ACCI Task Force on Campus Bridging

- Read current draft of final report (1 Mar 2011 v4.1)
- Submit a comment or position paper
- Read submitted position papers

Information about NSF-sponsored workshops

- Campus Bridging: Networking & Data-centric Challenges Workshop
- Campus Bridging: Software & Service Issues Workshop
- Campus Bridging: Campus Leadership

Campus Bridging community needs assessment

In early 2009 **National Science Foundation's (NSF) Advisory Committee for Cyberinfrastructure (ACCI)** charged six different task forces to make strategic recommendations to the NSF in strategic areas of cyberinfrastructure: Campus Bridging; Data; Grand Challenges and Virtual Organizations; High Performance Computing; Software and Tools; and Work Force

October 20, 2011

Education and Outreach Services

Steven Gordon (sgordon@osc.edu)

XSEDE

Extreme Science and Engineering Discovery Environment

Education and Outreach Goals

- Prepare the current and next generation of researchers, educators and practitioners.
- Create a significantly larger and more diverse workforce in STEM.
- Inculcate the use of digital services as part of their routine practice for advancing scientific discovery.

User/Admin View – logged into XSEDE Portal

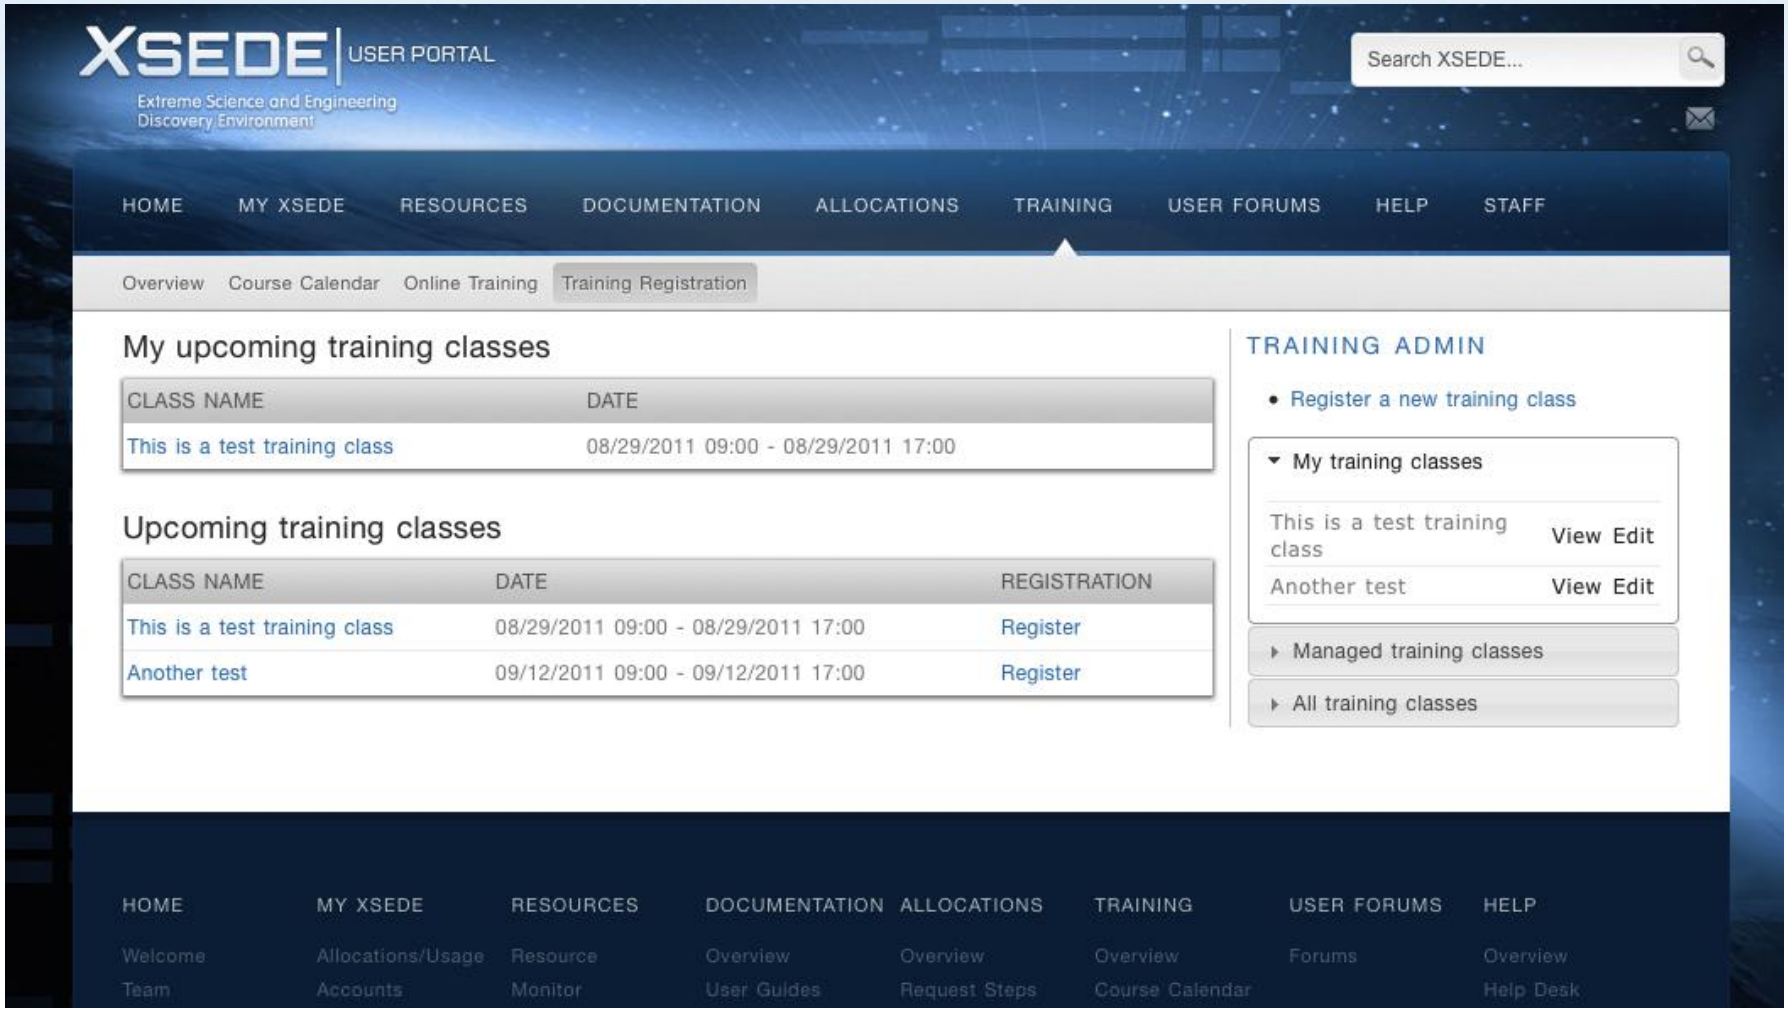

Training Goals

- Expand the scope/scale of training through expanded use of distance learning, new topics, etc.
- Create a simpler user experience through a single portal for all training at all sites
  - Updated portal coming soon
- Create an internal repository to promote sharing of materials indexed at a single site.

XSEDE Education Workshops

- Workshops for faculty
  - Focus on tools and pedagogy for teaching computational science
  - Workshops in various disciplines being planned for summer 2012 (chemistry, biology, computational thinking
  - Visits to campuses to encourage faculty interest in computational science
- UC Berkeley Par Lab Boot Camp on parallel programming
  - Given in August of each year and available online
  - http://parlab.eecs.berkeley.edu/2011bootcampagenda
- SC11 Education and Broader Engagement Workshops

Certificate and Degree Programs

- Creation of competency based model programs in computational science
- Recruiting campuses interested in starting programs
  - Assistance in starting new programs
  - Campus visits and faculty professional development
  - Programs in science and engineering
  - Teacher educator programs

Student Engagement

Components

- Students
  - Undergraduate and graduate
  - Drawn from contacts within and outside of XSEDE
  - 3-12 month appointments
- Projects
  - Provided and supervised by XSEDE researchers and staff

Outcomes

- Student presentations (papers, posters, etc)
- Case studies of successful and unsuccessful experiences
- More experienced practitioners entering STEM workforce

Process:

- Students and projects recruited and paired throughout the year.
- Researcher/staff supervises student work to complete project.
- Student develops and submits presentation material to relevant venue(s).

Underrepresented Engagement

Minority Institutions

SURA

- Identify established and emerging programs and researchers
- Expand awareness of XSEDE via campus visits, professional conferences
- Build a community promoting collaboration and peer support
- Target deep engagement that connect researchers with XSEDE expertise

SDSC

Work with SURA and Rice/ELA to implement their plans nationally, esp. in the southwest US and among Hispanic and Tribal communities

Minority Students at Research Institutions

Rice University/Empowering Leadership Alliance (ELA)

- Increase awareness and knowledge among underrepresented communities
- Identify and recruit minority students and mentors, leveraging ELA
- Provide education and professional development to participants
- XSEDE scholars program (http://bit.ly/xsede_2011)

Campus Champions

- “Champion” is a staff or faculty member on a campus that provides information on XSEDE to his/her colleagues
- Currently 100+ institutions represented by champions
- Receive training and support from XSEDE staff

Campus Champion Institutions

September 20, 2011

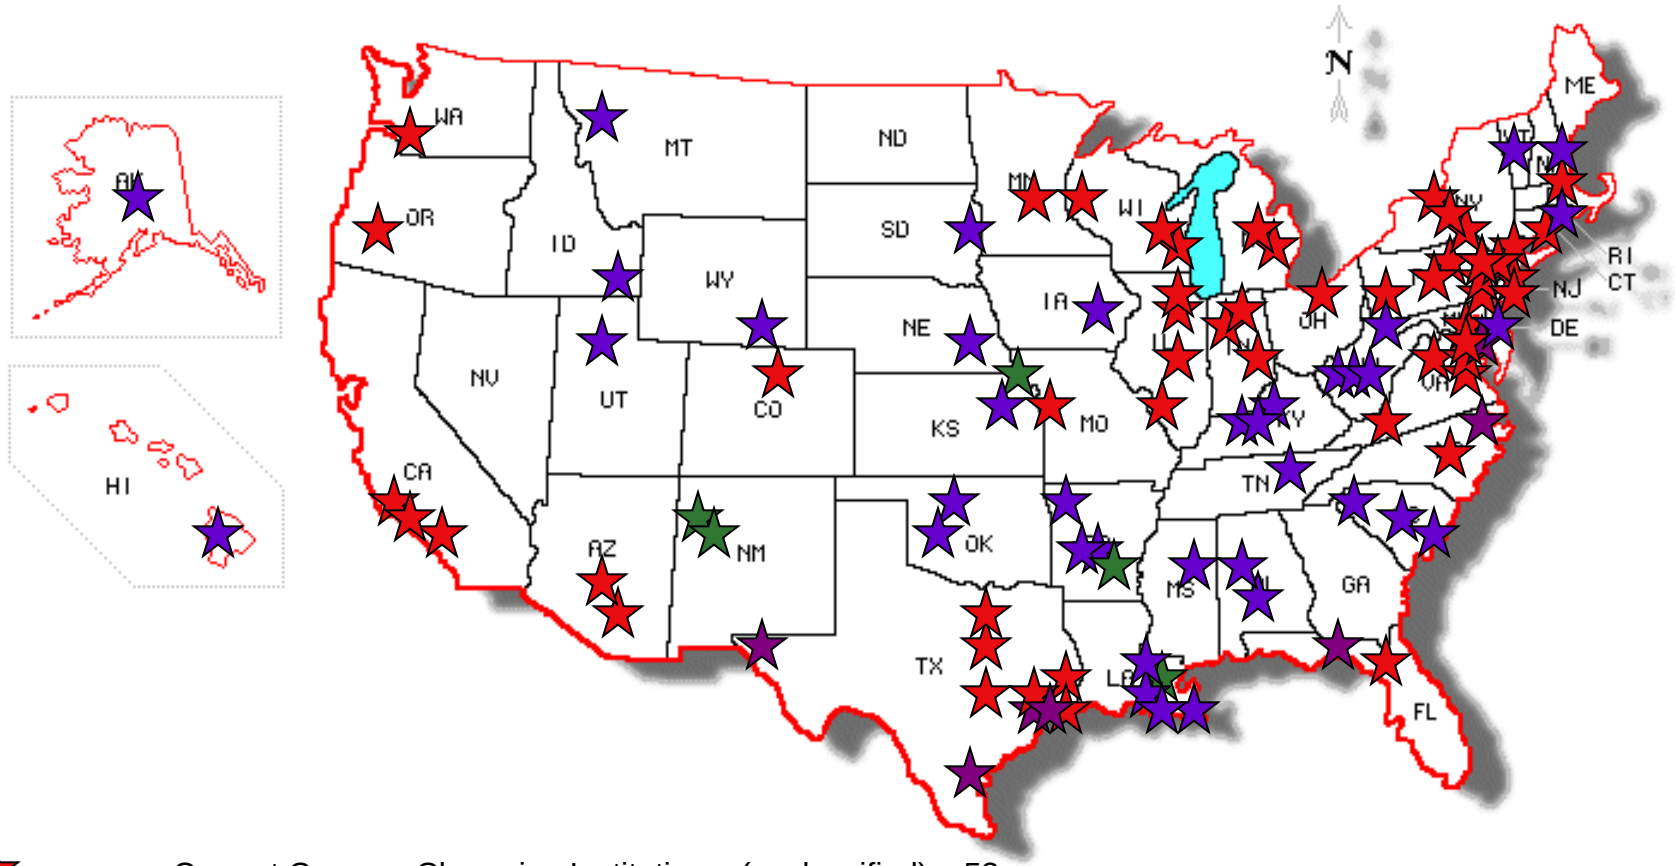

Current Campus Champion Institutions (unclassified) – 53

Current Campus Champion Institutions (EPSCoR state

Current Campus Champion Institutions (Minority Serving Institutions)-- 7

Current Campus Champion Institutions (both EPSCoR and MSI) – 5

Total Number of Campus Champion Institutions Overall -- 102

Speakers' Bureau

- **Audience**: Enhance XSEDE user diversity…
  - Demographically
  - Across disciplines
- **Venues**
  - Conferences and professional society meetings (as an exhibitor)
  - Campus visits
  - Presenter support
- **Criteria** for Venue selection:
  - Cost
  - Impact
- **Process**
  - Identify potential opportunities
  - <Go/No-Go> decision based on expected costs and impact
  - Execute selected events
  - Distribute contacts to XSEDE services for followup

For Further Information

- Website – www.xsede.org
- XSEDE Project
  - John Towns <jtowns@ncsa.illinois.edu>
- XSEDE Architecture
  - John Towns <jtowns@ncsa.illinois.edu>
  - Andrew Grimshaw <grimshaw@virginia.edu>
  - Ian Foster <foster@anl.gov>
- Campus Bridging
  - Craig Stewart <stewart@indiana.edu >
  - Rich Knepper <rknepper@indiana.edu>
- Education and Outreach
  - Steve Gordon <sgordon@osc.edu>
  - Scott Lathrop <scott@ncsa.uiuc.edu>


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