

# XSEDE

Extreme Science and Engineering  
Discovery Environment

# XSEDE Overview, Training Focus

Midwest Big Data Hub Meeting  
March 9 – 10, 2018

Susan Mehringer  
CAC Associate Director for Consulting  
XSEDE Training Lead  
[shm7@cornell.edu](mailto:shm7@cornell.edu)

# XSEDE Overview

## XSEDE:

- Is an NSF-funded virtual organization
- Integrates and coordinates the sharing of advanced digital services (e.g. supercomputers, visualization, data analysis resources)
- Serves researchers nationally to support science
- Provides users with seamless integration to NSF's high-performance computing and data resources

<https://www.xsede.org/about/what-we-do>

# XSEDE Resources

Resources include

- Multi-core high performance computing (HPC)
- Many-core HPC
- Distributed high-throughput computing (HTC)
- Visualization
- Data analysis
- Large-memory
- Data storage
- Cloud

<https://www.xsede.org/ecosystem/resources>

Resource	Org	Type
HP/NVIDIA Interactive Visualization and Data Analytics System (Maverick)	TACC	vis
IU/TACC (Jetstream)	TACC	compute
IU/TACC Storage (Jetstream Storage)	UT Austin	storage
LSU Cluster (superMIC)	LSU CCT	compute
Open Science Grid (OSG)	OSG	compute
PSC Bridges GPU (Bridges GPU)	PSC	compute
PSC Large Memory Nodes (Bridges Large)	PSC	compute
PSC Regular Memory (Bridges)	PSC	compute
PSC Storage (Bridges Pylon)	PSC	storage
SDSC Comet GPU Nodes (Comet GPU)	SDSC	compute
SDSC Dell Cluster with Intel Haswell Processors (Comet)	SDSC	compute
SDSC Medium-term disk storage (Data Oasis)	SDSC	storage
Stanford University GPU Cluster (XStream)	Stanford U	compute
TACC Data Analytics System (Wrangler)	TACC	compute
TACC Dell/Intel Knights Landing, Skylake System (Stampede2)	UT Austin	compute
TACC Long-term Storage (Wrangler Storage)	TACC	storage
TACC Long-term tape Archival Storage (Ranch)	TACC	storage

# Software on XSEDE Resources

The screenshot shows the XSEDE User Portal website. The browser address bar displays <https://portal.xsede.org/software#/>. The page header includes the XSEDE logo and the text "XSEDE | USER PORTAL" and "Extreme Science and Engineering Discovery Environment". A navigation menu contains links for MY XSEDE, RESOURCES (selected), DOCUMENTATION, ALLOCATIONS, TRAINING, USER FORUMS, HELP, ECSS, and ABOUT. A secondary menu lists various services: Systems Monitor, Storage, Remote Visualization, Software (selected), Queue Prediction, SU Converter, Science Gateways, and Scheduled Downtimes. The main content area is titled "Software" and contains the following text:

Find software available on XSEDE Service Provider sites. You can view by Resource, Site or Software type and then search for name, version, URL and more. To view details about a software package, click on the software name to see available versions. For more details click on the version to find out more about the software including how to access the software package.

Are you looking for software that is accessible via a science gateway? [Visit the Science Gateways Application List](#)

Help us gauge interest in potential future installations: if there's a software package you'd find useful, [submit a ticket](#) to let us know.

Below the text is a search bar with the placeholder "Search software" and a "Clear Search" button. To the right of the search bar is a "Collapse All" button. Under the search bar, there are two columns of filters. The left column is titled "View by:" and contains three buttons: "Resources" (selected), "Science Category/Domain", and "Sites". The right column is titled "Bridges" and contains a grid of software package names in buttons:

Bridges		
abaqus	Abinit	abyss
AIPS	allpaths-lg	anaconda
anaconda5	annovar	ansys

# Science Gateways

<https://www.xsede.org/ecosystem/science-gateways>

## Science Gateways Listing

This page lists all of the current science gateways. The list membership changes occasionally, as new projects join the community. Find the related science domain and links to the gateway home pages within the table below.

Below is a complete list of current science gateways, to see a detailed project description please click on the name of the science gateway.

To update the information contained in this table, please contact [help@xsede.org](mailto:help@xsede.org). To register your gateway, please complete the [Gateway Registration Form](#).

Title	Field of Science	Portal Homepage
<a href="#">Astero seismic Modeling Portal</a>	Stellar Astronomy and Astrophysics	<a href="#">Visit Portal</a>
<a href="#">Chem Compute</a>	Chemistry	<a href="#">Visit Portal</a>
<a href="#">CIPRES Portal for inference of large phylogenetic trees</a>	Systematic and Population Biology	<a href="#">Visit Portal</a>
<a href="#">Computational Anatomy</a>	Neuroscience Biology	<a href="#">Visit Portal</a>
<a href="#">Computational Chemistry Grid</a>	Chemistry	<a href="#">Visit Portal</a>

# XSEDE Allocations

The screenshot shows the XSEDE User Portal website. The browser address bar displays <https://portal.xsede.org/allocations/sta>. The page header includes the XSEDE logo, the text "Extreme Science and Engineering Discovery Environment", a search bar with the placeholder "Search XSEDE...", and a "SIGN IN" button. The main navigation menu is located below the header and includes links for "MY XSEDE", "RESOURCES", "DOCUMENTATION", "ALLOCATIONS" (which is highlighted), "TRAINING", "USER FORUMS", "HELP", "ECSS", and "ABOUT". A secondary navigation bar below the main menu contains links for "Announcements", "Resource Info", "Startup" (highlighted), "Education", "Research", "Submit/Review Request", "Manage Allocations", "Policies", and "About XRAS".

The main content area is titled "Startup Allocations". It features a sidebar on the left with a "Top of page" button and a list of links: "Submission Schedule", "Eligibility", "Required Components", "Startup Allocation Limits", "Example Startup Requests", "Campus Champion Allocations", and "Trial Allocations". The main text area contains the following content:

## Startup Allocations

Startup allocations, along with [Campus Champion](#) and [Trial](#) allocations, are one of the fastest ways to gain access to and start using XSEDE-allocated resources. We recommend that all new XSEDE users begin by requesting Startup allocations.

Appropriate uses for Startup allocations include:

- Small-scale computational activities that require the unique capabilities of resources allocated through XSEDE
- Application development by researchers and research teams
- Benchmarking, evaluation and experimentation on the various resources
- Developing a science gateway or other infrastructure software components

Users may request a single resource or a combination of resources. Each request is subject to [Startup Limits](#).

On the right side of the page, there is a vertical "FEEDBACK" button.



# Serve your community & help XSEDE via XRAC

## XSEDE Resource Allocations Committee (XRAC)

### Who

About 40 computational experts from academia, research labs, and industry

### What

Review large-scale requests to use XSEDE-allocated resources  
No more than 10 requests per meeting

### Why

Serve the community, see how the system works  
help ensure NSF resource portfolio used efficiently

### Where

Meetings held at seasonally advantageous locations around the country

### When

Four meetings per year, you're asked to attend at least three  
First week of March, June, December, and week before Labor Day

### How

All travel expenses paid by XSEDE • All-electronic submission and review system  
Meetings start Sunday evening and end Monday afternoon



Interested? Contact

Ken Hackworth • [hackworth@psc.edu](mailto:hackworth@psc.edu)

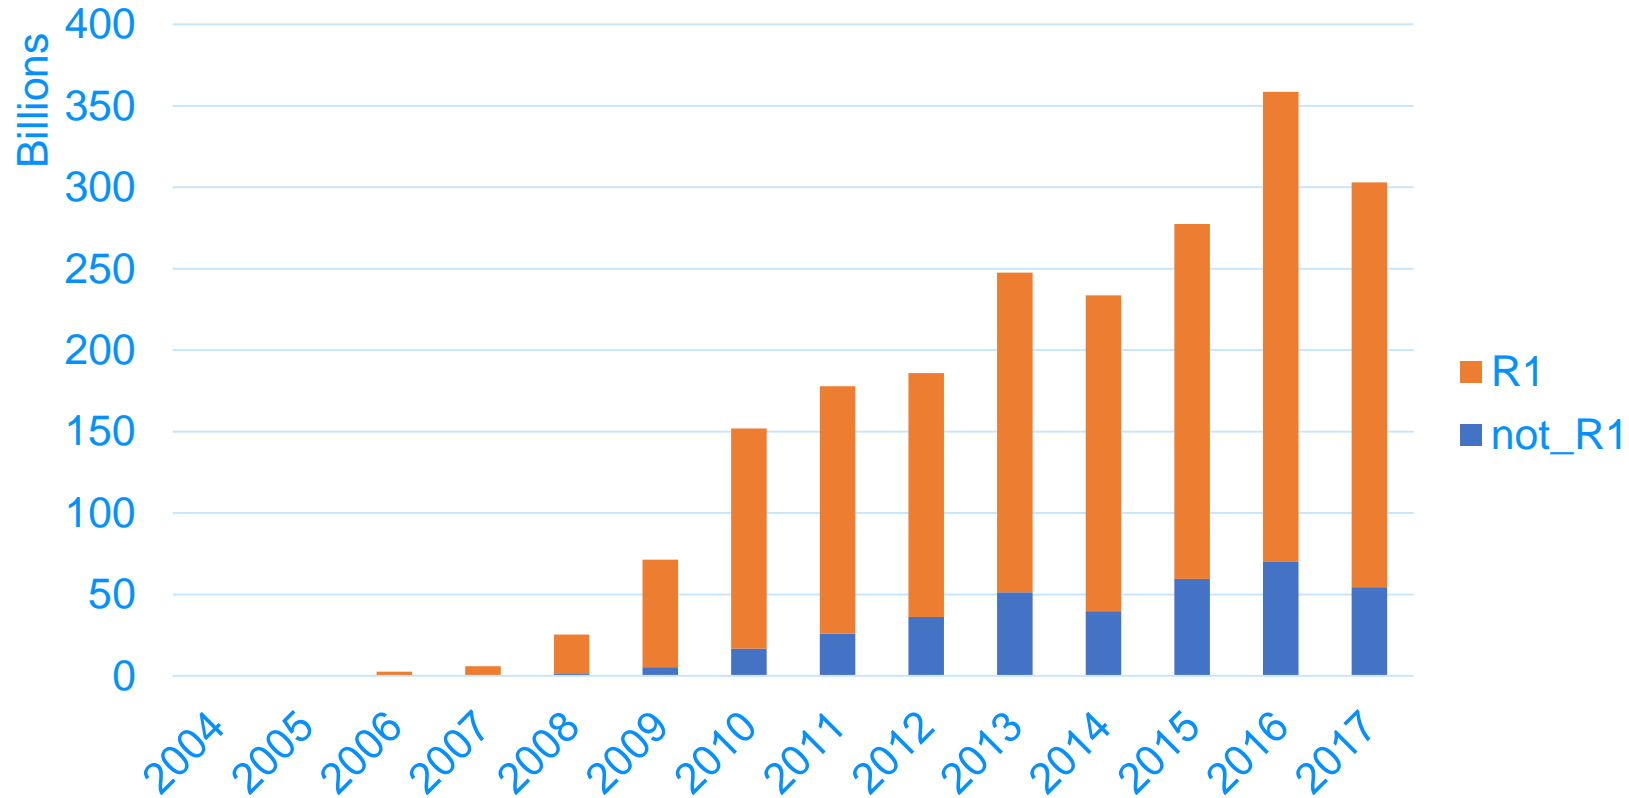


For more on XSEDE

[www.xsede.org](http://www.xsede.org) • [portal.xsede.org](http://portal.xsede.org)

# XSEDE

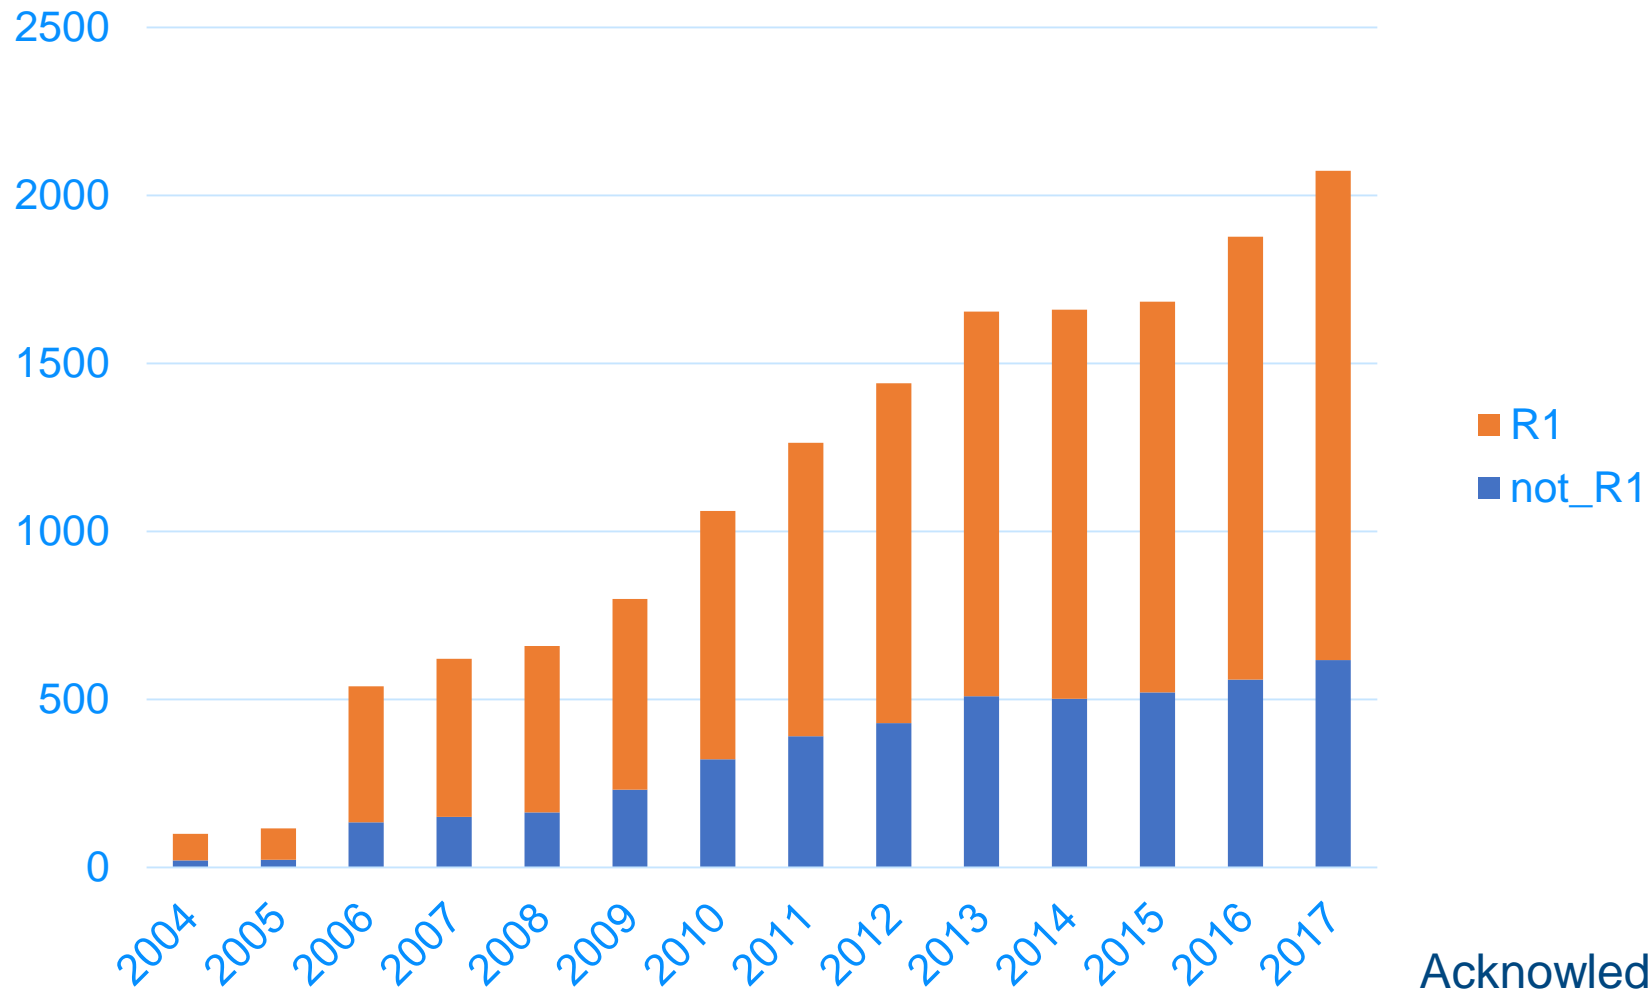
# Usage by R1s and non-R1s, in Normalized Units



Acknowledgement: Dave Hart

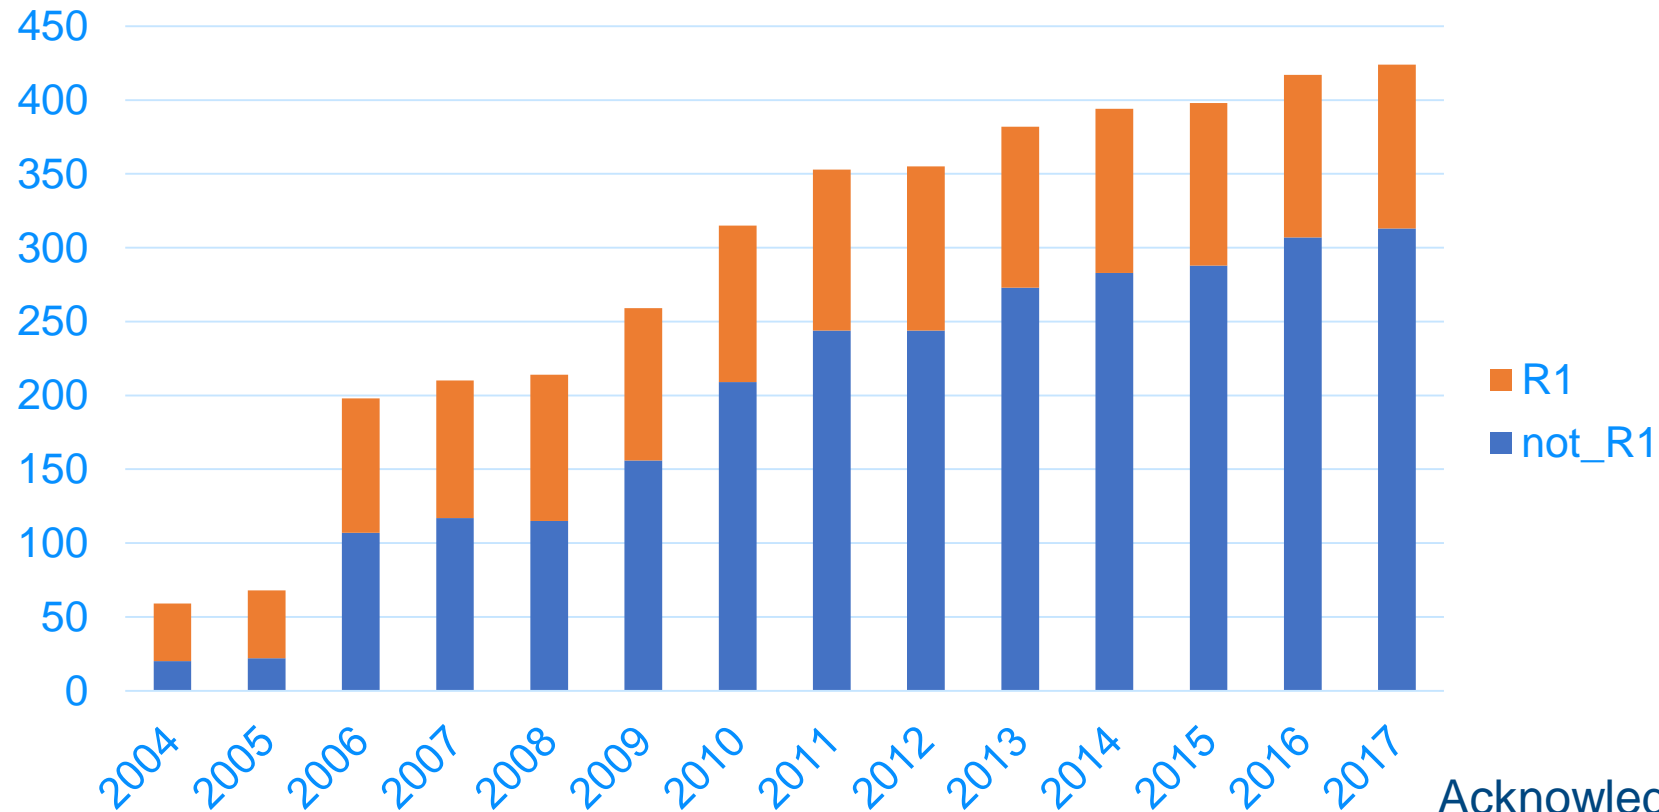


# Allocated PIs from R1s and non-R1s



Acknowledgement: Dave Hart

# Number of R1 & non-R1 Institutions Represented in Allocations



Acknowledgement: Dave Hart

# XSEDE Programs and Activities

Provides the expertise to ensure that researchers can make the most of the supercomputers and tools, including:

- [Extended Collaborative Support Service](#) (ECSS) – team experts with researchers
- [XSEDE Cyberinfrastructure Integration](#) (XCI) - provides an advanced hardware and software architecture that allows for individualized user experiences
- [XSEDE User Portal](#) web interface allows users to monitor and access XSEDE resources, manage jobs on those resources, report issues, and analyze and visualize results.
- Coordinated allocations of NSF's high-end resources and digital services, through the XSEDE Resource Allocation Service (RAS) team.
- A powerful and extensible network, maintained by the [XSEDE Operations](#) team, with fast connections to the XSEDE [Service Providers](#)
- Training, Education, and Outreach through the Community Engagement & Enrichment (CEE) team that expand participation in XSEDE-based projects, curriculum development, and traditional training opportunities.
- Advanced support for novel and innovative projects.
- A fellowship program that brings [Campus Champions](#) to work closely with XSEDE advanced user support staff.

<https://www.xsede.org/about/what-we-do>

# XSEDE Training

Training is available in a variety of formats, including multicast, webinars, online training, and in person workshops. Suggestions for new topics are encouraged via the [feedback](#) form. For more information, see:

- [XSEDE Training Overview](#) for a summary guide of materials available
- [XSEDE Training Course Catalog](#) including listings across formats and sites
- [Course Calendar](#) with registration for upcoming training courses
- [Online Training](#) on materials relevant to XSEDE users
- Badges are available
- Roadmaps are in development

Training materials focus on systems and software supported by the XSEDE Service Providers, covering programming principles and techniques for using resources and services. Training classes are offered in high performance computing, visualization, data management, distributed and grid computing, science gateways, and more.



# XSEDE Monthly Workshop Series

- Rotating (Nuts and Bolts) Topics
  - MPI, OpenMP, OpenACC, Big Data, Summer Boot Camp
- Up to 25 satellite sites per session
  - 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](mailto:tmaiden@psc.edu)

Acknowledgement: Tom Maiden

**XSEDE**

# XSEDE Monthly Workshop: Big Data

Day 1		Day 2
11:00	Welcome	Machine Learning: Recommender System with Spark
11:25	Intro to Big Data	
12:00	Hadoop	
12:30	Intro to Spark	
1:00	Lunch	Lunch
2:00	Spark	Deep Learning with Tensorflow
3:30	Spark Exercises	
4:30	Spark	Bridges: A Big Data Platform
5:00	Adjourn	Adjourn

Acknowledgement: Tom Maiden

# XSEDE Training Providers

This group	Offers	On	To	Approx
Training	Async online modules	Many HPC topics	Everyone	Ongoing
Training, SP	Webcast	Getting started	Everyone	Quarterly
Training	Multicast live training	HPC Topics	Satellite sites	Monthly
Broadening Participation	On-site training	HPC and XSEDE	Underrepresented institutions	Quarterly (academic year focus)
ECSS ESTEO	Webcast and live training	New resources, new capabilities	XSEDE Staff	Annual series
ECSS ESTEO	All formats	Many HPC topics	Everyone	Ongoing
Service Providers	On-site and webcast	Local resource	Everyone	
ECSS, SPs	Webcast	Varied	Campus Champions	Quarterly

Events and registrations: <https://portal.xsede.org/course-calendar>



# ESTEO

- Mentor Campus Champion Fellows
- Deliver live training events especially in collaboration with CEE Broadening Participation
- Orchestrate ECSS internal staff training seminars  
[https://www.xsede.org/documents/10165/1589367/RDMA\\_Spark\\_Hadoop.pdf](https://www.xsede.org/documents/10165/1589367/RDMA_Spark_Hadoop.pdf)
- Respond to requests for service  
<http://www.calendar.gatech.edu/event/595496>
- Review training modules, Retire training modules
- Support Meetings and BOFS
- Mentoring
- Review Education Allocation Proposals
- Support Campus Champions
- Manage US Participation in International HPC Summer School

Acknowledgement: Jay Alameda

# Education

<https://www.xsede.org/community-engagement/educator-programs>

## Curriculum and Educator Programs

XSEDE pursues innovation and collaboration in computational science education.

### Campus Visits

XSEDE campus visits emphasize the need for education and offer guidance concerning course

Campus visits bring together faculty, students discuss the importance of having a workforce

### Participating in Collaborative Online Courses

The XSEDE courses consist of recorded lectures that can be watched by students independently or in their own local classrooms. Each lecture comes with built-in quizzes that are used as part of the grading for the course. In addition, several computer exercises are typically available that students can run on XSEDE computational resources to gain practical experience and have credit for the work recorded in their class grade.

The capstone assignment for a course often is a final project supervised by the local faculty members.

The first of these courses is Applications of Parallel Computers, taught by Jim Demmel at the University of California, Berkeley. View the [course content](#).

### How Faculty Can Participate

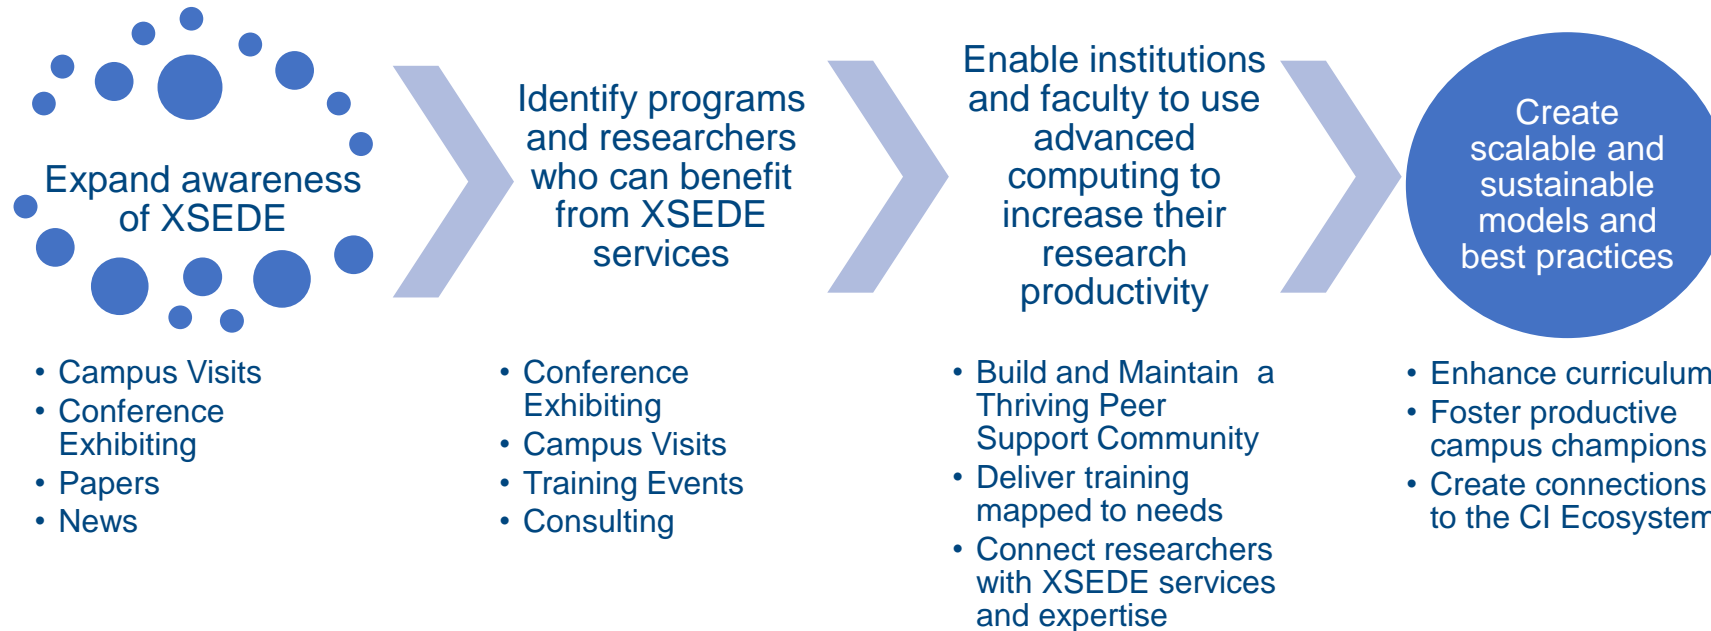
If you are a faculty member interested in collaborating with XSEDE in this program, you will need to create a course in your own academic schedule that your local students can register for and receive credit. You and your students will then use the online materials and XSEDE resources to complete the course.

Local faculty who participate in the program meet periodically with the XSEDE instructors and staff to discuss schedules, suggestions for course improvement, and any questions related to operations. The local faculties are responsible for assigning final grades to all of their own students.

*XSEDE sponsors full-semester online courses taught through collaborations with faculty at participating institutions.*

**XSEDE**

# CEE: Broadening Participation Program



Acknowledgement: Linda Akli

**XSEDE**

# Cyberinfrastructure Resource Integration

- Software toolkits, consulting services, provider coordination
  - Cluster distribution, scientific software, XSEDE-like environment
  - Site visits to help install
  - Information and support for joining the XSEDE federation
- Impacts
  - 7 clusters and more than 700 TeraFLOPS of computing on CRI software
  - New data analytics program supported at Bentley University, short video at <http://bit.ly/xsede-ba>

Acknowledgement: Rich Knepper

**XSEDE**

# Campus Champions

The Campus Champions Program is a group of 400+ Champions at 200+ US colleges, universities, and other research-focused institutions, whose role is to **help researchers at their institutions to use research computing**, especially (but not exclusively) large scale and high end computing.

The Campus Champions started as a way to drive users to XSEDE TeraGrid machines, but over time we've become a **national community of practice in research computing facilitation**.

## What is a Campus Champion?

A Campus Champion is an **employee of, or affiliated with, a college or university (or other institution engaged in research)**, whose role includes helping their institution's researchers, educators and scholars (faculty, postdocs, graduate students, undergraduates, and professionals) with their computing-intensive and data-intensive research, education, scholarship and/or creative activity, including but not limited to helping them to use advanced digital capabilities to improve, grow and/or accelerate these achievements.

## Who are the Campus Champions?

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

<https://www.hpc.msstate.edu/publications/docs/2017/01/1517907816979.pdf>

Acknowledgement: Marisa Brazil

**XSEDE**

# Questions?

Final comments:

- Student programs at PEARC:  
<https://www.pearc18.pearc.org/student-program>
- Many opportunities announced through the XUP; sign up for an account at <portal.xsede.org>

Acknowledgements:

- On individual slides
- [www.xsede.org](http://www.xsede.org)
- <portal.xsede.org>