

XSEDE For Extreme Computing in Science

Mats Rynge

XSEDE Extended Collaborative Support (ECSS)

Workflow Applications Team

XSEDE

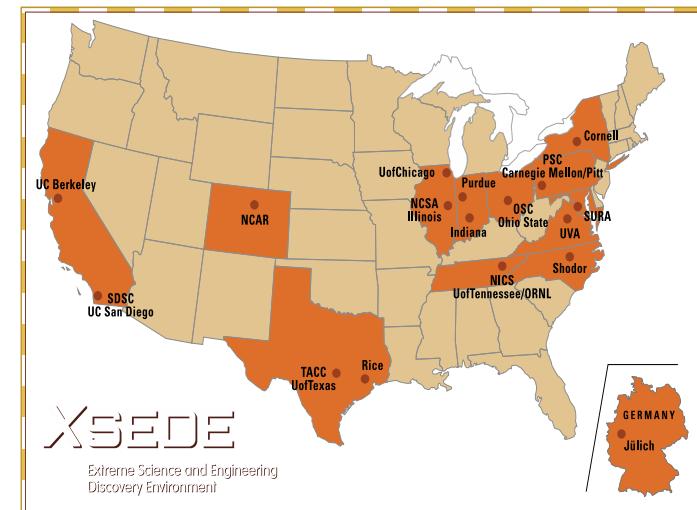
Extreme Science and Engineering
Discovery Environment



What is XSEDE?

eXtreme Science and Engineering Discovery Environment

- XSEDE is a NSF-funded, advanced, distributed open cyberinfrastructure, consisting of 9 supercomputers, 1 high throughput computing system, 3 high-end visualization systems, and 9 storage systems provided by 17 partner institutions. XSEDE also consists of fast networking, data collections, and science gateways. A defining property of XSEDE is a set of unified policies and programs.
- A single virtual system that scientists can use to interactively share computing resources, data, and expertise.
- Access is allocations based



XSEDE

HPC Resources



Blacklight - Blacklight is a SGI shared memory system intended for applications that require a large shared memory for computational tasks.



Gordon - Gordon is a unique, a flash-based supercomputer designed for data-intensive applications.



Keeneland - Keeneland is a balanced hybrid CPU/GPGPU system for use with codes that can take advantage of accelerator performance. Keeneland has over 200 nodes with 3 GPUs per node making it an excellent resource for scaling highly parallel codes.



Mason - Mason at Indiana University is a large memory computer cluster configured to support data-intensive, high-performance computing tasks using genome assembly software.



Stampede – A 10 PFLOPS Dell Linux Cluster based on 6,400+ Dell PowerEdge server nodes, each outfitted with 2 Intel Xeon E5 (Sandy Bridge) processors and an Intel Xeon Phi Coprocessor (MIC Architecture).



Trestles - Trestles employs flash-based memory and is designed for modest-scale research providing very fast turnaround time. It is intended for moderately scalable parallel applications with an emphasis on improving productivity for a broad spectrum of users. Trestles is ideal for applications with fast local I/O requirements that can benefit from the flash memory available on each compute node.

HTC Resources



Open Science Grid - The Open Science Grid (OSG) advances science through open distributed computing. OSG supports XSEDE users by providing a Virtual Cluster that forms an abstraction layer to access the opportunistic cores in the distributed OSG infrastructure.

Visualization Resources



Maverick - The Maverick Hewlett-Packard/NVIDIA Interactive Visualization and Data Analytics System is intended primarily for interactive visualization and data analysis jobs. Maverick's fast data movement and advanced statistical analysis capabilities allow for interactive query of large-scale data sets

Storage / Networking / Services / ...

<http://info.xsede.org> – Information services

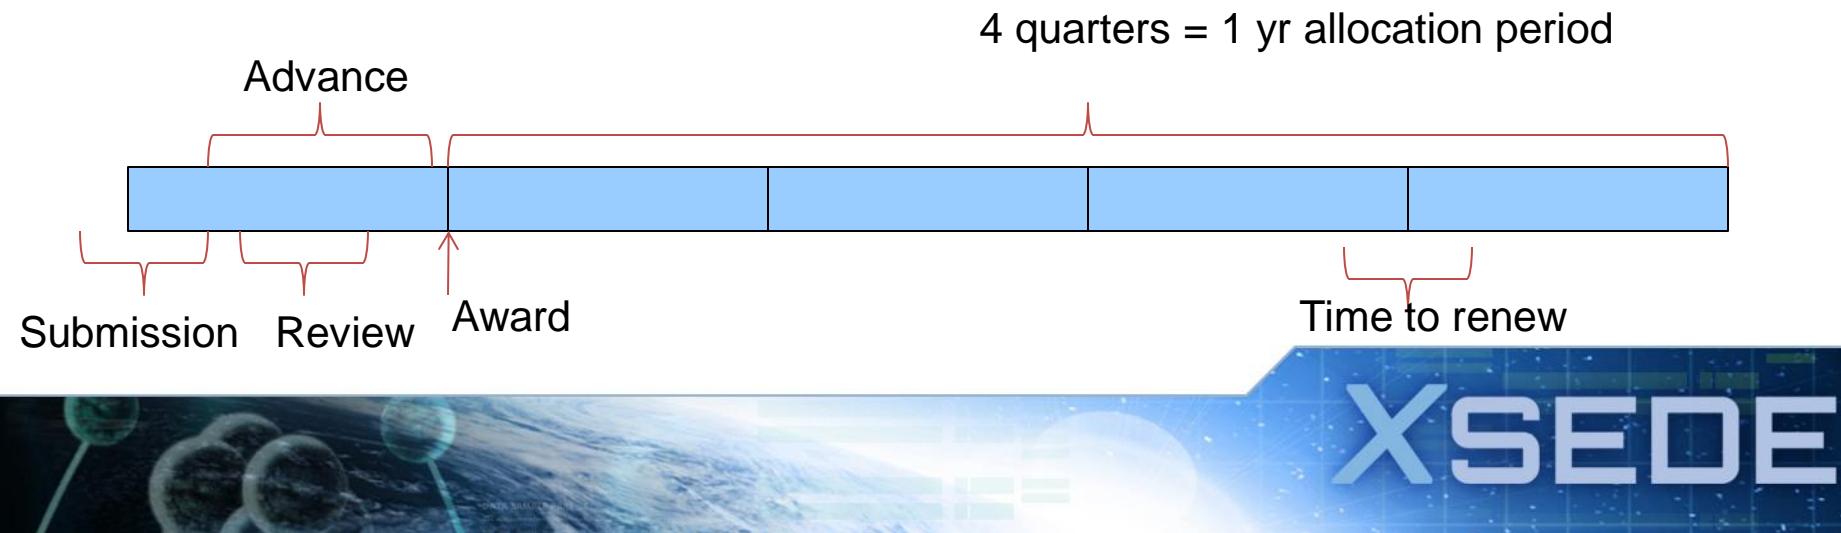
GRAM, UNICORE – Remote job execution

GridFTP, Globus Online – Data management

XSEDE

Access / Allocations

- Open to any US open science researcher (or collaborators of US researchers) regardless of funding source.
- To get started using XSEDE a researcher needs to:
 - apply for an XSEDE allocation (startup / research allocation)
 - request to be added to an existing one.
- One per PI (generally)
- 1-year duration
- Unused SUs are forfeited at the end of an award period
- Progress report required for renewal requests.
- Add users to a grant via XSEDE User Portal



Extended Collaborative Support Services

- Get expert staff as well as computer allocations
 - 37 FTEs, spread over ~80 people at almost a dozen sites.
 - Many PhDs or equivalent experience
- Wide range of areas,
 - Performance analysis
 - Visualizations
 - Petascale optimization techniques
 - Novel and Innovative Projects
 - Building Science Gateways.
 - Supporting Scientific Workflows



ECSS Workflows Team

To better enable challenging scientific research on XSEDE through the use of scientific workflow technologies

The XSEDE Workflow Community Applications Team's charter is to assist researchers to use scientific workflow technologies on XSEDE to solve challenging **scientific problems** involving parameter sweeps, multiple applications combined in dependency chains, tightly coupled applications, and similar execution patterns that **require multiple applications and multiple XSEDE resources**. The workflow team accomplishes its mission through the use of **third party workflow software** in collaboration with the workflow developers, service providers and XSEDE Extended Collaborative Support Services.

Science Gateways

A Science Gateway is a community-developed set of tools, applications, and data that are integrated via a portal or a suite of applications, usually in a graphical user interface, that is further customized to meet the needs of a specific community. Gateways enable entire communities of users associated with a common discipline to use national resources through a common interface.

<https://www.xsede.org/gateways-overview>



Thanks for listening!

Our reach will forever
exceed our grasp, but,
in stretching our horizon,
we forever improve our world.

XSEDE

Extreme Science and Engineering
Discovery Environment