


Cornell Center for Advanced Computing (CAC)



Services for Cornell faculty, staff & student researchers

- Grant proposal preparation
- Technical consulting
 - Software development
 - DB development and data management
 - Optimization and workflow
 - Cloud and container help
- Computational support
 - Cloud computing
 - Cluster management
- Leased and archival storage services
- Training



Who, what


- Susan Mehringer – shm7@cornell.edu
 - CAC Associate Director, Consulting
- Funded roles in national supercomputing
 - XSEDE Training Lead
 - Frontera training partner
 - Stampede2 training partner
 - Jetstream training partner
- Today's talk
 - CAC, XSEDE, Frontera resources
 - Pointers to RDMSG, public clouds via CIT
 - Won't cover many other compute resources available from specific CU depts. and public agencies: <https://research.cornell.edu/content/centers-institutes-programs-laboratories>



Jetstream

CAC resources

- Red Cloud
 - Instances up to 28 cores & 240GB RAM, 4 NVIDIA Tesla v100 GPUs and 4 NVIDIA T4 GPUs
 - 10 Gigabit network interconnect
 - Persistent disk storage (volumes) backed by Ceph storage with more than 1 petabyte raw capacity
 - Subscription model > Exploratory accounts available
- Private cluster management
 - House your HPC cluster in Rhodes Hall
 - CAC provides systems maintenance, i.e., software updates, server and network maintenance, power, cooling, etc.
 - Can help faculty architect the optimal cluster, servers, storage, and/or systems software, and, if procuring a cluster, to secure favorable pricing at no fee. Cluster installation is a fee-based consulting service.
- Computational consulting
 - Computation on resources above or elsewhere
 - Software development, DB development & data management, optimization & workflow, cloud computing, containerization ...



CAC rates


Rates for CAC computing and consulting services requiring non-trivial effort or resource utilization are based on the requirement of full cost recovery with a Provost subsidy, reviewed and adjusted annually, to ensure equitable recovery of CAC costs, and no over-recovery.

No fees for the following services

- Initial strategy meeting to discuss which computing and consulting services will best meet the research, education, or outreach components of your grant proposal or current project
- General help to get you started or answer the occasional question
- Exploratory accounts to try out Red Cloud; an exploratory account includes one hour of consulting, 165 compute hours on Red Cloud, and 50GB of storage
- Access to high-speed networks from CAC-supported services
- Up to 20 hours consulting to architect the optimal HPC cluster and storage system to be housed in our Rhodes Hall machine room and maintained by our staff. We will also use our experience, expertise, and vendor relationships to secure favorable quotes for your cluster procurement.

RDMSG

The Research Data Management Service Group (RDMSG) is a collaborative, campus-wide organization that assists with creating and implementing data management plans, applying best practices for managing data, and finding data management services at any stage of the research process.



Cornell access to public cloud services

- Cornell has agreements with Amazon Web Services (AWS) and Microsoft Azure; many infrastructure services have been moved to this platform
- For added convenience, the service is integrated into Cornell's system
 - Use your NetID password to log in
 - Collaborate with other community members using a common service
- Cornell negotiates institution-wide prices and terms (legal, policy, supply management, audit, and security) to provide the most benefit to our community

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/>

<https://portal.xsede.org/>

Research computing resources available

Location	Name	Description	CPUs	Peak Tflops	File Space (TB)
TACC	Stampede2	Dell/Intel Knights Landing, Skylake System	368,280	12,800.0	
TACC	Wrangler	Long-term Storage			10,000
TACC	Ranch	Long-term tape Archival Storage			61,440
SDSC	Comet	Dell Cluster with Intel Haswell Processors	46,752	2,000.0	
SDSC	Comet GPU	GPU Nodes	1,728	104.0	
SDSC	Data Oasis	Medium-term disk storage			4,000
PSC	Bridges	Regular Memory	21,056	894.6	
PSC	Bridges GPU	Bridges GPU	1,344	894.6	
PSC	Bridges GPU-AI	Bridges GPU-Artificial Intelligence	16	9.9	
PSC	Bridges Large	Large Memory Nodes	160	894.6	
PSC	Bridges Pylon	Storage			10,000
IU/TACC	Jetstream	Cloud, VMs, HTC	15,360	516.1	
IU/TACC	Jetstream Storage	Storage			960
OSG	OSG	Open Science Grid, HTC	60,000	50.0	
LSU	superMIC	Cluster	7,200	925.0	

Which is the right resource for your work?

- Software availability: <https://portal.xsede.org/software#/>
- Resource details & user guide: <https://portal.xsede.org/resource-monitor>
- Recommended use: <https://portal.xsede.org/allocations/resource-info>

PSC Regular Memory (Bridges)



PSC

Production Dates: 2016-01-01 - 2020-11-30

Startup Allocation Limit: 50000 Service Units (1 Service Unit = 1 Core Hour)

Description: Regular Shared Memory nodes each consist of two Intel Xeon E5-2695 v3 CPUs and 128GB of 2133 MHz DDR4 RAM configured as 8 DIMMs with 16GB per DIMM.

Recommended Use: Regular Shared Memory nodes provide substantial capacity, especially for jobs such as optimizations and parameter sweeps that require vast numbers of separate runs or threads, support interactivity and dedicated access, power complex workflows, and support the Hadoop ecosystem. MPI codes are well-supported on the RM nodes, with the best performance up to 1176 cores.

Allocations

Eligible: Almost all U.S.-based university and non-profit researchers

Cost: none

Available to:

- Experienced HPC users with large-scale needs
- New users in computational science, visualization, or data analysis
- Those deploying a science gateway
- Support academic courses and training activities

Types of allocations:

- Campus Champion
- Startup
- Research
- Education

<https://portal.xsede.org/allocations/announcements>

Support

Training:

<https://www.xsede.org/for-users/training>

Curriculum and Educator Programs:

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

Extended Collaborative Support Services:

<https://www.xsede.org/for-users/training>

Cyberinfrastructure Integration:

<https://www.xsede.org/ecosystem/ci-integration>

Info/Contact

General info

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

User docs

<https://portal.xsede.org/>

Allocations

<https://portal.xsede.org/allocations/announcements>

Submit a ticket

<https://portal.xsede.org/web/xup/help-desk>

Campus Champions

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



FRONTERA SYSTEM --- PROJECT

- ▶ A new, NSF supported project to do 3 things:
- ▶ Deploy a system in 2019 for the largest problems scientists and engineers currently face.
- ▶ Support and operate this system for 5 years.
- ▶ Plan a potential phase 2 system, with 10x the capabilities, for the future challenges scientists will face.



FRONTERA SYSTEM --- HARDWARE

- ▶ Primary compute system: DellEMC and Intel
 - ▶ 35-40 PetaFlops Peak Performance
- ▶ Interconnect: Mellanox HDR and HDR-100 links.
 - ▶ Fat Tree topology, 200Gb/s links between switches.
- ▶ Storage: DataDirect Networks
 - ▶ 50+ PB disk, 3PB of Flash, 1.5TB/sec peak I/O rate.
- ▶ Single Precision Compute Subsystem: Nvidia
- ▶ Front end for data movers, workflow, API



WHO CAN USE FRONTERA?

- ▶ You are eligible to apply for an allocation on Frontera if:
 - ▶ You are doing unclassified research at a US university.
 - ▶ Preferences to funded, and NSF-funded work.
 - ▶ A smaller amount of time is available for private sector or FFRDC users.
 - ▶ Allocation proposals should come from “PI-eligible” individuals.
- ▶ You may be a user of Frontera if:
 - ▶ You are at a US University, Company, or FFRDC and are added to an allocation by the PI.
 - ▶ You are at a non-US institution, actively collaborating with a US researcher with an allocation, and not in a country subject to US export control restrictions on supercomputing technology.



GETTING ACCESS TO FRONTERA

<https://fronteraweb.tacc.utexas.edu/allocations/> ← more info and submission windows

In general, proposers must show that multi-petascale compute resources are required and can be effectively exploited. Most require current peer-reviewed research funding to support the activities conducted on Frontera. Four allocation tracks:

Leadership Resource Allocation (LRAC) –

- Large allocations, require strong scientific justification, enable research that would otherwise not be possible
- Must be ready to use the allocated cycles
- Submissions 1-2 times per year, depending on availability of compute time

Pathways

- Small allocations to science teams with a strong scientific justification
- Use to scale codes using existing research funding for code and algorithm development
- Early career scientists and engineers encouraged to apply
- Submissions quarterly

Large-Scale Community Partnerships (LSCP)

- Extended allocations of up to three years to support long-lived S&E experiments

Director Discretionary Allocations (DD)

- For projects that don't fit well into the three tracks described above, such as: areas of urgent need, educational usage, and industrial collaborations/research.
- Submissions will be accepted on a rolling basis.

Additionally, users may request archive space for their data on TACC's tape storage system.



SYSTEM SUPPORT ACTIVITIES

Support

- Extended scientific consulting support
- Training (multiple formats), online documentation
- 24x7 ticket support
- Multiple queues, reservations, etc
- Shared Work Filesystem, Archive access
- Comprehensive software stack

Support +

- Auto-tuned MPI stacks
- Automated Performance Monitoring, with data mining to drive consulting
- Slack channels for user support

Support + +

- Full Containerization support
- Support for Controlled Unclassified Information (i.e. Protected Data)
- Application servers for persistent VMs to support services for automation



Info/Contact

General Info

<https://www.tacc.utexas.edu/systems/frontera>

“Welcome to Frontera March 20, 2019” on YouTube

User docs

<https://portal.tacc.utexas.edu/user-guides/frontera>

Allocations

<https://fronteraweb.tacc.utexas.edu/allocations/>

Submit a ticket

<https://portal.tacc.utexas.edu/> (login to see link)

