

New Faculty
Flyer

Research Computing

**Cornell Center for Advanced
Computing (CAC)**

Research Computing Services

We enable your success

Rhodes Hall

Table of Contents

Getting Started 1


Technology Leadership 2

Faculty Benefits 2

Process 3

Services 4

Grants 5



Cornell CAC


www.cac.cornell.edu

Welcome to Cornell

World-class experts like you are what makes Cornell a premier research institution. We're glad you're here! This flyer will help you understand what research computing services are available for faculty at Cornell.

Getting Started


Research—it begins with a big idea, a fleeting insight, or perhaps even serendipity. Whatever the spark, it's up to you to build a world-class collaborative team and find the services you need to make your idea a reality. The Cornell Center for Advanced Computing (CAC) provides professional research computing services to Cornell faculty in the sciences, engineering, business, arts, and humanities. Our expert systems and consulting staff include PhD-level computational consultants with expertise in many fields, from astronomy to biology, computer science, informatics, physics, and more.



Free CAC services include assessing your research needs, grant proposal strategy meetings, general help to get you started or answer the occasional question, exploratory accounts to try out Red Cloud, access to high-speed networks, and up to 20 hours consulting to architect your optimal HPC cluster and storage system that will be housed in Cornell's professionally managed machine room.

Computing and consulting services that require non-trivial effort are fee-based and require a Cornell account for charge-back based on the amount of services used. New faculty may use start-up packages to procure CAC services, request departmental funds, or include CAC services in grant proposals (for example, under the budget line "Consulting Services" in NSF proposals).

Technology Leadership



CAC has a history of computational firsts. We deployed the first IBM SP supercomputer, the first Dell Top500 supercomputer, the first federated cloud funded by the NSF in collaboration with researchers such as the College of Engineering’s Pat Reed (*left*), and the first large-scale MATLAB cluster that achieved a 175x faster computation of the CDC’s Hepatitis C model (*right*).

Today, we continue to test and deploy the latest computing technologies, software, and tools to keep Cornell faculty at the cutting edge. We’re containerizing applications so that researchers can run them on any platform they choose—supercomputers, Cornell’s Red Cloud, or public clouds. And, we’re housing and maintaining more than 18 faculty clusters with nearly 5 petabytes data storage capacity.

Faculty Benefits





By using CAC services, you can focus on your research goals rather than on keeping up with the latest and ever-changing computing technologies.

Our consultants will work with you and your research team to write programs, improve code, design databases, and perform any other short- or long-term services you need to enhance your team’s capabilities and achieve project milestones faster.

CAC data management and storage services, and well-designed pipelines and websites enable us to explore large data sets fast.

Shami Chatterjee, Astronomy

Process

Analyze	Plan	Deliver	Maintain
			
CAC professionals work closely with Cornell faculty and research staff to analyze project requirements.	Faculty or research staff select levels of computing, storage, and/or hourly consulting services desired.	Consulting or systems services are delivered. All services are accounted for monthly.	Quality, turnkey maintenance is available for HPC clusters, storage systems, websites, and databases.

Featured Service: Red Cloud



One of the more popular services we offer is Red Cloud, the University's on-premise cloud. Red Cloud is subscription-based to protect Cornell faculty from billing surprises that can occur with public clouds. Users can request instances with up to 128 cores and 240GB RAM. NVIDIA GPUs are also available. Persistent disk storage volumes are backed by Ceph storage with up to 1.9 petabytes of raw capacity. Unlike public clouds, CPU cores and RAM are not oversubscribed or shared with other users. Users have full access to the underlying physical processor and memory without competition, resulting in fast, consistent performance. Users can manage their cloud service using a Web console, command line clients, or any development library supporting the OpenStack API. Virtual clusters are available.

"Our goal is to increase your research productivity and accelerate discovery."

Rich Knepper, CAC

We're eager to meet you and learn more about your research computing needs. We offer a full range of services which are highlighted on the next page. Detailed descriptions of our services along with our technical documentation Wiki are available at www.cac.cornell.

Services

HPC & SERVER MAINTENANCE

CAC's professional systems staff will architect, house and maintain your private HPC cluster or servers so you can focus on your research.

**Turnkey Maintenance
Software Updates
Networking, Power & Cooling**

CLOUD COMPUTING

Cornell's Red Cloud is a subscription-based cloud that provides root access to virtual servers and storage on-demand.

**Up to 128 Cores and 240GB RAM per instance; plus NVIDIA GPUs
Virtual Clusters & Ceph Storage**

CLOUD APPLICATIONS

We'll build your ready-to-use cloud image and containerize your application for research efficiency and portability.

**Docker & Singularity
Portability to Clouds or HPC
Cloud-Based Web Applications**

DATA STORAGE

CAC offers storage services for research data with fast, no fee transfers in and out.

**Object Storage & Leased Storage
Globus Online Data Transfer
Archival Storage**

RESEARCH WEBSITES

Websites and research portals with custom capabilities such as databases, custom tools, and large-scale storage systems.

**Research Website Design
Hosting
Maintenance**

DATABASES

We design and implement research databases and deploy and operate database servers with robust performance.

**Database Server Capacity Planning
Database & Workflow Design
Data Visualization & Management**

PROGRAMMING & AI

We program and fine tune codes in C/C+, C#, Java, MATLAB, MPI, OpenMP, Perl, Python, R, etc. and can help you get started with AI.

**Parallel Performance Tuning
Code & Cache Usage Optimization
Performance Analysis & Debugging**

EDUCATION & OUTREACH



We produce online training on any subject for broader impact (you provide the expert, we do the rest)

**Guest Lectures & Courses
Webinars & Training Events
eCornell: Tableau, Python, DBs**

Grants

We collaborate with Cornell faculty in Ithaca, NYC and Qatar on grant proposals serving as PI, co-PI, senior personnel, and/or as a service provider. Contact us to explore how we might make your grant proposals more competitive and milestones easier to achieve. We also collaborate with faculty on journal publications.

Sample Awards

- CAC senior researcher Steve Lantz is providing investigative services on multi-core and many-core technologies to Peter Wittich, PI for the *Institute for Research and Innovation in Software for High-Energy Physics (NSF)*.

- CAC senior research associate Chris Myers was co-PI on a Scott McArt grant that investigated how pathogens infect bees and identified interventions to mitigate disease risk (*NIH grant; research continues with new USDA and NSF grants*).
- CAC was PI for an \$8.2.M *Data Analysis and Building Blocks for Multi-Campus Cyberinfrastructure through Cloud Federation* grant with multiple Cornell use case scientists, including Sara C. Pryor and Patrick Reed; finance and digital ag researchers from UB and UCSB; and, 15 REUs (*NSF*).
- CAC associate director Susan Mehringer and her consulting team are developing *Frontera Virtual Workshops* to provide online training to national supercomputing users (*Texas Advanced Computing Center*).

- CAC computational scientist Adam Brazier is Cyberinfrastructure Lead for the *North American Nanohertz Observatory for Gravitational Waves (NANOGrav)* consortium and Jim Cordes is co-PI. The goal of the consortium is to detect long-wavelength gravitational waves and individual binary supermassive black holes (*NSF*).

We have two easy ways to learn more:

Visit us on the web: www.cac.cornell.edu

*Call or email CAC director Rich Knepper
to discuss your needs*



*CAC is providing data workflows
and managing data for the NANOGrav
Physics Frontiers Center through 2026*

Cornell Center for Advanced Computing

722 Rhodes Hall
Ithaca, NY 14853-3801

Rich Knepper, PhD
Cell (812) 361-0690
Email: rich.knepper@cornell.edu

www.cac.cornell.edu



© 2023 Cornell University