

New Faculty  
Flyer

# Research Computing

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

We enable your success

*Professional research computing and software services*

*Rhodes Hall*

# Table of Contents

**Getting Started** ..... 1

**Technology Leadership** ..... 2

**Faculty Benefits** ..... 2

**Process** ..... 3

**Services** ..... 4

**Grants** ..... 5

**Cornell CAC**

We enable your success


*[www.cac.cornell.edu](http://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 Center for Advanced Computing (CAC) provides professional research computing and software 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 research needs assessments, 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 housed in our machine room.

Professional 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 faculty at the cutting edge. For example, 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 maintaining 31 clusters in the University’s machine room for faculty research teams with 4.5 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 computing technologies that enable it.

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

*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 online. | 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 28 cores and 224GB RAM. NVIDIA GPUs are also available. Persistent disk storage volumes are backed by Ceph storage and offer more than 1 petabyte raw capacity. Linux and Windows Server operating systems are supported. Unlike many cloud platforms, CPU cores and RAM are not oversubscribed—an Intel processor core is behind each core on the virtual server for fast and consistent performance. Users can manage their cloud service using a Web console, command line clients, or any development library supporting the OpenStack API.

*“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](http://www.cac.cornell).

# Services

---

## HPC & SERVER MAINTENANCE

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

**Server & Network Maintenance**  
**Software Updates**  
**Power & Cooling**

## RESEARCH WEBSITES

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

**Website Design**  
**Hosting**  
**Maintenance**

## CLOUD COMPUTING

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

**Up to 28 Cores & 224GB RAM**  
**NVIDIA GPUs**  
**Ceph Storage**

## 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**

## CLOUD APPLICATIONS

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

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

## PROGRAMMING

We program in C/C+, C#, Java, MATLAB, MPI, OpenMP, Perl, Python, R, Scala, etc. and fine tune codes.

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

## DATA STORAGE

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

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

## 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 service provider. Contact us to explore how we might make your grant proposals more competitive and achieving grant milestones easier. We also publish with faculty.

## 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 is co-PI on Jonathan Butcher's *Engineering Research Center for the Engineering of Emergent Biocomplexity Planning Grant (NSF)*.
- CAC director and Cornell Vice President and CIO David Lifka is PI for the *Data Analysis and Building Blocks for Multi-Campus Cyberinfrastructure through Cloud Federation* grant with Cornell use case scientists James Cordes, Angela Douglas, Sara C. Pryor, Patrick Reed, and finance and digital agriculture faculty at 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 technical lead for the *Scalable Cyberinfrastructure to support Multi-Messenger Astrophysics* project. MMA is an exciting new field of science that combines traditional astronomy with the brand-new ability to measure phenomena such as gravitational waves and high-energy neutrino particles that originate from celestial objects (NSF).





**We have two easy ways to learn more:**

*Visit us on the web: [www.cac.cornell.edu](http://www.cac.cornell.edu)*

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

*Analyzing Fast Radio Bursts  
(case study available)*

## **Cornell Center for Advanced Computing**

528 Rhodes Hall  
Ithaca, NY 14853-3801

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

[www.cac.cornell.edu](http://www.cac.cornell.edu)

