



RESEARCH COMPUTING AND DATA SERVICES

HIGH PERFORMANCE COMPUTING
AT THE HIGHEST PROVEN VALUE



Rosen Center for Advanced Computing

HIGH PERFORMANCE COMPUTING

The Rosen Center for Advanced Computing (RCAC) operates a significant shared cluster computing infrastructure developed over several years through focused acquisitions using funds from grants, faculty startup packages, and institutional sources. These "community clusters" are now at the foundation of Purdue's research cyberinfrastructure.

We strongly encourage any Purdue faculty or staff with computational needs to join this growing community and enjoy the enormous benefits this shared infrastructure provides.

Peace of Mind

Purdue IT's system administrators take care of security patches, attempted hacks, operating system upgrades, and hardware repair so faculty and graduate students can concentrate on research. Research support staff are available to assist you by providing consultation and software.

Low Overhead

Our data centers provide infrastructure, including networking, racks, floor space, cooling, storage, and power. Each cluster is built with a lifespan of five to six years, with free support for the life of the cluster. All of this is included with the Community Cluster Program and is available to you at no extra charge.

Cost Effective

We work with vendors to obtain the best price for computing resources by pooling funds from different disciplines to leverage greater group purchasing power. Partners have invested several million dollars in computational and storage resources through the Community Cluster Program since 2006.

Flexible

As a partner in a community cluster, you always have ready access to the capacity you purchase, with potential access to much more. The Community Cluster Program shares compute nodes among cluster partners when the nodes are idle. This way, you get more computational value per dollar than is feasible as an individual.

Additional Benefits

- **Parallel Filesystem:** Access to large-scale, high-performance, parallel scratch for running jobs
- **Archive:** Access to the high-performance HPSS Archive system "Fortress" for long-term storage of research data
- **Research Data Depot:** High-performance, expandable space is available to any research group to:
 - Share data and results among your group or with collaborators using Globus transfer service
 - Centrally install and manage the group's applications
 - Define and manage access to custom UNIX groups for easy project-based collaboration
- **Cloud Lab Folders:** Centralize your lab's documents and collaborate in a managed folder utilizing **Box.com**.
- **Version Control:** Self-managed Purdue-hosted Github repositories for documents and source code
- **Remote Desktops:** Access community cluster systems via user-friendly Thinlinc Remote Desktop connections.
- **Notebooks:** Work in Python notebooks on cluster resources for reproducible, shareable data analysis.
- **Open OnDemand:** Access clusters, files, and applications from your browser using Open OnDemand.

COMMUNITY CLUSTER RESOURCES

1

GAUTSCHI

Gautschi is a Community Cluster optimized for communities running traditional, tightly-coupled science and engineering applications. It consists of Dell compute nodes with two 96-core AMD Epyc "Genoa" processors (192 cores per node) and 384 GB of memory. All nodes have 200 Gbps NDR Infiniband interconnect and service through 2030.

2

GILBRETH

Gilbreth is a throughput-optimized Community Cluster designed for AI applications such as machine learning. Gilbreth consists of Dell compute nodes with Intel Xeon processors and Nvidia GPUs.

3

GEDDES

Geddes is a Community Composable Platform optimized for composable, cloud-like workflows that are complementary to the batch applications run on Community Clusters. Geddes consists of Dell Compute nodes with two 64-core AMD Epyc 'Rome' processors (128 cores per node).

ADDITIONAL RESOURCES

▪ Anvil: National resource provider

Anvil is a NSF-funded national advanced cyberinfrastructure resource operated by RCAC serving thousands of researchers across the U.S. Anvil consists of 1,000 nodes with two 64-core AMD Epyc "Milan" processors each, and delivers over 1 billion CPU core hours through ACCESS (formerly XSEDE) each year, with a peak performance of 5.1 petaflops. Anvil also has GPUs, large memory nodes, and a cloud subsystem. Access to Anvil is free and coordinated by the ACCESS allocation process.

▪ Negishi Interactive:

The interactive tier on our Negishi cluster provides entry-level access to high performance computing. This includes login to the system, data storage on our high-performance scratch filesystem, and a small allocation that allows jobs submitted to an "interactive" account limited to a few cores. This subscription is useful for getting workloads off your personal machine, integrated with more robust research computing and data systems, and a platform for smaller workloads. Transitioning to a larger allocation with priority scheduling is easy and simple.

▪ Scholar:

The Scholar cluster is open to Purdue instructors from any field whose classes include assignments that could make use of supercomputing for modeling or data science, from high-end graphics rendering, weather modeling, simulation of millions of molecules, and exploring masses of data to understand the dynamics of social networks.

Coffee Hour Consultations

Coffee Hour Consultations are excellent opportunities in a casual setting to consult and discuss computing questions with RCAC computational scientists.



RESEARCH DATA SOLUTIONS

RCAC maintains several different storage resources to accompany computational systems.




Data Depot: The Data Depot is a high-capacity, fast, reliable, and secure data storage service designed, configured, and operated for the needs of Purdue researchers in any field and shareable with both on-campus and off-campus collaborators.



Fortress: The Fortress system is a large, long-term, multi-tiered file caching and storage system that utilizes both online disk and robotic tape drives. Fortress has an impressive storage capacity of over 10PB and writes two copies of every file — either to two tapes or to disk and a tape — to protect against medium errors.



Scratch Storage: Each cluster is assigned a default parallel scratch filesystem, providing work-area storage optimized for a variety of job types. They are designed to perform optimally with data-intensive computations while scaling well to large numbers of simultaneous connections.



Software Defined Storage: Our central software defined storage resource provides access to on-demand provisioning of block, filesystem and object storage for modern, cloud native workflows on the Geddes Composable Platform and community clusters as well as readily accessible cold storage for large scale datasets.



Globus: Globus is a powerful and easy-to-use file transfer service for sharing files virtually anywhere. Using Globus, you can easily share data with collaborators worldwide with dropbox-like simplicity!



Purdue Data Resources: Purdue offers several storage solutions outside of Research Computing. One such offering is PURR, a research collaboration and data management solution for Purdue researchers and their collaborators.

Controlled Data



Export controlled data: RCAC provides resources for data and computation in support of projects with heightened security requirements. Research requiring protection for Export Control (EAR, ITAR), or Controlled Unclassified Information can all be performed.



REED Folder: A REED Folder is a managed storage solution built on top of the Box.com cloud platform for research projects requiring compliance with regulations or heightened security.



Protected data Filesystem: To support biomedical applications, RCAC provides data storage capabilities configured to hold external datasets such as the database of Genotypes and Phenotypes (dbGaP) and The Cancer Genome Atlas (TCGA), and the UK BioBank."

RESEARCH SOFTWARE ENGINEERING (RSE)

RCAC offers a wide range of research software engineering solutions to help advance scientific discovery, education, enable greater research reproducibility, and generate lasting impact at Purdue and beyond. RCAC's **Scientific Solutions**, **Envision Center** and web developers has comprehensive experience in software design, programming languages, and deployment platforms gives us the tools to help you create or modify software to take advantage of the latest technologies in advanced computation, web frameworks, data analysis, visualization, and management.

Find out more: <https://www.rcac.purdue.edu/rse>

Or ask us a question: RCAC-Help@Purdue.edu

HOW WE CAN HELP

Virtual Simulation

Digital training tools that allow for immersive experiences and learning platforms by utilizing Virtual Reality and Augmented Reality. These tools are great for use in research and the classroom, helping to bridge the gap between theoretical knowledge and practical application.

Multimedia Production

Graphics, animated visualizations, interactive web deployments, and mobile app development are a few of the many ways we can help improve your project. Multimedia productions suit a variety of output, from promotional and publication materials to classroom and research content.

Data Visualization & Analysis

From geospatial to tomographic, the Envision Center can assist in data handling, production, analysis, and tool creation for visualizing complex 3D data.

Website and Web Application Development

Our team of software engineers, project managers, and computational scientists can build science gateways, websites and applications for bringing your research to the web.

Research Software

With expert knowledge in a broad range of programming languages such as Python, C, R, and Matlab, our software engineers can assist with designing, developing, and optimizing research software.

Partner on Proposals

The mission of Purdue's ITaP Research Computing staff is to support and advance research discoveries at Purdue through partnerships with faculty, research groups, and centers. Our research scientists collaborate with faculty on grant proposals and support them by providing cyberinfrastructure solutions to funded projects.

TESTIMONIALS ABOUT OUR SERVICES



"The efficiency, power, and reliability of Bell and Negishi have been paramount. Their role in our success cannot be overstated, these clusters will definitely continue to play a central role in my research for many years to come."" - **Daniel Dawson, an associate professor of earth, atmospheric, and planetary sciences**, who used the community clusters to conduct high-resolution simulations of severe thunderstorms and tornadoes.



"When you're trying to do science, you don't want to spend a lot of time resolving technical issues. Having someone to help us with that was a tremendous benefit." - **Asif ud-Doula, associate professor of physics at Penn State Scranton**, who used Anvil to develop numerical tools to model the stellar winds of massive stars.



"For years, RCAC clusters have faithfully provide the advanced computing power necessary for cutting edge research in time domain astronomy. As our needs expand and as we increasingly employ the latest in machine learning techniques to analyze massive astrophysical data sets, so has RCAC's portfolio of computing options and stellar support." - **Danny Milisavljevic, an Associate Professor of Physics and Astronomy**, who uses RCAC Community Clusters to model and study the births and deaths of stars, amongst other astrological phenomenon.

SELF SERVICE TOOLS


Computational Science Expertise

In addition to the peace of mind gained from professional systems engineering staff, community cluster partners can draw from the expertise of Purdue IT's experienced computational scientists, software engineers, and visualization experts.

Purdue IT's computational scientists are experienced users of computational resources, with advanced degrees in Engineering, Big Data, Bioinformatics, Biology, Chemistry, and Physics. Computational science staff can help with a wide range of issues: from answering user questions and providing training, code development, software installation, designing effective workflows, and performance analysis. Additionally, research solutions engineers are available to consult on applying new technology solutions for science problems.

Track Your Usage

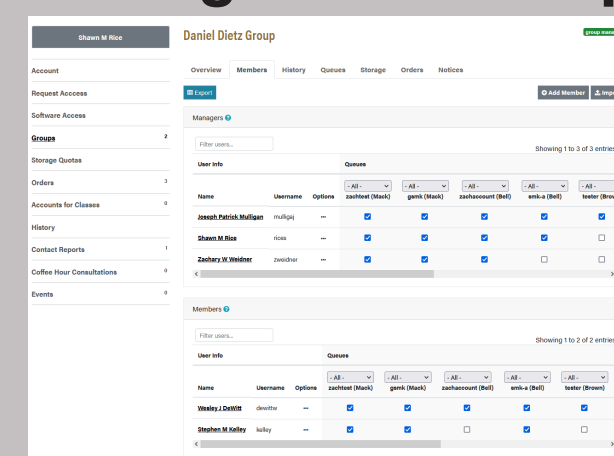
Track which students use the most computing, generate reports for sponsors, and monitor trends in your group's resource usage.



Faculty Partners by Cluster

Negishi 30 Departments 167 Partners	Gilbreth 36 Departments 111 Partners	Bell 45 Departments 161 Partners	Brown 49 Departments 173 Partners
Halstead 36 Departments 102 Partners	Rice 23 Departments 69 Partners	Conte 26 Departments 62 Partners	

Manage Your Group





Want to learn more? Contact us here:

Purchase Options

rcac-cluster-purchase@lists.purdue.edu

Support Questions

rcac-help@purdue.edu



Rosen Center for
Advanced Computing

101 Foundry Drive, suite 3314
West Lafayette, IN 47907
rcac@purdue.edu