

Contents

Overview ..... 2

CRCF HPC Cluster Specifications ..... 2

    Alipi & Clusty ..... 4

    Key Personnel (AY 2024) ..... 4

Strategic Plan (5-year, 2023-2028) ..... 5

Accomplishments ..... 6

    Administrative ..... 6

    Technical ..... 6

Metrics ..... 7

Work in Progress ..... 10

User Survey ..... 10

Final Points ..... 11

    Good Citizenship ..... 11

    Interested in Helping CRCF? ..... 11

Acknowledgements ..... 11

## Overview

The Centralized Research Computing Facility (CRCF) is Villanova's first core facility (established November 2021) and provides support for computational research. Most research support is in the form of high performance computing (HPC), but other forms of support are available as well. We encourage all Villanova researchers to contact us if they are interested in research computing but do not yet have a strategy towards achieving their goals.

CRCF supports three HPC clusters and their users:

- Augie: this is CRCF's largest and newest cluster. It was procured through an NSF Campus Cyberinfrastructure (CC\*) grant in Summer 2020. The cluster was installed in March 2021. Augie is open to all Villanova researchers and Villanova students.
- Alipi: this cluster is limited to College of Engineering researchers. Alipi is a couple of years older than Augie. Alipi is also smaller than Augie. **Alipi is to be retired in Summer 2025.**
- Clusty: this is the Astrophysics cluster that has been opened up to all users. The cluster is managed by Andrej Prsa in the Astrophysics and Planetary Science department.

CRCF emphasizes that *centralized computing is more than just HPC equipment*. CRCF currently provides the following to the Villanova community:

- State-of-the-art HPC resources
- A central knowledge base/pool of expertise for research computing resources
- A partner to help faculty integrate HPC into their coursework and proposals
- A resource for developing computer programming and software applications skills for students, staff, and faculty. This resource follows Augustinian values of helping others (akin to service learning)

In summary, *CRCF is a partnership of faculty, technical staff and students that empowers Villanova to maximize its computational research potential.*

CRCF has three thrusts as part of this mission/vision:

**Thrust #1: HPC Support through a Collaboration of Faculty and IT Staff.** This thrust involves the management of Augie, Alipi, and Clusty.

**Thrust #2: Research Computing Center of Expertise.** CRCF provides training opportunities including on-demand new user training, webinars for research computing topics of interest, and a point of contact for interaction with experts at other institutions. All researchers interested in research computing or HPC should approach CRCF to help develop a strategy to meet their goals.

**Thrust #3: HPC Peer User Support.** A wealth of research computing expertise exists on campus, and CRCF strives to facilitate communication between users seeking help and the experts that can help them. These volunteer experts personify Villanova's Augustinian tradition, providing support for commonly used programming languages and applications.

## CRCF HPC Cluster Specifications

### Augie

Item	Values
Compute Nodes/Cores	
RAM	
Disk Space	
GPU	
Interconnect	Mellanox 200 Gbps
Software	C/OpenMPI/MPICH, Fortran/OpenMPI, MATLAB, Python, PyTorch, R, COMSOL, CUDA, OpenFOAM, QuantumEspresso, VASP, LAMMPS, CP2K, IQTREE, PyMT, TensorFlow, Intel compilers, AMD compilers, PAML, RAXML

The status of Augie's applications is as follows:

1. The following software applications have been installed and performance tested, with sample scripts available to users: C/OpenMPI/MPICH, Fortran/OpenMPI, MATLAB, Python, PyTorch, R, COMSOL, CUDA, OpenFOAM, QuantumEspresso, VASP.
2. The following software applications has been installed, sample scripts have been provided in some cases, and performance testing is in progress: LAMMPS, CP2K, IQTREE, PyMT, TensorFlow, Intel compilers, AMD compilers, PAML, and RAXML.
3. The following software applications are slated for installation: Gaussian, Abaqus, Singularity, and sage.
4. The following software applications need further exploration or discussion prior to initiating an installation effort: Mathematica, ANSYS, and Anjuta.

Only those software packages listed in Items 1-3 above are currently supported by CRCF. Users are encouraged to install their own software packages in their home directories for their own use.

Users are encouraged to visit the [Augie App Notes](#) folder on the [AugieUsers SharePoint site](#), which contains lessons learned and names of power users for many of the software applications listed above.

## Alipi & Clusty

Cluster	Alipi	Clusty
Cores	208	212
RAM	1,408 GB	208 GB
Disk Space	26 TB	15 TB
GPU	N/A	N/A
Interconnect	InfiniBand	10 Gbps
Software	Python, LAMMPS, <u>QuantumEspresso</u> , VASP, MATLAB Compilers: <u>Akantu</u> , Intel	Ubuntu server, Python, Perl, R, sqlite3, octave

## Key Personnel (AY 2024)

CC\* Committee: Matt Morrissey (UNIT), Chris Washburn (Research Computing Administrator), Aaron Wemhoff (COE, Administrative Director)

Key UNIT & CLAS IT Support (Augie): Jonathan Hardy, Jonathan Graziola, Leo Nelson, Peter Palladino, Gavin Printz

HPC Advisory Board: David Cereceda (COE), Arup Das (VSB), Justin DiBenedetto (CLAS), Andrej Prsa (CLAS), Michael Robson (Smith College), Jason Simms (Swarthmore College), Bill Wagner (VSB)

## Strategic Plan (5-year, 2023-2028)

Mission: Support and promote computational research for Villanova researchers and their collaborators

Vision: CRCF, as a center of expertise in computational research, is one of the top contributors to Villanova's research enterprise.

Goals & Metrics (5-year goal in parentheses, AY2024 value after colon):

1. CRCF is a significant enabler of research productivity
  - a. Number of grant proposals submitted that use CRCF resources (10/yr.): 13
  - b. Number of peer-reviewed publications produced using CRCF resources (10/yr.): 11
2. CRCF is the center of expertise for computational research
  - a. Number of workshops or tutorials offered (4/yr.): 2 (JetStream2, Google Colab)
  - b. Number of researchers using external services (OSG, PATH, XSEDE/ACCESS) (5/yr.): 1 user achieved research using ACCESS with JetStream2.
  - c. Number of CI grant proposals stemming from the center (1/yr.): 1 in progress
3. CRCF has a robust multidisciplinary community of computational researchers
  - a. Number of total users (170): 236
  - b. Number of active users (40/mo.): average of 20 per month, max 26
  - c. Number of departments represented by active CRCF researchers (12): 12
  - d. Number of power users per software application (3): 0 or 1
  - e. Number of software applications with power users (12): 6
  - f. Number of classes that use Augie for instruction (8/yr.): 3 in the past year. One additional class had access to JetStream2 for HPC calculations.
4. CRCF's resources grow to match demand
  - a. Low idle time (< 15%): 29%
  - b. Small average job wait time (< 12 hours): we are still looking on ways to track this
  - c. Number of new condoers per year (2): 1
  - d. Acting as a host site for REDCap: pushed to data storage investigation
  - e. Dedicated HPC IT Admin: done
5. CRCF operates for the good of all humankind beyond research
  - a. Request that hosting data center must meet energy efficiency requirements (PUE < 1.5): Netrix does not track PUE; request made for them to do so
6. Request that data center have PPA in place for renewable energy: Netrix does not do this; request has been made for them to do so

## Accomplishments

CRCF's accomplishments for the AY have been split into administrative and technical categories.

### Administrative

The team spent the AY continuing to grow CRCF to ensure long-term use and management of the facility. Highlights of the facility's administrative accomplishments include:

1. Ran CRCF "AugieFest" annual all-hands meeting.
2. Ran webinars on JetStream2 and Google Colab
3. Achieved NSF ACCESS allocation, used primarily for JetStream2 GPU
4. Hosted first intern – Yin Wah Yip. A second intern, Arielle Korn, will start in June.
5. Created policies on condo pre-emption and long qos jobs
6. Renewed COMSOL and ANSYS/Fluent annual research software licenses
7. Surveyed researchers and developed writeup on data storage issues and recommendations & submitted for integration in IT strategic planning
8. Developed Augie financial and housing plan & submitted for integration in IT strategic planning
9. Arranged for the procurement of Gaussian computational chemistry software
10. Onboarded students in computational chemistry course
11. Recruited a couple of power users
12. Cleaned up information repository for users
13. Attended TechNova IT Strategic Planning steering committee meetings
14. Provided detailed documents on (1) financial issues surrounding CRCF, (2) data management issues on campus, and (3) thoughts regarding the HPC transition to a new phase.
15. Continued to provide monthly updates to users

### Technical

1. Augie continues to evolve technically to meet the needs of the user community. Some highlights:
2. Learned how to use spack for quick installation of some software packages
3. Installed OpenFOAM, tinker and sage
4. Upgraded SLURM to current version
5. Initiated process for migration to Ubuntu OS and made significant progress
6. Installed 6 condo nodes
7. Updated user provisioning script
8. Learned how to spawn MATLAB jobs on Augie by running MATLAB locally

## Metrics

### Total number of users (as of 3/14/25)

- Augie: 231
- Alipi: 55
- Clusty: 127 (note: all students in astrophysics & planetary science are on by default)

### Active Departments

The following departments contain active users since March 2023:

- CLAS: Biology, Chemistry, Computing Sciences, Mathematics and Statistics, Physics, Psychological and Brain Sciences, Sociology
- COE: Civil and Environmental Engineering, Chemical and Biological Engineering, Electrical and Computer Engineering, Mechanical Engineering
- VSU: Finance

### Publications using CRCF resources in the past year

Bazurto, Eric Daniel. "Computational Modeling of Li and Na Electrolyte Systems: From all-Atom to Coarse-Grained." Master thesis, Department of Chemistry, Villanova University, 2024.  
<http://ezproxy.villanova.edu/login?url=https://www.proquest.com/pqdtlocal1006301/dissertations-theses/computational-modeling-li-na-electrolyte-systems/docview/3098439025/sem-2?accountid=14853>

Fathi, Mohamed M. "Forecasting Long-Term Fluvial Geomorphology Behavior Under various Climatic Conditions." PhD dissertation, College of Engineering, Villanova University, 2024.  
<http://ezproxy.villanova.edu/login?url=https://www.proquest.com/dissertations-theses/forecasting-long-term-fluvial-geomorphology/docview/3156041738/se-2>.

Fathi, Mohamed M., Zihan Liu, Anjali M. Fernandes, Michael T. Hren, Dennis O. Terry, C. Nataraj, and Virginia Smith. "Spatiotemporal Flood Depth and Velocity Dynamics Using a Convolutional Neural Network Within a Sequential Deep-learning Framework." *Environmental Modelling & Software* 185 (2025): 106307. <https://doi.org/10.1016/j.envsoft.2024.106307>.

Fathi, Mohamed M, Zihan Liu, Anjali M. Fernandes, Michael T. Hren, Dennis O. Terry, C. Nataraj, and Virginia Smith. "Continuous Modeling of Flood Dynamics Using a Convolutional Neural Network within a Sequential Deep-Learning Framework."

Gruppi, Maurício, Soham Dan, Keerthiram Murugesan, and Subhajit Chaudhury. "On the effects of fine-tuning language models for text-based reinforcement learning." *arXiv preprint arXiv:2404.10174* (2024). <https://doi.org/10.48550/arXiv.2404.10174>

Musa, Md Rajib Khan, Yichen Qian, Jie Peng, and David Cereceda. "Accelerating the discovery of low-energy structure configurations: A computational approach that integrates first-principles calculations, Monte Carlo sampling, and Machine Learning." *Scripta Materialia* 259 (2025): 116535. <https://doi.org/10.1016/j.scriptamat.2024.116535>

Nguyen, Hiep, Lynn Yip, and Justin DeBenedetto. "Automatic Quality Estimation for Data Selection and Curriculum Learning." In The 2nd BabyLM Challenge at the 28th Conference on Computational Natural Language Learning, pp. 212-220. 2024. <https://aclanthology.org/2024.conll-babylm.18/>

Peng, Jie, Yichen Qian, and David Cereceda. "A First-Principles Study of the Structural and Thermo-Mechanical Properties of Tungsten-Based Plasma-Facing Materials." *Metals* 14, no. 10 (2024): 1197. <https://www.mdpi.com/2075-4701/14/10/1197>

Qian, Yichen, Mark R. Gilbert, Lucile Dezerard, Duc Nguyen-Manh, and David Cereceda. "First-principles study of the energetics and the local chemical ordering of tungsten-based alloys." arXiv preprint arXiv:2410.03998 (2024). <https://arxiv.org/pdf/2410.03998>

Sankar, Harini. "Encoding Context in Spoken Word Recognition: A Computational Model." Master's thesis, Department of Psychological and Brain Sciences, Villanova University, 2024.

Scott, Reese, and Robert Styer. "Number of solutions to  $a^x + b^y = c^z$  with  $\gcd(a, b) > 1$ ." arXiv preprint arXiv:2401.04197 (2024). <https://doi.org/10.48550/arXiv.2401.04197>.

Styer, Robert. "At most one solution to  $a^x + b^y = c^z$  of  $(a), (b), (c)$ ." *Glasnik matematički* 59, no. 2 (2024): 277-298. <https://doi.org/10.3336/gm.59.2.02>.

Styer, Robert. "SERIJA III [www. math. hr/glasnik](http://www.math.hr/glasnik)." (2024).

#### **Proposals submitted in the past year that would utilize CRCF resources**

- Nersesov, Sergey (1 proposal)
- Ural, Ani (3 proposals)
- Juretus, Kyle (3 proposals)
- Nataraj, C. (5 proposals)
- Smith, V. (1 proposal)

#### **Miscellaneous**

- Many researchers have journal papers under review for work that used Augie.

#### **Courses where Augie, Alipi, or clusty have been integrated**

- BL 2149 TOP: Cyber Law
- CHM 9947 Computational Chemistry
- CSC 2405 Computer Systems II
- CGS 5990 Fairness in AI
- CGS 5900/PSY 8900 Cognitive Science Seminar
- ECE 8487 - Advanced Machine Learning
- Mathematics Senior Seminar
- ME 3600 Fluid Mechanics
- ME 3950 Heat Transfer I
- ME 7030: Numerical Methods in Engineering Simulation
- SPA 3200 Introduction to Spanish Translation




### CRCF workshops offered in the past year

- JetStream2 (Yin Wah Yip)
- Google Colab (Justin Debenedetto)


### Augie Usage

The usage goals for Augie are (1) less than 20% idle time, and (2) an average of 20% usage by the Open Science Grid (OSG). In the last 12 months, the idle time is slightly high (29%), and the OSG usage is under the limit (18%). These are both acceptable values for purposes of ensuring that Augie's growth matches demand. While CLAS and COE demand vary widely, the average Augie usage per college is 28% for CLAS and 24% for COE, with small usage by VSB and zero usage by FCON.



### Number of Active Users

Active users are defined as those that submitted jobs to Augie within a specified month. The annual average is 20.3 users, which is up from 16.2 users in the previous 12 months.



### Work in Progress

The following efforts are currently ongoing:

- Administrative
  - Work with the Provost's Office and UNIT on the new phase of research computing through the TechNova IT Strategic Plan
  - Provide better means to track integration of Augie into proposals
  - Hold workshops on AI code generation and using MATLAB-based batch submissions
  - Recruit power users
  - Track average job wait time
  - Mentor a summer intern to provide documentation and a webinar on Globus
  - Procure license renewals for COMSOL and the Academic Research license for ANSYS-CFD
- Technical
  - Complete OS transition to Ubuntu
  - Install Gaussian, Abaqus, and SageMath
  - Consider data backup to cloud storage

### User Survey

A survey was provided to the Augie Users in March 2025. No feedback was garnered outside of indications of satisfaction with the current state of affairs. We will continue to provide an annual survey in March 2026.

## Final Points

### Good Citizenship

#### Some reminders:

- No running jobs on the head node – submit batch jobs instead
- Don't use the debug queue for production runs
- Don't submit lots of jobs to occupy a large percentage of cores on the cluster
- Be sure to acknowledge use of the HPC clusters – see user terms and conditions documents:
  - Augie: Augie Users SharePoint, file [Administrative/Augie-HPC-TC 20210507.pdf](#)
  - Alipi: Villanova HPC Team → Alipi Users Channel, file Alipi Usage Policies.docx
  - Clusty: contact Andrej Prsa
- Need help on Augie/Alipi/Clusty?
  - Don't contact UNIT directly or put in a ticket
  - Augie:
    - Find appropriate document in the [Augie App Notes](#) documents. See if the document specific to your app has the information you need.
    - Many documents have listed superusers; contact the superuser for help.
    - If the above two items don't work, then email [ccstarcommittee@villanova.edu](mailto:ccstarcommittee@villanova.edu).
  - Alipi: post issue in Teams channel. If no response in 24 hours, then email [engineering-hpc@villanova.onmicrosoft.com](mailto:engineering-hpc@villanova.onmicrosoft.com).
  - Clusty: contact Andrej Prsa

### Interested in Helping CRCF?

Contact [ccstarcommittee@villanova.edu](mailto:ccstarcommittee@villanova.edu) if you are interested in...

- Supporting other Villanova researchers as a power user
- Collaborating with a researcher at a small local college
- Providing additional feedback on CRCF operations
- Working as an undergrad support software technician during the summer
- Providing your thoughts about how to make AugieFest a popular, fun event
- Providing ideas for tutorials

## Acknowledgements

Funding and support for CRCF from the Office of the Provost, the College of Liberal Arts and Sciences, the College of Engineering, and UNIT are gratefully acknowledged.