COLUMBIA | RESEARCH
Shared Research Computing
Policy Advisory Committee

Shared Research Computing Policy Advisory Committee (SRCPAC)
Spring 2025 Meeting

Date: Tuesday, March 7, 2025
Time: 11:00 AM - 12:30 PM
Chair: Alex Urban, Assistant Professor of Chemical Engineering

Meeting Minutes

Welcome & Introductions – Professor Urban, Chair of SRCPAC

Professor Urban called the meeting to order and outlined the agenda. He provided an overview of shared research computing at Columbia, detailing its key components, respective charges, and associated university units, including SRCPAC, the Shared Research Computing Facility, Research Computing Services (RCS), and Foundations for Research Computing (FoRC).

Empire AI Alpha Phase – Professor Urban

Professor Urban gave an update on Empire AI, a major GPU initiative funded by New York State. Announced during last year's meeting by EVP Jeannette Wing, the initiative is now operational in its ALPHA stage, with Columbia utilizing approximately 17% of available compute time.

Halayn Hescock, Senior Director of CUIT Research Services, noted the upcoming BETA phase, expected to launch in the summer, which will introduce broader access and a new structure, including a fee structure that is being determined. ALPHA access remains limited but can be requested via the EVPR.

High-Performance Computing Updates – Professor Urban & Halayn Hescock, Senior Director of CUIT Research Services

Current and Proposed Hardware Retirement Policy
Professor Urban reviewed current HPC systems:

- Terremoto Phases 1-2: Retiring
- Ginsburg Phases 1-3: Phase 1 retirement is December 2026
- Insomnia: Not subject to traditional retirement
- Manitou GPU Cluster: Being integrated into Insomnia

He explained that constraints such as space, electricity, cooling, and software costs influence retirement policies. A revised model will extend hardware operation beyond five years until failure or

required space. Hescock noted that this change incurs additional costs, including user-funded cables, networking, and storage.

Professor Urban acknowledged the complexity of providing access to computer nodes and emphasized the need for flexibility. A new scheme, developed by the HPC Manager, Max Shortte, will be piloted and refined. An Operating Committee will be convened to review policies. Interested participants should reach out. Additionally, Manitou GPUs will soon be available to Insomnia users.

HPC Statistics

Professor Urban outlined efforts to optimize cluster utilization. Initially, Ginsburg faced underuse due to policy revisions and technical issues. Newly revised policies, including increased job allowances, aim to eliminate idle hardware. Further policy refinements are expected to enhance usage for Ginsburg and Insomnia.

Free / Edu Tier Updates

A new model dedicates hardware to free and educational tiers, with new nodes contributed annually by A&S, SEAS, and EVPR. State-of-the-art resources are available to all campus users who apply, running at lower priority. These tiers have supported 56 educational and 160 free-tier accounts, completing nearly 900,000 jobs since last fall. Going forward, these tiers will no longer rely on retired equipment.

High-Performance Computing Updates – Axinia Radeva, Senior Manager, Research Computing Services, CUIT

SRCPAC Website Enhancements

The updated website provides details on:

- HPC options (buying vs. renting nodes)
- Free and education tier information
- Recommended acknowledgment language for proposals and publications

Additional resources:

- www.cuit.columbia.edu/research
- www.cuit.columbia.edu/rcs/training (includes RCS video library)

Training and Workshops

Since April 2024, 15 workshops (six in-person) have hosted over 180 attendees. Sessions, led by RCS/HPC or vendors (Google Cloud, Intel, NVIDIA), primarily attract graduate students, researchers, and IT staff. SEAS and the Graduate School of Arts and Sciences have the highest participation.

Globus: High-Speed Data Transfer Service

Radeva highlighted Globus milestones:

- 2021: Globus Standard launched with Google Drive Connector
- 2023: Globus Open Access added
- 2024: Globus BAA signed, enabling sensitive data transfer (PHI, PII, RHI)
- NEW: Globus as a service

She noted ongoing efforts to expand Globus beyond data transfer into automation and computation. Columbia users now have free access to Globus High Assurance via CUIT and CUIMC.

ACCESS Roadmap at Columbia

Beyond Columbia's HPC clusters, RCS facilitates access to national high-performance computing systems via NSF's ACCESS program. A new request page and intake form launched in August to streamline Discover Allocations.

Radeva clarified that Columbia does not pay for ACCESS; NSF provides access through a proposal process, which requires PIs to justify their computing needs. RCS can support application preparation.

Foundations for Research Computing Update – Professor Marc Spiegelman, Foundations Advisory Committee

Foundations Overview

Professor Spiegelman outlined the original objectives of Foundations: 1) to provide the investment in people and computational skills required to complement our investment in hardware, software and systems administration, and 2) to provide a structured pathway for onboarding students and researchers into modern research computing at Columbia.

Initially, Foundations was designed with three levels (novice, intermediate, advanced), with the latter focused on faculty. He emphasized that Columbia offers broader computational training beyond the biannual Foundational Workshops, including CUIT and Library-led RCS/HPC services. Demand remains high for these labor intensive workshops—649 applicants last year, with only 166 accepted and 100 attending the first-day session. Due to capacity constraints, Libraries must selectively admit participants based on skill level.

Key challenges identified:

- Scalability – Foundational Workshops are labor-intensive and difficult to sustain at scale.
- Faculty Coordination – Declining faculty involvement.

- Shifting Scope – Foundations has evolved beyond its original purpose, making its relationship to SRCPAC unclear.
- Outdated Curriculum – Training content has not kept pace with computational advancements.

Given these challenges, Professor Spiegelman proposed rethinking Columbia's approach to informal computational training.

Foundations Framework

To address sustainability and increasing demand, Professor Spiegelman proposed:

- Re-engaging the Foundations Faculty Advisory Group.
- Revising the introductory curriculum and models (asynchronous learning resources, AI tutors, peer mentoring).
- Research Computing Services (CUIT) continues to provide training for support of research computing resources.
- Centering computational literacy as a structuring focus in Libraries.

He presented two potential directions: 1) discontinuing Foundational Workshops and transforming Foundations into a training repository while redistributing responsibilities to CUIT and Libraries, or 2) restoring a unified Foundations model, reestablishing it as a centralized computational training resource, incorporating scalable beginner training and AI-enhanced curricula within Libraries.

Professor Spiegelman and Sharon Sputz, Associate Vice President for Research Initiatives and Development, encouraged faculty involvement, including serving as SRCPAC Co-Chair or joining a temporary subcommittee on entry-level training. Sputz emphasized the need for faculty input for what computational skills faculty believe students should develop as they enter labs and classrooms currently.

Closing Remarks and Discussion Points

Faculty Governance, Engagement, and Communication

There was a request for increased communication regarding workshop notifications: while they are currently available via the CUIT research page listserv, the RCS training calendar, Foundation for Research Computing site, and library sites, participants recommended distributing the notifications via additional platforms.

Discussions ensued on ways to increase faculty participation due to reduced committees that met more frequently in the past. There was a realization that the operating committee no longer convenes and this has caused communication and participation challenges.

The recommendation was to increase faculty engagement in a more meaningful and consistent way with governance structures, transparent decision-making, and assurance that contributions lead to tangible outcomes.

Additional discussions on aligning SRCPAC's direction with the evolving research and learning landscape followed, particularly in the context of emerging technologies like LLMs. This discussion related to Professor Spiegelman's question on the future of the Foundations initiative, and underscored the potential risks of losing collaborative learning opportunities amid these technological changes. Additionally, discussions took place on the training needs across campus given the advancements in computing since Foundation for Research Computing was created. For example, answering: How can faculty, researchers, and administrators collectively identify the most critical priorities for advancing research computing within Columbia's institutional and technological landscape? Participants recognized the Foundations initiative's need for a faculty-driven mandate that aligns research computing training with current faculty needs.

Professor Urban concluded by affirming the importance of faculty engagement and governance. In addition, he emphasised the improvements that have been made to address many faculty requested actions, such as addressing hardware pricing, seasonal purchasing restrictions, hardware retirement, and centralizing information on SRCPAC's website.

Slides provided on: https://research.columbia.edu/content/srcpac

See a list of attendees below.

Attendees

Maneesha Aggarwal (Zoom)
Suzanne Bakken
Thomas Boag
Jonathan Cain
Robert Cartolano
Rémi Cousin
Jessica Eaton
Farshid Fathi
Michael Faucher (Zoom)
R. Yamir Gomez Carrasco
Halayn Hescock
Owen Keith
Rob Lane
Walt Mankowski
Razvan Popescu
Axinia Radeva
Michele Simoncelli
Marc Spiegelman
Sharon Sputz
Alexander Urban
Anthony Vanky
Eric Vlach
Daniel Woulfin
Tian Zheng
+ 3 more in-person attendees
+ 8 more virtual attendees