



**Updates from START Research
Computing Working Group**
- part of the START Committee on
Increasing Our Research Excellence

Chen-Nee Chuah
(on behalf of START RCWG)
Feb 4, 2025

START Research Computing Working Group



- Chen-Nee Chuah, Professor, Electrical & Computer Engineering (Chair)
- Simon Atkinson, Vice Chancellor for Research, Office of Research
- Rachael Callcut, Associate Dean of Data Science and Innovation, School of Medicine
- Titus Brown, Professor, Population Health & Reproduction [2023-24]
- Ian Korf, Professor, Molecular and Cellular Biology
- Jawdat Al-Bassam, Associate Professor, Molecular and Cellular Biology
- Viji Murali, CIO and Vice Provost, Information and Educational Technology
- Neeraj Chauhan, Associate Chief Information Officer, Enterprise Infrastructure Services, IET
- Jeremy Phillips, Director of Information Technology, College of Letters and Science
- Steve Pigg, Executive Director of Information Technology, College of Engineering
- Kent Anderson,
- Charles Bookman, Technical Director, HPC-CF
- Jeffrey Ross-Ibarra, Professor, Evolution and Ecology / Faculty Director of HPC-CF



High Priority Issues

- General research computing services, including clouds
 - UC Davis is the only UC campus without a Research IT unit!
- **HPC current unmet needs & future expansion**
 - Funding for HPC-CF is negotiated annually under ACCD model
- **Research Data & IT Security**
 - Increasingly stricter requirements from funding agencies
- Data Storage/Backup/Archival
 - Digital birthright: free-tier for all faculty
- Network infrastructure
 - Digital birthright: Good wireless connectivity in all buildings
- *Use of ChatGPT & AI services for research*

Winter 2024



➤ SWOT analysis

- Identify critical areas for research computing needs
- Main challenge: totally uncoordinated across campus with no central oversight, leading to redundancy & inefficiency
- Unmet HPC needs

➤ Working group and HPC-CF faculty advisory group independently issue request for a Design & Construction Management (DCM) project to assess power/space/cooling capacity on campus

➤ Consult

- UCDH IT
- Library & Datalab leadership
- HPC-CF staff and faculty advisory group
- Denise Ehlen (OR)
- Cheryl Washington (IET SISO)



➤ Submit first two recommendations

➤ **Rec#1: Research Computing Oversight Committee**

Establish a Strategic Administration of Research Computing (STARCC) Oversight Committee to centralize the coordination of research computing facilities and services, including but not limited to: HPCCF, data centers, network infrastructure, web services, and software, with direct report line to decision makers. We recommend that the committee convene under the direction of the Vice Chancellor of Research. The committee should have representations from the following constituencies:

- HPC-CF
- IET leadership
- IT leadership (college/school level)
- Library (e.g., Online Strategy & DataLab)
- OCVR research data security team
- Faculty members spanning different disciplines

Reference model: UCDH system-wide Health Data Oversight Committee (HDOC), especially its subcommittee on Advanced Computing Committee, and Data Center of Excellence (CoE).



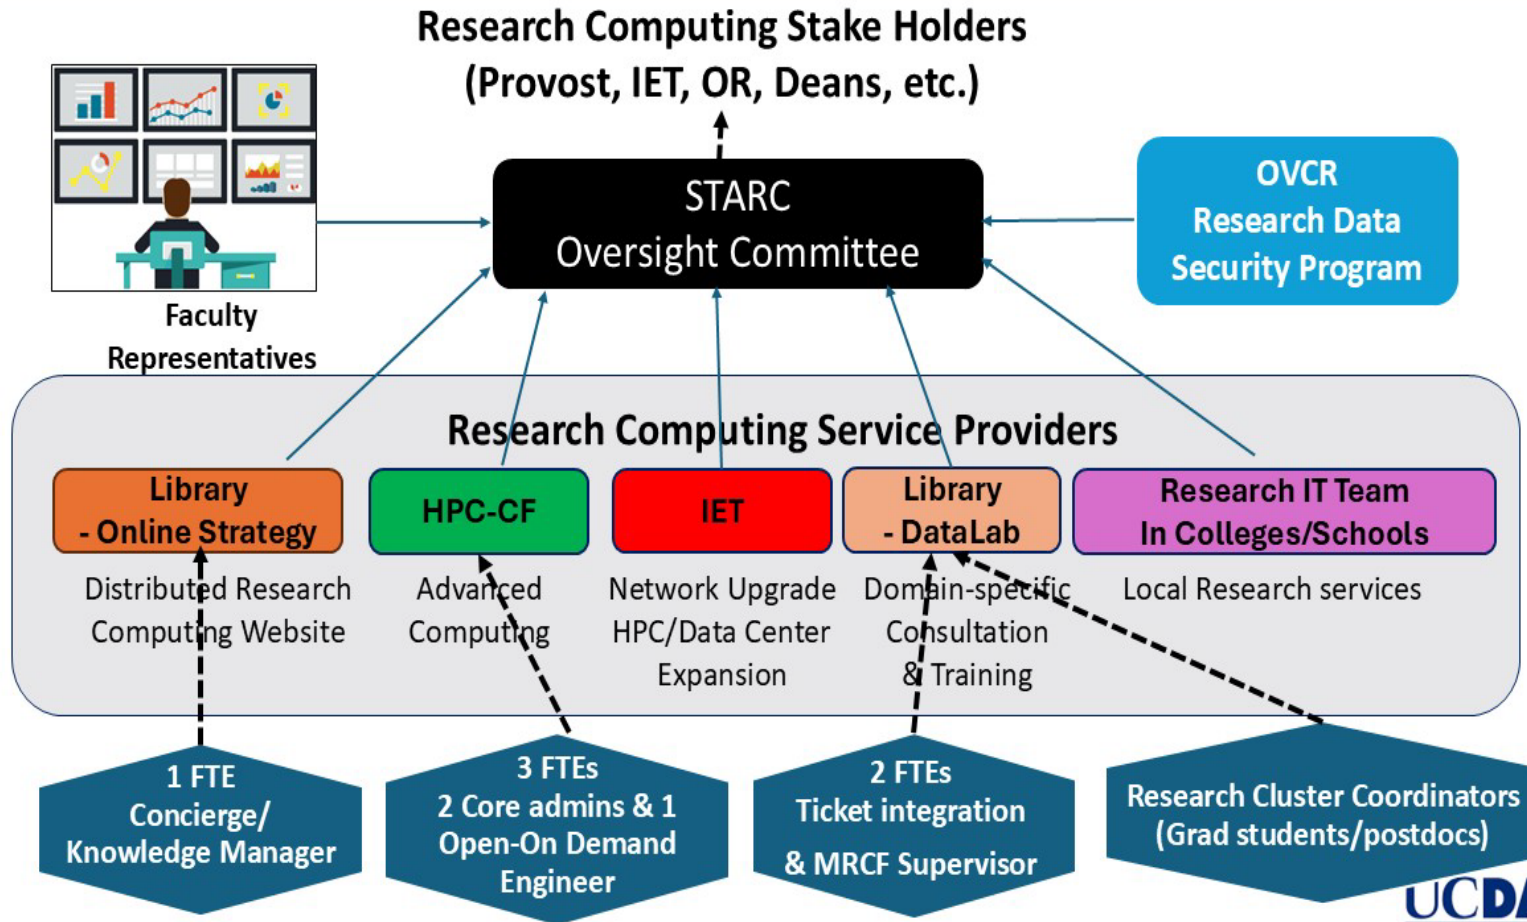
➤ Rec#2: FTE for Unified Research Computing Services and Community Building Effort

- Concierge/Knowledge Manager to maintain the content on a single Research Computing Website that compiles and consolidates research IT services from many different campus units into a “one stop shopping” experience to aid in resource discovery, particularly for new and prospective researchers and faculty.
- System Integration Engineer to be housed in HPC-CF reporting to Charles Bookman and his team responsible for:
 - Implementing application integrations as requested by Research Cluster Facilitators and the Distributed Research Computing Group
 - Investigate, assess, and validate public research computing resources (e.g., NSF ACCESS program or Google TPU program)
- Multi-Disciplinary Research Cluster Facilitators housed in the DataLab who will function as method- and domain-specific facilitators with the goal of providing consultation to new and existing researchers to identify customized research computing solutions, relevant training resources, and further build and grow a communities of practice where researchers can engage and share knowledge/skills

Proposed Research Computing Core



Strategic Administration of Research Computing (STARC) Oversight Committee





Summer 2024: HPC needs

- Spring quarter: consulted a group of CoE faculty to identify needs for HPC expansion, summarized in a document prepared by Yubei Chen
 - Discussion for short-term/medium-term/long-term solutions for HPC expansion
 - Info gathering from IET, HPC-CF, Office of Research
 - Current HPC-CF infrastructure houses computing and storage in four locations: the Campus Data Center, Academic Surge, Watershed, and in the Genome Center
 - Requested DCM report on power/space/cooling was scoped to assess opportunities to add 40 racks at 35kW each; ~1,400kW of computing and represents resources needed to fulfill current needs and very modest growth in the next 2 years
- **Missed opportunity: Faculty w/ technical expertise (solution space) not consulted**
 - Yubei Chen now serves on HPC-CF faculty advisory group
 - Vinod Narayanan (director of Energy Efficiency Center) is consulted for modern data center and liquid cooling technologies

DCM Executive Summary



- Came out end of Aug
- Recommendations: **None of the existing space can accommodate high-performance computing (HPC) with existing infrastructure.** Although some of the existing spaces may have physical space for some portion of HPC, **the cost of the required infrastructure outweighs the benefits of the marginal increase in HPC**, which will have little impact on the long-term growth of HPC on the UCD campus. Furthermore, to meet the requirements for HPC, a major capital project would be required to bring the required infrastructure to the building.



Addressing CoE HPC needs

- Worked with HPC-CF to identify short-term solutions
- Identified two possible space: Library Server Room & Academic Surge
 - Pending MOU between library & CoE to be revised by ECE-IT committee
 - Neither room has enough power → need to estimate expansion cost
- **HIVE HPC Cluster: <https://hpc.ucdavis.edu>**
 - **Does not meet CoE PI's needs for hosting dedicated GPU development clusters**
- **More STEM faculty need to get involved!**
 - **Articular unique HPC needs & usage models**
 - **Share technical expertise, best practices, and brainstorm solutions with HPC-CF and CoE-IT staff (e.g., incorporating cutting edge cooling technologies or AI accelerator hardware)**

START Rec#3 on HPC (under review)



- **Stable and continuous central staff funding** (replacing ACCD model)
 - **HPC-CF staff:** (i) core system administration team plus additional FTEs for (ii) central file systems and (iii) domain-specific HPC architects to provide support for diverse HPC-needs on campus to continue to main and expand different areas of research excellence
 - **HPC-CF Equipment maintenance**
 - “Public Tier” hardware: the largest expenditure is to replace compute nodes and storage that are made available at no direct cost to UC Davis researchers.
 - Compute infrastructure not covered by recharge model includes racks, cooling, power supplies, etc.
 - Software licensing
 - Network infrastructure
 - Staff training

START Rec#3 on HPC (under review)



- **Stable and continuous central staff funding** (continued)
 - **HPC-CF Equipment Expansion**
 - Purchase new hardware & storage to accommodate an expanding user base for both Public Tier and new research projects under recharge model
- Sustainable HPC service and funding models need to address the mismatch between funding mechanisms (external funding agencies) and campus recharge model
 - A specific allocation of F&A to invest in research computing, especially to cover the intrinsic cost for servers (cooling for the machine room even though the rest of the building is empty)

START Rec#3 on HPC (under review)



- **Request one-time bridge funding to buy-out ancient hardware and move to new cluster models (e.g., HIVE)**
 - This will free up data center space for newer and more energy efficient equipment, leading to better utilization of existing HPC space.
- **Funding to investigate and implement short-term HPC expansion**
 - Commission a new DCM project to evaluate the expansion cost of using Academic Surge and Shields Library Server Rooms for housing additional racks
- **Evaluate multiple options for meeting future HPC needs**
 - Thoroughly investigate liquid cooling and other modern data center technologies
 - Partner with other campuses and industry partners in co-location and fundraising efforts for new data centers
 - Negotiate better contracts with co-location and cloud service providers

Action Items



- Review the two START Research Computing Recommendations and use the Feedback Form to provide your endorsement
 - <https://leadership.ucdavis.edu/news/messages/provost-messages/START-Research-Recommendations-Dec-2024>
 - <https://leadership.ucdavis.edu/about/provost/initiatives/start/feedback>
- Request clarification on F&R (indirect cost) distribution
 - A fixed portion of indirect cost recovery (ICR) should be redirected to research computing infrastructure and admin support (e.g., HPC-CF), as assumed by most funding agencies.
 - *Reference model: Indiana University retains 30% of ICR starting July 1, 2023 (“applicable grants”), out of which 65% is allocated to IU Research and 35% is allocated to University IT Services:*
<https://research.iu.edu/funding-proposals/proposals/budgets/rates.html>

Other Recommendations (on-going)



- Secure Research Computing and Data Storage
 - To comply with research security mandates of funding agencies
 - On-premise secure computing services (IET) vs. outsourcing (Sherlock @ San Diego Super Computing)
- Network Infrastructure
 - Goal: Ubiquitous & robust wireless connectivity across campus
 - Specific research needs: upgrade horizontal wirings