

UC Berkeley

Publications

Title

Research Facilitation on a Budget

Permalink

<https://escholarship.org/uc/item/0zq9850x>

Authors

Christopher, Jason
Culich, Aaron
Dombrowski, Quinn
et al.

Publication Date

2018-07-01

Peer reviewed

Research Facilitation on a Budget

Jason Christopher
UC Berkeley
jchristo@berkeley.edu

Aaron Culich
UC Berkeley
aculich@berkeley.edu

Quinn Dombrowski
UC Berkeley
quinnd@berkeley.edu

Debra McCaffrey
University of Michigan
debmccaf@med.umich.edu

Amy Neeser
UC Berkeley
aneeser@berkeley.edu

Andrew Wiedlea
Lawrence Berkeley Lab
AWiedlea@lbl.gov

ABSTRACT

Achieving broad uptake of research computing services is a tremendous challenge when funding for staff positions is constrained. Outreach and ongoing engagement with researchers is both essential and time-consuming, leading to a tension between supporting day-to-day operations and building the kinds of partnerships that ensure ongoing support for the program.

Since 2015, Berkeley Research Computing (BRC) has been hiring primarily graduate students into a part-time “domain consultant” role (influenced by ACI-REF job descriptions and Campus Champion activities) that addresses the program’s staffing needs in an affordable way, while providing those graduate students with the technical training and professional work experience required for professional research facilitator positions. The domain consultant program has evolved from an hourly student position into a codified set of practices informed by IT service management, addressing needs including: in-person consulting, tier 2 triage of HPC troubleshooting tickets, support for cloud computing and compute in virtualized analytics environments, and user training. In addition, consulting engagements are reviewed and discussed regularly, both in team meetings and in one-on-one meetings with the service manager, to provide opportunities for consultants to hone their skills.

This paper will also highlight the value of programs such as BRC’s for addressing gaps in graduate education practices that can hinder PhD recipients’ success when applying for research facilitator positions. It will also illustrate the value of thinking broadly about partnerships when developing a consulting program, by describing the program’s recent expansion into research data management, and at the Lawrence Berkeley Lab.

CCS CONCEPTS

• **Management of computing and information systems** → **Project and people management** → Project Staffing;

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

PEARC '18, July 22–26, 2018, Pittsburgh, PA, USA
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-6446-1/18/07...\$15.00
<https://doi.org/10.1145/3219104.3219137>

KEYWORDS

Research facilitation

ACM Reference format:

J. Christopher; A. Culich; Q. Dombrowski; D. McCaffrey; A. Neeser; A. Wiedlea. 2018. Research Facilitation on a Budget. In *Proceedings of PEARC conference, Pittsburgh, Pennsylvania USA, July 2018 (PEARC'18)*, 5 pages.

1 INTRODUCTION

One of the most significant challenges currently facing research facilitation as a career path is the pipeline problem. Graduate students, even those with extensive disciplinary and computational knowledge in the context of their own research, have few opportunities to develop the interpersonal professional skills around consultation that are essential for research facilitation. Another major challenge for research computing organizations is funding constraints. Berkeley Research Computing has taken steps to simultaneously address both sets of issues by implementing a program of hiring graduate students as computational “domain consultants”. Domain consultants have some existing familiarity with research computing, but receive additional training as part of their appointment, and develop their expertise through meeting with and advising researchers from a wide range of campus domains.

2 BACKGROUND

The Berkeley Research Computing (BRC) program launched in 2014, as an expansion of the University of California (UC) Berkeley’s reconceptualized Research IT group. Unusually, the group already had a long track record of working closely with humanities researchers and museum staff at the time research computing was introduced, which influenced BRC’s emphasis on consulting (i.e. research facilitation) as a core component of the overall program.

From a technical perspective, the emerging BRC program aspired to offer a local high-performance computing service to Berkeley researchers, including both no-cost access through a “faculty computing allowance”, and a condo purchase option. BRC would also develop \AEoD (Analytics Environments on Demand), a virtual Windows-based research desktop service leveraging existing investment in VMware and Citrix by

administrative infrastructure. In addition, BRC would offer consulting around, and facilitated access to, cloud computing resources.

The budget for the program overall was modest in the context of its scope: \$1.5 million over 3 years, jointly funded by the university chancellor, vice-chancellor for research, and CIO. In order to leverage existing expertise and avoid the temporal and financial costs of starting a high performance computing (HPC) program from scratch, BRC established an agreement with the HPC system administration and user services team at Lawrence Berkeley National Lab (LBNL), located less than a mile away from UC Berkeley and part of the University of California system, to manage the HPC aspect of the program through buying time from the existing LBNL team. The cloud and virtual desktop aspects would be directly managed by Research IT professional staff.

Unlike the LBNL HPC team, where some members held PhDs in the sciences, the varied backgrounds of Research ITs professional staff did not include experience with the hard sciences. This limitation was less problematic than one might assume, since the demand for consultation came disproportionately from researchers in domains outside those traditionally supported by HPC. Nonetheless, expanding the pool of research computing consultants to support a broad range of researchers -- including those in the hard sciences -- remained a high priority.

3 DEVELOPING THE DOMAIN CONSULTANT PROGRAM

Funding restrictions meant that hiring multiple full-time BRC research facilitators would not be feasible. Choosing one domain to disproportionately support through a research facilitator position was antithetical to BRC's commitment to supporting the greatest possible breadth of campus research, even when factoring in existing staff's experience in the humanities, social sciences, and professional school programs. Reconceptualizing the core of the research facilitator role into a 25% appointment (8-10 hours/week) for graduate students would enable BRC to offer tailored consultation for a larger range of disciplines, and provide technical and professional-skills training that would make Berkeley graduate students more competitive for professional research facilitator positions, if they chose not to pursue a traditional academic career path.

BRC staff had already established contacts with graduate students with proficiency in at least one of BRC's areas of technical focus (e.g. HPC, cloud, or virtual desktops), through participating in graduate student oriented "working groups" run by other organizations on campus, such as the Berkeley Institute for Data Science (e.g. "The Hacker Within", a support group for graduate students who are new to computational tools and methodologies), and the social sciences-based D-Lab (e.g. Cloud Working Group, Text Analysis Working Group, and Machine Learning Working Group). BRC's first domain consultant hire in November 2014 grew out of BRC's emerging partnership with UC Berkeley's Statistical Computing Facility (SCF). This first

domain consultant hire was a person who has the dual role of an adjunct faculty member and statistical computing consultant in the Department of Statistics. The initial appointment was for 5% of his time, intending to expand up to 25% as the program grew.

In summer 2015, BRC's lead Research Computing Architect led the launch of the cloud consulting aspect of the BRC consulting program to go beyond traditional cluster computing environments to facilitate access to and effective use of cloud computing resources from commercial cloud providers (such as Amazon AWS, Google Cloud Platform, or Microsoft Azure), as well as national centers (such as the then-new XSEDE Jetstream cloud).

As a complement to the consulting service, we also established a Cloud Working Group (CloudWG) in Fall 2014. The CloudWG served as a place to provide training on cloud computing, and as a way to work directly with researchers together to better understand in a first-hand way the challenges of researchers learning to use cloud computing for the first time. It also helped and as a way to help scale the consulting capabilities by connecting together researchers with similar use cases together, rather than relying solely on the limits of one-on-one consulting. The working group was built as a partnership between BRC and the Social Sciences D-Lab which offered a set of services (training, working groups, physical collaborative space, and consulting) for social science and humanities researchers. D-Lab had a flourishing consulting model based on hiring graduate students, which was influential in BRC's decision to pursue a similar approach.

BRC also developed a partnership with the newly launched Berkeley Institute for Data Science (BIDS) in Summer 2014. Rather than create a new working group with BIDS, a consultant/architect embedded with the collaborative space and participated in BIDS Reproducibility Working Group and Career Paths Working Group. The working groups at D-Lab and BIDS eventually also became a forum to identify graduate students with the required skills and interest in research computing.

As both the consulting service and CloudWG co-evolved, additional planning went into the next phase of the BRC Consulting Program, BRC established a formal job description for the next round of Domain Consultant hiring. BRC looked to the broader research computing landscape to better understand the state of emerging career paths. The models that stood out as applicable to our campus-focused challenges were two NSF-funded models: the XSEDE ECSS and Campus Champions[2] and ACI-REF (Advanced CyberInfrastructure Research & Education Facilitators)[3]. The BRC Domain Consultant model was directly influenced by job listings from Clemson University (ACI-REF), University of Wisconsin - Madison (ACI-REF), Indiana University (XSEDE Campus Bridging Engineer), University of California, Santa Barbara (Supercomputing Consultant), and Carnegie Mellon University (HPC Specialist).

BRC's partnerships with campus organizations became valuable sources of potential domain consultant candidates, as the collaborative physical spaces provided a chance to interact with members of the research community. In fall 2015, BRC staff used

a lightweight, informal interview process to hire a third year graduate student in nuclear engineering who was active in The Hacker Within working group at BIDS.

In many ways, this graduate student's background aligned perfectly with what BRC staff had envisioned for a Domain Consultant. She was herself a proficient HPC user as a result of her graduate training in nuclear engineering, but more significantly, she had experience supporting other researchers, as the person responsible for managing her group's private cluster. Those experiences, combined with her attention to detail and interest in improving BRC's processes and documentation, led to significant contributions to the program over her first year at BRC, even as her involvement with individual consultations with researchers remained quite low.

In 2016, BRC restructured how it organized the virtual desktops (AEoD Service), cloud, and consulting aspects of the BRC program to break down the original silos that divided them. Together, two BRC architects, including an architect previously focused exclusively on the AEoD Service, co-led further development of the consulting program. Together they introduced new management practices and processes that laid the foundation for the expansion of BRC consulting.

4 CONSULTING AS A PRACTICE

The BRC Consulting program was re-formulated as a set of shared practices, documented by an extensive "consulting handbook", made increasingly more useful for the consultants through ongoing and detailed review, discussion and re-write of sections of the handbook. The practices of: shared record-keeping, consulting manager review of open engagements, regular team meetings, consultant trainings and oversight established the role and value of taking an IT service management approach to managing the program.

4.1 Consulting log, "backgrounders", and Asana

The BRC consulting practice codified and expanded upon an existing informal practice of shared record keeping for consulting engagements. Each consultation (defined as a meeting with a researcher, or extensive email exchange) is recorded in the "consulting log", a Google spreadsheet viewable and editable by everyone affiliated with BRC. The spreadsheet captured the following, with required fields denoted by an asterisk:

- Date*
- Consultant(s)*
- Client* -- who the consultant actually met with; often a graduate student
- PI -- whose project was discussed
- Department*
- Position* (e.g. graduate student, faculty, staff, etc.)
- Service* -- what area of BRC was relevant (e.g. HPC, cloud, etc.)
- Topic* -- keywords for what was discussed
- Hand-off or referral -- who the consultation was referred to, if applicable

- Notes

The consulting log served as a key source of data for BRC's annual report, particularly as a way of providing a different lens onto uptake of computational resources rather than simply looking at utilization of HPC cycles. The Consulting program added to BRC's disciplinary breadth across academic disciplines -- many of the disciplines that showed only trivial usage of the HPC cluster benefited substantially from consulting engagements, becoming "repeat customers".

Actively keeping track of consultations also had immediate value, enabling the consulting service manager to monitor how much time was being spent with each group. The BRC Consulting Handbook established a guideline of 5 hours of consulting per quarter, per research group. Further consultation in a given quarter, while possible, had to first be approved by the consulting service manager, to ensure that the engagement remained within the scope of "consulting" as understood by BRC, rather than spiraling into a more time-and resource-intensive technical advising or development project.

Another piece of shared documentation established by the Consulting Handbook was "backgrounders", which served as both a script of prompts for a consulting engagement, and a way of collecting notes from the engagements in a standardized way that would make it easier to hand off engagements between consultants as needed. Backgrounders were based on a Google Docs template that captured next steps, and completed steps for the consulting engagement at the top. Below that were the sections that could serve as a script to guide the domain consultant's meeting with the researcher, or at least help structure their notes:

- Overview
- Use case
 - Research statement
 - Problem statement
 - BRC recommendations / solutions
- Issues
- Other options
- Referral process
- Initial emails
- Notes
- Additional reference materials

Ensuring consistent uptake of backgrounders was much more challenging than doing the same for the consulting log, with domain consultants in particular resisting the "paperwork-heavy" nature of the backgrounders, even after stating that many of the fields were optional.

As time passed, BRC became increasingly reliant on Asana, a web-based project management platform, for monitoring and tracking both tasks and consultations. Using Asana, every engagement (a set of activities, including but not limited to meetings, with a given research group around a given topic) became a "task", and follow-up steps became sub-tasks that could be assigned to anyone in BRC, given due dates, and tracked

separately. Within Asana, engagements could be grouped by consultant, and tagged as “Open - Active”, “Open - Inactive”, “Open - Waiting on client”, and “Closed”, giving the consulting service manager easier visibility on the status of engagements, as well as their distribution among the consulting team. Given the primacy of Asana for work tracking, backgrounders were ultimately deprecated in favor of simply adding comments to the Asana task for a given engagement, and the consulting log was retired once BRC obtained a paid Asana subscription that supported the use of custom fields[1] for tasks and data export.

4.2 Consulting Meetings

Another important aspect of the consulting practice was the set of meetings that aimed to foster domain consultants’ engagement with the larger program, and sense of teamwork among the group. Domain consultants were encouraged to attend biweekly BRC staff meetings that provide visibility into all aspects of the program, including those that the domain consultants may have less exposure to. Attendance was also required at a biweekly BRC consulting meeting, which provided an opportunity to plan trainings and events, discuss current and upcoming issues with the HPC cluster (e.g. downtimes, upgrades), and for the service manager to review open consultations with the group. These reviews of open consultations allowed the service manager to identify trends, and for the group to share successful approaches to working with researchers. In addition, the consulting meetings were a chance to re-balance the workload between consultants, to accommodate the domain consultants’ primary role as grad students (e.g. in the lead-up to comprehensive exams or dissertation chapter due dates).

The weeks alternating with the consulting meetings have varied in their use, ranging from planned internal trainings (often in anticipation of the domain consultants leading a public training on the same topic) to one-on-one check-ins with the domain consultants and various members of the BRC staff.

5 CONSULTING PROGRAM EXPANSION, PART ONE

In summer 2016, we hired two additional domain consultants: a graduate student in the Environmental Science, Policy & Management (ESPM) Department, and a recent UCB graduate with a PhD in chemistry.

5.1 Career Placement Success

In fall 2017, one year after her hire, our domain consultant who had recently graduated with a PhD in chemistry, was offered a position as a research facilitator at the University of Michigan medical school. As with many PhDs in the sciences, she developed her technical proficiency largely outside the context of coursework, through a combination of self-teaching, MOOCs, and D-Lab Cloud, Text Analysis, and Machine Learning working groups focused on technical topics. Nonetheless, her year of experience as a BRC domain consultant put her at a significant

advantage when applying for full-time research facilitation jobs, and left her well-prepared for such a position.

6 CONSULTING PROGRAM EXPANSION, PART TWO

During summer 2017, the architect who was leading the consulting service turned the consulting service manager role over to a new BRC technologist, who had recently become more involved with BRC after working closely with the campus digital humanities program for a number of years. That fall, new leadership in the Research Data Management (RDM) program (a sister program to BRC under the umbrella of UC Berkeley’s Research IT group) and an interest in partnership with Lawrence Berkeley Lab’s IT Division led to consideration of possible expansions to the BRC Consulting program.

While the Lab’s IT division had extensive experience with HPC user services and desktop support, department head, aspired to develop a “science accelerator” program that could more directly engage with researchers through consultation with domain consultants drawn from a pool of graduate students. To ensure sufficient availability of consultants across both UC Berkeley and the Lab, the Lab offered to fund three additional domain consultant positions. Filling these positions has been very difficult. At the time of publication of this paper, these three positions are still unfilled.

In contrast, RDM already offered consulting, but as an increasing number of RDM consultations overlapped with BRC due to a compute-related facet, it made sense to standardize and largely consolidate the consulting practices between the two groups to make consultations more seamless for researchers and reduce the administrative overhead of managing two separate consulting programs. The Consulting Program service manager reviewed the consulting handbook with RDM program staff, making updates and changes both to reflect the reality of the consulting practice as implemented, and to accommodate RDM’s different staffing model. This in turn allowed the RDM program to expand its consulting pool to include a diverse group of librarians who may otherwise have limited involvement with the program and experience working with BRC. In addition to the expansion with RDM, BRC also attempted to hire several more domain consultants, as mentioned above. After interviewing several candidates— graduates students as well as research staff — we discovered the challenge of finding the right combination of incentives to hire non-graduate students. There were several candidates that we wished to hire, but due to several factors only one person accepted the offer to join, a graduate student in Plant & Microbial Biology.

7 CHALLENGES FOR THE BRC DOMAIN CONSULTANT MODEL

In many ways, the graduate student domain consultant model used by BRC is a tremendous success. It has reduced staffing costs and increased disciplinary diversity, and has even had 50% representation by women. Furthermore, to date, all domain

consultants have stayed at least one year. Nevertheless, successful outreach, and maintaining a steady pipeline of potential applicants, have been ongoing challenges.

Three of the four current and former domain consultants were identified through campus working groups, and personally encouraged to apply. BRC staff have established relationships with the campus graduate professional development center, as well as student-run organizations for post-PhD employment such as Beyond Academia, and have recently made efforts to work through those groups to identify students who are interested in holding jobs. However, the fact remains that for many advisors, holding employment other than a graduate instructor or graduate research assistant is a distraction from a student's primary responsibilities, and out of concern for BRC maintaining a good relationship with researchers, approval from a graduate student's advisor is a requirement for employment as a domain consultant.

Additionally, even when finding a qualified candidate whose advisor approves of the request for the graduate student's time, there are sometimes other challenges, such as the right combination of incentives. During the hiring process of one very viable candidate who was very interested in working with us, the incentive structures we could provide did not meet key requirements, so the candidate declined in the end. The candidate was seeking tuition fee remission as part of the compensation package, however our unit— an administrative unit at the time— was not able to offer fee remissions as incentive, whereas other organizations such as D-Lab, an academic unit rather than an administrative unit, were able to offer it.

Other BRC outreach efforts to support greater campus awareness of the program in general (e.g. developing relationships with departmental and divisional chief academic officers, creating flyers and handouts tailored for specific disciplines and presenting those at departmental meetings) may have a positive impact on domain consultant recruitment. Larger-scale campus efforts to destigmatize holding internships or jobs while completing a graduate program may also have a positive effect. But the difficulties BRC has experienced in attempting to rapidly scale up its consultant pool in response to interest from the Lab should serve as a cautionary tale for others considering this model, to treat outreach and promotion of domain consultant positions as an ongoing activity, in order to ensure a ready pool of interested and eligible candidates who can step into the role as turnover or new funding opportunities emerge.

8 CONCLUSION

Hiring graduate students in a fractional appointment (20-25%) with disciplinary knowledge and some degree of existing technical skills, and providing them with training and mentoring to help address gaps in their technical knowledge, is a viable cost-saving approach for doing research facilitation with constrained resources. Success with such a model is likely due to the active ongoing implementation of a shared set of consulting practices and structures, as well as continual outreach efforts to establish and maintain contact with potential future candidates. Benefits of this model accrue not only to the campus community where the

graduate students consult with researchers, but also to the larger ecosystem of research computing organizations that can draw from these graduate students when looking to hire early-career research facilitators.

REFERENCES

- [1] Devon Watts and Nick Fassler, 2016. Now rolling out in beta: Track anything in Asana, with custom fields, <https://blog.asana.com/2016/07/custom-fields-beta/>
- [2] John Towns, Timothy Cockerill, Maytal Dahan, Ian Foster, Kelly Gaither, Andrew Grimshaw, Victor Hazlewood, Scott Lathrop, Dave Lifka, Gregory D. Peterson, Ralph Roskies, J. Ray Scott, Nancy Wilkins-Diehr. 2014. "XSEDE: Accelerating Scientific Discovery", *Computing in Science & Engineering*, vol.16, no. 5, pp. 62-74, Sept.-Oct. 2014, doi:10.1109/MCSE.2014.80
- [3] David A. Anisi. 2003. *Optimal Motion Control of a Ground Vehicle*. Master's thesis. Royal Institute of Technology (KTH), Stockholm, Sweden.
- [4] James Bottum, Maureen Dougherty, Gwen Jacobs, Scott Yockel, Lauren Michael, James Cuff, Paul Wilson, Advanced Cyberinfrastructure - Research and Educational Facilitation: Campus-Based Computational Research Support, <https://aciref.org/facilitators/facilitation/>
- [5] Lauren Michael and Bruce Maas. 2016. "Research Computing Facilitators: The Missing Human Link in Needs-Based Research Cyberinfrastructure". *ECAR Research Bulletin*. May 16, 2016. <https://library.educause.edu/~media/files/library/2016/5/erb1602.pdf>