

# An Expansion of the User Support Services for the Research Computing Group at the University of Colorado Boulder

Shelley L. Knuth<sup>†</sup>  
Research Computing  
University of Colorado Boulder  
Boulder, CO USA  
[Shelley.Knuth@colorado.edu](mailto:Shelley.Knuth@colorado.edu)

Daniel Trahan  
Research Computing  
University of Colorado Boulder  
Boulder, CO USA  
[Daniel.Trahan@colorado.edu](mailto:Daniel.Trahan@colorado.edu)

Timothy Dunn  
Research Computing  
University of Colorado Boulder  
Boulder, CO USA  
[Timothy.Dunn@colorado.edu](mailto:Timothy.Dunn@colorado.edu)

Joel Frahm  
Research Computing  
University of Colorado Boulder  
Boulder, CO USA  
[Joel.Frahm@colorado.edu](mailto:Joel.Frahm@colorado.edu)

Zebula Sampedro  
Research Computing  
University of Colorado Boulder  
Boulder, CO USA  
[Zebula.Sampedro@colorado.edu](mailto:Zebula.Sampedro@colorado.edu)

Thomas Hauser  
Research Computing  
University of Colorado Boulder  
Boulder, CO USA  
[Thomas.Hauser@colorado.edu](mailto:Thomas.Hauser@colorado.edu)

Andrew J. Monaghan  
Research Computing  
University of Colorado Boulder  
Boulder, CO USA  
[Andrew.Monaghan@colorado.edu](mailto:Andrew.Monaghan@colorado.edu)

Aaron Holt  
Research Computing  
University of Colorado Boulder  
Boulder, CO USA  
[Aaron.Holt@colorado.edu](mailto:Aaron.Holt@colorado.edu)

Patrick Burns  
Information Technology  
Colorado State University  
Fort Collins, CO USA  
[Patrick.Burns@ColoState.edu](mailto:Patrick.Burns@ColoState.edu)

## ABSTRACT

In any research computing (RC) environment, the role of the user support group is a combination of education, consulting, and outreach. As this role is the most public-facing of a research computing group's team, it is important to ensure an excellent level of support is provided to users. The Research Computing group at the University of Colorado (CURC) Boulder maintains a large-scale computing cluster with several hundred active users, among other services, and has done so since 2011. The user support group at CURC provides a variety of services intended to support these users. This paper describes those services the CURC group provides, as well as explores the various ways that these services have been improved in 2018-19. The impact on users is also assessed.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

PEARC '19, July 28-August 1, 2019, Chicago, IL, USA  
© 2019 Copyright is held by the owner/author(s).  
ACM ISBN 978-1-4503-7227-5/19/07.  
<https://doi.org/10.1145/3332186.3332229>

## CCS CONCEPTS

- Human centered computing • Human-computer Interaction (HCI) • HCI design and evaluation methods

## KEYWORDS

Research Computing, User Support, Trainings, Education, Allocations, Case Management

## ACM Reference format:

Shelley Knuth, Joel Frahm, Andrew Monaghan, Daniel Trahan, Zebula Sampedro, Aaron Holt, Timothy Dunn, Thomas Hauser, Patrick Burns. 2019. An Expansion of the User Support Services for the Research Computing Group at the University of Colorado Boulder. In *Proceedings of PEARC19 (PEARC'19)*. ACM, New York, NY, USA, 6 pages.

## 1 Introduction

The University of Colorado (CU) Boulder Research Computing group (CURC) manages and maintains several large-scale computing resources, including a 480+ node supercomputer called the Rocky Mountain Advanced Computing Consortium (RMACC) Summit, as well as a 2.6 PB data storage system called the PetaLibrary. In addition to these resources, the CURC group offers several user-facing services that utilize the underlying hardware, including a JupyterHub server and an EnginFrame portal to an underlying visualization cluster. The user support branch of the CURC group is responsible for ensuring user access to these resources is expedient and as effortless as possible. Several user support services have been in place to enhance the CURC user experience. In 2018-19 these services

were expanded upon and improved to include additional trainings, better access to trainings, a new documentation portal, streamlined account and allocation services, changes to consultations, new partnerships with on- and off-campus groups, and an initiative to collect user support data with the purpose of improving the user experience. This paper describes these enhanced services and how they have improved the user experience.

## 2 Improvements on general user support

The CURC group supports users from three disparate groups 0 CU, Colorado State University (CSU), and RMACC. In the past year, approximately 730 users were active on the RMACC Summit system. To ensure users understand how to use the available resources, CURC offers several avenues by which users can receive support: trainings, an email-based ticketing system (ServiceNow), and consultations [1]. CURC trainings are always offered in person in classrooms on campus; however, instructors broadcast most trainings online as well. In 2018 and 2019 CURC offered over fifty trainings on a variety of topics both high-performance computing (HPC) and non-HPC related, such as "Introduction to Python", "Containerization for Research and Development Applications", and "Fundamentals of HPC". Over 400 people attended these trainings offered through Research Computing from 2018 through March 2019. Often more users attend online rather than in person, indicating that online offerings are an important way of broadening participation. CURC advertises these workshops via a number of avenues, including a user email list, partner email lists, and through a Meetup group. The training topics offered by CURC are selected based upon user requests or selectively by user support staff based on tickets, consultations, or common errors encountered while running jobs on the system.

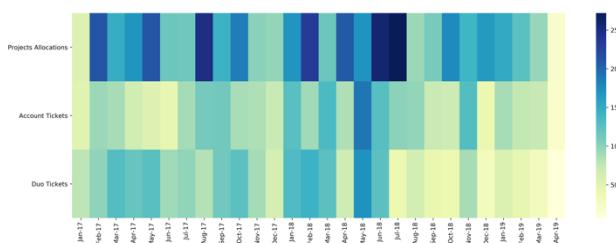
In addition to trainings, CURC began offering office hours in mid-2018 as part of a collaborative effort with several other partner organizations on campus (see Section 3). This effort began with only one staff member attending office hours for one hour per week, but the number of people attending (approximately 150 people since 2018) motivated increased staff support. We now have 2-3 staff members attend for two hours per week for several weeks of the year. Attendance at office hours varies greatly, and at times the CURC staff only sends one person for the two-hour time period. The office hours, along with the standard one-on-one consults, have served approximately 150 users since CURC began office hours. The office hours themselves account for approximately 20-25% of these users.

Several user-related processes were improved in April 2018 through automation and improvement to the user portal. These processes have included the account and allocation request procedures. The account creation process was previously a manual procedure where users would request an account, the user support group would receive the request via ServiceNow, and respond by manually sending an email to the user within 24 hours (usually much less). The email from CURC user support would then ask the user to sign up for two factor authentication and ask the user to send in another request for a startup, or "general" allocation of core hours. This email was long and difficult for users to understand, especially for new research computing users. Often users would get lost in this process and not request an allocation or two factor authentication, resulting in more tickets and a loss of productive time. In addition to a cumbersome back end, the front-end procedure to request an account on the CURC website was also challenging for users. There were several questions asking users' reasons for requesting CURC resources, which were particularly difficult for new users to answer. The new process has the following steps: first, the user verifies their identity by authenticating against a trusted identity server for their organization. Next, the user answers a minimal number of questions relevant for reporting. The user then submits this request via the website. Once the user is verified, an account is created, and the user is automatically added to the CURC general allocation. An email (via ServiceNow) is sent confirming that the account has been created, and providing basic information about CURC resources in a much more compact format compared to the original email. Unfortunately, setting up two-factor authentication still requires manual intervention due to external issues. However, in recent months alternate procedures have likely made automation of this process possible as well, and will be soon implemented. The entire process, while originally intended to provide the user support group with a way to communicate with end users early in the process, proved tedious and time consuming for both staff and users.

In addition to streamlining the account request process, the allocations process was also improved. All users on CURC resources receive a "general" allocation upon setup, which allows them to share 20% of CU's allocation on RMACC Summit (CU's allocation is approximately 67 M core hours per year) via FairShare. Users can move to the remaining 80% of CU's RMACC Summit allocation by submitting a longer allocation request wherein the user has performed tasks such as benchmarking and ensuring code has been optimized for RMACC Summit. Previously, the higher priority allocation process meant that users would have to read through and decipher a twelve-page document using complex terminology.

Complaints about submitting allocation requests were high among users, leading to numerous tickets.

To examine the impact of these changes a heatmap (Figure 1) was generated showing the number of tickets received in ServiceNow by month using certain associated keywords. The three topics chosen are: Project allocations, Account Tickets, and Duo Tickets. Recall that in the prior account creation process three emails were sent - one to create the account, one to ask to be set up with two-factor authentication (Duo), and one to ask for a general allocation. As well, several emails were sent back and forth between the user and user support requesting clarification on larger allocation requests. After the April 2018 transition account requests were done via an automated email and no additional emails about two-factor authentication or general allocations were needed.



**Figure 1. Heatmap showing some of the cases generated by users in the ServiceNow ticketing system by keyword from 1 January 2017 through 31 March 2019.**

Examining the Account Tickets panel of Figure 1 indicates that there were a minimal number of account requests early in RMACC Summit's implementation, but overall the number of account creation ticket requests was steady over the remaining time period, with the exception of the period immediately near the procedure transition. This is an expected result because tickets were still automatically generated in the account creation process even after the transition. However, Duo Tickets noticeably decreased after June 2018. Also in June 2018, CURC removed VASCO tokens as an option for two factor authentication, and a number of tickets were initiated by CURC in May and June requesting a change to Duo for those using Vasco. But the number of tickets asking for assistance from Duo is decreased from before the transition.

Allocations are more difficult to assess, as general allocation requests and the large "Projects Allocations" are also included. Overall, there appears to be a decrease in the number of tickets after July 2018, but the tickets are still high. However, CURC

has seen a 31% increase in allocation requests since the new procedure has been implemented, so the reduction in tickets is not due to a reduction in allocation requests. An examination of the average time to approval of allocation requests before the transition was 31 days, compared to 6 days after the transition. Co-author Frahm, the head of the CURC allocations committee, has noted that since the online template has been utilized it is much easier to decipher users' requests.

Finally, one last improvement to CURC user support services in 2018 was expanded and revamped documentation. Previously, CURC's documentation had been outdated with no revisioning policy in place. The pages were written directly on the CURC website, and it was an archaic process to update each page. Multiple people would edit pages without a procedure in place, leading to disjointed pages and incoherency. In addition, users would frequently submit tickets pointing out outdated documentation referring to old systems or code that did not work properly. The CURC user support team overhauled the documents pages, transferred them to Read the Docs (<https://readthedocs.org/>), and implemented a continuous improvement process to ensure the documents stay current. One user support staff (co-author Trahan) is now responsible for reviewing the documentation monthly, and a user support committee reviews and approves any changes to the pages.

The documentation changes also included an overhaul of the website. Many of the pages were out-of-date, too long, or suffered from staff making simultaneous changes. The old website, which is maintained in a different system than the documentation, was also difficult to edit. After a significant review process, a new website was implemented that was a significant improvement over the old site. These changes will be assessed via a usability study conducted in early summer 2019.

### 3 Collaborations

Local and regional partnerships have been central to CURC's endeavors since its creation. For example, NCAR and CU Anschutz were partners on Janus, the first CURC supercomputer. Additional institutional collaborations have begun since 2011. These collaborative efforts have created mutually beneficial cyberinfrastructure efforts.

#### 3.1 Colorado State University

A joint CU/CSU proposal submitted to the National Science Foundation (NSF) was awarded in 2015 to procure a new HPC system (grant #1532236; [2]). This system, RMACC Summit, first came online in January 2017 to a subset of CURC users, and became fully operational in May 2017.

During 2017 CSU staff members were partially redirected to make RMACC Summit available to the CSU community during the system's initial deployment. In January 2018, as part of the NSF grant objectives, CURC user support staff began working closely with CSU to develop a new user support structure. Initially, a former CURC user support staff member was hired to assist with setting up a case management system, answering questions already being sent in by users, and offering consultations to some of the early RMACC Summit users. In January 2018 a "Summit on Summit" was held at CSU to introduce the broader community to the resource. By May 2018, a temporary CSU user support staff member was hired and was dedicated to addressing CSU user needs. The staff member was housed within the CURC user support group at CU Boulder to learn first-hand how to implement user support offerings for RMACC Summit. The staff member offered online office hours, online consults, and answered cases. The benefit of remaining at CU in Boulder (approximately 45 minutes south of CSU) was that the CSU user support member could gain experience from and collaboratively address problems with CU user support and system administration staff. The CSU user support person also made monthly visits to CSU and offered in-person consultations, training, and office hours. The user support team leads at both CURC and CSU, in addition to this staff member, would also meet bi-monthly to discuss strategies to develop and enhance the user support program at CSU.

This collaboration offered an interesting opportunity to examine the importance of user support in HPC operations. From January 2017 until January 2018, RMACC Summit was online and both CURC and CSU users were able to access and utilize the system. However, without a permanent user support structure in place at CSU only a few users were aware of the system. Only 36 CSU users utilized the CSU allocation between 1 January and 30 June 2017. By contrast, 236 users utilized CSU's allocation between 1 January 2017 and 5 April 2019, including 71 in the most recent month (5 March-5 April 2019). The influence of user support can be seen in the CSU allocation consumption over time. From 1 January to 31 December 2017, before CURC began assisting CSU with user support, only 31% of CSU's allocation was being used. In 2018, significant staff time was dedicated to assisting CSU users, and allocation usage increased to 84%. In January 2019, CSU's primary user support staff member left the university, leaving

a gap in the services until a new hire was made at the end of March 2019. During the gap user support was limited to answering tickets, and use of the CSU share dropped to 48%.

An important goal is increasing the use of the CSU share. It is speculated that there are several reasons why a larger increase was not realized during 2018 when the CSU staff member was in place. First, the new hire was not physically at CSU; it was eventually determined that having someone onsite would provide more expedient user support services. Second, the CSU user support staff member was only 50% time, and wasn't able to do as much promotion of computing and consulting services across the CSU campus. A new hire is now in place for CSU, is housed on the CSU campus, and is a full-time staff member. CSU and CU user support staff still meet bi-weekly to ensure coherence between the two groups. An updated report on the share of RMACC Summit utilized by CSU will be presented at PEARC, but it is expected to increase through the remainder of 2019 as user services come back online.

The CSU/CURC partnership thus far has resulted in establishing a consultation service, various trainings, and the establishment of a formal ticketing system for CSU. In addition to benefits to CSU users, CURC users and the CURC user support staff have benefitted as well. First, CURC user support staff, in assisting with building out the CSU user support structure, reviewed policies and ideas about the current implementation of user support, resulting in changes to CURC user support procedures. Second, the CU users benefitted from some of the new services that were spun up for CSU users. One example of this was containerization. CURC user support staff learned software containerization in order to assist bioinformatics users at CSU in deploying complex software pipelines on RMACC Summit. This additional capacity was subsequently used to provide containerized solutions for numerous CU users, as well as trainings on Singularity containerization available to both universities.

### 3.2 Rocky Mountain Advanced Computing Consortium

CU Boulder leads the Rocky Mountain Advanced Computing Consortium (RMACC). RMACC is a group of institutions in the Rocky Mountain West that either have HPC systems onsite or have researchers or students requiring access to large-scale computing resources. Over time, RMACC has led to numerous collaborative activities among members that have benefitted the CURC user community, as well as the users at other RMACC institutions. One of these is the NSF CyberTeam project. The

CURC group has collaborated with CSU and the University of Utah on an NSF Cyberteam project assisting users on each campus with establishing better data workflows. Workflow facilitators with complementary knowledge meet with users to ensure personnel on each campus are managing data in the most beneficial way for their projects. While specific to data related topics, this project (funded through NSF grant #1659425) has had an impact on the user support offerings available to all RMACC HPC users, as well as providing access to a HIPAA compliant data storage system. A detailed description of this project will be given in an oral presentation by co-author Monaghan entitled, "Improving regional cyberinfrastructure services through collaboration: Cyberteam for the Rocky Mountain Advanced Computational Consortium".

RMACC Summit is also available to RMACC members, with CURC providing user support. Ten percent of the user cycles go to RMACC members. There are presently about 40 RMACC users across six institutions outside of CU/CSU. Efforts are ongoing to continue reaching out to users at the RMACC institutions, but bringing in more users from these regions has been challenging. The CURC group has held office hours, town halls, and offered trainings online – which have been advertised through email lists and at the annual RMACC conference – but less participation than expected has occurred. This may be partly due to a lack of dedicated RMACC user support staff; rather, it is part of the normal duties of the CURC staff to reach out to RMACC users. Further efforts on engaging RMACC users are being explored.

### 3.3 Center for Research Data and Digital Scholarship

Perhaps of most direct benefit to CURC users has been our collaboration with the CU Boulder Libraries through the Center for Research Data and Digital Scholarship (CRDDS) [3]. The mission of the center is to be a hub for data-related resources on campus and to foster collaboration, learning, and development on data and digital humanities related topics. The center accomplishes this in several ways – trainings, consultations, resource offerings, and outreach. Four initiatives that are branches of the center – Cyberinfrastructure, Research Data Management, Education and Training, and Digital Scholarship – work collaboratively to achieve the center goals. Collaborations with the Libraries and other groups across campus through CRDDS has allowed CURC to greatly expand training offerings and consultation expertise. This partnership has also allowed for exposure of other topics and events to our user base.

Since its inception in 2017, CRDDS, along with a partnership with the Laboratory for Interdisciplinary Statistical Analysis (LISA), has offered trainings such as, "Introduction to GIS", "Parallel R", "Statistics in R", "OpenRefine", "Digital Humanities Fundamentals – Esri Story Maps", "Introduction to Web Scraping", and "Introduction to Geocoding". These trainings were advertised among each of the collaborators' email lists and websites, as well as through a collaborative list and website, reaching several entities across campus that might not otherwise have known about the workshops. In 2017 and 2018 alone there were over 1,000 attendees at CRDDS offered workshops. In addition to the trainings, CRDDS offers events, such as meet-and-greets, and consultations, including office hours. The office hours are offered with personnel from the Libraries, Research Computing, LISA, and another campus group – Earth Lab, giving attendees a wide variety of expertise to draw from when seeking assistance. In 2017 and 2018 over 400 people attended these office hours or sought out one-on-one consultations from the group. Resources, such as small- and large-scale data storage, assistance with writing data management plans, or assistance with data publishing are also available through CRDDS. The CURC group contributed extensively to the trainings and consultations, and has several separate offerings as noted above.

## 4 User Support Data Analytics

As with any service effort, an analysis of the benefits to the users should be made. In 2018 and 2019, CURC underwent a study to explore the types of users taking advantage of CURC resources, how they were using these resources, and if any modifications should be made. To do this, CURC began a project to analyze data from several sources - the ticketing system, Slurm job statistics, trainings, and consultations. CURC is also using this data to examine who is using the system, and what pain points they may experience. This data collection is still in the preliminary stages, but thus far we have been able to begin evaluating the quality of the user experience. A few early results are presented below.

CURC's primary user base (in terms of jobs run on RMACC Summit) are researchers and students in the Physics and Engineering departments. These users comprise of 53% and 24% of the total users, respectively (from 1 January 2017-1 January 2019). CURC does have users in other departments that are not as traditionally HPC-centric, such as Integrative Physiology, Political Science, and Athletics. Unfortunately, challenges to our data collection system prevents a breakdown of personnel by status (faculty, etc.) for who is running jobs on

RMACC Summit. However, data shows that graduate students are the largest group attending trainings and workshops, as well as receiving assistance through consultations and office hours.

After an assessment was completed of *who* is running on the system one was completed of *what* is being run on the system; in other words, what types of jobs are being submitted. An examination of Slurm job statistics from 1 January 2017 through 1 January 2019 was performed. Table 1 shows Slurm statistics for certain types of job sizes. "Did not run" are jobs that failed; "Single node, single core" are jobs that were allocated one node and one CPU; "Single node, multi-core" are single-node jobs with multiple allocated CPUs; "Midsize" jobs are 1-4 node jobs; "Large" jobs are 4-16 node jobs; and "XL" jobs are 16+ node jobs. These sizes were subjectively selected after a cursory examination of the data.

This data shows that most of the jobs run on RMACC Summit over this time period were single node, single core jobs (52%), while the least were XL jobs. In fact, the vast majority (89%) were single node jobs. This is was not surprising to the user support team given anecdotal knowledge of incoming tickets and discussions during consultations. To that end, a new Slurm Quality of Service (QOS) was setup in 2018 designed for users to run short-duration, single-node jobs. The user support team wanted to improve the user experience by providing a separate QOS for these jobs so that users might be able to complete jobs quicker, rather than waiting in QOSes that allowed for longer wait times. In total, over 5 million jobs were run in 2017 and 2018 accounting for a total of 204 million CPU hours over this time period.

	Jobs submitted	Total Wall time (days)
Did not run	78289	0
Single node, single core	2892577	392088
Single node, multi-core	2020370	313587
Midsize	235800	25884
Large	256513	36512
XL	36515	8345
Total	5520064	776416

**Table 1. Slurm statistics showing the number of jobs submitted and the total wall time in days from 1 January 2017 through 1 January 2019.**

Table 1 also shows the total wall time in days for jobs in each of the six categories listed above. The total wall time is how many total days jobs ran in each category over the course of these two years. Again, single node jobs occupied the most time over the course of the two years. This data would indicate that ensuring good support around the smaller jobs would be the beneficial to the most users.

However, while most users may be running jobs that occupy less infrastructure, the users that are running larger jobs consume the largest share of resources. Table 2 shows the average number of nodes and CPUs used broken out by job size. While larger jobs may only represent 10% of the total jobs being run, these jobs still use 95% more compute resources than the single node jobs. The amount of resources being used by these larger-sized jobs was much higher in 2018 than 2017, as shown in Table 2. This is likely because the overall number of jobs was much higher in 2018 than 2017 as the number of people creating accounts on the system increased. An exploration of how efficient these larger jobs are being run on the system is forthcoming. As well, the amount of memory being used by these jobs was two to three orders of magnitude higher than (in MB) for the larger jobs as compared to the single node jobs.

	Average Num. Nodes		Average Num. CPUs	
	2017	2018	2017	2018
Did not run	0	0	0	0
Single node, single core	1	1	1	1
Single node, multi-core	1	1	15	18
Midsize	2	2	37	56
Large	6	5	193	143
XL	27	28	619	559

**Table 2. The average number of nodes and CPUs used for each job size from 1 January 2017 through 1 January 2019.**

Next, CURC explored the amount of time users spend running jobs, and spend waiting in the queue. Table 3 shows the average runtime of jobs in 2017 and 2018, as well as the “average overestimation”. The “overestimation” refers to how long the user requested the job to run compared to how much time the job actually needed to run. Users can either opt to use the default time for the QOS and partition to which they are submitting, or they can set this number explicitly. The results of this analysis are interesting and a bit challenging to explore. In 2017, the average elapsed runtime of jobs showed a fairly steady increase as the number of nodes increased. This is not unexpected because larger jobs are likely to be more complex. However, the results are not as straightforward for 2018. While the longest elapsed time for jobs occurs in the XL size category, matching what is shown in 2017, large jobs on average ran in a shorter amount of time than the single node jobs. Furthermore, the average runtime for a single job was generally shorter for 2018 than 2017.

	Average job runtime (hours)		Average over estimation, default & explicit (hours)	
	2017	2018	2017	2018
Did not run	0	0	0	0
Single node, single core	2.93	1.72	14.4 15.0	6.3 14.8
Single node, multi-core	3.09	1.77	10.8 11.3	4.6 8.0
Midsize	3.15	2.16	6.5 8	4.8 10.4
Large	4.61	1.60	7.4 7.9	4.7 8.0
XL	6.20	3.95	6.9 7.0	5.3 8.1

**Table 3. Slurm statistics showing the average amount of elapsed time in each job size category.**

Further examination of this required. However, there are a few interesting items that occurred between these two years. First, more jobs ran in 2018 than in 2017. This may be because there were more users on the system in 2018 than 2017 as RMACC Summit became more well known. It might be assumed, but

not directly correlated, that the increase in the number of trainings offered as well as more outreach activities such as office hours contributed to the increase in accounts created, and therefore jobs. As well, nearly two-thirds more jobs were run in the single node categories in 2018 compared to 2017. There was a nearly even amount ran in the largest job categories between the two years, and in the Midsize category two-thirds more jobs were run in 2017 than in 2018. One possible reason for this is that in RMACC Summit’s early adoption phase, unit testers were hand-picked by the CURC team to do testing on the system. These users continued to run as RMACC Summit came online, likely dwarfing the smaller jobs submitted. As more users created accounts on RMACC Summit in 2018 it is possible that these users were running shorter-duration jobs than their predecessors.

An additional interesting statistic from Slurm is the average overestimation of jobs. This is an important statistic to note because the more wall time a user requests from Slurm the longer (typically) they will wait in the queue before their jobs run. Table 3 shows how many hours users are typically overestimating by job size in both 2017 and 2018. One item to notice is that users allowing the default time to determine how much wall-time was requested usually fared worse in 2017 than 2018. One possible reason for this is that changes to the default wall time were made in mid-2017 when this number was reduced from 24 hours to 4 hours. Therefore, the user was only requesting four hours of wall time if using the default rather than the maximum allowed of 24 hours, leading to a large decrease in the amount of possible requested wall time. This dramatic change is not seen when comparing the explicit requests between 2017 and 2018. Another item to note is that, generally speaking, users overestimated more hours on average when explicitly requesting a wall time rather than allowing the default to be used.

CURC user support also examined the types of software being used by the users. This is useful because it can help determine what software might be of most use to maintain as part of our software stack. Data has shown that the Intel and GCC compilers are most readily used. It is then difficult to ascertain which software are most used compared to each other; however, NetCDF, Python, and R software is frequently loaded by users on RMACC Summit.

## 5 Summary

This paper presents the current status of user support services offered by the University of Colorado Boulder Research Computing group. Several CURC user support services

underwent a major overhaul recently. These included an improvement upon the account creation and allocation request processes. Previously, both processes required a significant amount of manual intervention from CURC user support staff and users, costing time and productivity. A thorough review of these procedures was completed in 2018. Improvements to documentation and the website were made as well.

CURC users have benefitted from several collaborations, both external and internal, throughout 2018. Some of these collaborations are long-standing, while others have been developed within the past year. A collaboration between CU and CSU has resulted in a new implementation of user support services at CSU, where previously there had been few services offered. This collaboration has benefitted users at CU as well - by motivating CURC user support staff to review and improve current processes, and by teaching staff members new tools, such as containerization. This has resulted in a knowledge transfer between CURC staff and users. Other collaborations with RMACC and CRDDS have resulted in benefits to the CURC user community. Particularly, CRDDS has given CURC the opportunity to provide more learning opportunities to the greater CU community, has provided a greater selection of trainings for CURC users, and has resulted in more direct knowledge transfer to users via avenues such as office hours.

Finally, a preliminary assessment of user services was conducted using results from a user data analytics project that is currently underway. An exploration of the types of users and the types of jobs run was completed. CURC also examined the amount of resources used by jobs, including wall times and average runtimes. In the future, CURC would like to explore other pain points for users such as the failure rates of jobs and why they are unsuccessful, and the average wait time for jobs in the various QOSes. A usability study with users wherein in-person feedback is solicited is also planned, in order to acquire first-hand knowledge of problems from the perspectives of users.

## ACKNOWLEDGMENTS

The authors would like to acknowledge The National Science Foundation (grant nos. #1532236 and #1659425), for their support.

## REFERENCES

- [1] S. Papich, S. Knuth, P. Ruprecht, T. Hauser, 2017. The impacts of outreach efforts by Research Computing at the University of Colorado Boulder. *ACM*, <http://dx.doi.org/10.1145/3093338.3093377>.
- [2] J. Anderson, P.J. Burns, D. Milroy, P. Ruprecht, T. Hauser, and H.J. Siegel, 2018. Deploying RMACC Summit: An HPC Resource for the Rocky Mountain Region. *Proceedings of PEARC17, New Orleans, LA, USA, July 09-13, 2017*. DOI: 10.1145/3093338.3093379
- [3] S.L. Knuth, A. Johnson, T. Lindquist, D. Weiss, D. Hamrick, T. Hauser, and L. Reynolds, 2016. The Center for Research Data and Digital Scholarship at the University of Colorado-Boulder. *Bull. Assoc. Information Sci. and Technology*, 2, 46-48.