

Ask.Cyberinfrastructure.org: Creating a Platform for Self-Service Learning and Collaboration in the Rapidly Changing Environment of Research Computing

Julie Ma MGHPCC Holyoke, MA jma@mghpcc.org	Torey Battelle Colorado School of Mines Boulder, CO battelle@mines.edu	Katia Bulekova Boston University Boston, MA ktrn@bu.edu
Aaron Culich UC Berkeley Berkeley, CA aculich@berkeley.edu	John Goodhue MGHPCC Holyoke, MA jtgoodhue@mghpcc.org	Jacob Pessin Boston University Boston, MA jpessin@bu.edu
Vanessa Sochat Stanford University Stanford, CA vsochat@stanford.edu	Dana Brunson Internet2 Ann Arbor, MI dbrunson@internet2.edu	Tom Cheatham University of Utah Salt Lake City, UT tec3@utah.edu
Sia Najafi Worcester Polytechnic Institute Worcester, MA snajafi@wpi.edu	Chris Hill Massachusetts Institute of Technology Cambridge, MA cnh@mit.edu	Adrian Del Maestro University of Vermont Burlington, VT adrian.delmaestro@uvm.edu
Bruce Segee University of Maine Orono, ME segee@umaine.edu	Ralph Zottola University of Alabama Birmingham, AL rzottola@uab.edu	Scott Valcourt University of New Hampshire Durham, NH sav@cs.unh.edu
Zoe Braiterman OWASP New York, NY Zoe.braiterman@owasp.org	Raminder Singh Harvard University Boston, MA R_singh@g.harvard.edu	Robert Thoelen Pratt & Whitney Rthoelen@ieee.org
		Jack Smith West Virginia Research Jack.smith@wvr.org

ABSTRACT

Ask.CI [3], the Q&A site for Research Computing, was launched at PEARC18 with the goal of aggregating answers to a broad spectrum of questions that are commonly asked by the research computing community. As researchers, facilitators, staff, students, and others ask and answer questions on Ask.CI, they create a

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. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Copyright ©JOCSE, a supported publication of the Shodor Education Foundation Inc.

© 2021 Journal of Computational Science Education
DOI: <https://doi.org/10.22369/issn.2153-4136/12/2/9>

shared knowledge base for the larger community.

For smaller institutions, the knowledge base provided by Ask.CI provides a wealth of knowledge that was previously not readily available to scientists and educators in an easily searchable Q&A format. For larger institutions, this self-service model frees up time for facilitators and cyberinfrastructure engineers to focus on more advanced subject matter. Recognizing that answers evolve rapidly with new technology and discovery, Ask.CI has built in voting mechanisms that utilize crowdsourcing to ensure that information stays up to date.

Establishing a Q&A site of this nature requires some tenacity. In partnership with the Campus Champions, Ask.CI has gained traction and continues to engage the broader community to establish the platform as a powerful tool for research computing. Since launch, Ask.CI has attracted over 250,000 page views (currently averaging nearly 5,000 per week), more than 400

contributors, hundreds of topics, and a broad audience that spans the US and parts of Europe and Asia.

Ask.CI has shown steady growth in both contributions and audience since it was launched in 2018 and is still evolving. In the past year, we introduced *Locales*, which allow institutions to create subcategories on Ask.CI where they can experiment with posting institution-specific content and use of the site as a component of their user support strategy.

CCS CONCEPTS

- General and reference → Document types → General literature •
- Human-centered computing → Human computer interaction (HCI) → Interaction paradigms → Web-based interaction

Keywords

Q&A, Research computing, Self-service learning

1. BACKGROUND

Ask.cyberinfrastructure.org is a collaborative, crowd-sourced Q&A site specifically curated for the research computing community. The project began in September 2017 with the vision of constructing a resource that allows the research computing community to more quickly find answers to commonly asked questions by way of a shared, public knowledge base, resulting in better/faster science results.

The goal of the project is to aggregate answers to a broad spectrum of questions that are commonly asked by researchers and educators as they utilize advanced computing and data resources. The result is a self-service knowledge base for domain researchers, facilitators, cyberinfrastructure (CI) engineers, and others who comprise the research computing community. The hope is that this site will become the go-to platform for sharing frequently asked questions, comparing solutions, and building on previous work pertaining to research computing. Making this knowledge readily available in the public domain will free up time for facilitators and CI engineers to focus on more advanced subject matter, thereby elevating the research computing practice. Simultaneously, the platform allows users to apply an andragogic approach to explore the information on the site at their own pace.

As the project launched, it rapidly drew the enthusiastic support of the XSEDE Campus Champions leadership team. We have collaborated since December 2017 to build the site and launched it together in July 2018 at PEARC18. A naming contest at the conference yielded “ask.cyberinfrastructure.org”, which we then nicknamed “Ask.CI” (available at <https://ask.ci/>). The site has been well-received, with nearly 5,000 page views per week, over 400 registered contributors, and hundreds of topics.

Establishing a Q&A site of this nature requires some tenacity. We have gained some traction and hope to continue to engage the broader community to firmly establish this platform as a tool for the global research computing community. Ask.CI has further inspired thinking and awareness about the importance of the subject matter. Throughout the development process, we have been thinking frequently about what defines “research computing” in relation to other computing disciplines. The hope is that not only will Ask.CI become a great resource for the community, but that it will also provide public testimony of the importance of research computing and how it exists in relation to enterprise IT, computer science, and domain research.

1.1 StackExchange and Discourse

As we investigated possible technologies upon which to build the site, there was consensus about using a platform that supports a voting mechanism that enables crowd-sourced monitoring and pushes the best answers to the top, with moderation tools to manage spam/trolling. This led us to StackExchange [2], the gold standard for Q&A platforms. StackExchange is the platform behind Stack Overflow and many other widely used Q&A sites. While it is not easy to establish a StackExchange site, it offers several advantages, including search engine visibility, built-in backup and maintenance, security and resistance to spam and trolling, voting, and a clear question-and-answer syntax that yields definitive answers. Launching a StackExchange site involves a rigorous, four-phase process, including a restart if any phase exceeds a specified time limit. In our first iteration, which closed in May of 2018, we reached the second phase and attracted a working group of volunteers who developed questions and answers to post on the site in preparation for when it became functional. Since our StackExchange site was not yet functional, we curated these on another platform called Discourse [1], an open source Internet forum, and we developed a methodology for culling nuggets of information from ad hoc user questions and adapting them for use by a general audience. While Discourse is not as ubiquitous as StackExchange, it is a very flexible platform with a low startup threshold.

When our StackExchange effort was closed, we decided to formalize the Discourse content, add a voting mechanism to mimic Stack Exchange functionality, and launch our Q&A site on the Discourse platform. Subsequently, we discovered several benefits of using the Discourse platform over StackExchange. These include: 1) having flexibility in the question format to include discussion topics as well as Q&A and 2) having the ability to set up categories, which we are using to create institution-specific “locales,” described below.

1.2 Broadening Participation

The idea for Ask.CI originated with the NSF-sponsored Northeast Cyberteam Program, which aims to make research computing more accessible to small/medium sized institutions in northern New England. Ask.CI became an integral part of the strategy. By creating a shared, public knowledge base composed of content which is often found behind the firewall at large institutions, Ask.CI enables researchers at smaller institutions to become more self-sufficient, reducing the need for Research Computing Facilitation support.

2. Ask.CI 2020

Since launching at the 2018 Conference on the Practice & Experience in Advanced Research Computing (PEARC18), we have nurtured Ask.CI into an active, growing site managed by a dedicated group of volunteer site moderators who meet weekly via Zoom. The primary purpose of the meeting is to actively curate the site and discuss outreach activities. Active curation includes reviewing new posts, checking for unanswered topics, considering new subject matter areas to cover, and planning weekly marketing activity. While Ask.CI has shown steady growth in both contributions and audience, finding methods to continue to grow audience participation is an ongoing focus of our attention, as the expert research computing knowledge that we seek to gather is widely distributed among the community. We have employed several methods to do this, described below.

2.1 Events, Meetings and Conference Calls

We actively seek opportunities to talk about Ask.CI with audiences at venues where the research computing community has congregated. Since launch at PEARC18, we have conducted Birds of a Feather (BoF) sessions at all subsequent PEARC and Super Computing (SC) conferences. We have also presented on the CI Brown Bag, XSEDE Campus Champions, Campus Research Computing Consortium (CaRCC), and EDUCAUSE Research Computing (RCD) calls, and at regional gatherings whenever possible.

2.2 Question of the Week

Each week at our site moderators' Meeting, we review new content that has been posted on the site and content that has not been answered. If we find a topic that has been unanswered for over a week, we will likely mark it a question of the week (QoW). QoWs are emailed to the Campus Champions mailing list and tweeted to the Ask.CI twitterverse. A goal for this year is to expand the recipient list beyond the Campus Champions, to other groups in the Research Computing ecosystem that might be able to answer the questions.

2.3 Friday Factoid and Sunday Science

In addition to the QoW, we occasionally tweet Friday Factoids, interesting tidbits of relevant material about the research computing world, particularly if they are timely with calendar or current events. We also post Sunday Science stories, which are more domain-specific deeper dives. All three of these methods function as reminders that attract the community back to the site, in part to see if new topics have been posted, or simply as a reminder that Ask.CI exists and that new content is always welcome.

3. INFRASTRUCTURE

3.1 Q&A and Discussion Zone

One of the most stringent requirements of building a site using StackExchange is that questions must be written in a manner that there can be a well-defined, best answer. This is necessary so that the voting mechanism, which is required to ensure that content stays up to date, can work correctly.

In research computing, where work is frequently in unexplored territory, sometimes it is not clear if there is a best answer, or how this can be evaluated. Often, a dynamic and relevant discussion among subject matter experts can help to formulate an answer or a response to a particular situation. For this reason, we created the Discussion Zone, where discourse in response to a question can take place. We distinguish topics that go in the discussion zone from those that have a clear-cut best answer, which are categorized as Q&A. At present, for the most part, we have not found a need to further categorize questions by subject matter or domain, and in fact believe it to be useful to have topics on wide-ranging subject matter in a single category, as this will encourage cross-pollination of solutions among different domains.

3.2 Tags

To facilitate searching for answers that pertain to particular subject matter on the site, questions are tagged with labels that identify subject matter and other characteristics of the content. Users can also search by tag to obtain a listing of all topics tagged with this particular label. This mechanism is used both to delineate content, and also to aggregate it when appropriate. There are currently 216 active tags on the site, many of which are only used a few times. There are also tags to indicate the audience for a

particular topic. Following the CaRCC model of audience delineation, these are researcher, system, and data, representing topics of interest to researcher-facing, systems-facing, and data-facing facilitators.

3.3 Voting

Faced with ever-changing technology advances, one of the key strategies for keeping content up to date on the site is a voting mechanism modeled after the voting function on StackExchange sites. Users can vote for the "best" answer to a given question, and they can also vote on "best" questions. The software then rearranges topics so that topics and questions with the highest number of votes appear first. Unlike the StackExchange model, our system only allows "up" votes, which creates a more convivial environment for participants. The intended result is that over time when new content is posted on a topic that renders other answers obsolete, voting will ensure that the most relevant answer appears first.

4. EXPANSION VIA LOCALES

In spring 2019, one of the Ask.CI moderators observed that there could be significant benefits to having institution-specific content on the site. We introduced a program that allows institutions to create subcategories on Ask.CI, dubbed "locales," where they can experiment with posting institution-specific FAQs and using the site as a component of their user support strategy. The intent is that by sending users to Ask.CI via the institution-specific sandbox, it will encourage them to start down a path of self-service learning, simultaneously encouraging user-to-user collaboration both within the institution's own user community and across the research computing community as a whole. One of the key benefits of having this exploration occur under the umbrella of Ask.CI is the simplicity of migrating content from the main site to a locale and vice versa.

We piloted the locale concept from April to November 2019 with six institutions and formally announced the Locales program at SC19. As of April 2020, there are 13 locales in service: Brown University, Colorado School of Mines, the Computing against Covid-19 Project, Harvard University, MGHPCC, the Northeast Cyberteam, Northeastern University, Ohio Supercomputer Center, the ResearchSOC Cybersecurity community of practice, Stanford University, Tufts University, University of Maine, and Yale University. Other institutions are in the process of starting theirs up: MIT, University of Alabama, University of Missouri, University of New Hampshire, University of Vermont and several other institutions. The CaRCC consortium and US-RSE have expressed interest in building locales later in 2020. Locale moderators join our Ask.CI site moderators' call once a month to check in and exchange ideas. We are also developing a toolkit to facilitate integration with existing support platforms (websites and ticketing systems) that are often behind institutional firewalls.

5. EVALUATION AND METRICS

A standing item on the site moderators' weekly meeting agenda is to review certain statistics to monitor the health of the site, including page views (daily, monthly, and aggregate), users, and return visits. We also periodically count the number of unique institutions represented, the topics by audience type (researcher-facing, systems-facing, data-facing, end-user), and the total number of locales.

6. LESSONS LEARNED AND FUTURE PLANS

Ask.CI is entering its third year of existence and is growing at a steady rate. We have learned a great deal about the effort required to establish and maintain a site of this nature. A few key observations are noted below.

6.1 Creating a Workflow Shift Takes Time

The process of creating a public record of content that is typically disclosed on a private mailing list or behind a firewall requires a shift in mindset and habits. Although most people understand the value of creating this knowledge base, it is an extra step, and not all individuals with the knowledge are readily able to take the time to post topics or answers on a regular basis. In the upcoming year, we hope to introduce a program which creates incentives to contribute, by recognizing the effort and the value of the information already posted.

6.2 Discourse Flexibility Allows Creativity in Outreach Not Possible with StackExchange

As described in Section 1, we began this effort thinking that we would build this site on the StackExchange platform. While we were initially disappointed that our first effort was terminated, we have found the flexibility of being able to support a discussion zone and locales has been a happy outcome of shifting to the Discourse platform. This year, we hope to put mechanisms in place to address one of the key benefits that StackExchange platforms enjoy, which is the high Google ranking that they inherit from StackExchange. The inherited ranking can put content from other sites earlier in search results than Ask.CI content, even if the Ask.CI content is more relevant to the topic being searched.

The methods that we have in mind for this year would put Ask.CI on a more equal footing with sites that currently have this advantage.

6.3 Convivial Weekly Meetings Have Yielded a Dedicated Group of Moderators and Allow for Ebb and Flow of Individual Workloads

The weekly site moderators' meetings comprise a group of seven volunteers who have been the backbone of Ask.CI. The site would not have matured to the level that it has reached without their tireless efforts. In a given week, the number of people at the meeting will vary, so it has been very beneficial to have a large enough group to be able to make forward progress each week.

6.4 Outreach is Key

The Question of the Week, Friday Factoid, and Sunday Science communications have had a noticeable impact on our page view statistics each week. We hope to reach out to other communities and mailing lists this year to expand our presence and welcome any suggestions/recommendations from the community.

6.5 Cross-Posting to Specialized Communities of Experts for Specific Topics

We have identified certain specialized boards, both in the community and at vendors, that have been great sources of one-off content when the need has arisen. Finding a systematic way to keep those groups engaged with our site will have long term positive results.

6.6 Locales Have Great Potential When Institution/Community of Interest is Ready

Although we are in early stages with many of our locale partners, we have had very promising initial results. We hope to capitalize on this concept with further outreach to the community to find other institutions that are ready to create a locale. In addition, building the integration tools mentioned above will facilitate the onboarding of other institutions and create a more seamless experience for the end users.

6.7 Lessons Learned and Future Plans

We welcome and encourage feedback and participation on the Ask.CI platform. Participation can take many forms, from simply reading content on the site and posting a reply occasionally, to full participation on the site through a locale.

7. ACKNOWLEDGEMENTS

This work has been partially supported by the National Science Foundation under grant award ACI-1659377. The authors also thank the many the Ask.CI site moderators and Locale moderators and volunteers whose efforts to build and maintain the site are beginning to yield significant results, and the hundreds of contributors who have generously shared knowledge and experience on Ask.CI.

8. REFERENCES

- [1] "Discourse — Civilized Discussion." 2020. Discourse — Civilized Discussion. Discourse. Accessed June 16, 2020. <https://www.discourse.org/>.
- [2] "Stack Exchange." 2020. Accessed June 16, 2020. <https://stackexchange.com/>.
- [3] "Ask.Cyberinfrastructure." 2020. Ask.Cyberinfrastructure. Accessed June 16, 2020. <https://ask.cyberinfrastructure.org/>.
- [4] Goodhue, J., Ma, J., Del Maestro, A., Najafi, S., Segee, B., Valcourt, S., Zottola, R. Northeast Cyberteam:Workforce Development for Research Computing at Small and Mid-sized Institutions. Proceedings of the Conference on Practice and Experience in Advanced Research Computing (PEARC20), Portland, OR, July 26–30, 2020, doi:10.1145/3311790.3396662.
- [5] Goodhue, John, et al. "Northeast Cyberteam Program – A Workforce Development Strategy for Research Computing." The Journal of Computational Science Education, vol. 11, no. 1, 2020, pp. 8–11., doi:10.22369/issn.2153-4136/11/1/2.