

What do we (not) know about Research Software Engineering?

Authors:

* These authors co-led this work and are the corresponding authors.

Name	Affiliation	ORCID
Anna-Lena Lamprecht*	Utrecht University	0000-0003-1953-5606
Carlos Martinez-Ortiz*	Netherlands eScience Center	0000-0001-5565-7577
Michelle Barker*	Research Software Alliance	0000-0002-3623-172X
Sadie L. Bartholomew	National Centre for Atmospheric Science and University of Reading	0000-0002-6180-3603
Justin Barton	University of London	0000-0002-1244-9484
Neil Chue Hong	University of Edinburgh	0000-0002-8876-7606
Jeremy Cohen	Imperial College London	0000-0003-4312-2537
Stephan Druskat	German Aerospace Center (DLR)	0000-0003-4925-7248
Jeremy Forest		0000-0002-8486-5559
Jean-Noël Grad	Institute for Computational Physics, University of Stuttgart	0000-0002-5821-4912
Daniel S. Katz	University of Illinois	0000-0001-5934-7525
Robin Richardson	Netherlands eScience Center	0000-0002-9984-2720
Robert Rosca	European XFEL	0000-0001-5796-7587
Douwe Schulte	Utrecht University	0000-0003-0594-0993
Alexander Struck	Cluster of Excellence Matters of Activity	0000-0002-1173-9228

Marion Weinzierl	Durham University, UK	0000-0002-2302-5476
------------------	-----------------------	---------------------

Abstract

As recognition of the vital importance of software for contemporary research is increasing, Research Software Engineering (RSE) is emerging as a discipline in its own right. We present an inventory of relevant research questions about RSE as a basis for future research and initiatives to advance the field, highlighting selected literature and initiatives. This work is the outcome of a RSE community workshop held as part of the 2020 International Series of Online Research Software Events (SORSE) which identified and prioritized key questions across three overlapping themes: people, policy and infrastructure. Almost half of the questions focus on the people theme, which addresses issues related to career paths, recognition and motivation; recruitment and retention; skills; and diversity, equity and inclusion. However, the people and policy themes have the same number of prioritized questions. We recommend that different types of stakeholders, such as RSE employers and policy makers, take responsibility for supporting or encouraging answering of these questions by organizations that have an interest. Initiatives such as the International Council of RSE Associations should also be engaged in this work.

1. Introduction

Research Software Engineering (RSE) is emerging as a discipline in its own right. The term RSE was coined around a decade ago to recognize the vital importance of software for contemporary research, and the role of people, policy and infrastructure in its development, support and maintenance. Research Software Engineering is now an increasingly recognized term and substantiated knowledge about its aspects and fields of activity is essential for its further development. Early RSE initiatives often relied on personal experiences and anecdotal evidence to gain support. Now there is increasing (empirical) research being undertaken into RSE, which provides substantiated evidence and insights to support the advances in the field.

To get a better understanding of the current body of knowledge and open research questions about RSE, we organized a community workshop titled “What do we (not) know about RSE?” (1) within the 2020 Series of Online Research Software Events (SORSE) (2). The workshop aimed to bring together members from the international RSE community to collect research questions and available scientific literature. In this article we present the outcomes of the workshop, including the crowd-sourced inventory of relevant research questions, including pertinent literature and initiatives, that is available for reference and reuse (see table in the Supporting Information).

The remainder of this article is structured as follows. In Section 2 we describe the workshop setup and the method/process of collecting and treating research questions and available literature. In Section 3 we summarize the results and discuss selected observations and insights. In Section 4 we formulate recommendations on the utilization of this work and next steps. Finally, Section 5 concludes the paper with a summary and perspectives for future work.

The materials produced for, during and following the workshop are available at <https://docs.google.com/document/d/1fACmxmxEJYjPWWdHIABV3FhIBg9UjMI-1AFKfRo3F4s/edit?usp=sharing>.

2. Method/Process

We organized the workshop as part of SORSE (2020), an initiative of international RSE associations to provide an opportunity for the Research Software Engineers (RSEs) to develop and grow their skills, build new collaborations and engage with RSEs worldwide during the Covid19 pandemic. The event was promoted via the communication channels of the international RSE communities, consequently, participants most likely had RSE backgrounds. Participation was free of charge and open to anyone interested. 27 people from at least seven countries took part in one of the two online workshop sessions on 28th and 30th October 2020 (repeated to cater for different time zones).

Three invited short talks (pre-recorded) set the scene for the workshop discussions: Simon Hettrick (Software Sustainability Institute) presented an estimate on “How many RSEs?” there are in the world; Daniel S. Katz (University of Illinois) reported on experiences with “Forming and Supporting RSE Groups and Communities”; and Zhian Kamvar, Toby Hodges and Serah Rono (The Carpentries) discussed “Research Software Engineers and The Carpentries”.

The next part of the workshop was a World Cafe (3) to collect answers to the question of “What we do (not) know about RSE”. Participants were randomly split into three breakout groups, each focussing on one of the three themes: people, policy, infrastructure. These themes were chosen as they are key themes for other initiatives in the sector, including the Research Software Alliance and the Software Evidence Bank (Software Sustainability Institute, n.d.). The challenges of theory-software translation utilized somewhat similar themes, with questions identified in the areas of design, infrastructure and culture (4). They can be defined as follows:

- **People:** RSE personnel, whether explicitly employed as RSEs or not even aware of the title, but performing a similar role. This category will be addressing issues including RSE career paths, recognition and motivation; recruitment and retention; skills; and diversity, equity and inclusion.
- **Policy:** The policies surrounding RSE, both internally within organizations, and externally from national bodies and funding organizations. This category will be addressing challenges related to recognition, funding, and demonstrating the importance and impact of RSE.
- **Infrastructure:** The infrastructure used by RSEs, including software tools, (shared) hardware platforms, and code sharing platforms. This category will encompass topics including barriers to RSE reuse of code; identifying commonly used productivity tools and code sharing platforms; and constraints in carrying out RSE tasks.

After 20 minutes, the groups changed to another of the topics, and again after another 20 minutes, so that all participants had a chance to contribute to all areas. Participants were asked

to work together to brainstorm and record interesting research questions in the three themes, the motivation for asking them, and if applicable, to provide links to existing research that (at least partially) address the questions. The collection happened collaboratively and in real time in a set of shared online documents (1).

Following a coffee break, participants were again randomly split into three breakout groups, this time to discuss a prioritization of the collected questions for one of the three themes. After 30 minutes the groups reported back in the plenary session to share their results. After a brief discussion of the planned follow-ups, the workshop was closed.

After the workshop, we aggregated the notes from the two workshop sessions into a single document per priority area. For two weeks we shared these with the participants and asked them to revise and comment, also allowing for adding of further questions and references to existing literature. Following this community consultation period, we clustered and synthesized the collected questions, to integrate similar issues and remove overlaps. The resulting 65 questions for all themes were also divided into high priority and low priority questions.

From these we generated an online form and asked the participants to up/down vote the (prioritized) questions from the different areas. The survey was available to participants for three weeks, and 15 participants responded. Based on these responses, we further reassigned questions as “high priority” or “low priority” where there was a high level of agreement that they should be re-assigned.

3. Results and Discussion

This section details the final listing of the questions, analysis of these, and challenges to validity of the outcomes.

3.1 Final questions list

In addition to a series of working documents that record a lively discussion, the final outcome of the workshop is a table containing 65 questions (provided in the table in the supplementary material), divided into the three themes of people, policy and infrastructure, and further classified as high or low priority. A high level overview of the final breakdown of questions is provided in Table 1.

Themes	Prioritization	Total
People	High	10
	Low	20
Policy	High	10
	Low	5

Infrastructure	High	6
	Low	14

Table 1: Breakdown of prioritized questions across the three themes

The following questions, or groups of closely-related questions, were prioritized:

3.1.1. People

1. Why do people become RSEs? How do we highlight the importance of RSEs and of selecting an RSE career to potential candidates for RSE roles? What makes these careers attractive? What are the considerations RSEs take into account to decide to take a role (in academia vs outside of academia)? Salaries? Progression opportunities?
2. What background do RSEs have? Where do they come from?
3. How are RSEs recruited? How are they found?
4. Why do RSEs quit and what do they do after they are RSEs? Do available career progression paths help keep people in the career?
5. How can an RSE career track have branching (e.g., not just go from senior to manager type roles). How do we create enough structure to allow for professional and personal growth, and recognition of that work?
6. What are the disadvantages to being labeled RSE?
7. What would RSEs have liked to learn before starting their RSE careers? What skills will someone starting an RSE career in five years need?
8. Should RSE training be part of the curriculum of every student?
9. How do we prevent waste of resources and effort from people moving on, or to other projects, such that software becomes unmaintained and unusable?
10. How much of the RSE work needed is done by staff (e.g. postdocs) in unstable positions?

3.1.2 Policy

1. What proportion of funding should be dedicated to the infrastructure required for generic software projects?
2. Which funders (if any) support grants for long-term sustainability for software infrastructure?
3. How are RSE groups funded?
4. Does stating in the funding guidance that RSEs can be put on a research funding proposal increase the likelihood of RSEs being put on fundable/successful grants?
5. How crucial are RSEs to research projects, short and long term? What is the value of RSEs to an institution? What contribution do RSEs provide to research? What research results evaluation schemes exist that 'acknowledge' RSE work? Are there any alternative

ways of defining success for RSEs? What is the return on investment for hiring an RSE? How could this be quantified?

6. What are suitable merit evaluation schemes/metrics for RSEs?
7. How many RSE groups are in each country (in relation to the number of universities/research institutes)?
8. How do we get policy to support re-use and sustainability rather than “novel innovative” re-inventing the wheel?
9. What are the financial/reputational costs of NOT providing career paths to RSEs?
10. How can we get more RSE expertise in the project proposal review process? Software management plans should be reviewed by experienced RSEs.

3.1.3 Infrastructure

1. What software should be preserved and/or maintained? What criteria should be used to decide on this?
2. Do we have the right infrastructure to publish research software (together with the relevant data)? What needs are (not yet) covered?
3. Do we have the right infrastructure to find research software once it has been published? How do people make their software findable at the moment? Do they?
4. What would motivate people to stop re-implementing established solutions (in terms of writing software packages for which there are existing implementations)?
5. How do the infrastructure needs of RSEs differ from those of software engineers/developers outside of research, and those of standard researchers?
6. What are the bounding constraints RSEs encounter most often? Which one is most critical (and what is the spread over different types of resources)?

3.2 Analysis of questions

This section analyses patterns in the final list of questions, particularly the prioritized questions, and links findings to related work identifying issues that need further investigation of relevance to RSE.

3.2.1 Emphasis on people theme

The first thing to observe is that the *people* theme is the largest, having almost double the number of questions of the other two themes (30/65). The number might indicate that this is the theme area where most research is required; however, the *people* and *policy* themes also have the same number of prioritized questions. We also note that two-thirds of the *policy* questions are considered priorities, in comparison to a third for each of the *infrastructure* and *people* questions.

Going deeper into the questions themselves, it can be seen that all of the people-themed questions are centered around RSE career paths, with training, recruitment and retention of talent being identified as an issue. RSE careers are also central to a majority of the policy

questions (and also play a role in infrastructure). To some extent this may reflect that many of the survey participants were probably RSEs, because this is the target audience of SORSE events.

There has been an increasing research emphasis on the people undertaking software development, and this is not surprising as the evolution of the RSE community has focused on RSE recognition and career paths (5–9). A similar workshop on building the research innovation workforce identified 12 thematic challenges in problem areas involving: diversity and inclusivity; fostering the development and support of the workforce ecosystem and talent pipeline; establishing viable career paths and normative role descriptions in the workforce; enhancing internal and external communications and education for stakeholders; compensation; workforce sustainability; the establishment of an identity of the field as a discipline; the position of research computing within institutional organizations; and the need for continuing training and education for professionals (10).

Equity, diversity and inclusion is another important aspect within RSE. Work such as that undertaken by Chue Hong, et al. (11)—highlights evidence for a lack of diversity within the RSE community. This work also highlights potential interventions and examples of approaches that can contribute towards supporting enhancements in equity, diversity and inclusion for research software engineering. One of the prioritized questions from the people theme focused explicitly on this topic, and others on topics such as RSE recruitment, retention and community development could be argued to consider it implicitly.

Five of the prioritized questions in the people theme also reference community, one of the four pillars of RSE identified by Cohen et al. (12). The other three pillars of RSE have some alignment with the themes used here; training is encompassed within the people theme used here, policy is the same, and software development aligns somewhat with the infrastructure theme. Community is often highlighted as particularly important to open source software, to enable innovation and sustainability. It could be argued that community fits under the people theme, noting that community can also contribute to the development of policy and how infrastructure is provisioned and used.

3.2.2 Overlaps across the three themes

Our approach to gathering questions was to classify them into three separate themes. However, there is inevitable overlap between the themes, as illustrated by Figure 1.

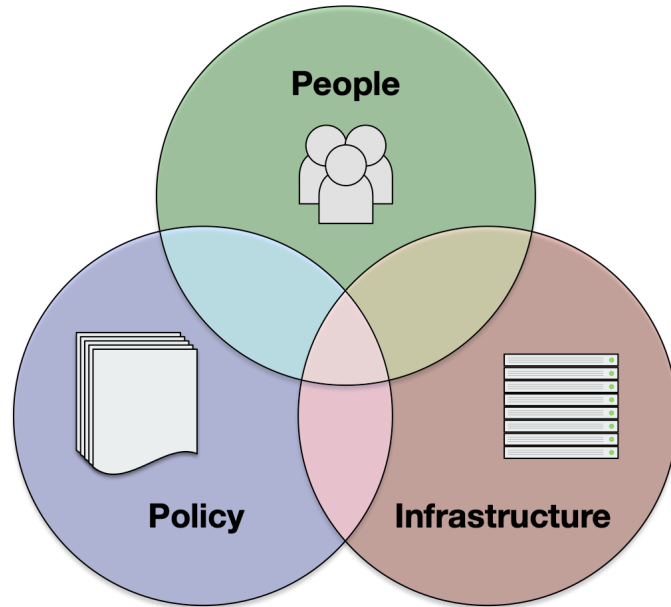


Figure 1: Interrelationship of the three themes

There are strong links and overlaps between the three themes given that infrastructure exists to support people in undertaking their work and policies exist to help ensure that people and infrastructure can operate safely, securely and effectively. For example, software sustainability is affected by the skills and motivation of the RSEs that develop it; the software’s sustainability may be incentivized by, or evaluated against, relevant policy, and may then be included in relevant infrastructure such as a repository.

At the centre of the diagram, where all three themes intersect, we have the combination of individuals, the software they produce or use, the infrastructure that they work with and the policies that guide the way the individuals and the infrastructure work. While the high-level overlaps between these areas are clear, the effects they have on the RSE landscape in individual domains, communities or institutions are much more complex to predict or understand. As such, the approach of considering the groups separately for the purpose of crowdsourcing questions that highlight what we still need to know about RSE is the most practical approach and we can see from the wide array of questions raised that there is already much to understand within the individual themes. Looking at how these questions affect other themes, through overlaps with them, would be useful work for future analysis.

3.2.3 Relevance of existing literature

For many of the questions (24 out of 65, or 37%), the table in the supplementary material references existing literature that touches upon or contextualizes the issues raised here, or looks at some aspects of the questions that may be relevant to answering the broader question. 10 out of 26 (or 38%) of the prioritized questions point to existing literature; however, we

observe that preliminary work towards answering the most pressing issues is no more advanced than for all the questions as a whole.

3.3 Challenges

The major threat to the validity of this analysis is the sampling of participants, as the workshop was held as part of the SORSE events which naturally results in a high proportion of active RSEs. For a better representation of the current issues this workshop would need to be updated and repeated, possibly with a broader audience, including policy makers from government, funders and research organizations. As a consequence, the results of this study should be seen as major questions about RSE from the viewpoint of RSEs.

It should also be noted that the list of existing resources included in the table in the supplementary material is not exhaustive, as it was based on crowd-sourcing rather than a formal literature review. An outcome of this work will be the future inclusion of identified literature in the Software Sustainability Institute's Open Evidence Bank (13), a curated collection of articles and data that contribute to understanding of the research software landscape. The Open Evidence Bank's aims are to create an open registry of relevant research, ensure that research is easily discoverable and accessible by the community, and provide evidence to underpin policy and best practice. It is therefore an ideal place to deposit literature collections such as those that emerged from this workshop.

4. Recommendations

This section provides recommendations on how RSE stakeholders could utilize this work, and suggested next steps.

4.1 Relevance for RSE stakeholders

To encourage answering of at least the prioritized questions, it would be useful to categorize them further by identifying which stakeholders are best positioned to facilitate this. High-level analysis suggests that the organizations that employ RSEs would gain the most from insights related to the people theme questions, as the first seven of the ten prioritized questions focus on RSE recruitment, upskilling, recognition and retention. Whilst these could be addressed by individual organizations, it would be significantly more beneficial if considered at national, international or disciplinary levels, and are thus potentially relevant to governments, disciplinary consortiums and/or university associations. This information would also be advantageous to the funders and policy makers who could make use of it to incentivize changes in how the system works, based on the resulting understanding of what change is needed.

The policy themed questions are most relevant to policy makers and funders by their very nature, but vary considerably in focus. The first five of the ten prioritized questions relate to aspects of funding, including broader questions on how to demonstrate the value of investing in RSE roles to maximise research impacts. Some of the prioritized policy questions focus on the

policy aspects, such as recognition, motivation and funding for RSEs, whilst others highlight the need for information on demographics. Three of the six prioritised infrastructure questions for the infrastructure theme identify questions related to infrastructure to enable reproducibility of software. Another suggests the need for better understanding of the differences between the infrastructure needs of RSEs and software engineers outside academia, pointing to the need for comparison with other sectors.

4.2 Recommendations on next steps

It is recommended that relevant stakeholders consider addressing the priority questions that have been identified by the workshop participants as a first step towards enhancing the capabilities of the RSE community to improve research outcomes. It would be valuable to involve the recently formed International Council of RSE Associations, and the (currently seven) national RSE associations that it encompasses, to engage with this analysis. It should be noted that there are also a range of other institutions, communities or initiatives that already have research projects of relevance to some of the priority areas (which the list of relevant research assists in illuminating) that could be supported or encouraged to focus specifically on some of these issues.

5. Conclusion

The process of crowd-sourcing a prioritized inventory of research questions about RSE and the resulting analysis has yielded valuable results as a basis for future research and initiatives to advance the field. Classification into the three overlapping themes of people, policy and infrastructure proved useful for enabling initial observations, such as a strong emphasis on people-themed questions relating to career paths, training, recruitment and retention of RSEs. This exercise has also facilitated identification of literature that provides context to the identified questions and/or begins to address these questions. However, it is clear there is still much to learn in this field, and there are a range of stakeholders who would benefit from addressing these questions. These include the organisations employing RSEs, and the policy makers incentivising change in the sector. We recommend that further work is undertaken by relevant stakeholders to advance addressal of these questions.

Acknowledgements

Input was also provided by the following workshop participants: Mary Chester-Kadwell, University of Cambridge; Adrian D'Alessandro, Imperial College London; Julia Damerow, Arizona State University; Niels Drost, Netherlands eScience Center; Teri Forey, Freelance RSE; Krishnakumar Gopalakrishnan, University College London; Simon Hettrick, Software Sustainability Institute; Christine Rogers, McGill University; Malin Sandström, International Neuroinformatics Coordinating Facility; Vanessa Sochat, Stanford University; Thomas White, Deutsches Elektronen-Synchrotron Hamburg;

Funding Acknowledgements

Michelle Barker: Wellcome Trust grant 222436/Z/21/Z; Jeremy Cohen: UKRI EPSRC grant EP/R025460/1; Jean-Noël Grad: Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) Research Software Sustainability grant 391126171 (PI: Holm); Alexander Struck: Cluster of Excellence »Matters of Activity. Image Space Material« funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2025 – 390648296.

References

1. Lamprecht A-L, Martinez C, Barker M. What do we (not) know about RSE?: SORSE Workshop October 28 & 30, 2020 [Internet]. 2020. Available from: <https://tinyurl.com/rseresearch>
2. SORSE [Internet]. SORSE. 2021 [cited 2021 Mar 26]. Available from: <https://sorse.github.io/>
3. Hurley TJ, Brown J. Conversational leadership: Thinking together for a change. *The Systems Thinker*. 2009;20(9):2–7.
4. Jay C, Haines R, Katz DS, Carver JC, Gesing S, Brandt SR, et al. The challenges of theory-software translation. *F1000Res*. 2020 Oct 2;9:1192.
5. Akhmerov A, Cruz M, Drost N, Hof C, Knapen T, Kuzak M, et al. Raising the Profile of Research Software. 2019 Aug 27 [cited 2021 Jun 1]; Available from: <https://zenodo.org/record/3378572>
6. Anzt H, Bach F, Druskat S, Löffler F, Loewe A, Renard BY, et al. An environment for sustainable research software in Germany and beyond: current state, open challenges, and call for action. *F1000Res*. 2021 Jan 26;9:295.
7. Brett A, Croucher M, Haines R, Hettrick S, Hetherington J, Stillwell M, et al. Research Software Engineers: State Of The Nation Report 2017 [Internet]. 2017 Apr [cited 2021 Mar 26]. Available from: <https://zenodo.org/record/495360>
8. Cohen J, Woodbridge M. RSEs in Research? RSEs in IT?: Finding a suitable home for RSEs. *arXiv:201010477 [cs]* [Internet]. 2020 Oct 20 [cited 2020 Oct 24]; Available from: <http://arxiv.org/abs/2010.10477>
9. Hardey, Leng. Linking Professional Skills to RSE Career Paths. 2020 Oct 8 [cited 2021 Jun 1]; Available from: <https://zenodo.org/record/4073299>
10. Arafune L, Brunson D, Hacker T, Smith P. Building the Research Innovation Workforce: A workshop to identify new insights and directions to advance the research computing community. [Internet]. 2020. Available from: <https://www.rcac.purdue.edu/ciworkforce2020/report/report.pdf>
11. Chue Hong N, Cohen J, Jay C. Understanding Equity, Diversity and Inclusion Challenges Within the Research Software Community. *SE4Science, ICCS21: International Conference on Computational Science* [Internet]. 2021; Available from: Forthcoming
12. Cohen J, Katz DS, Barker M, Chue Hong N, Haines R, Jay C. The Four Pillars of Research Software Engineering. *IEEE Softw*. 2021 Jan;38(1):97–105.
13. Software Sustainability Institute. Open Evidence Bank [Internet]. [cited 2021 Mar 26]. Available from: <https://www.software.ac.uk/open-evidence-bank>

Supplementary Data

List of questions

<https://docs.google.com/spreadsheets/d/1uHi0t7foLBzY53FitPqbQZ-q44AnHyxFwchJx2ArKDQ/edit?usp=sharing>