

OpenStack Gender Diversity Report

Daniel Izquierdo, Nicole Huesman, Alexander Serebrenik, Gregorio Robles

The OpenStack Foundation's goal is to *"promote the global development, distribution and adoption of open infrastructure."* As in many other open source communities, and in the technology industry as a whole, the OpenStack community reflected a lack of representation of females and underrepresented minorities, a fact that should be supported with evidence. Intel and Bitergia have conducted research to assess the current state of gender diversity within the community, examining both code and non-code contributions, including leadership, governance, and event representation, among other elements. In this paper, the authors summarize the results of this research. They discuss the importance of using data to increase awareness of this issue and to guide subsequent actions to improve the diversity and inclusion within the OpenStack community. This is the first comprehensive analysis of gender diversity within the OpenStack community, and has opened the possibility of extending insights from this research to other communities.

Why Gender?

Engagement of people of different genders in software engineering is heavily skewed towards men; multiple studies show that the situation in open source is even worse [11, 12]. There are several reasons why diversity, in all of its forms, is essential in open source communities, and the larger technology industry. Indeed, in addition to the ethical arguments favoring diversity, recent studies show that gender diversity leads to better balance between the traits of feminine and masculine roles [2], makes teams more productive [12] and benefits both industrial and academic teams [5]. Moreover, a more diverse development team comes closer to the inherently diverse population of the software users. Therefore, a more diverse development team is more likely to understand and represent the users' needs, contributing to the better alignment between software and its intended customers [9].

Not surprisingly, multiple open source projects have tried to increase the diversity of their communities, specifically gender diversity, as evident in the formation of many "Women in" groups established in Linux, GNOME, Ubuntu or Debian. Despite these efforts, the situation has not changed much in the last 15 years, resulting in what has been called a 'lost decade' in the integration of women in open source [10]. Understanding the success or failure of these and similar activities starts with measuring the participation of females - in terms of both population and level of activity – within a specific open source project. Below we report on our experience with measuring gender diversity.

While we focus on OpenStack, a well-known open source project involving thousands of contributors, many of them employed by companies, the proposed analysis and methodology have been developed with the aim of increasing the understanding of diversity and inclusion across the open source ecosystem as a whole.

Why Focus on OpenStack?

The OpenStack community has long recognized the importance of gender diversity, reflected in ongoing discussions and activities hosted during OpenStack Summits and the creation of the Women of OpenStack working group, which aims to *"... increase the diversity of the OpenStack community by overcoming OpenStack's barrier to entry through educational sessions, professional networking,*

mentorship, social inclusion, and enhanced resource access". Measuring the effectiveness of activities aimed to increase gender diversity requires repeated measurement and examination of the results of these activities before (baseline), during, and after their implementation. Moreover, understanding how people interact and perform is essential to assure transparency and fairness [6]. However, the OpenStack Foundation did not have baseline information about the involvement of women. Indeed, obtaining such baseline information is not easy for open source foundations: unlike traditional companies, open source foundations do not typically have a human resources division with detailed records of participants, their activities and their gender.

Hence, Intel and Bitergia, a software analytics company, decided to conduct research to produce these initial datasets related to participation of different genders, and share the results with the community in an effort to create broader awareness of the current state of diversity, and in turn, narrow this gap. The results of this research have been communicated during OpenStack Summits in 2016, 2017, and 2018, and welcomed by the community. Based on the feedback received, subsequent diversity reports have been prepared, incorporating this feedback to address the needs and expectations of the community. These reports measure participation by gender across both code contributions and non-code contributions. While non-code contributions do not directly affect the OpenStack project's code base, they are considered essential to the organization and overall well-being of the project. Non-code contributions included leadership and governance, and event representation and participation, among other factors.

Challenges of Measuring Gender Diversity

Measuring gender diversity based on digital traces such as logs of developers' communication and collaboration, or records of their attendance and participation in events, involves multiple challenges. First, gender is a complex social construct: since digital trace data rarely records gender explicitly, most studies reduce gender to binary and employ heuristics to decide whether the individual is a woman or a man [7]. Risks implied by this strategy are related to the inherent imprecision of such heuristics [7] and to the marginalization of non-binary individuals. In our study, the data, and the gender of the developers, has been updated using a manual approach up to a certain threshold. The OpenStack project selection is defined by the Governance project. For these projects, the developers with the highest level of activity were manually inspected and their gender updated if needed. This was done in the same way, but working at the level of companies. However, given the size of the community and the total number of different identities (more than 7,000 across all Git repositories), the remaining number of 'unknown' profiles is still high (around 25%), although their activity is lower (around 15%).

Another challenge was the difference between academia and the industry. While academia has a long experience in measuring in software engineering, industry sees some of these as potentially innovative, and therefore, is reluctant to share the methodology or results within the community and with other potential competitors. Thus, lack of common methodology and process may lead to discrepancies between internal company studies and the current work. This risk has been minimized stating the definition of each metric, using open source software for reproducibility, and choosing the official projects and data sources (defined by the OpenStack Governance project).

Measuring diversity in OpenStack is challenged by inherent imprecision of gender identification

approaches (whether manual or automatic), lack of data, and lack of methodology and process common to academia and industry.




Figure 1: Percentage of women (blue) in governance and leadership positions. The numbers in brackets are the total members of each group.

Figure 1 shows percentage of women in the governance structures of the OpenStack Foundation (October/November 2017). One of the lowest percentages is at the technical position of the Project Team Leaders. The Technical Committee is another governance technical position, displaying 15%, while in previous analysis this was 0%. These roles are filled either by community elections or by direct representation of the companies.

When considering percentage of women among attendees and speakers of the OpenStack Summits and Project Team Gatherings in 2016-2017, we observe a consistent pattern of participation around 10-12%. We do not see differences in percentages of speakers and attendees.

Percentages of women among code and non-code contributors remain in the same range of 10-12%. Percentages of their *contributions* (Figure 2) are usually lower. In particular, the lowest share of contributions is in the developer mailing list *openstack-dev*, while “other mailing lists” have the highest percentage of contributions by women. This observation concurs with the study of Drupal mailing lists [11]. Adding the time aspect to the picture reveals that participation of women increased in all of the code and non-code related contributions over time: in particular, women move from developers to maintainers and act as code reviewers.




Figure 2: Percentages of code- and non-code-related artifacts contributed by women (blue), men (red) and individuals whose gender could not be identified (green).

Percentage of women is 10-12% depending on the data source and the analysis; participation at the governance and leadership level increased remarkably; participation in all of the code- and non-code-related contributions also increased.

Discussion and Lessons Learned

The authors have presented the results of the OpenStack Gender Diversity reports at several OpenStack and Linux community events—specifically, OpenStack Summits, Open Source Summits, CHAOSSCon and the Open Source Leadership Summit—in 2017 and 2018. The feedback received from the open source communities is overwhelmingly positive. The ability to measure and report on the level of female and minority representation in open source communities, as was done in the case of the OpenStack community, is key to having well-grounded discussions about this issue. Data aids in increasing awareness and understanding. During this time, discussions focused on this representation gap—and approaches to resolving it—have gathered momentum across the open source community as a whole.

Thanks to the dissemination of the results, the authors were able to collect feedback with respect to the initial results that enabled them to evolve the parameters of the research over time to address the needs and expectations of the community. This feedback can be categorized into two main areas: feedback that can be used to improve the current set of metrics and datasets, and feedback with respect to how those metrics and datasets can be used to increase gender diversity. In regards to the second area, the focus was on providing the community with actionable recommendations, and extending the benefit of insights gained through this data, experience and metrics beyond OpenStack.

Recommendations identified in the reports aimed to increase diversity and inclusion included: increased understanding around the importance of diverse, inclusive teams; continued tracking of female participation; enhanced collaboration with Project Team Leads and other leaders across the community to identify, document and communicate best practices; continued support of working groups; and improved onboarding guidelines and mentoring programs.

The research community has recently worked on identification of barriers for participation of women in software development [3, 4, 8], and we plan to explore whether the recommendations derived from these studies, such as the use of GenderMag [1] to reveal gender-bias issues in software, can benefit open source communities.

The authors recognized the importance of being inclusive in these analyses. After all, diversity is not a mere measurement of gender, but of other underrepresented minorities as well, with multiple facets and dimensions. In addition, participation in open source communities is more than writing code; non-code contributions must also be considered and accounted for within the overall health of a community. The authors view this analysis as a starting point that they have always intended to evolve beyond. Aspects such as, but not limited to, project leadership, events inclusivity, and code of conduct creation and reinforcement have helped to shed light on the sensitivities within a community about this topic; this is part of further work on improving inclusion of the analysis undertaken.

The privacy of the data was another critical consideration. The datasets are provided under requirement for this reason, and presentations use aggregated information where there is not a direct match between developer and gender. However, it is hard to curate, improve and in general manage a set of thousands of developers. Other ways to participate and help community members should be explored to keep improving the community dataset as more and more people enter the community each month.

This research has helped pinpoint areas where further research is needed. Further research, using qualitative methods, offers deeper insight, interpretation and understanding of these numbers. For example, the quantitative data can be used to identify projects that attract a greater level of diverse contributors, and equally important, those that do not. Qualitative research is then needed to understand why this is the case.

There are still open questions in the results that are focused on providing more advanced datasets that may help provide greater insight into the idiosyncrasies of open source projects and communities. For instance, is there any difference between the development activities by gender? Is there unconscious bias towards contributions depending on the gender of the developers? At the communication level, is the collaboration taking place primarily between developers of the same gender? Are code reviews fair? Does the presence of a greater number of female participants within a project help attract more diverse participants to that specific project, when compared to other projects?

The expectation of the authors is that the OpenStack community enters the next phase now, and uses the numbers to drive deeper qualitative research, learn from the most—and least—diverse projects and organizations, and deliver data-driven strategies to guide actions and activities to resolve this issue. The

goal is to create a welcoming, inclusive environment that encourages diverse contributions across all constituents of a community.

With respect to extending the benefit of insights gained through this data, experience and metrics beyond the OpenStack project, the authors have fostered collaboration with other projects, such as the CHAOSS (Community Health Analytics for Open Source Software) project. This project, announced by the Linux Foundation in September 2017, was created to establish standard, implementation-agnostic metrics to gauge the overall health of open source communities. In December 2017, the CHAOSS project's Diversity and Inclusion Working Group was formed with the goal of integrating the methods and processes of the OpenStack Gender Diversity Reports, among other initiatives, to create a proper body of knowledge about the topic of diversity and inclusion.

Open Source communities should continue monitoring participation of women, discover and learn best practices for attracting and retaining women, and broaden the focus to include other underrepresented minorities.

References

[1] Margaret M. Burnet, Simone Stumpf, Jamie Macbeth, Stephann Makri, Laura Beckwith, Irwin Kwan, Anicia Peters, and Will Jernigan. GenderMag: A method for evaluating software's gender inclusiveness. *Interacting with Computers*, 28(6):760-787, 2016.

[2] Luiz Fernando Capretz and Faheem Ahmed. Why do we need personality diversity in software engineering? *SIGSOFT Sofw. Eng. Notes*, 35(2):1-11, March 2010.

[3] Dena Ford, Alisse Harkins, and Chris Parnin. Someone like me: How does peer parity influence participation of women on Stack Overflow? *VL/HCC*, pages 239-243, 2017.

[4] Dena Ford, Justin Smith, Philip J. Guo, and Chris Parnin. Paradise unplugged: identifying barriers for female participation on Stack Overflow. *FSE*, pages 846-857, 2016.

[5] Abdul Rehman Gila, Jafreezal Jaafa, Mazni Omar, and Muhammad Zahid Tunio. Impact of personality and gender diversity on software development teams' performance. In *International Conference on Computer, Communications, and Control Technology*, pages 261-265, Sept 2014.

[6] Jesus Gonzalez-Barahona, Daniel Izquierdo, Stefano Maffulli, and Gregorio Robles. Understanding how companies interact with free software communities. *IEEE Sofw.*, 30(5):38-45, September 2013.

[7] Bin Lin and Alexander Serebrenik. Recognizing gender of Stack Overflow users. *MSR* pages 425-429, 2016.

[8] Christopher J. Mendez, Hema Susmita Padala, Zoe Steine-Hanson, Claudia Hilderbrand, Amber Horvath, Charles Hill, Logan Simpson, Nupoor Patl, Anita Sarma, and Margaret M. Burnet. Open source barriers to entry, revisited: a sociotechnical perspective. *ICSE*, pages 1004-1015. 2018.

[9] Michael J Muller and Sarah Kuhn. Participatory design. *Communications of the ACM*, 36(6):24-28, 1993.

[10] Gregorio Robles, Laura Arjona Reina, Jesús M González-Barahona, and Santago Dueñas Domínguez. Women in free/libre/open source software: The situation in the 2010s. In *IFIP International Conference on Open Source Systems*, pages 163-173. Springer, 2016.

[11] Bogdan Vasilescu, Andrea Capiluppi, and Alexander Serebrenik. Gender, representation and online participation: A quantitative study. *Interacting with Computers*, 26(5):488-511, 2013.

[12] Bogdan Vasilescu, Daryl Posnet, Baishakhi Ray, Mark GJ van den Brand, Alexander Serebrenik, Premkumar Devanbu, and Vladimir Filkov. Gender and tenure diversity in GitHub teams. *CHI*, pages 3789-3798. ACM, 2015.