Making Open Source Project Health Transparent

Sean P. Goggins, University of Missouri

Matt Germonprez and Kevin Lumbard, University of Nebraska Omaha

We explore the Community Health Analytics for Open Source Software project and how it plays an integral role in the automation of key measures to make the state of open source readily observable.

We are long past discussions about the benefits of proprietary software versus open source software (OSS). The world we live in is built on OSS. The software exists in complex infrastructures and supply chains, often alongside proprietary programs. It is more ubiquitous and complex than ever, and it is continuing to grow. Corporations have embraced OSS in a way few could have imagined 40 years ago. For example, automotive-grade Linux is deployed in dozens of vehicle models, and Kubernetes enables massive on-demand scaling for a wide range of online firms. All this begins with a project or Git repository, where bugs are fixed, features are added, and discussions about the inclusion of contributions take place. When OSS projects are used by others, they become part of a larger supply chain composed of many released versions that delivers value to software consumers. The projects and products depend on yet another layer of OSS: the dozens to thousands of libraries imported by each software component. OSS's ubiquity is making it more visible, and its complexity is making it more difficult to manage. In this article, we illustrate how a five-year project, Community Health Analytics for Open Source Software (CHAOSS), helps improve the effectiveness of the people who build, maintain, contribute to, and consume OSS in an interconnected world.

As Gonzalez-Barahona[1] describes in an earlier column, the history of free and open source software (FOSS) is largely the story of computer scientists laboring to

Digital Object Identifier 10.1109/MC.2021.3084015
Date of current version: 30 July 2021

EDITOR **DIRK RIEHLE**
Friedrich Alexander-University of Erlangen Nürnberg;
dirk.riehle@fau.de

make core computing functions available across an expanding number of architectures. The growth of contemporary software infrastructure is the cousin of early FOSS work, and it persists with increasing corporate sponsorship through paid contributors and organizations, such as the Linux Foundation. The motivation for OSS infrastructure work is moving beyond the core value of open production, increasingly centering on the open sourcing of corporate intellectual property that a firm, or a collection of firms, determines is essential and nonmarket differentiating. If we all need it, why not share the cost of maintenance and evolution?

Individuals and organizations from a significantly more diverse set of domains than the early days are also working on OSS projects and delivering OSS products. The United Nations Children's Emergency Fund maintains a portfolio of more than 800 OSS projects aimed at attaining policy goals, social good, and greater diversity in the OSS contributor pool. Scientists researching diseases, biological and plant genomes, and pharmacological treatments depend, to a growing extent, on a substantial OSS ecosystem. Journalists are using OSS to close information and skill gaps within individual, increasingly resource-starved newsrooms to fulfill their roles, and even video game designers are starting to build their work using OSS engines, such as Godot (https://godotengine.org/). If you use a computer, drive a car, or purchase groceries, it is a virtual certainty that a number of OSS products interact to ensure your success. The realization that our world is constructed on OSS illuminates our need to understand how our built environment shapes our lives[4] and how we can continue to structure and maintain a world we want to inhabit.

FROM THE EDITOR

Hello everyone, and welcome back to the "Open Source Expanded" column! We are well into the open source community theme now. In this instance, Sean Goggins and colleagues from the Linux Foundation's Community Health Analytics Open Source Software project discuss how to identify healthy open source project communities and determine if something is going wrong. Practical metrics are important for any open source project leader, and our experts have a story to tell. Enjoy, and as always, stay safe and healthy, and keep on hacking. — *Dirk Riehle*

Diversity, ubiquity, and complexity within each OSS application adds a responsibility for computer scientists and others to be aware of the health and sustainability of their projects and of those they depend on. One result of our collective recognition of OSS complexity awareness is the formation and development of the Linux Foundation's CHAOSS project (https://chaoss.community), which is the product of active engagement from hundreds of OSS maintainers, contributors, corporations, and domains of construction and use. In the remainder of this article, we describe the CHAOSS approach for addressing OS health and sustainability, the project's core focus areas, and how the project's software (Augur) plays an integral, ethically grounded role in the automation of key measures that make project growth, risk, value, and potential transparent at today's OSS scale.

OSS HEALTH

CHAOSS develops tools to support consistency for OSS maintainers and other stakeholders in their individual assessments of projects and ecosystem health and sustainability. While mundane in appearance, critical historical gaps have been closed through the project. For example, OSS health and sustainability metrics originate from earlier tools focused on measuring commit activity, which the authors outline in a recent review of the literature.[2] Activity metrics are helpful but incomplete for understanding contemporary health and sustainability questions that organizations, individuals, and foundations ask about OSS project portfolios that often number in the thousands, usually include dependencies on many other OS projects, and influence corporate valuations.[3] As OSS began experiencing something of a Cambrian explosion, the need for CHAOSS reached critical mass in 2017.

CHAOSS recognizes that the growing complexity associated with OS work, when mitigated through research that boosts visibility beyond activity metrics, is likely to accelerate innovation by increasing the adoption of shared, essential resources in a larger number of cases. For example, within organizational boundaries, key contributors are known. However, in OS projects, the dynamics of contributor turnover can create uncertainty about incorporating work into commercial products. The effects of that turnover and the limits to its visibility heighten perceptions of risk. The lack of visibility, like other issues, can have direct and lasting impacts on the work found in OS projects.

Most organizations (for example, corporations, nonprofits, collaboratives, and universities) that engage in OS projects rely on a small number of experts who use heuristics to assess

opportunities and approximate the value and risk of participation. Prior to the CHAOSS project, metric definitions were idiosyncratic; the tools were a bricolage of homegrown and small-scale OSS projects. Organizations had difficulty consistently understanding the return on their OS investments, especially more strategic ones that crossed ecosystems and included competitors. Metrics and tooling that reach beyond activity measures make OS projects work and the subsequent evolution of OSS more visible. The growing uptake of standardized metrics and software from CHAOSS is helping organizations assess risk and value in ways that overcome the useful but haphazard assessments that are commonplace.

In a drive to make OS projects more sustainable, the CHAOSS project has published more than 55 metrics and tools aimed at lifting the veil of complex interdependencies that encumber sustainable growth. In the sections that follow, we define the scope of five core working groups and their focus areas that increase the visibility of multiple dimensions related to project health. While metrics are organized within the working groups that develop them, you may recognize alternate, potentially more useful structural presentations, and we welcome those suggestions on our mailing list (https://lists.linuxfoundation.org/mailman/listinfo/chaoss). Our working group structure and the developed metrics are illustrated in Figure 1.

Five core working groups

CHAOSS metrics for code development activity and quality as well as issue resolution, efficiency, and community growth are devised by the Evolution Working Group (https://github.com/chaoss/wg-evolution). Long-standing metrics focused on commit activity are developed and maintained primarily within this group. Most of the activity metrics that are not produced there are part of the Common Working Group (https://github.com/chaoss/wg-common), which also creates measurements that are of interest to multiple other groups. The Risk Working Group (https://github.com/chaoss/wg-risk) maintains metrics focused on license coverage, Construction Industry Institute best practices, and hazards pertaining to maintainer diversity. More recently, it has defined metrics, measures, and resource lists for understanding the increasing complexity of dependencies between projects. Dependency concerns are especially prominent in the work of OS program offices, community managers, and project maintainers.

The Value and the Diversity, Equity, and Inclusion (DEI) Working Groups develop and maintain metrics that are more difficult to derive solely from Git platforms, issue trackers, and electronic project communication. For example, the Value Working Group (https://github.com/chaoss/wg-value) develops metrics for assessing project popularity, labor investment, and OSS as a social good. The DEI Working Group (https://github.com/chaoss/wg-diversity-inclusion/) crafts metrics focused on event inclusivity and that aim to raise awareness of project practices, such as mentorship and managing burnout, and other factors that have been shown to foster or erode inclusiveness. These two groups help CHAOSS move beyond measuring health and to advance the sustainability of OSS as a whole.

GROWING OSS CONTRIBUTION

CHAOSS has generated concrete metrics for more than 10,000 OS projects and implemented reporting systems tailored to the needs of several dozen corporate organizations. The most prominent types of analysis desired by project maintainers relate to the retention of contributors and the responsiveness of maintainers to contributions. In some cases, projects have focused on competitive analysis along the lines of maintainer responsiveness, recognizing that faster replies are more likely to keep contributors engaged. Figure 2 illustrates a competitive analysis of maintainer responsiveness, and Figure 3 describes the integration of several CHAOSS metrics focused on contributor retention.

To make OS projects more sustainable, CHAOSS focuses on making work visible. Our members apply theories of organizational development and analytical tools that advance visibility to support the exponential growth and increasing interdependency that OSS is undergoing. As CHAOSS cofounders, maintainers, and board members, we apply our deep, embedded fieldwork in concert with machine learning and network science to facilitate an energetic response to these changes. CHAOSS recognizes that software engineering is part of adapting to this phase shift. In that context, one risk to future software engineering practices is that social, organizational, and technical responses to change will unintentionally replicate existing approaches that are incommensurate to a problem. For example, more of the work is performed by a less diverse collection of people, who build OSS for pay, than we find in other professions. Advancing inclusivity in OS work and the now ubiquitous impact OSS has on society are growing focuses of the CHAOSS project.

DEI event badging

The reshaped world of OSS demands that CHAOSS not only standardize metric definitions and tool kits but innovate with programs that strive to increase the workforce through greater diversity and inclusion. Our DEI event badging program (https://chaoss.community/diversity-and-inclusion-badging/) has already recognized 21 major initiatives in the Kubernetes, cloud-native, and other OSS communities (https://github.com/badging/event-diversity-and-inclusion), following an open peer-review process modeled after *Journal of Open*

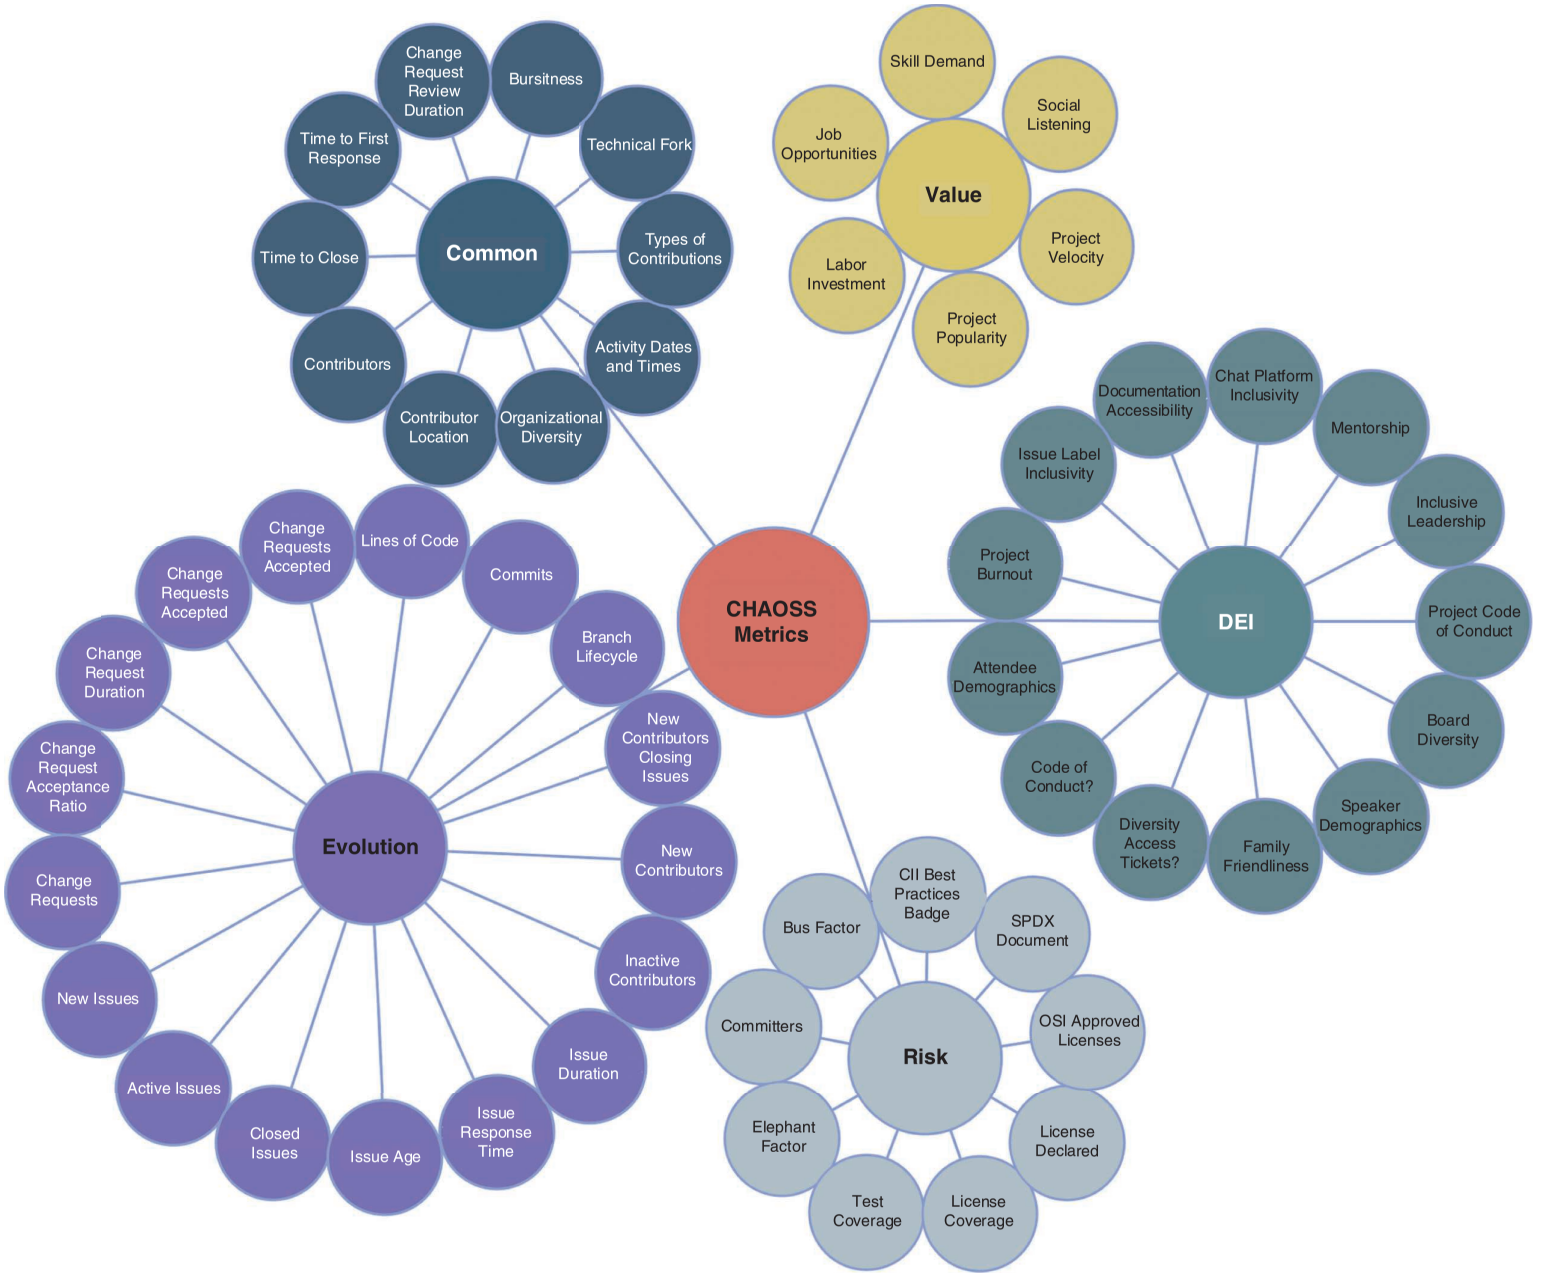

FIGURE 1. The 57 CHAOSS health and sustainability metrics organized by five working groups: Risk; Evolution; Value; Common; and Diversity, Equity, and Inclusion (DEI). CII: Construction Industry Institute; SPDX: Software Package Data Exchange; OSI: Open Source Initiative.

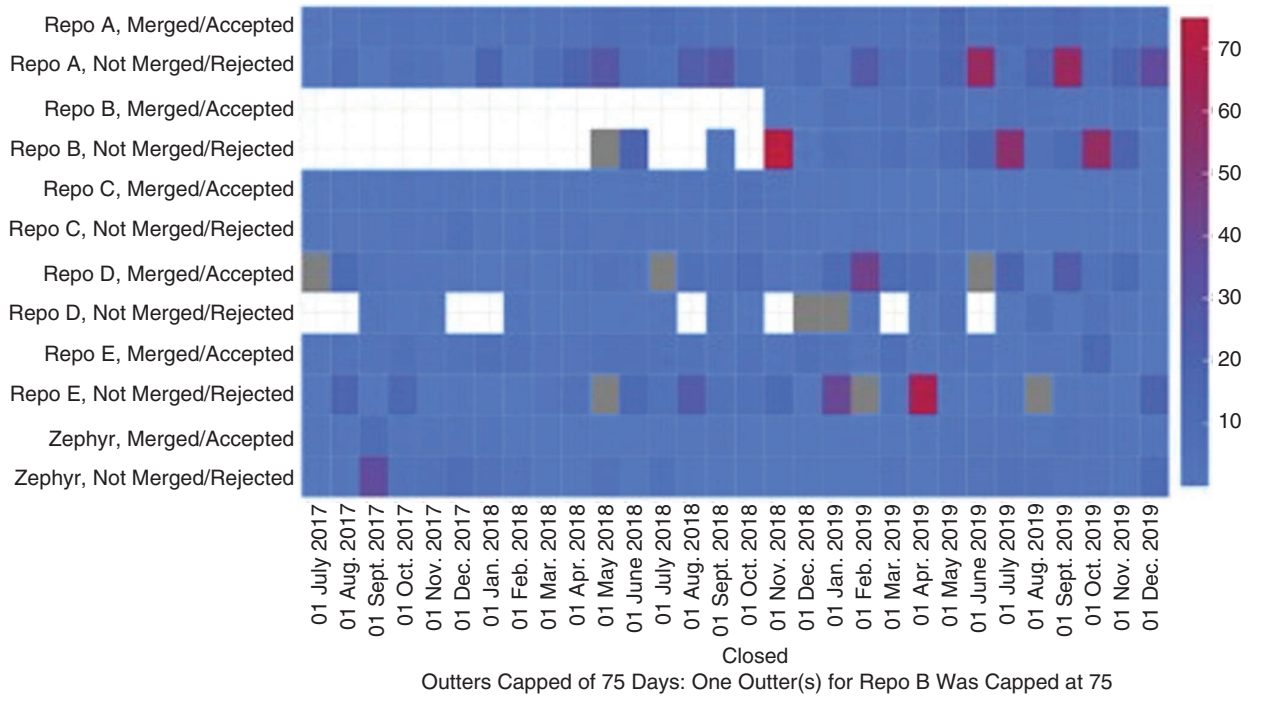

FIGURE 2. Zephyr is a real-time operating system for critical infrastructure. This analysis compare's Zephyr's pull request responsiveness to that of competing products through time, using a heat map–style visualization. The competitors have been anonymized to maintain a focus on Zephyr.

Source Software (https://joss.theoj.org/). Such rapid success illustrates a wider recognition of DEI as central to managing the changed nature of OSS. The badging program has several tiers based on the number of CHAOSS DEI metrics attained, in much the same way the Core Infrastructure Initiative (https://bestpractices.coreinfrastructure.org/en) issues badges, as shown in Figure 4.

TOOLS FOR SCALING HEALTH AND SUSTAINABILITY AWARENESS

Developing and distributing common definitions for OSS health metrics includes collaboration among hundreds of OSS organizations that recognize the urgent need to move beyond activity metrics, yet a dictionary of sorts is of limited use without implementation in software. CHAOSS tools, such as GrimoireLab (https://github.com/chaoss/grimoirelab) and Augur (https://github.com/chaoss/augur; https://github.com/chaoss/augur-community-reports), implement CHAOSS metrics and present them in ways that enable maintainers, contributors, and other stakeholders to draw inferences about the relative health and sustainability of their projects by using indicators whose consistency, if not perfection, can be trusted.

Analyzing information within ecosystems

Prior to the CHAOSS project's introduction of Augur, OSS metrics collection and persistence focused on the analysis of individual efforts and predefined collections of initiatives, using definitions that were specific to each tool. Augur's design supports the phase shift in the number of projects and the dependencies between them by collecting lists of repositories where dependencies are managed and where contributors and maintainers work, beyond the predefined scope of analysis. This type of "snowball collection" of basic information makes it possible for OS program offices, community managers, scientists, and other stakeholders to take a peek at parts of their infrastructure that were not visible before CHAOSS consistently defined metrics and developed tools to address contemporary challenges.

CHAOSS recognizes that ecosystems are defined by the goals and questions of each OSS stakeholder. A science OS ecosystem is typically bounded by a particular field that develops and uses OSS. Corporations often participate in and manage numerous interconnected ecosystems. With a consistent taxonomy of metrics, each stakeholder is enabled, through the flexible and well-defined data structures Augur implements, to combine and present the most critical information for decision making at any point in time. CHAOSS and Augur can answer questions such as, "Where are the most vulnerable dependencies across 11,000 OSS projects in one program office?" and "What are the most vulnerable dependencies in each project?"

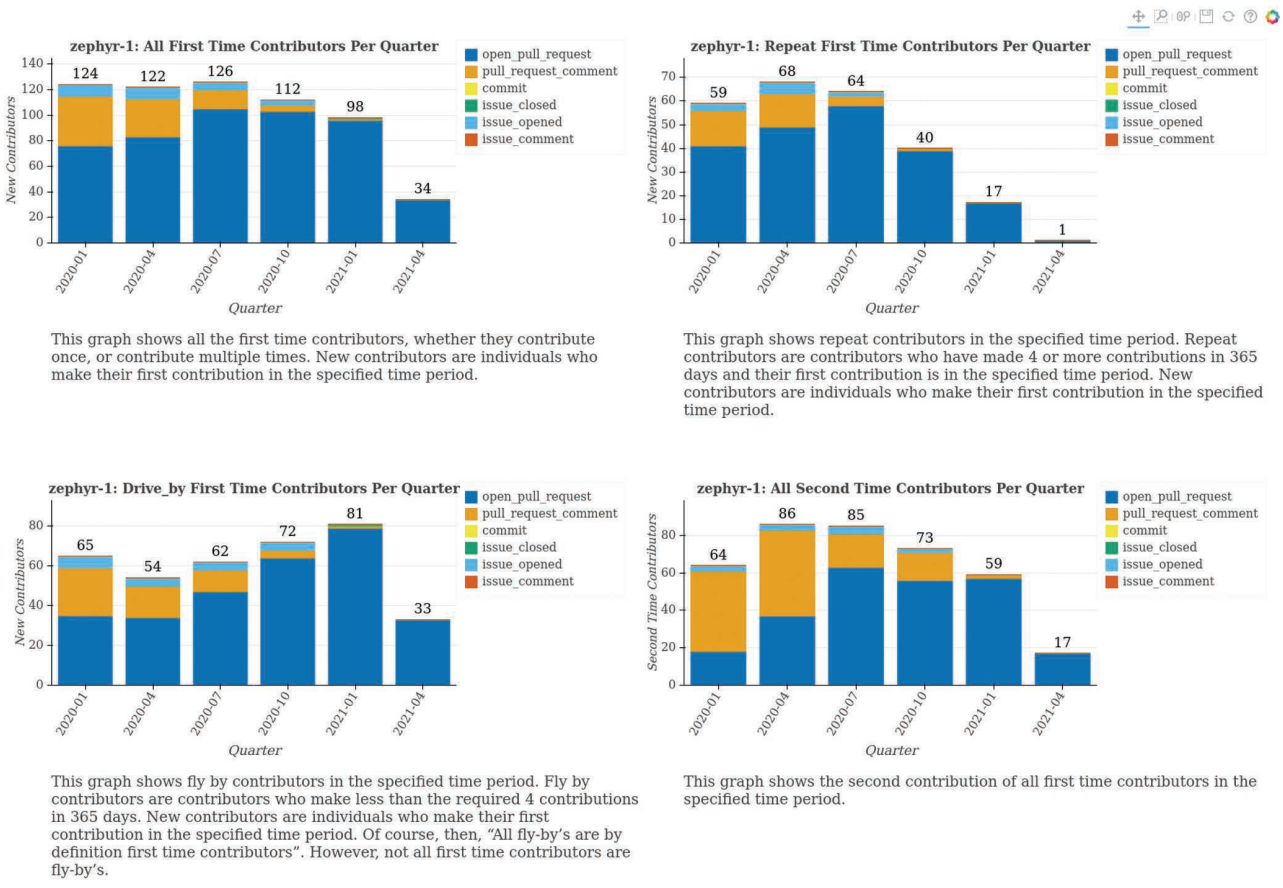

FIGURE 3. New contributors can be understood from many perspectives. Zephyr maintains a constantly updating instance of Augur that can generate a report (http://zephyr.osshealth.io:5222/api/unstable/contributor_reports/new_contributors_stacked_bar/?repo_id=26222) at any time.

Linking ecosystems for corporatized open

Source. Rapid OSS growth also means fast changes to the ways corporations bound the different ecosystems in their spheres. Applying CHAOSS metrics enables OSS leaders' real-time awareness of software that may become part of an ecosystem they are engaged in or soon will be. Risk awareness, in contrast, focuses on the OSS ecosystem as it is now, with the software bill of materials Augur provides, using the Software Package Data Exchange standard and FOSSology scanners; details about file-level license declaration completeness and diversity; and reasonably clear information about the organizations and contributors that are most essential for a project's health and sustainability. Through support from programs, such as the Google Summer of Code, and growing partnerships with the Open Source Security Foundation, CHAOSS and Augur are on the leading edge of defining essential indicators of OSS health: software security, software bills of material, development time, and runtime dependencies.

Augur: Making subtle project changes Transparent, with ethical artificial intelligence.

Dependencies, licensing, emerging ecosystems, and software bills of material, while complex, are being incorporated into CHAOSS metrics and Augur by using discrete, discoverable data. There are, of course, a number of OSS projects that have fractured or declined despite their critical importance. Often in these cases, the reasons are subtle and, with the growth of OSS, more challenging to identify early enough for intervention. Typically, early warning signs can be found in project communications on mailing lists, issue trackers, and merge request discussions. The volume of those communications enables "trouble in paradise" issues that go unnoticed for too long.

Machine learning, artificial intelligence, computational linguistics, topic modeling, clustering, and statistical anomaly detection are useful for identifying project health concerns earlier. In our experience working with OSS teams using Augur's seven computational models, applying these technologies is most effective if two conditions are mutually understood. First, no one model is sufficient for identifying

projects experiencing challenges or in the early signs of remarkable success. An integrated analysis using all seven models provides more useful information. Second, human judgment, clear communication, and the protection of individual privacy are paramount if CHAOSS metrics emerging from these tools are to be useful and accepted by the maintainer and contributor communities. In our view, metrics and tools leveraging these technologies are unethical if humans are removed from the interpretation of results.

USEFUL SIGNALS FOR A HIGH-VELOCITY OSS WORLD

The issues of OSS health and sustainability are nearly as old as OSS itself. In our four years focused on CHAOSS, we have observed nearly daily a range of customs, communication patterns, and contribution cycles, and we have seen the effects that changing alliances have on goals and the questions OSS leaders need answered. In some cases of applying CHAOSS metrics, within weeks, questions that were never asked before emerge and require answers. At times, these sudden shifts have financial motivations; in other cases, legal rulings and the unexpected growth of new technologies drive them. Candidly, many discussions about the shortage of OSS engineers reference Ostrom's "overgrazed commons" metaphor to illustrate the constraint. Yet, activity metrics alone show how the OS contributor community remains insufficient for the work at hand, without suggesting remedies.

Based on emerging needs, the CHAOSS project will monthly identify and develop new signals. Often, Augur and other CHAOSS tools respond rapidly to implement those metrics to support decision making. To sustain our efforts, the partnerships we have formed help reseed our commons, but the volume of work continues to outpace the availability of OSS engineers. To continue developing useful metrics and tools, and for OS in general, it is becoming increasingly clear that one of our necessities is to identify and remove shared obstacles, which take many forms, such as barriers to careers in OSS. There are known impediments of context and distance that are basic features of OSS and that will likely remain in the future. However, others may be more manageable.

We think part of the upcoming contribution of CHAOSS can involve working with current and future projects to classify goals and related activities into categories of obstacles that provide shared utility. Not every project is at the same point in its evolution or focuses its questions on the same CHAOSS metrics. In fact, there is a good deal of organization-, ecosystem-, and maintainer-focused curiosity that drives the application of CHAOSS metrics through a "context" filter. In each case, prior work and lessons inform how we identify and classify obstacles. Briefly, during activities when CHAOSS contributors seek and explore obstacles in a general sense, they attempt to map risks to project health. Some activities in this category include health and sustainability metrics, including software licensing and regulation, and the increasing number of dependencies, such as when an OSS project imports a library from another.

Reflecting on our work with partners on DEI in OSS, we have observed and been told of project communication patterns that welcome newcomers and others that do not. We have gathered a number of oral narratives on the CHAOSS podcast that illustrate how efforts to help people recognize relationships between their motivations is an effective tool to make newcomers recognize, frankly, that OSS exists and that they can be paid for it. The increasing interest in understanding obstacles that impact DEI health is a reflection of what we should all know: contributors are the lifeblood of every OSS project, and OS health and sustainability require creating a diverse, equitable, and inclusive environment.

Level	Badge	Percentage of Requirements Met
Pending	D&I Pending	Less Than 40%
Passing	D&I Passing	Greater Than or Equal to 40% and Less Than 60%
Silver	D&I Silver	Greater Than or Equal to 60% and Less Than 80%
Gold	D&I GOLD	Greater Than or Equal to 80%

D&I: Diversity and Inclusion

FIGURE 4. There are four levels in the CHAOSS DEI event program. Each represents a progressively higher percentage of the DEI metrics that have been met. To date, 21 major OS conferences have achieved one of the badge levels.

The future of OS metrics, as found in CHAOSS and through advanced tooling, such as Augur, seems virtually assured by the incredible growth of OSS. Through our work, we have assisted a number of OS projects, talked with hundreds of contributors and maintainers, and occasionally ventured out of that bubble. Health and sustainability for OSS is tied very clearly to many aspects of our lives as people: fishing, home improvement, travel, close relationships, and all the things that bring us joy. One CHAOSS member summed up this occasionally overlooked interconnection well, and we close with that thought: "Open source helps power virtually every piece of technology in our lives. The only way open source technology will equitably serve all of us is if we center DEI in the design and development of that technology."

ACKNOWLEDGMENTS

This work was supported through funding from the Alfred P. Sloan Foundation as well as Mozilla. Jesus Gonzalez-Barahona and the GrimoireLab initiatives are founders of the CHAOSS project, and their software remains the most widely used for OS metrics.

REFERENCES

1. J. M. Gonzalez-Barahona, "A brief history of free, open source software and its communities," *Computer*, vol. 54, no. 2, pp. 75–79, 2021. doi: 10.1109/MC.2020.3041887.
2. S. Goggins, K. Lumbard, and M. Germonprez, "Augur: Architecture for capturing, reshaping, and socially contextualizing open source software communities," in *Proc. ACM SoHeal Conf. Int. Conf. Softw. Eng.*, 2021. doi: 10.5281/zenodo.4627236.
3. M. Germonprez, G. J. P. Link, K. Lumbard, and S. Goggins, "Eight observations and 24 research questions about open source projects: Illuminating new realities," *Proc. ACM Human-Comput. Interact.*, vol. 2, pp. 1–22, Nov. 2018. doi: 10.1145/3274326.
4. S. W. Goldhagen and A. Gallo, *Welcome to Your World: How the Built Environment Shapes our Lives*. New York: Harper, 2017.

SEAN P. GOGGINS is a professor of computer science at the University of Missouri, Columbia, 65211, Missouri, USA; cofounder of the Linux Foundation's Community Health Analytics Open Source Software project; and creator of the University of Missouri's master's program in data science. Contact him at outdoors@acm.org.

MATT GERMONPREZ is the Mutual of Omaha Distinguished Chair of Information Science and Technology and a professor of information systems and quantitative analysis in the College of Information Science and Technology, University of Nebraska Omaha, Omaha, 68182, Nebraska, USA. Contact him at mgermonprez@unomaha.edu.

KEVIN LUMBARD is a Ph.D. candidate at the University of Nebraska Omaha, Omaha, 68182, Nebraska, USA. Contact him at klumbard@unomaha.edu.

Digital Object Identifier 10.1109/MC.2021.3094709