moz://a

MOSS

2016 Review and 2017 Strategic Plan

April 25, 2017

Table of Contents

Introduction

Mozilla Open Source Support (MOSS) is Mozilla's program for supporting the free and open source community of which we are a part, and from which we get much of our strength and our asymmetric advantage. The budget for 2016 (not including those awards made in 2015 but paid in 2016) was US$3M. This document reviews what was achieved or not achieved in 2016, and proposes ways to take the program forward in 2017.

2016 Review

Foundational Technology

The Foundational Technology track is open to free/open source projects whose code is incorporated into Mozilla projects, runs in our infrastructure, or which Mozillians use to get work done. It is our way of thanking those projects for the value we have already received from their hard work - there is no requirement that the work funded be of direct benefit to Mozilla or its mission. It launched in late 2015, with the first 7 awardees announced at the Mozilla All Hands in Orlando in December of that year. Those awards totaled $533,000.

In 2016, we have made another 6 awards, totaling almost exactly the same amount of money - $535,000. Here are brief status reports from some of these 13 projects, several of which are now completed:

Bro network traffic analysis system ($200k). Goal: create bro-pkg system, allowing simpler community sharing of scripts and plugins. Done.

CodeMirror in-browser code editor ($20k). Goal: modularizing the core library (done) and improving bidirectional text support (WIP).

Discourse web-based discussion platform ($25k). Goal: achieving feature parity with the most popular mailing list systems. Nearly done.

Django web development framework ($150k). Goal: rewrite internals to be event-driven and more suitable for implementing WebSocket-based servers. Nearly done.

Django REST Framework web API framework for Django ($50k). Goals: Python client, command line client, schema support (done); JavaScript client, docs tooling, WebSockets support (WIP).

Kea DHCP server ($100k). Goal: implement a remote management API. WIP.

Mercurial distributed source code management system ($75k). Goals: improve history search in the web interface; build JSON API; more usable client default settings. WIP.

Redash data visualisation system ($100k). Goals: improve setup, maintainability, localizability and extensibility. WIP.

The Intern website and web app testing system ($35k). Goal: implement automated performance benchmarking, visual regression, and accessibility testing capabilities. Done.

This is a collection of significant improvements to important open source packages, and Mozilla should be proud of making them possible.

Mission Partners

The Mission Partners track is open to any free/open source project, but the funded activity must meaningfully further Mozilla's stated priorities or mission - to ensure the Internet is a global public resource, open and accessible to all. It launched in mid-2016, with the first 8 awardees announced around the time of the Mozilla All Hands in London in June. It was able to leverage the hard work done for Foundational Technology in terms of reusing much of the infrastructure and committee, and in working out the complexities of paying awards. Those awards totaled $385,000; we since and, since then, 1 further award has been made of $56,000. We are also aiming to make another $300,000 in awards by the end of the year, for a total of $741,000. This is against a notional budget of $1.25M for this track.

Here are brief status reports from some of the projects funded under this track:

Caddy secure-by-default web server ($50k). Goals: automated plugin deployment; API and GUI to avoid need for use of command line. WIP.

NVDA screen reader ($15k). Goal: Improve Firefox e10s Windows accessibility and make necessary changes in NVDA so screen reader users can benefit from multi-process Firefox. WIP.

PeARS distributed web search engine ($15.5k). Goal: complete and release beta version. WIP.

Tor anonymity network ($152.5k). Goal: Improve metrics for Tor network - usage, tool use, browser downloads etc. WIP.

Each of these projects, in the opinion of the MOSS committee, is doing work which significantly furthers Mozilla's mission or other stated priorities.

Secure Open Source

Secure Open Source aims to make the Internet more secure by providing manual source code audits by professionals for key pieces of free/open source software. Pilots began in 2015, but it became an official MOSS track in mid-2016 and over the course of the year (as of mid-November 2016) has completed 7 audits, with another 4 in progress. Unlike the other two tracks, while there is a nomination process whereby anyone can make suggestions, projects are approached by Mozilla with an offer of audit rather than applying to the program themselves.

Highlights of the year include finding a couple of denial of service flaws in the JPEG standard itself during an audit of the libjpeg-turbo JPEG library, and finding critical vulnerabilities both in the PCRE regular expression library and the curl HTTP client. A highlight of a different sort was auditing the dovecot IMAP server, a widely deployed piece of technology, and finding almost nothing wrong with it. Assuming the auditor is trusted to have actually done the work (which they are), this is the kind of result we want to see more of!

SOS has three inputs - money, projects to audit, and auditors. In 2016, the limiting factors were auditors and projects. So we focused on establishing a stable of auditors with whom we have working relationships. We now have 6 on the books, which should set us up for greater throughput in 2017.

Publicity and Outreach

As Foundational Technology and Mission Partners rely on quality applications for awards to be made, publicity and outreach needs to be an important part of what MOSS does.

During the year, our committee promoted MOSS to their contacts, and we also emailed specific key mailing lists with encouragements to apply. On the web, as well as blog posts about and promotion of specific track launches, we have produced quarterly updates about the program, published on the main Mozilla blog - Q1, Q2, Q3.

Towards the end of the year, we launched a proper MOSS website on the mozilla.org site. We have also built a "Friends" page listing other funding organizations, and some of those organizations have in turn promoted us to their contacts.

In the middle of the year, we moved to a quarterly batches application model to try and galvanize applications. We also hosted a Twitter chat about funding for open source on 21st November, with many of the existing awardees participating, which we hope will raise awareness in advance of the end of the month deadline.

SWOT Analysis

Strengths

28 open source projects have been significantly helped in what they do by a MOSS award. One particular strength of MOSS is our ability to give awards to individuals, rather than requiring that we fund a 501(c)3 or company.

We have had positive press mentions of the project in, for example, The Register, Silicon, Naked Security and Betanews.

The MOSS Committee has done an excellent job analyzing the proposals, asking difficult questions, giving good feedback to applicants about how to improve their applications, and choosing those which fit well with the laid-down criteria. We have chosen a group with a wide range of experience and perspectives, and appropriate skills for making the evaluations.

Many Mozillians see MOSS as a great thing, and something we should have done a long time ago. One said it made them proud to be a Mozillian.

Weaknesses

MOSS has had some difficulty generating applications outside of periods of publicity associated with track launches. Foundational Technology had 12 non-junk applications in the last quarter of 2015, and has only had another 14 since, which includes 5, which were re-applications from organizations whose first versions were declined. Mission Partners had 47 applications (of varying quality) for the initial round, but prior to the Twitter chat, had had only 4 since, of which the one that was awarded was a re-application. (Post-chat, we have had another 5, which is encouraging.)

SOS Fund takes a reasonably large amount of administration - much more than the other two tracks. In the case of all 3 tracks, assistance from a contract manager has proved invaluable in lightening the load, but SOS Fund still requires a lot of coordination to push projects through the process, and that load will scale linearly as we increase the amount of audits the program runs.

The two weaknesses above have combined to mean that MOSS did not spent its full allocated budget for 2016. At time of writing (end-November), US$1.22M has been spent; we expect it to be close to $1.6M by the end of the year.

The relatively small amount of time the MOSS administrator is able to devote to MOSS means that the oversight of projects and what they do with the money is fairly light. This is less of an issue with track 1, where we have Mozilla champions and where we are not as fussed what they do with the money as long as it’s used to develop and improve their software, and somewhat more of a weakness related to track 2.

Opportunities

Particularly with the Mission Partners track, MOSS offers the opportunity to further our mission by contributing to work in areas of software and security that, for whatever reason, our own resources are not best placed to do or replicate. Not all the smart people work for us, and providing grants to the projects they do work on allows us to direct their efforts in support of our goals.

A successful and well-known MOSS program is a great opportunity for Mozilla both to rebuild and strengthen its reputation in the open source community (damaged by things like DRM, UI churn and product churn) and to help retain both volunteer and paid talent by demonstrating that we care about the communities they come from.

Threats

It is important that MOSS demonstrate value to Mozilla as well as to the projects we support. If we fail to maximize the opportunities for pushing forward the mission and for positive publicity, there is a risk that the program will be seen as not having sufficient impact.

If we are unable to generate quality applications on an ongoing basis, the program will be unable to make many awards for Foundational Technology or Mission Partners.

There is a tension between being both a fundraiser and a grant giver under the same Mozilla brand. Some may find it strange that they are giving money to Mozilla while Mozilla gives money to someone else. This tension is broader than just MOSS, of course.

2017 Strategic Plan

Publicity and Outreach

We need to address the lack of publicity surrounding MOSS in 2017, which will involve working with the PR, press and social media teams to draft a publicity plan. Placement in the standard tech press is not expected to be particularly effective; we need to reach out through specific developer-focused channels. It would be good to talk to the DevRel team about what options they have. We had 15 applications subsequent to the November 2016 Twitter chat, 4 of which specifically mentioned social media as the source, which shows that we can move the needle if we publicize this in the right places.

We are planning to produce a set of interview videos with previous recipients early next year. These will be done by Gerv, using recipients in and around the Bay Area, and can be used to promote MOSS in the run-up to future deadlines. They will also be usable more permanently on the MOSS website.

Now that more awardees are finishing up the work they started, we need to encourage them to do a blog post or other report on how it worked, speaking of MOSS in a positive light and encouraging others to apply. This will become more of a publicity stream as MOSS becomes established.

In the end, though, it may be that cultivating leads through the personal contacts of people involved in MOSS and other Mozillians is the slow-but-sure way to generate interest.

Foundational Technology

The criteria for Foundational Technology means that the pool of applicants is limited, and the growth potential is also limited. This track only makes awards to projects that Mozilla or Mozillians use or deploy, which is not a short list in absolute terms, but is a small group compared to the universe of open source software.

We have an incomplete list of projects in use by Mozilla. We have attempted to reach out to those projects via their Mozilla contacts, which had led to a small number of applications. One option would be to reach out to some or all of these projects directly. This merely needs someone to research the points of contact, and could be done in an afternoon. It might also be worth doing a mail to all Mozillians, both in an effort to expand the list, and also to get them to consider whether other open source communities of which they are members could apply. We should point out that if a significant number of Mozillians use some particular piece of desktop software for their work, which also falls under the Foundational Technology umbrella.

One outcome of the Twitter chat in November was that participants listed aspects of open source, which were hard to get funding for. These tended to be those which are hard to measure crisp outcomes for, including maintenance, standardization, metrics, usability, user support, community building, tooling improvements, performance and documentation. Perhaps, because Foundational Technology is a "thank you" program and so we are not so concerned about the exact work funded, there is an opportunity here? Perhaps we should specifically flag the validity of, and make it easier for people to apply for, less tangible goals such as the ones listed above. At the moment, we still require a fairly clear and crisp project plan and for the applicants to work on something tangible and measurable. We should discuss whether that's as necessary for this track as we thought it was. A greater degree of trust in the maintainers is required for this sort of award, but perhaps we can use our deep open source community links to do a reputational evaluation, which could establish such trust.

Mission Partners

The potential applicant pool for Mission Partners is the whole of the world of open source. Therefore, the challenge is making sure that this group hears about the opportunity - see Publicity and Outreach, above.

We don't anticipate making any change to the scope of or criteria for this track, but we will continue experimenting with ways to broaden awareness. Because the potential applicant pool for this track is wide, this is a good place to invest in relationships with other funders, so projects are directed to the organization best placed to help them with a particular funding issue.

One extension we would like to try is to see how we can bring mission partners to non-Western geographies and non-English linguistic communities. We envisage a pilot with a small, local awards committee (perhaps containing one or two members of the main committee for experience) and its own, smaller, budget, which can make awards which make sense for that geography or locale. Options include Spanish-speaking South America, Brazil, Africa or South Asia. This would promote decentralization, a current Mozilla priority.

Secure Open Source

We feel the concept for SOS has been proved, and the next step is to ramp up the level of output. If the MOSS budget continues at the 2016 level of US$0.5M, we expect to be able to spend 100% of it without too much difficulty, now that we have six auditors on the books.

If SOS were to continue at its current scale, of the three inputs to SOS - money, auditors and projects - it would probably be the project input which needs most attention in 2016. We would need to devote some time to researching potential audit targets, and perhaps developing some more objective evaluation criteria to make sure we were targeting our money most effectively. So far, we have been using common sense and criteria produced by others, but our focus is not the same and so we will need to develop our own capabilities here.

In order to scale above US$0.5M, we would need two things. The first is to seek funding from additional sources. This was always the plan for SOS Fund - that, once we had proved the concept, we would seek support from non-profits, companies and governments who believe this work is valuable but do not wish to establish their own program or center of excellence. Mozilla's SOS contributions would then become just one among many as the program grows.

The second thing we would need would be a part-time administrator. SOS Fund requires much more administration per project than the other two tracks, and the expense per audit is normally about US$25K, compared to an average MOSS award on the other two tracks of US$50K-75K, so there are more projects for the same money. Coordinating with projects and auditors, and guiding a project through the process, is a high-touch activity. Also, if we seek external funding, we would also need to manage donor relations. Therefore, we would need to budget for and hire an administrator, probably part-time, with an understanding of the open source community and good personal and organizational skills.

Rust

In 2017, we also hope to investigate ways in which the security of the open source ecosystem can be improved other than by audit. One avenue we have been exploring is leveraging Rust. Several projects we were asked to audit are very large, but they consist of a relatively small core and then a large number of protocol or codec-specific plugins. Examples of such projects are wire shark (network protocols), libav (audio/video codecs) and imagemagick (graphics formats). Often, security vulnerabilities in such applications are found in particular plugins for old, unmaintained or unpopular formats. In 2016, we have done initial investigations into what it would take to build the capability for plugins for these applications to be written in or converted to Rust, to provide the memory safety and concurrency guarantees which would prevent those plugins being used as attack vectors. We believe that both vlc/libav and wireshark are promising routes for this, and the next step is opening a conversation with the maintainers about it.

Another route to consider would be funding green field Rust replacements for key pieces of open source software which have a historically poor track record of security.

Conclusion

MOSS has enjoyed a reasonable level of success in 2016, but changes and additional effort are necessary to make sure it fulfills its potential to benefit Mozilla, Mozilla's mission and the open source community in 2017 and beyond.