



# Tor Metrics Final Report

Project name: Tor Measurement and Metrics Improvement

July 31, 2017

Karsten Loesing and Karin Herm  
The Tor Project

[karsten@torproject.org](mailto:karsten@torproject.org)

0CC8 157C 2578 EEE9 2B07  
45D6 4EFD 4FDC 3F46 D41E

[iwakeh@torproject.org](mailto:iwakeh@torproject.org)

B788 B367 68B6 9B4D D8E1  
F69E 7525 6C22 E303 259C

# Table of Contents

<b>Table of Contents</b>	<b>2</b>
Background (previous state of code)	3
What you applied to do, and how much for	3
What you did, and whether it was exactly the same as what you planned	4
Who worked on the project	4
How long it took, and whether that was what you expected	5
What positive outcomes you have seen from the work being done (uses of the code, people appreciating new features, community growth etc.) and what the impact has been	5
What you would do differently if you did it all again	6
What you plan to do next	6

## Background (previous state of code)

Tor Metrics has been a part of Tor since 2009 and was originally funded by the National Science Foundation. In the beginning, Tor Metrics only had one full-time developer who added new features and requests from other projects. While this work showed the importance of the work for the rest of the Tor community, having only one person working on the project limited our ability to enhance core parts of Tor Metrics.

The MOSS award helped us address this problem, allowing us to add another member and giving developers the chance to discuss design choices. This is obviously a better scenario than having one developer reviewing his own code. Additionally, the MOSS award allowed us to make crucial improvements to the Tor Metrics code base, including refactoring the code and adding more tests. We kept our development cycle healthy and were able to build the necessary foundation to scale and become more mature.

## What you applied to do, and how much for

Work done under this grant sought to improve Tor Metrics in six core areas, which became our six smaller projects.

1. Strengthen our data collection infrastructure to increase the number of instances and operators to create a more reliable system and expand our observational capacity. Budget: \$24,500.
2. Strengthen our performance measurement infrastructure by building and providing an observation “kit” that contains: Software (Tor DescripTor) that helps use and analyze CollecTor outputs, plus user-friendly tutorials on using DescripTor to analyze large amounts of network data. Budget: \$28,500.
3. Strengthen our network status monitoring infrastructure. We want to encourage third-party operators to run Tor network monitoring services (Tor Onionoo) by building and providing a service implementation kit that contains related software tools and tutorials. We would also resolve some of the most pressing usability issues of Tor's Atlas network status website, which publishes network monitoring results. Budget: \$27,000.
4. Improve the security of metrics reported by individual Tor network relays by reducing the amount of sensitive, potentially personally identifying, data we store for measurement and further obfuscating that data in the reporting process. Budget: \$24,000.
5. Improve the accuracy and depth of performance measurements by developing better user models and by using an improved traffic generator. Budget: \$25,000.

6. Make it easier for people to find, compare, and interpret Tor usage and performance metrics by redesigning the primary Tor Metrics website, reflecting Mozilla's 8th Principle. This includes adding metrics, such as Tor Browser download and update statistics requested by the community. Budget: \$23,500.

## What you did, and whether it was exactly the same as what you planned

We had six objectives in all, and we largely stuck to this plan. We divided the 12 months of this project into increments of approximately two months. There was some overlap in these smaller projects, but we gave each the time which was required.

More information is available in the documents related to individual milestones, as well as in invoice correspondence. Monthly reports are available for [July](#), [August](#), [September](#), [October](#), [November](#), [December](#), [January](#), [February](#), [March](#), [April](#), [May](#), and [June](#).

Despite not being explicitly mentioned in our milestone documents, another crucial improvement to Tor Metrics was that the overall quality of the code in the Metrics code base improved significantly.

Two objectives for Milestone 4 were to reduce the amount of personally-identifying information and to obfuscate data stored by Tor relays and bridges to the directory authorities. While we completed all objectives listed under this milestone, we decided not to merge the code written for these two objectives. This was not our plan, but we realized that to do so would risk making the statistics we collect less useful. We did, however, complete the code, run it in a separate simulation, and [publish](#) the results of this simulation. More time and resources than we originally estimated will be needed to address this issue further. We hope to return to this work in a future proposal.

## Who worked on the project

Karsten Loesing and Karin Herm carried out most of the work of this grant. Some deliverables, such as the reorganization and redesign of the Tor Metrics website, involved other people at the Tor Project, such as UX team lead Linda Lee, project manager Isabela Bagueros, web designer Raphael Bergmann, and former communications director Kate Krauss. Developers Georg Koppen and Nicolas Vigier of the Tor Applications Team provided insights into Tor Browser download statistics. Another deliverable involved web developer Iain Learmonth, who made usability improvements to the Tor Atlas website and who was assisted by several anonymous volunteers in the Tor community. The OnionPerf deployment deliverable depended on the work of researcher Rob Jansen and services admin Silvia Puglisi. Researchers Aaron Johnson and Paul Syverson as well as Tim Wilson-Brown of the Tor

Network Team provided valuable technical feedback on the Tor daemon deliverable. Tommy Collison edited this final report.

## **How long it took, and whether that was what you expected**

As described above, this project had a bigger workload than anticipated. We still completed all deliverables by the end of the 12 months, although at times we had to involve two people rather than one.

## **What positive outcomes you have seen from the work being done (uses of the code, people appreciating new features, community growth etc.) and what the impact has been**

Initial funding from the National Science Foundation allowed us to hire a developer to work full-time on Tor Metrics. The MOSS award allowed us, for the first time, to focus exclusively on making Tor Metrics better, including making important improvements which were not visible to the user, such as refactoring the code. These improvements were compounded by the fact that we put out several releases in the 12 months of this award.

Our redesign of the Tor Metrics website was not simply cosmetic: it involved changing how we organized the site's content, making it more accessible to our diverse user base. This reorganization comes just before a major planned redesign of the Tor website and support portal, and so will provide guidance as to the best way to carry out such work. We also made metrics-lib documentation easily accessible on the website, making it easy for people to get involved.

A valuable impact of this work is that Tor network data no longer depends on a single host. Just a few weeks after setting up a mirror for this data collection, the main instance broke with no warning or indication. We would have lost a couple hours of data if this had happened before we made our improvements, but we were able to restart the primary instance and obtain the missing data from our backup host.

Finally, our improvements to Atlas were a benefit to the wider community, particularly volunteers who operate relays and others who look out for inconsistencies or misconfigurations in the network.

## What you would do differently if you did it all again

Broadly speaking, we would take a similar approach if we were to do it all again, but we have identified two areas of improvement:

We would reconsider how we divide our efforts. Instead of having six small projects, we may organize the work into four slightly larger projects. It takes time after switching between tasks and contexts to work productively, and having a smaller number of tasks would increase productivity. A downside of this approach, however, is that there would be greater impact if we failed to deliver on one of our larger four projects.

Also, we would be more conservative when planning deliverables and outcomes based on analyses. Because we underestimated the complexity involved, we had to adapt the deliverables for Milestone 4.

## What you plan to do next

The MOSS award allowed us to work on the most pressing issues in the most relevant code bases. As mentioned above, this work also helped us identify the next steps of Tor Metrics development. We intend to focus on documentation and specification, to be more transparent in what we do and to encourage others to get involved. Additionally, we want to consolidate services and increase our reliance on existing frameworks. This will allow us to reduce the amount of code we have to maintain, and make it easy for others to contribute to Tor Metrics.