



Tor Metrics Roadmap 2017/18

November 17, 2017
Karsten Loesing, Karin Herm, Iain R. Learmonth

The Tor Project

karsten@torproject.org
0CC8 157C 2578 EEE9 2B07
45D6 4EFD 4FDC 3F46 D41E

iwakeh@torproject.org
B788 B367 68B6 9B4D D8E1
F69E 7525 6C22 E303 259C

irl@torproject.org
A8F7 BA50 41E1 3333 9CBA
1696 76D5 8093 F540 ABCD

Table of Contents

Table of Contents	2
Background	2
High-level goals	3
Effort and capacity estimation	4
Short-term goals (Q4/2017)	4
Medium-term goals (until Q3/2018)	5
Long-term goals (Q4/2018 or later)	6

Background

This is what we, the Tor Metrics Team, would like to work on in the 12 months between October 2017 and September 2018.

When writing this roadmap, we took the following aspects into account:

- We approached this task of writing a roadmap from two sides: top-down by thinking about the vision of what we want to achieve and bottom-up by going through the list of open issues.
- We asked people in the Tor community for input on this roadmap in order to learn if some goals may be less relevant to them than we expected and to identify goals we have overlooked.
- We received funding for two important documentation tasks where we document how we aggregate data and why we designed our data pipeline the way we did. We mention these deliverables here, though they were already fixed and not subject to change at the time of writing this roadmap.
- We need to remain open to make enhancements for other teams to fulfill their existing or upcoming sponsor obligations. We need to reserve some part of our time for these unforeseen tasks.
- We will need to keep services running and fix any bugs that show up during the roadmapped interval. These tasks are highly unpredictable and not explicitly mentioned in this roadmap, but we need to do them anyway.

High-level goals

In the process of writing this roadmap we identified 10 high-level goals that we want to achieve in the next 12 months. We care about all ten of these high-level goals, though some are more important to us and what we do than others. The following list is ordered by priority for this roadmap, going from highest priority to lowest:

- **Research:** Enable researchers to write better papers about the Tor network. Our two existing documentation deliverables, in particular the one where we document how we're aggregating data, will be a big step into this direction. These two external deliverables are already funded and will use all the resources we can dedicate to this goal during the next 12 months.
- **Code:** Make our code cleaner, more efficient, and more robust (without adding specific features). This is not just because we love clean code, it's because we still identified a lot of technical debt that we need to pay off before making more enhancements and inviting more volunteers to contribute to our code base.
- **Operation:** Make our services easy enough to operate to not rely on a single person anymore. Like the previous high-level goal, this is a sort of technical debt that we need to eliminate in order to reduce the time in the future to run services and do emergency responses.
- **Sources:** Provide a complete archive of Tor network data and closely related data. This is one of the core reasons why Tor Metrics exists, and it's one of the most important high-level goals in this roadmap. This high-level goal includes making existing data sources more robust, and it also includes adding new data sources.
- **News:** Enable others to learn about and collaboratively analyze network-wide events. As we have realized over the past months, others in the Tor community (with dcf leading the way) are better at analyzing network-wide events than we are. Which is great, because whoever does the job best should do it. We want to support these friendly people by providing them tools and a platform to share their findings about events in the Tor network.
- **Services:** Provide a complete 'toolbox' for relay operators, directory operators, and "network angels" who keep an eye on the wider network (like nusenu) to keep the Tor network healthy.
- **Visualizations:** Provide customizable graphs and tables that cover most questions that website visitors may have and that are understood by most of them. We realize that these visualizations are the most visible part of our work. However, we decided to give them less priority in the next 12 months, because it will be easier to improve these visualizations when things are more consolidated.
- **About:** Make it obvious for team-external people what we do, what is part of our work and what is not, and how they can contribute. In the past it has not always been obvious what belongs to Tor Metrics and what does not. During the last months we improved the situation a lot by working towards making the Tor Metrics website the primary landing site for everything

related to Tor Metrics. Some work remains, but not as much as on other high-level goals mentioned above.

- **Development:** Enable team-external developers to write code that uses our data and tools. We put quite some effort on this high-level goal in the recent past by attempting to make the Tor Metrics Library more widely known. This goal was a main goal during the year 2016/2017 and thus effort for it will be limited during the next months to very important on-demand support, if these don't interfere with higher prioritized goals.
- **Team:** During the last months we agreed upon processes and guidelines to make working in the team as smooth as possible. This already enabled us to avoid duplicating efforts in the team by agreeing on how we do things before doing them. Thus, we already made a lot of progress here over the past year, so there does not remain as much work to do and we can focus on other goals.

Effort and capacity estimation

When writing this roadmap we performed a quick estimation of effort required to complete any given goal. These estimates are mainly based on having completed similar projects in the past.

For each task, we're assigning one out of the following estimates: M for tasks that would very roughly require the current team 1 week to complete and L for 1 month. For ease of effort computation, we use M as base unit and define: $4 M \approx 1 L$.

The underlying assumption is that there are no other distractions during those times, which is rarely the case. All in all we assume that in any given calendar month we'd be able to finish 2 M tasks.

We hope that we can grow the team to also include Iain as core contributor beginning in July 2018. This would grow the team capacity to roughly 1.5 times the current capacity.

In sum, the team capacity in the next 12 months will be around 27 M tasks.

Short-term goals (Q4/2017)

The high-level goals above give an outlook on where we'd like to be in 12 months from now. In the process of creating this roadmap we went one step further and wrote down possible (lower-level) goals for achieving our high-level goals. The following list contains all goals that we'd like to start or even complete within the remaining months of 2017.

- Code
 - Update all code to use Java 8 features (M)

- Switch ExoneraTor and metrics-web from Tomcat to embedded Jetty, and harmonize directory structures (M)
- Operation
 - Deploy better notification system for operational issues (M)
 - Make sure that each service has at least two operators (M)
- Sources
 - Add web server logs as new data source (M)
- News
 - Provide metrics timeline events as both a table on Tor Metrics pages and as an RSS/Atom feed that is also syndicated via Twitter to increase community engagement (M)
- Services
 - Make Atlas part of the Tor Metrics website (M)
 - Make Compass part of Atlas in order to provide aggregated search results and details pages on Atlas and to finally shut down Compass (M)
 - Resolve ExoneraTor database slowness (M)
- Visualizations
 - Add OONI graphs to Tor Metrics based on pre-aggregated data from vanilla Tor test and bridge reachability test (M)
 - Add IPv6 relay graphs to Tor Metrics based on existing descriptor contents (M)
- About
 - Rename Tor Metrics components to make them easier to understand for people outside of the team (M)

In sum, we're planning to either start or complete 12 M tasks until the end of 2017 which corresponds to 44% of our capacity.

Medium-term goals (until Q3/2018)

In addition to the goals in the previous section we came up with another set of goals that we'd like to complete within the 12 months of this roadmap, but which we don't necessarily plan to start working on before the beginning of 2018. We realize that some goals are harder to plan than others, in particular goals that depend on people outside of the team. These have the potential of keeping us busy in the next 12 months which we need to take into account.

- Research
 - Document aggregated statistics files better (Sponsor 13) (L)
 - Put external research data on the website, including Rob's phantomtrain OnionPerf data and Yixin's BGP monitoring data (M)
- Sources

- Re-process bridge descriptor tarballs to retain “contact” lines in server descriptors and “fingerprint” lines in statuses (M)
- Support PrivCount development and integrate its data as successor to statistics reported by relays and bridges (L)
- Help the Core Tor and Tor Browser teams with ongoing measurements (M)
- About
 - Document the Tor Metrics system/pipeline, including explicitly writing down our requirements (Sponsor 13) (L)

In total, the goals above amount for 15 M tasks which we’re planning to start in 2018 and complete by the end of September 2018 or which are ongoing or on-demand tasks. Together with the the previous section this corresponds to 100% of our capacity.

Long-term goals (Q4/2018 or later)

Finally, we’re listing all goals which seemed like too much for the upcoming 12 months. These may become part of the next roadmap, assuming that they’ll still be relevant in late 2018 or 2019.

- Research
 - Add and maintain a list of relevant research papers on the deployed, public Tor network (M)
- Code
 - Make descriptor parsing multi-threaded (M)
 - Make code more testable and increase test coverage (L)
 - Support parsing of arbitrary large descriptor files, possibly using Java NIO (L)
 - Evaluate possible web frameworks rather than our own servlets and JSPs (L)
 - Turn Onionoo, metrics-web's modules, and ExoneraTor into a single service (L)
 - Use more off-the-shelf components and frameworks to reduce the amount and complexity of our code (L)
- Operation
 - Use a database for news and other content on the website, and automate updating news (L)
- Sources
 - Research IP2Location as alternative source for geolocation data (L)
 - Add geolocation data as new data set to the archive and use it in our own tools (L)
 - Improve (directory-request and possibly other) statistics reported by relays and bridges (L)
 - Extend tor’s directory protocol to make collection of relay descriptors more robust (L)
 - Add metrics on core Tor development from Trac, mailing list archives (L)
 - Resume gathering bridge distribution statistics from BridgeDB (L)
 - Prepare for RSA IDs being replaced with Ed25519 IDs some time after 2018 (L)

- Services
 - Bring back Tor Weather using Onionoo as data source (L)
 - Create a service to search for raw descriptors in the archive (L)
 - Extend Onionoo to process vote documents in order to include bandwidth measurements which may include switching to an efficient database (L)
 - Make consensus health (website and mailing list) an “official” part of Tor Metrics again (M)
- Visualizations
 - Add sparkcharts to the website (M)
 - Move either metrics-web’s modules or website to a separate code repository (M)
 - Switch to a graphing framework and underlying database that support interactive graphs (L)
 - Provide a dashboard with visualizations that users care most about and that they can customize themselves (L)
 - Provide Tor network data in a time series database like Prometheus or Influx DB (L)
 - Evaluate existing visualizations and suggested additions, decide which to remove or add (L)
 - Extend the aggregation module that produces server and bandwidth statistics to include IPv6 capability of relays, to provide similar statistics for bridges as for relays, and to output consensus weights and path selection probabilities as metric rather than absolute relay or bridge numbers (L)
 - Provide graphs showing network diversity in terms of relays being part of the same family or being run in the same AS or country (L)
- About
 - Put more testimonials on the website (M)