Collecting BridgeDB usage statistics

Philipp Winter
phw@torproject.org

April 25, 2019

1 Preliminary remarks

Note that the context for this proposal is not a research project, but rather an engineering task that will hopefully result in Tor Metrics publishing the data we seek to collect. As a result, the outcome of this project would be data, code, and visualizations; but probably no research paper.

2 What are you trying to learn? What are the benefits?

Tor users can use *bridges* to connect to the Tor network if the publicly-documented relays are blocked by their ISP. The distribution of these bridges is handled by a tool called BridgeDB. Users can request bridges from BridgeDB over email (by writing to `bridges@torproject.org`), over HTTPS (by visiting `https://bridges.torproject.org`), and directly in Tor Browser, after solving a CAPTCHA.

At the moment, we have no understanding of how BridgeDB is used. This prevents The Tor Project from making sensible design decisions. For example, it is difficult to estimate the effect of fixing the following proposal: `https://bugs.torproject.org/28496` Worse, we don't even know how many people use BridgeDB, raising the question of how we should prioritize its development. Below is a set of concrete questions that we would like to answer in this project:

- How many requests does BridgeDB see per day?
- What obfuscation protocols are the most popular?
- What bridge distribution mechanisms are the most popular?
- From what countries do we see the most bridge requests?
- How many BridgeDB requests fail and succeed, respectively?
- How many requests does BridgeDB see from Yahoo/Gmail/Riseup?
- How many HTTPS requests are coming from proxies?
- How many requests are suspicious, and likely issued by bots?

We believe that answers to these questions benefit The Tor Project because it allows us to make sensible design decisions and prioritize our development tasks. It also benefits Tor users because it will allow us to react to bridge enumeration attacks, improve bridge distribution to minimize the number of failed requests, and infer what bridge types may be blocked in a given country. Finally, Tor researchers benefit because our published metrics would provide concrete answers to questions such as "how many bridge requests come from bots?". These numbers are generally based on unfounded assumptions.

3 What exactly is your plan?

Each client request to BridgeDB carries with it some information that allows us to answer the above questions. We plan to extend BridgeDB's source code to collect the following information:

- The distribution mechanism over which bridges were requested. Currently, this is HTTPS, email, or Moat.
- The requested transport protocol. Currently this is vanilla, FTE, obfs3, obfs4, or scramblesuit.
- The request's origin. For Moat and HTTPS, it's the user's two-letter country code, e.g., IT for Italy. For email, it's the user's email domain (Gmail, Yahoo, or Riseup).
- Whether the request was successful or unsuccessful, i.e., resulted in BridgeDB handing out bridges or not.
- Heuristics that allow us to infer if a request came from a bot, and to estimate the total number of bots. On our bug tracker, we discussed a number of methods to estimate if a request came from a bot: `https://bugs.torproject.org/9316#comment:19` One example is to count the subnets (e.g., /16 networks) from which BridgeDB receives the most requests.

We intend to keep the metrics safe by (*i*) limiting their granularity, by (*ii*) only collecting what we consider safe to publish, by (*iii*) minimizing the data we collect, and by (*iv*) making sure that the benefits outweigh the risks. In particular, to protect users whose request is the only one in a given bucket (e.g., there may be only one user in Turkmenistan who successfully requested an FTE bridge over HTTPS on 2019-04-02), we would bin the statistics by rounding them up to the next multiple of ten. We further plan to aggregate metrics at a granularity of once per day. Note that our bin size and aggregation period are uninformed guesses. We intend to adjust these parameters if they prove problematic in practice. Taking into account our privacy-preserving methods, our data would reveal information like:

- 41–50 users successfully requested an obfs4 bridge over Moat in Morocco on 2019-04-25.
- 1–10 users unsuccessfully requested a scramblesuit bridge from a Riseup email account on 2019-04-01.
- BridgeDB saw 101–110 successful requests on 2019-04-10.
- The HTTPS distribution mechanism sees more unsuccessful requests on average than the Moat and email mechanism.

4 What attacks did we facilitate?

An adversary would get to learn information like "11–20 users successfully requested an obfs4 bridge over HTTPS in Turkmenistan" or "1–10 users unsuccessfully requested an FTE bridge from Gmail." We believe that in specific cases, this information could help an adversary confirm assumptions about users. Still, we believe that the risk is low considering that our bin size does not provide *exact* numbers, leaving the adversary with uncertainty. We further believe that powerful adversaries already are able to obtain the information we seek to publish.

We believe that our heuristics for bot detection are potentially problematic. For example, prefixes as short as a /8 can map to a single city according to the NetAcuity geo-location provider. If Tor Metrics were to publish this data, it could reveal the physical location of Tor users with too high of a granularity. To mitigate this issue, we are considering collapsing our heuristics into a scalar "bot score," i.e., a score of 0 would signal that a request is unlikely to come from a bot while 1 would signal high probability. Such a score would eliminate privacy concerns.

5 Why do the benefits outweigh the risks?

We believe that the benefits are tangible, clear, and significant. The risks however (apart from some of our heuristics) are minor and clearly outweighed by the benefits. As for our heuristics, we are still in the process of refining them to strike a reasonable balance between use and harm.