

Towards an Evaluation Framework for Tor Load Balancing

Andre Greubel
University of Wuerzburg
Wuerzburg, Germany

andre.greubel@uni-wuerzburg.de

Samuel Kounev
University of Wuerzburg
Wuerzburg, Germany

samuel.kounev@uni-wuerzburg.de

ABSTRACT

Tor is a widely used anonymization network. Traffic is routed over different relay nodes to conceal the communication partner. The Tor Load Balancing Mechanism (TLBM) is responsible for assigning traffic to relays. However, up to this point, there are no agreed-upon quality and security metrics for evaluating the quality of the TLBM. In practice, multiple papers provide their own evaluation using custom metrics incomparable to related work. This paper aims to start bridging this gap by arguing for the necessity of a unified framework to evaluate TLBMs and by giving an overview of aspects that should be included in such a framework.

1 INTRODUCTION

In Tor, clients can build privacy-preserving paths (*circuits*) consisting of relay nodes (*relays*) that forward encrypted TCP packets to other relays in a circuit or destination server on the internet. This way, only the communication from the client to Tor and from the last relay to the destination can be observed. As the available bandwidth of each relay is limited, the *Tor Load Balancing Mechanism* (TLBM) has to ensure that no relay attracts more traffic than it can handle. Additionally, circuits in Tor are anonymous, only if the attacker is not able to simultaneously surveil the first and last node of the circuit. Otherwise, a powerful traffic correlation attack potentially enables complete de-anonymization of that circuit [4, 6].

To decrease feasibility of this and similar (cf. [4]) attacks, traffic should be distributed to as many entities as possible. However, recent research has shown that the current measurement method determining the load distribution in Tor lacks in both security and quality [1, 3, 4, 12]. Furthermore, the current load distribution highly favors high-bandwidth relays [2, 4, 7, 8]. This leads to a reduction of resources necessary to perform several large-scale de-anonymization attacks by more than 80% [4]. Based on these shortcomings, several projects proposed replacements or improvements of the current TLBM, both in security and quality [3, 5, 9, 11].

2 GOALS OF A FRAMEWORK

Yet, despite the importance of the TLBM on the security of Tor and its known shortcomings, there currently is no consensus on how to evaluate a TLBM. In practice, all evaluations of the existing TLBM and all proposals for replacing or improving it provide their own evaluation method and evaluation metrics.

However, different terminology and evaluation metrics prevent scientists and Tor developers alike to compare approaches and identify the best one. In fact, this might be one reason why none of the proposed replacements or improvements were ever included in Tor, despite the community’s awareness (cf. [10]) of these issues. Furthermore, separating the discussion about the metrics from the following evaluation also enables proposals on new designs to be more focused on the actual design. Hence, a unified conceptual framework containing guidelines and established metrics could help to solve these practical issues with the TLBM.

3 BLUEPRINT OF A FRAMEWORK

Such a framework should provide (1) a classification of TLBM approaches, (2) unified terminology, (3) metrics and criteria for evaluating a given TLBM.

The classification should include the general ideas used to design the TLBM. Most notably, it should be able to clearly communicate the security assumptions used to design the TLBM and its design approach (decentralized measurements, trusted measurement entities, semi-trusted measurement entities, ...). Furthermore, there should be exact definitions of the terms used. Most importantly, the term “bandwidth” and its various interpretations during measurements, aggregation, and load distribution in the TLBM should be standardized. An example of such an attempt (including nine different definitions of the term bandwidth) is presented in [4]. However, these definitions are limited to aspects relevant to a centralized measurement process (like the current TLBM) and are not directly applicable to decentralized approaches.

Lastly, such a framework should also include metrics and criteria that a secure TLBM of high quality should fulfill. In the remainder of this section, we will describe such criteria and metrics of the quality and security of the TLBM that we consider to be relevant for an evaluation framework. For better readability, we loosely categorized them into three categories (“security”, “load distribution”, and “other”).

3.1 Security Aspects

In this section, we describe security aspects that a TLBM should fulfill to reduce its attack surface.

Trust-Less The TLBM should not require trust in self-reported data from a non-trusted entity. This is desirable as it prevents relays from influencing the load distribution by lying about self-reported information.

Fair The TLBM is fair if no untrusted entity can be assigned more traffic by other means than actually providing more bandwidth in real-world communication. This is desirable as it prevents relays from influencing the TLBM by other

means than lying about its bandwidth. For example, if it is possible to detect bandwidth measurements in a TLBM, relays might artificially inflate their traffic by providing additional bandwidth during measurement only.

Anonymous The TLBM should not rely on collection, verification, or publication of real-world communication data – or information that could be used to reconstruct this data. This is desirable as Tor is designed to conceal communication information. Giving as few as possible entities the possibility to save and collect this information raises the effort to access this information for an attacker.

Autonomous A TLBM is autonomous if it limits the necessary up-time of entities required to be able to use the network. In the current TLBM, this most notably includes the Directory Authorities and relays. Note that this also implies that the network should be *stable*, i.e., that entities can be expected to behave similarly as before. This enables all entities to re-use old information (like topology information) instead of being required to constantly acquire up-to-date information. This is desirable, as it makes denial-of-service attacks on Tor more difficult.

3.2 Load Distribution Aspects

In this section, we list metrics and criteria that describe whether the load distribution in the network (the “output” of the TLBM) has desirable properties.

Balanced In a perfectly balanced network, each relay has the same utilization (defined as the ratio between its used and available bandwidth). Some balance is necessary, as the available bandwidth of each relay is limited. Furthermore, the assignment of traffic is not necessarily equal to the output of the TLBM [4]. To limit the probability of a relay attracting too much traffic because of these differences, it is desirable that the network is as balanced as possible.

Decentralized The network is perfectly decentralized if $x\%$ of the relays serve in $x\%$ of the circuits. If all circuits handle the same amount of traffic, this is equivalent to $x\%$ of the relays handling $x\%$ of the traffic. This is desirable, as, this way, attackers cannot get more information (surveilled traffic) per invested resources (compromised relays) by focusing on popular relays.

3.3 Further Aspects

Low Overhead A TLBM has a low overhead if its entities require little resources like bandwidth or computing power. This is desirable, as Tor and its entities are run by volunteers who should not be de-incentivized in their work by requiring cost-intensive resources from them.

Simple Design A TLBM has a simple design if the behavior of entities and their trust assumptions are clearly stated, reasonable, and have limited dependencies with each other. This reduces attack surface and simplifies security analysis, leading to a more secure and trustworthy system.

User Performance The performance of a Tor circuit is, among other things, limited by the maximum bandwidth of each relay in a circuit and the latency between the relays.

It is also dependent on client parameters like geolocation and relays already chosen for the circuit. As many of the current deficiencies are caused by deliberate choices aimed at increasing the user performance [4], this aspect should be included in a framework.

4 FUTURE WORK AND CONCLUSION

The most important step for building such a framework includes proposals on how to exactly define, quantify and measure these aspects. Note that this is not trivial: For example, there are currently at least three different approaches to measure decentralization [2, 4, 8]. However, none of the approaches are able to measure a “ground truth”, and all three metrics proposed allow for various interpretations themselves.

But even with a consensus on how to define, measure, and interpret those metrics, one has to discuss trade-offs: Some aspects, even in our preliminary list, directly contradict each other: For example, if Tor has differently sized entities, it cannot simultaneously be perfectly balanced and decentralized. Up to this date, there is no consensus on how these trade-offs should be designed and which aspects are most important.

As such, a complete framework most likely cannot be built by a single research group but rather has to be a project of the TLBM research community at large.

Overall, we hope that, in this work-in-progress paper, we showed the need for such a project and offered fellow researchers a quick overview of aspects of interest for it.

REFERENCES

- [1] Kevin Bauer, Damon McCoy, Dirk Grunwald, Tadayoshi Kohno, and Douglas Sicker. 2007. Low-resource Routing Attacks Against Tor. In *Proceedings of the 2007 ACM Workshop on Privacy in Electronic Society*.
- [2] Tao Chen, Weiqi Cui, and Eric Chan-Tin. 2019. Measuring Tor Relay Popularity. In *Security and Privacy in Communication Networks*.
- [3] André Greubel, Alexandra Dmitrienko, and Samuel Kounev. 2018. SmarTor: Smarter Tor with Smart Contracts. In *Proceedings of the 34th Annual Computer Security Applications Conference*.
- [4] Andre Greubel, Steffen Pohl, and Samuel Kounev. 2020. Quantifying measurement quality and load distribution in Tor. In *Proceedings of the 36th Annual Computer Security Applications Conference (ACSAC 2020)*.
- [5] Aaron Johnson, Rob Jansen, Nicholas Hopper, Aaron Segal, and Paul Syverson. 2017. PeerFlow: Secure load balancing in Tor. In *Proceedings on Privacy Enhancing Technologies*.
- [6] Aaron Johnson, Chris Wacek, Rob Jansen, Micah Sherr, and Paul Syverson. 2013. Users Get Routed: Traffic Correlation on Tor by Realistic Adversaries. In *Proceedings of the 2013 ACM Conference on Computer and Communications Security*.
- [7] Lei Yang and Fengjun Li. 2015. mTor: A multipath Tor routing beyond bandwidth throttling. In *2015 IEEE Conference on Communications and Network Security*.
- [8] Damon McCoy, Kevin Bauer, Dirk Grunwald, Tadayoshi Kohno, and Douglas Sicker. 2008. Shining Light in Dark Places: Understanding the Tor Network. In *Privacy Enhancing Technologies*.
- [9] Asya Mitseva, Thomas Engel, and Andriy Panchenko. 2020. Analyzing PeerFlow. In *Sicherheit 2020*.
- [10] Mike Perry. 2018. Tor’s Open Research Topics: 2018 Edition. <https://blog.torproject.org/tors-open-research-topics-2018-edition> (accessed 25.08.2020).
- [11] Robin Snader and Nikita Borisov. 2009. EigenSpeed: secure peer-to-peer bandwidth evaluation.. In *Proceedings of the 8th international conference on Peer-to-peer systems*.
- [12] Fabrice Thill. 2014. *Hidden Service Tracking Detection and Bandwidth Cheating in Tor Anonymity Network*. Ph.D. Dissertation. University of Luxembourg.