A Wide-Area Testbed for Tor

Roger Dingledine
The Tor Project
arma@mit.edu

David Goulet
The Tor Project
dgoulet@torproject.org

Prateek Mittal
Princeton University
pmittal@princeton.edu

Nicholas Feamster
Princeton University
feamster@cs.princeton.edu

Rob Jansen
U.S. Naval Research Laboratory
rob.g.jansen@nrl.navy.mil

Matthew Wright
Rochester Institute of Technology
Matthew.Wright@rit.edu

1. INTRODUCTION

In this talk, we will describe a plan for the development of a wide-area testbed for the Tor anonymity system. Our proposed testbed would have the ability to directly run real-world software for anonymous communication, including support for clients, relays, and other critical components of Tor, such as directory services and bandwidth authorities. Researchers could select the experimental parameters, including user models, upload modified versions of Tor and associated scripts, and run an instance for a period of time.

In the rest of this talk abstract, we lay out the motivations for having a testbed, the characteristics and use cases we anticipate, and the goals for the HotPETs talk.

2. MOTIVATION FOR A TESTBED

The Tor testbed will enable improved research efforts in a range of areas that are relevant to anonymity in general and Tor in particular. It will allow multiple research groups to experiment with anonymity protocols, such as those used for defenses against traffic analysis [3, 14], AS-aware or trusted path selection [11], faster path selection [5], transport design [8], load balancing [12], and safe data collection [9]. It will enable research that evaluates attacks against Tor [10, 16], as well as defenses against these attacks in a way that does not endanger the privacy of real Tor users.

The proposed testbed can have a transformative impact on anonymity research, especially involving topics that have a strong dependency on the network layer of communications. In most existing evaluation frameworks, such as Shadow [1, 7] and ExperimenTor [2], the network layer is partially abstracted away. This means that attacks that rely heavily on timing information, such as end-to-end correlation [13], congestion-based attacks [4], and latency- and bandwidth-based attacks [6, 15], cannot be reliably studied in these frameworks. Many of these attacks have thus been tested in the live Tor network, which can harm performance and security of Tor users. Also, any system promising performance gains has only limited basis to estimate delays due to the underlying network.

Beyond its uses in research, this testbed will be critical to bringing ideas from the research world into deployment in Tor. Currently, there is a gap between research-based evaluations of an idea in a simulation platform like Shadow and the kind of extensive testing that would be necessary before deployment in a live network. This gap slows innovation, since Tor is rightfully loathe to deploy a new idea when it may lead to issues that could adversely affect users. The testbed we envision would bridge that gap and enable Tor to have more confidence that new ideas would work as intended. This helps not only the research areas that need detailed networking experiments, but also changes to Tor for which simulation is sufficient in the research sense, such as guard selection.

3. TESTBED CHARACTERISTICS

In this section, we outline the salient characteristics and design goals of our proposed testbed.

- Real Tor code: To reduce the gap between evaluation of a research idea, and practical deployment in Tor, the proposed testbed should run real Tor code.
- Wide-area network: An important limitation in existing approaches for private Tor deployment is that they are often use co-located machines in a single organization. It is important for the testbed to be wide-area, to incorporate realistic network latency, throughput, routing and failure characteristics.
- Testbed configuration: The testbed should provide support for flexible configurations, such as for traffic generation and relay workload. This allows researchers to perform a "what-if" analysis and to reason about the Tor network at multiple scales.

4. USE CASES

In this section, we provide example use cases that currently have no outlet via existing experimentation techniques but can be enabled via our proposed testbed.

- Practical traffic analysis: A long-standing challenge for traffic analysis research is to quantify the vulnerability of the Tor network to correlation and fingerprinting attacks. Attacking the real Tor network, however, introduces significant ethical challenges. Our testbed introduces an important design space for understanding the security of Tor.
- Impact of inter-domain routing: A wide-area testbed network opens up opportunities to study the impact of inter-domain routing on Tor, including attacks such as RAPTOR [16].
- Fault tolerance and failure resilience: Will Tor successfully resist attempts to DDoS directory authorities? Bandwidth authorities?
- Performance analysis: A practical wide-area testbed

provides enhanced validation of research ideas that aim to enhance Tor's performance.

5. GOALS FOR A HOTPETS TALK

The primary goal for the HotPETs talk will be to gauge interest in the project. We intend to pose questions for the audience, including:

- Who would be interested to use such a testbed in their current and near-term research projects?
- Who would be interested to use such a testbed in their long-term research?
- What research directions would be enabled by such a testbed beyond what we've already mentioned?
- What requirements would researchers have for such a testbed to make the most effective use of it?
- What experiments would such an experiment enable that are not suitable for PlanetLab or DETER/Emulab environments?
- What existing proposals for Tor should be tested with the goal of near-term deployment?

Realistically, covering the motivation and discussing the above questions will more than fill the alloted time. If time unexpectedly permits, however, we would also like to elicit some thoughts about the challenges of building such a testbed, with questions such as:

- How do we accurately model users and user behavior?
- Since concurrent use of the testbed is likely to harm some of the experimental accuracy, how would you like to see time sharing work on the testbed?
- While leveraging some cloud and hosting providers together with academic institutions is a straightforward deployment path, is it important to include some other locations?
- Since the testbed will not perfectly replicate Tor, e.g. in size and network locations, what remains untestable even with such a testbed in place?

6. REFERENCES

[1] Shadow: Running Tor in a box for accurate and efficient experimentation.

[2] K. Bauer, M. Sherr, D. McCoy, and D. Grunwald. ExperimenTor: A testbed for safe and realistic Tor experimentation. In *Proceedings of the USENIX Workshop on Cyber Security Experimentation and Test (CSET 2011)*, August 2011.

[3] X. Cai, X. C. Zhang, B. Joshi, and R. Johnson. Touching from a distance: Website fingerprinting attacks and defenses. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, CCS '12, pages 605–616, New York, NY, USA, 2012. ACM.

[4] N. S. Evans, R. Dingledine, and C. Grothoff. A practical congestion attack on Tor using long paths. In *USENIX Security*, 2009.

[5] J. Geddes, R. Jansen, and N. Hopper. How low can you go: Balancing performance with anonymity in Tor. In *Proceedings of the 13th Privacy Enhancing Technologies Symposium (PETS 2013)*, July 2013.

[6] N. Hopper, E. Y. Vasserman, and E. Chan-Tin. How much anonymity does network latency leak? *ACM Transactions on Information and System Security*, 13(2), February 2010.

[7] R. Jansen, K. Bauer, N. Hopper, and R. Dingledine. Methodically modeling the Tor network. In *Proceedings of the USENIX Workshop on Cyber Security Experimentation and Test (CSET 2012)*, August 2012.

[8] R. Jansen, J. Geddes, C. Wacek, M. Sherr, and P. Syverson. Never been KIST: Tor's congestion management blossoms with kernel-informed socket transport. In *23rd USENIX Security Symposium (USENIX Security 14)*, pages 127–142, San Diego, CA, Aug. 2014. USENIX Association.

[9] R. Jansen and A. Johnson. Safely measuring tor. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 1553–1567. ACM, 2016.

[10] R. Jansen, F. Tschorsch, A. Johnson, and B. Scheuermann. The sniper attack: Anonymously deanonymizing and disabling the Tor network. In *Proceedings of the 21st Symposium on Network and Distributed System Security*, 2014.

[11] A. Johnson, R. Jansen, A. D. Jaggard, J. Feigenbaum, and P. Syverson. Avoiding the man on the wire: Improving tor's security with trust-aware path selection. *arXiv preprint arXiv:1511.05453*, 2015.

[12] A. Johnson, R. Jansen, A. Segal, N. Hopper, and P. Syverson. PeerFlow: Secure load balancing in Tor. *Proceedings on Privacy Enhancing Technologies (PoPETs 2017)*, 2017(2).

[13] A. Johnson, C. Wacek, R. Jansen, M. Sherr, and P. Syverson. Users get routed: Traffic correlation on Tor by realistic adversaries. In *Proceedings of the 20th ACM conference on Computer and Communications Security (CCS 2013)*, November 2013.

[14] M. Juarez, S. Afroz, G. Acar, C. Diaz, and R. Greenstadt. A critical evaluation of website fingerprinting attacks. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 263–274. ACM, 2014.

[15] P. Mittal, A. Khurshid, J. Juen, M. Caesar, and N. Borisov. Stealthy traffic analysis of low-latency anonymous communication using throughput fingerprinting. In *Proceedings of the 18th ACM conference on Computer and Communications Security (CCS 2011)*, October 2011.

[16] Y. Sun, A. Edmundson, L. Vanbever, O. Li, J. Rexford, M. Chiang, and P. Mittal. RAPTOR: Routing attacks on privacy in Tor. In *24th USENIX Security Symposium (USENIX Security 15)*, pages 271–286, Washington, D.C., Aug. 2015. USENIX Association.