

# One Fast Guard for Life (or 9 months)

Roger Dingledine<sup>1</sup>, Nicholas Hopper<sup>2\*</sup>, George Kadianakis<sup>1</sup>, and Nick Mathewson<sup>1</sup>

<sup>1</sup> The Tor Project, <https://torproject.org>  
{arma,asn,nickm}@torproject.org

<sup>2</sup> University of Minnesota, Minneapolis, MN USA  
hopper@cs.umn.edu

**Abstract.** “Entry Guards” in the Tor anonymity network mitigate against several traffic analysis attacks including the “predecessor” attack, statistical profiling, and passive AS-level correlation attacks. Several recent works have shown that the current design does not provide sufficient mitigation against these attacks and may also introduce new vulnerabilities. We propose a simple response to these results: Tor clients should move from using three entry guards to a single, fast entry guard, and rotate entry guards after 9 months rather than after 45 days. We measure the likely effect on anonymity and performance of these changes, and discuss some of the remaining problems with entry guards not addressed by this proposal.

## 1 Introduction

The Tor anonymity network allows *clients* to build anonymous connections by establishing nested, encrypted tunnels through *circuits* of three relays, chosen at random from a list of several thousand volunteer-operated hosts around the world. The *entry* relay knows the client and a *middle* relay; the *middle* relay knows an entry and an *exit* relay, and the *exit* relay knows the destination(s) visited through the circuit, but can only associate them with the middle relay. However, it is generally accepted that if an adversary controls both the entry and exit relays, then a timing attack that correlates the traffic on each end will allow this adversary to link clients with their destinations.

In Tor, each client selects a small set of relays at random – its *guards*, and chooses only from those relays when making the first hop of all of its circuits. This entry guard design mitigates several attacks, including the “predecessor attack” [14], the selective denial of service attack [3], and statistical profiling attacks, by decreasing the opportunities for an attacker to be the first hop of a particular user. An additional benefit is that to be chosen as an entry guard, a relay must contribute to the Tor network for several weeks, increasing the “startup cost” to an adversary of launching relay-based attacks.

**Guard Rotation Weaknesses.** In the current version of the Tor protocol, clients choose three guards, and each guard is discarded after it has been used

---

\* Work done while on sabbatical with the Tor Project

for 30-60 days. This rotation of guards serves several purposes. First, if a client has unluckily chosen a compromised guard, rotating her set of guards gives her a chance to regain some privacy. Additionally, if clients never rotated their guards, then guards would tend to accumulate more clients the longer they participated in the network, leading to poor load-balancing. Moreover, a client’s choice of guards would be strongly dependent on the time she joined the network, perhaps aiding attacks that partition clients by time of participation.

However, several recent works have suggested that the current parameters for guard rotation – the number of guards, criteria for selecting guards, and the rotation time – may expose users to more privacy loss than anticipated at the time of their selection.

In [5], Elahi *et al.*, using historical Tor network data, simulated a relatively weak adversary: after all clients choose an initial set of guards, the adversary can add a single corrupt guard to the network. Their study then determined, for several rotation policies, what fraction of clients eventually chose this guard. They found that if clients never rotated guards unless a current guard was unavailable, only 10% of clients used the corrupt guard by the end of the simulated 8 month period, whereas using the current rotation policy, 14% of clients would use the corrupt guard within the first three months.

Biryukov, Pustogarov and Weinmann [2] also considered guard rotation as part of a larger attack focused on hidden services running over Tor, finding that an adversary running 13.8% of the Tor network for 8 months would have a 90% chance to be chosen as a guard by a given client within that time frame.

Finally, Johnson *et al.* [6] proposed a new metric for security in anonymity networks that are vulnerable to end-to-end correlation attacks: time to first compromised circuit. They showed, using more recent historical data, that an adversary controlling 10% of the guard capacity (and a fast exit node) of the Tor network had an 80% chance to compromise at least one circuit of any given user within six months. They also showed that increasing the rotation period would partially mitigate the attack.

Taken together, these works suggest that, while entry guards do help to mitigate known attacks against Tor users, the current guard parameters are not strong enough to prevent large-scale attackers from de-anonymizing most users within a relatively short time frame. This motivates us to propose new parameters that improve resistance to these attacks.

## 2 Proposal

### Move to a single guard node

Currently, Tor clients select three relays to use as guards, and only select additional guards when less than two of these guards are reachable. We propose using a single relay as a guard, and only choosing a new guard if this relay becomes unreachable. If that Guard node ever becomes unusable, rather than replacing it, the client picks a new guard and adds it to the end of the list. When choosing

the first hop of a circuit, Tor considers all guard nodes from the top of the list sequentially until it finds a usable guard node.

### **Raise the guard bandwidth threshold**

The current Tor protocol specifies that a node can be chosen as a guard if it has a sufficient *weighted fractional uptime* and if its measured bandwidth is above the median (currently about 250 KB/s). However, as Elahi *et al.* [5] have shown, reducing the number of guards from three to one will leave some clients with relatively poor performance due to having low-throughput guards. Thus we propose increasing the bandwidth threshold from 250KB/s to 2MB/s; we analyze the effect of this change on anonymity and performance in the next section.

### **Increase the guard rotation period**

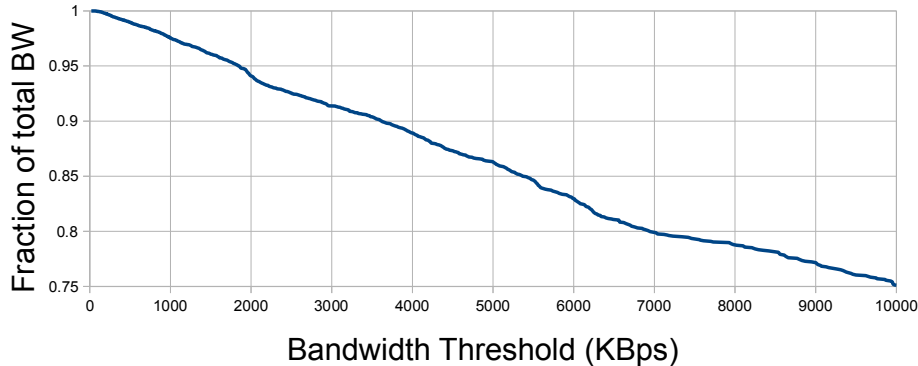
Since Johnson *et al.* found that a primary factor driving the exposure of nodes to end-to-end correlation attacks was rotation to compromised guards, we propose increasing the period for which users retain a guard from 60 days to 9 months. Analysis by Johnson *et al.* found that this significantly increases the time to first compromise. We study the impact on availability of guard nodes in the next section.

### **Fractionally weight guards**

In the current Tor protocol, the trusted directory authorities publish a consensus “network status” document each hour, listing information about all the relays, including public keys, *flags* indicating whether the relay can be used as a guard or exit node, the *advertised* bandwidth that a relay claims to provide, and the *measured* bandwidth as determined by a set of trusted “bandwidth authorities.” In order to balance the traffic load, clients choose relays proportionally to their measured bandwidth capacity; and since not all relays can act as guards or exits, directory authorities also calculate a set of “bandwidth weights” that specify what fraction of the bandwidth from each class of relays (exit, guard, middle, and guard/exit) should be used for each position in a circuit, based on relative scarcity (or surplus) in each class.

Currently, if relay  $R$  has recently become eligible for use as a guard it is heavily under-utilized, because the bandwidth weights allocate a large fraction (currently around 60%) of  $R$ 's bandwidth for use as a guard, but only a few clients have rotated their guards and had an opportunity to choose  $R$ . Increasing the guard rotation period will make this worse, since clients will rotate guards even more infrequently.

We can mitigate this phenomenon by treating these recent guards as “fractional” guards. To do so, when creating the network status entry for a guard, directory authorities will read the past 10 months’ historical network status documents to calculate the *visibility* of the guard; that is, in how many consensuses it



**Fig. 1.** Fraction of remaining guard bandwidth weight as guard cutoff bandwidth is increased, as of February 2014.

has had the guard flag. The authorities then include this visibility in the guard’s network status entry.

A guard  $N$  that has been visible as a guard for fraction  $F$  of the status documents for the last rotation period has had the opportunity to be chosen by approximately fraction  $F$  of the clients. Thus it should be treated by both clients (when choosing middle and exit nodes) and directory authorities (when generating bandwidth weights) as having fraction  $F$  probability of being a guard and  $1 - F$  probability of being a non-guard. For example, if a relay does not have the exit flag, and has been a guard for fraction 0.2 of the last rotation period, then instead of having 40% of its bandwidth available to act as a middle relay, it would have  $0.2 \times 40\% + 0.8 \times 100\% = 88\%$  of its bandwidth available for use as a middle relay.

### 3 Measurements

#### 3.1 Anonymity

In aggregate our changes greatly reduce both the fraction of users vulnerable to compromise at a given rotation and increase the expected time to first compromise. However, there are two conflicting effects to consider. First, reducing the number of guard nodes from three to one reduces the fraction of nodes that are vulnerable to compromise by a relay adversary that controls fraction  $g$  of the guard bandwidth from roughly  $3g$  to only  $g$ . In contrast, raising the bandwidth threshold for entry guards increases the fraction of guard bandwidth that can be compromised by an adversary with a fixed bandwidth budget;<sup>3</sup> we examine this effect. We also examine the *diversity* or concentration of routes [4].

<sup>3</sup> We assume here that the adversary can allocate this bandwidth among hosts capable of supporting the minimum guard bandwidth; the increased threshold does make it difficult to control a high fraction of guard bandwidth using, e.g. botnet nodes on DSL or slow cable modem connections.




Fig. 2. Number of guards remaining as bandwidth threshold increases

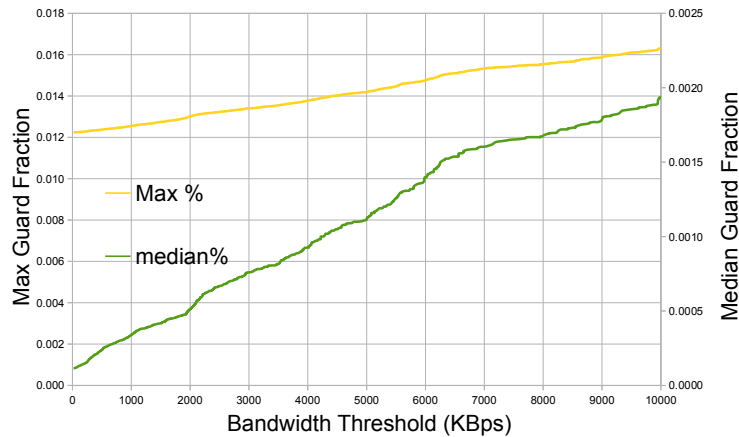
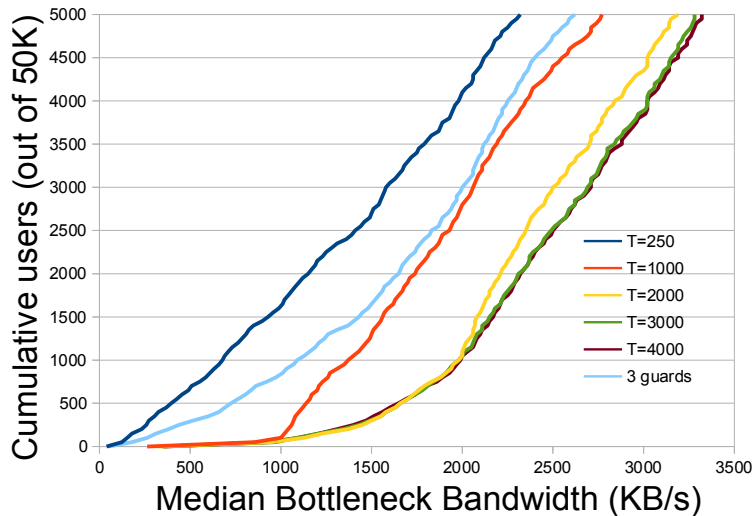


Fig. 3. Median and Maximum probability of guard selection as bandwidth threshold increases.

**Compromise** Figure 1 shows that, as of February 2014, increasing the bandwidth eligibility threshold from 250 KB/s steadily decreases the available bandwidth, but at a very low rate: increasing to 1MB/s only reduces the guard bandwidth by 3% and increasing to 2MB/s only reduces the guard bandwidth by approximately 7%. This means that an adversary that could compromise fraction  $g$  of the guard bandwidth at threshold 250 KB/s would control fraction  $1.07g$  of the reduced pool of guard bandwidth.

**Diversity** Another concern when increasing the bandwidth threshold for guard nodes is that the diversity of paths will decrease [4]: having fewer possible paths through the network could increase both the feasibility of attacks that enumerate possible paths and the possibility of single-node failures. Figures 2 and 3 show that while increasing the bandwidth threshold to 2000 KB/s will decrease the number of guards to just under 1000, the probability distribution on guards will only have a small change, with the median guard’s load increasing from 0.015% to 0.05%.



**Fig. 4.** For bottom 10% of simulated users, distribution of median bottleneck bandwidth under current parameters (3 guards), and single guard selection with varying bandwidth cutoff.

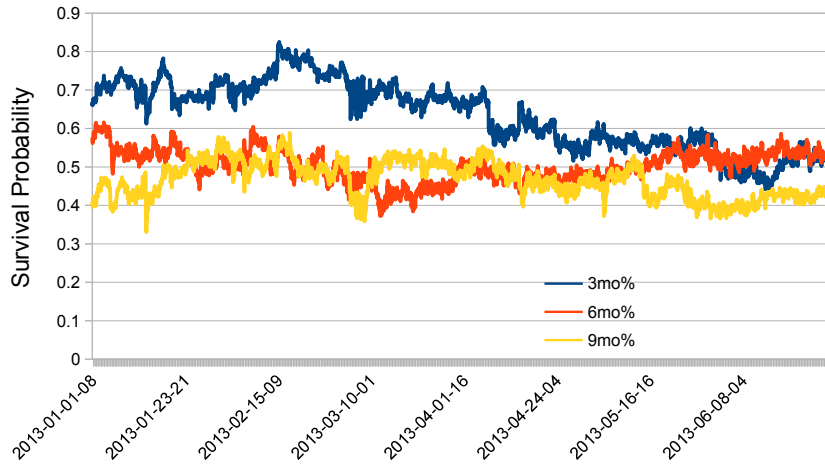
### 3.2 Performance

**Bottleneck bandwidth** To analyze the effect of reducing the number of guards and increasing the bandwidth threshold, we simulated 50 000 Tor clients building 600 circuits each, under several parameter sets: using 3 guards, using 1 guard with the current bandwidth cutoff, and using 1 guard with higher cutoffs: 1000, 2000, 3000, and 4000 KBps. For each circuit, we computed the “bottleneck” bandwidth: the minimum measured bandwidth of the guard, middle and exit relays. Figure 4 shows the median bandwidth available to the bottom decile of users in each simulation. We see that without raising the bandwidth cutoff for guards, reducing to a single guard would lead to some users experiencing heavy reductions in typical throughput rates. However, increasing the threshold to 1000KBps results in circuits with similar typical performance, and increasing the threshold to 2000KBps will slightly improve the typical performance seen by most clients. Increasing the guard bandwidth threshold beyond 2000KBps does little to further improve performance. Coupled with the results showing the decline in diversity above 2000 KBps, these results support an increase in the guard bandwidth cutoff to 2000 KBps.

### 3.3 Availability

Using historical network data<sup>4</sup>, we computed, for each hour starting on January 1, 2013, the probability that a guard chosen according to the consensus network


<sup>4</sup> <https://metrics.torproject.org>



**Fig. 5.** For each consensus network status between January 1 and June 29, 2013, the probability that a guard chosen from this document is present in the network status 3, 6, and 9 months later.

status would be present (useable as a guard) in the network 3, 6, and 9 months later. The results are shown in Figure 5; in general, the probability of a guard being available 6 or 9 months after being chosen was between 40 and 50%, and the conditional probability that a guard was available after 9 months given that it was available after 6 months was generally above 80%.

Additionally, we computed for each hour the expected time to the first failure of a guard not available after 9 months. The results are shown in Figure 6. In this case the average rotation time was typically around 40 days. In combination, these results suggest that reducing to a single guard will not result in more frequent rotations than the current parameters, and that a longer rotation period of 9 months will usefully extend the expected time to first compromise.



**Fig. 6.** Expected time until a guard not available after 9 months is first missing from the consensus network status.

## 4 Unsolved Problems

### 4.1 Picking alternative guards

Instead of picking a new guard when the old guard becomes unusable, an alternative would be to pick a number of guards in the beginning but only use the

top usable guard each time. When a client’s guard becomes unusable, it moves to the guard below it in the list.

This behavior would make some attacks harder; for example, an attacker who shoots down a client’s guard in the hope that the client will pick his guard next, is now forced to have evil guards in the network at the time the client first picked guards. However, there is also a significant probability that the alternative guards will be unavailable. Understanding the anonymity tradeoffs of this behavior remains an interesting open problem.


## 4.2 Multiple paths and performance

Several recent systems [1,12,13] have explored the performance benefits of having multiple, disjoint paths available when building circuits or, in the extreme case of Conflux [1], when routing traffic. Reducing the number of guards to a single node would clearly also reduce the performance benefits that these systems can offer. Measuring the effect of this proposal on these systems, and identifying methods to retain their effectiveness without losing the anonymity benefits from the move to a single guard, remains an interesting direction for future research.

## 4.3 Fingerprinting

**Fingerprinting** One common problem associated with Tor’s current guard node parameters is that each client’s guards serve as a “fingerprint” for the guard: Since the most likely set of three guards, as of April 2014, has probability approximately  $1.7 \times 10^{-6}$ , then assuming 1 million active Tor clients this set of guards has an expected 1.7 users; thus every user’s guard set is distinct with high probability. This could allow malicious exit nodes – in connection with other attacks – to link clients across destinations, and malicious network providers to link mobile clients across locations.

As Figure 7 shows, our proposed changes would partially mitigate this situation. Using the network status as of April 2014, the median user would have an anonymity set of just over 2000 users. However, 10% of users would have an anonymity set of 512 or fewer users, and 1% would have an anonymity set of fewer than 100 users. As a result, mobile clients may still be trackable across multiple internet access points, since if a particular guard has even 100 clients, it is unlikely that two of these clients will be in the same geographic region. A more thorough discussion of these problems and some directions for future work can be found at [7,9,11].



**Fig. 7.** Distribution of anonymity set size with single guards and 2000 KBps cutoff for guard bandwidth.




**Guard buckets** One of the proposed solutions to the fingerprinting problem, is for the directory authorities to split the set of guard nodes into equivalence classes (buckets), and have each client pick and use a bucket of guards. This way there are  $\frac{N}{k}$  potential guard sets instead of  $\binom{N}{k}$  guard sets, where  $N$  is the total number of guards and  $k$  is the number of guards per client. Since there are less potential guard sets with this scheme, more users hide behind each guard set, giving them a greater anonymity set against fingerprinting. Unfortunately, the tradeoffs of this scheme have not been carefully analyzed, and the engineering effort is substantial.

#### 4.4 Guard Enumeration

This proposal does not address the problem that attackers with the ability to cause many circuits to be built by a specific client or hidden service can discover the guards used by that client. Potential defenses proposed include using multiple guard layers [10] or reusing parts of a circuit if it needs to be rebuilt for the same purpose. Effectively mitigating this attack remains an open and interesting question for future research.

#### 4.5 Measured vs. Advertised Bandwidth

Currently, directory authorities use a relay’s advertised bandwidth to determine whether it is eligible for the guard flag, and clients use the measured bandwidth when choosing guards. This could lead to an “attack” wherein an adversary chooses to advertise bandwidths just above the cutoff for guard eligibility, even though they will be chosen with lower probability. Alternatively, we could consider using measured bandwidth in determining eligibility for use as a guard. This could potentially lead to undesirable dynamics for relays that have measured bandwidth close to the cutoff, especially if being a guard increases the probability of having a reduced bandwidth measurement.



**Fig. 8.** Distribution of relative change in measured bandwidth 60 days after first guard flag.

However, since some guards will have decreases in bandwidth measurements, it is unclear how to handle this dynamic, and whether the potential attack is serious enough to justify the additional complexity this might entail.

An initial analysis shows that since January 2013, this is not the case — on average, when a relay is first assigned the guard flag it is more likely to have an increase in measured bandwidth after a complete guard rotation period, as shown in Figure 8, even when compared to a random sample of high-uptime relays. However, since some guards will have decreases in bandwidth measurements, it is unclear how to handle this dynamic, and whether the potential attack is serious enough to justify the additional complexity this might entail.

## 4.6 Assigning Guard Flags

As mentioned in Section 2, Tor clients currently consider a relay eligible to be a guard if it has a weighted fractional uptime greater than the median over all relays, has been present in the network consensus for at least two weeks (or longer than 12.5% of the relays) and meets a minimal bandwidth threshold. Our results show that a sizable fraction of relays satisfying this criterion will still be present after 6 or 9 months, but many will also drop out of the network significantly earlier. Thus it remains an interesting question for future work whether there is a secure algorithm for selecting guards that will simultaneously improve the expected time to failure of guard nodes while maintaining the level of available guard bandwidth.

## 4.7 Privacy Cost and Benefit of Rotation

For users with trustworthy guards, rotation exposes users to the risk of choosing malicious guards. Under the assumption that users have a stable behavioral profile, the first exposure to a compromised guard results in the permanent loss of privacy for a user since statistical profiling techniques are generally very effective. However, if we assume that statistical profiles degrade in value over time or that users acquire new private behaviors over time, then rotation also has a privacy *benefit* to users when rotating away from a compromised guard. Potentially some choices of rotation parameters could then benefit the network as a whole by reducing the expected payoff of an attack to the point where it no longer exceeds the cost. Developing more accurate models of the changes in users' private behavior and the value of this information over time is an interesting direction for further research.

## 5 Conclusion

We have outlined a proposal that can address many of the weaknesses associated with the current guard parameters in the Tor network, as well as introduced a set of open problems that require further research. Furthermore, we authored a technical proposal outlining our suggestions and submitted it to the Tor development team [8]. We hope this document will foster a discussion about these problems, and welcome future work on their solution from the PETS community.

## Acknowledgements

This work was supported by the U.S. National Science Foundation under grants 1111734 and 1314637, and DARPA. The authors thank several members of the Tor developer community for useful comments on the evaluation of Proposal 236 [8], including Tariq Elahi, Ian Goldberg, Mike Perry and Leif Ryge.

## References

1. AlSabah, M., Bauer, K., Elahi, T., Goldberg, I.: The path less travelled: Overcoming Tor's bottlenecks with traffic splitting. In: Proceedings of the 13th Privacy Enhancing Technologies Symposium (PETS 2013). (July 2013)
2. Biryukov, A., Pustogarov, I., Weinmann, R.P.: Trawling for tor hidden services: Detection, measurement, deanonymization. In: Proceedings of the 2013 IEEE Symposium on Security and Privacy. (May 2013)
3. Borisov, N., Danezis, G., Mittal, P., Tabriz, P.: Denial of service or denial of security? How attacks on reliability can compromise anonymity. In: Proceedings of CCS 2007. (October 2007)
4. Dingledine, R.: Measuring the safety of the Tor network. Technical Report 2011-02-001, The Tor Project (February 2011) <https://research.torproject.org/techreports/measuring-safety-tor-network-2011-02-06.pdf>.
5. Elahi, T., Bauer, K., AlSabah, M., Dingledine, R., Goldberg, I.: Changing of the guards: A framework for understanding and improving entry guard selection in tor. In: Proceedings of the Workshop on Privacy in the Electronic Society (WPES 2012), ACM (October 2012)
6. Johnson, A., Wacek, C., Jansen, R., Sherr, M., Syverson, P.: 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)
7. Kadianakis, G., Hopper, N.: Set of guard nodes can act as a linkability fingerprint. <https://trac.torproject.org/projects/tor/ticket/10969> (February 2014)
8. Kadianakis, G., Hopper, N.: "The move to a single guard node" proposal (236-single-guard-node.txt). <https://gitweb.torproject.org/torspec.git/blob/HEAD:/proposals/236-single-guard-node.txt> (March 2014)
9. Mathewson, N.: Comment on "Brainstorm tradeoffs from moving to 2 (or even 1) guards". <https://trac.torproject.org/projects/tor/ticket/9273#comment:3> (July 2013)
10. Øverlier, L., Syverson, P.: Locating hidden servers. In: Proceedings of the 2006 IEEE Symposium on Security and Privacy, IEEE CS (May 2006)
11. Ryge, L., et al.: entry guards and linkability. Mailing list thread: <https://lists.torproject.org/pipermail/tor-dev/2013-September/005423.html> (September 2013)
12. Wacek, C., Tan, H., Bauer, K., Sherr, M.: An Empirical Evaluation of Relay Selection in Tor. In: Proceedings of the Network and Distributed System Security Symposium - NDSS'13, Internet Society (February 2013)
13. Wang, T., Bauer, K., Forero, C., Goldberg, I.: Congestion-aware Path Selection for Tor. In: Proceedings of Financial Cryptography and Data Security (FC'12). (February 2012)
14. Wright, M., Adler, M., Levine, B.N., Shields, C.: The Predecessor Attack: An Analysis of a Threat to Anonymous Communications Systems. ACM Transactions on Information and System Security (TISSEC) 4(7) (November 2004) 489–522