

Protecting Tor from botnet abuse in the long term

Nicholas Hopper

Tor Tech Report 2013-11-001
November 20, 2013

1 Introduction

Starting on August 20, 2013 the Tor network has seen a rapid spike in the number of directly connecting users. This spike is apparently due to the large “mevade” click-fraud botnet running its command and control (C&C) as a Tor Hidden Service. Figure 1 shows that estimated daily clients increased from under 1 million to nearly 6 million in three weeks. Figure 2a shows the effects on performance: measured downloading times for a 50 KiB file doubled, from 1.5 seconds to 3.0 seconds.

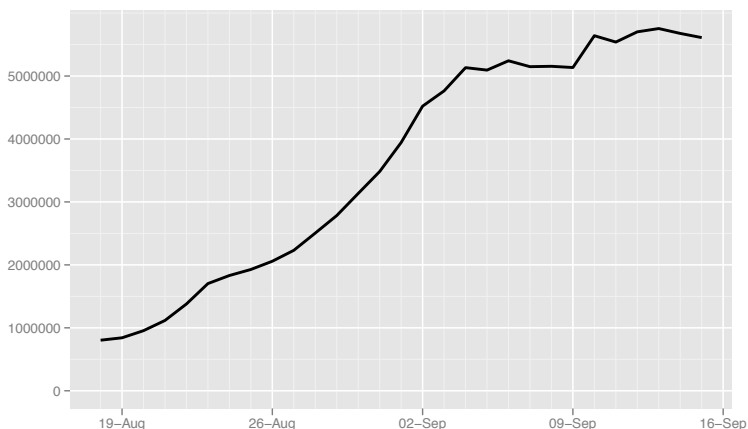
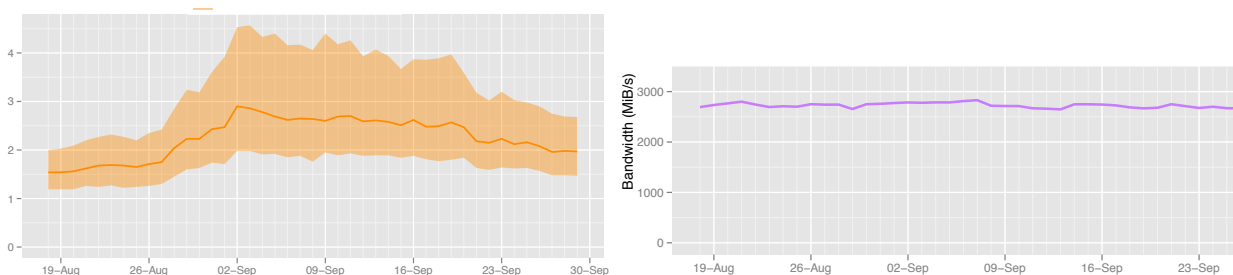


Figure 1: Estimated daily Tor users, 18 August to 13 September 2013

However, the amount of traffic being carried by the network did not change dramatically, as seen in Figure 2b. The primary cause of the problems seems to be the increased processing load on Tor relays caused by the large increase in key exchanges required to build anonymous encrypted tunnels, or *circuits*. When a Tor client - an *Onion Proxy* or OP - connects to the network, it sends a `CREATE` cell to a Tor node, called a *guard*, which contains the first message g^a in a Diffie-Hellman key exchange, called an “onion skin”; the node receiving the create cell



(a) Measured 50 KiB download times, in seconds (b) Reported Bandwidth Histories, all relays.

Figure 2: Download times and total bytes carried by the Tor network.

computes the shared key g^{ab} and replies with the second message g^b , creating a 1-hop circuit. After this, the client iteratively sends onion skins in `EXTEND` cells to the end of the circuit, which extracts the onion skins and sends them in `CREATE` cells to the next relay, until all three hops have exchanged keys.

- Extending a circuit – decrypting an “onion skin” and participating in a Diffie-Hellman key exchange – is sufficiently compute expensive that high-weight relays can become CPU-bound. The total bandwidth that high-weight relays can handle in onion-skins is significantly lower than the network bandwidth.
- The hidden service protocol – explained in section 2 – causes at least three circuits to be built every time a bot connects.
- When onion skins exceed the processing capacity of an OR, they wait in decryption queues, causing circuit building latencies to increase.
- Queued onion skins eventually time out either at the relay or the client, causing the entire partial circuit to fail, causing more onion skins to be injected to the network.

In response to this, the Tor Project modified release candidate 0.2.4.17-rc to prioritize processing of onionskins using the more efficient ntor [7] key exchange protocol. Adoption of this release helped the situation: as Figure 2a shows, measured download 50 KiB times as of late September decreased to roughly 2.0 seconds. Figure 3 shows that circuit extensions using tor version 0.2.4.17-rc range between 5% and 15%, while circuit extensions using the stable release, version 0.2.3.25, ranged between 5% and 30%. By November 2013, further efforts to find and remove the infection by anti-malware teams from companies including Microsoft have mitigated the immediate threat, though the several unmanaged hosts remaining could still revive the botnet.

In this document, we consider longer-term strategies to ease the load on the network and reduce the impact on clients. Full evaluation of the effectiveness of these strategies, impact on privacy and performance for regular users, and relative ease of deployment remains an ongoing challenge; the Tor Project welcomes the collaboration of the research and anonymity community in meeting this challenge.

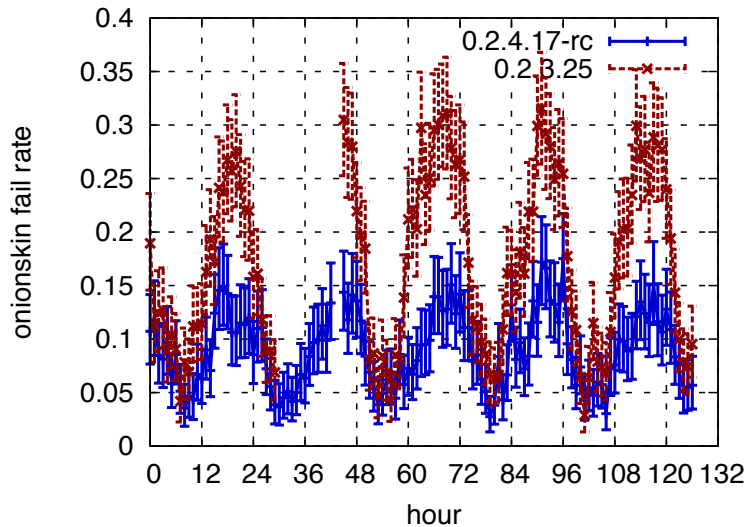


Figure 3: Hourly measured failure rates, starting 27 September 2013, of EXTEND cells, for latest stable Tor release (0.2.3.25) and ntor-prioritizing release candidate (0.2.4.17-rc).

We assess these strategies with the security goal of ensuring the availability of Tor under the threat of a botnet that uses hidden services as its primary C&C channel. A “medium-term” strategy is one which might be successful against a botnet that does not modify the behavior of basic Tor binaries and does not respond strategically to changes in Tor designed to mitigate botnet-induced load. On the other hand, a “long-term” strategy must contend with a botnet that could be deployed in response to such methods, where the behavior of both the botnet and the tor software can change adaptively to circumvent mitigation mechanisms. We note that some attacks are out of the scope of this document, in particular we do not consider a botnet that simply communicates with non-hidden servers through Tor, since such an attack must contain traditional network addresses, and we do not consider a botnet that simply seeks to conduct a denial of service attack on Tor by flooding the network with traffic in excess of its capacity.

The remainder of this manuscript describes the major categories of medium- and long-term responses that Tor could consider, along with the technical challenges that each would pose to the research community. Section 3 considers medium-term strategies aimed at eliminating a current botnet threat to Tor. Section 4 considers longer-term mechanisms to limit the rate of circuit-building requests by any botnet. Section 5 describes mechanisms to reduce the load from ordinary clients. Section 6 considers the idea of isolating hidden-service from regular Tor client traffic.

2 Background: Tor Hidden Services

The Tor network provides a mechanism for clients to anonymously provide services (e.g., websites) that can be accessed by other users through Tor. We briefly review the protocol for this mechanism:

1. The hidden service (HS) picks a public “*identity key*” PK_S and associated secret key SK_S . The HS then computes an “*onion identifier*” $o_S = H(PK_S)$ using a cryptographic hash function H . Currently, the hash function H is the output of SHA1, truncated to 80 bits. This 10-byte identifier is base32-encoded to produce a 16-byte `.onion` address that Tor users can use to connect to HS, such as `3g2upl4pq6kufc4m.onion`.
2. The HS constructs circuits terminating in at least three different relays, and requests these relays to act as its *introduction points* (IPs).
3. The HS then produces a “*descriptor*,” signed using the SK_S , that lists PK_S and its IPs. This descriptor is published through a distributed hash ring of Tor relays, using o_S and a time period τ as an index.
4. A client OP connects to the HS by retrieving the descriptor using o_S and τ , and building two circuits: one circuit terminates at an IP and the other terminates at a randomly-selected relay referred to as the *rendezvous point* (RP). The client asks the IP to send the identity of the RP to the HS.
5. The HS then builds a circuit to the RP, which connects the client and HS.

Since lookups to the distributed hash ring are performed through circuits as well, and each descriptor has three redundant copies, a client connecting to a hidden service could require building up to 6 circuits; to reduce this load, clients cache descriptors and reuse rendezvous circuits any time a request is made less than ten minutes after the previous connection.

3 Attack the Botnet

One set of possible medium-term responses would be to build defenses that protect Tor against abuse by the current botnet. We consider an escalating series of possible mechanisms, depending on the adaptivity of the botnet.

3.1 Descriptor Blacklisting

Discover the `.onion` address of the C&C hidden service using some variant of Biryokov *et al*'s “Trawling for Hidden Services” attack [4]. Essentially, the descriptor can be recognized by its popularity - the volume of requests is an order of magnitude larger than all other hidden services combined. Once the descriptor has been discovered, blacklisting the HS public key from the Hidden Service directory will prevent attempts to reach the blacklisted `.onion` address by clients.

Unfortunately, this approach is both technically and philosophically problematic. Existing Tor client software does not deal gracefully with failed descriptor lookups, leading to a situation where lookup failures increase the circuit load on the network. Furthermore, an adaptive or forward-looking botmaster can defeat identification through volume by multiplexing across

multiple `.onion` addresses. Blacklisting can be defeated using multiple keys or domain generation algorithms – in the hidden service case, generating a sequence of public keys using a fixed-seed pseudorandom sequence. Philosophically, a mechanism to blacklist certain hidden service keys could potentially be abused for censorship.

3.2 Deanonymize the server

Assuming the botnet is prepared for blacklisting, the Tor Project could instead attempt to discover the entry guards of the C&C server, repeatedly changing their availability (e.g. by rotating identity keys), until a colluding entry guard is chosen, eventually learning the IP address of the C&C server. Once the server is deanonymized, traditional anti-malware organizations could be alerted, hopefully taking the botnet offline. From a technical standpoint, coordinating an attack of this scale, including the repeated coercion or identity-rotation of entry guards, while not further disrupting the network, would seem to pose significant organizational and engineering challenges. Depending on the fraction of colluding entry-guard bandwidth, the time-scale of the attack could also be problematic.

3.3 Fingerprint and blacklist clients

A related, less costly, alternative is to look for connections to this hidden service (e.g. by waiting until colluding nodes are chosen as Introduction Points or Rendezvous points) and attempt to “fingerprint” a bot’s “phone home” connection [17, 6, 12, 15]. This fingerprint could then be used to temporarily blacklist bots and/or the C&C server at entry guards. Deploying the solution would involve technical challenges such as determining a stable fingerprint, and upgrading a significant fraction of relays to support fingerprint matching and blacklisting. Furthermore, an adaptive botnet could make this very difficult by employing many of the same traffic analysis countermeasures employed by obfuscated transport mechanisms, for example, “morphing” [18] C&C connections to look like interactions with various popular destinations (e.g. choosing at random an Alexa Top 1000 web site, or even running a relay and choosing at random another stream to emulate).

3.4 General Objections

Beyond the technical difficulties and the possibility for abuse of any of the mechanisms described in this section, another concern is that “attacking” a botnet with hundreds of thousands or even millions of nodes could lead to retaliation. So far the network has withstood the onslaught of new clients created by this bot, but a concerted attack by a botnet of this size could easily bring down the entire Tor network.

4 Throttling

Since the primary concern from the point of view of the other users of Tor is the *rate* at which botnet nodes consume the collective computing resources of the relays, another set of potential solutions is to attempt to throttle or otherwise limit the rate of requests from the botnet. Two key points to recall in evaluating solutions from this class are that (i) in many ways the botnet has more resources available than the set of all regular Tor clients and (ii) neither bots nor the C&C server are constrained to follow the standard Tor algorithms, although the current implementations may do so.

4.1 Can we throttle by cost?

One way to control the rate at which circuit building requests enter the network is by making it costly to send them. Tor could do this by requiring proof of the expenditure of a scarce resource, for example, human attention, processor time, bitcoins, and so on. If the cost to build a circuit or connect to a hidden service can be correctly allocated it could be the case that ordinary users and services can easily afford the cost while the price for a botnet becomes prohibitive. Depending on the resource used, correctly allocating the cost is an important research question; we consider the problem for Proof of Work (CPU-based) and CAPTCHA (human attention-based) systems below.

Besides the cost allocation problem, another technical challenge is ensuring that resources can't be double-spent, so that each resource expenditure in a given time period only authorizes a single circuit or hidden service connection. Several approaches exist, but each would require further investigation:

- Make the unit of pay cover a single circuit extension and have one of the relays extending the circuit issue a challenge back to the client, which then must be answered before the `CREATE` (or `EXTEND`) cell is processed, similar to the scheme described by Barbera *et al.* [2]. This has the unfortunate side effect of adding an extra round-trip time to every circuit-building request. Finding a way to hide this extra round-trip time could make it a viable alternative, for some resources.
- Relay descriptors could include “puzzle specifications” that describe what the challenge will be for a given time period, requiring a method to prevent “precomputing” a batch of payments before the time period; how to solve this problem is an open question.
- Another method would use an extra trusted server that verifies resource expenditures and issues relay- and time period-specific signed tokens, similar to *ripcoins* [16] or the tokens in *BRAIDS* [9]. Using blinded tokens would limit the trust required in the server so that it can't compromise anonymity, and relay-specificity would allow each relay to verify that tokens aren't double-spent. However, this adds an extra signature-verification to the task of onion-skin processing and another server and key that must be maintained.

All of these solutions also require adding extra verification to the onion-skin handshake, incurring the additional risk implied in changes to the core Tor protocol.

4.1.1 Proof of work (proves once more not to work?)

When the resource in question is processor time and challenges are, e.g. hashcash [1] targets, the cost allocation strategy should dictate that the hidden service must pay a cost for each connection, since bots clients and normal hidden service clients will have essentially identical use profiles (from the point of view of relays) and computational resources. On the other hand, the C&C hidden server(s) will collectively initiate many more circuits than any single “normal” hidden server.

The key security challenge when considering an adaptive botmaster’s response to this approach is the “chain-proving” attack (by analogy to chain voting [11]). In this attack, the C&C server solves the first challenge it receives when a bot contacts the hidden service, but then on each additional challenge, the *previous bot* is asked to solve the puzzle in time to allow the *next bot* to connect. In principle the difference in latencies (caused by the need to pass a puzzle to the bot through Tor) could potentially be detected, but an adaptive botmaster could well build shorter circuits, and employ multiple bots in an effort to reduce the time needed to solve a “proof of work” puzzle.

4.1.2 CAPTCHAs

If CAPTCHAs are used to verify expenditure of human attention, the relative cost allocation should change to favor the client: clients of most hidden services will have human users, while hidden servers will not. This raises additional technical problems, such as how CAPTCHAs can be served through Tor without a GUI interface, how a user’s solution can be transferred to the hidden service without violating privacy or allowing overspending, and how to deal with the needs of completely headless services where neither the HS client nor the HS server have a user’s attention to give.

Another technical challenge to deploying CAPTCHAs is how to defeat computationally expensive automated solvers. Most commercially-deployed web CAPTCHAs can be solved with success rates on the order of 1-10% per challenge, and the typical service mitigates this by temporarily blacklisting an IP address after a small number of attempts. With anonymous users, this becomes a more challenging problem to solve; without blacklisting a bot can simply attempt as many CAPTCHAs as necessary to obtain an automated solution.

4.1.3 Network Contributions

As in several Tor incentive schemes [14, 9, 13, 10] we could require proof of contributions to the network to receive high performance at volume, while providing a level of performance to non-relay users that would be acceptable for normal clients but would significantly impede the activity of a hidden C&C server. However, this would have significant impact on the security of regular Tor users since it might simply incentivize a botmaster to run many compromised relays.

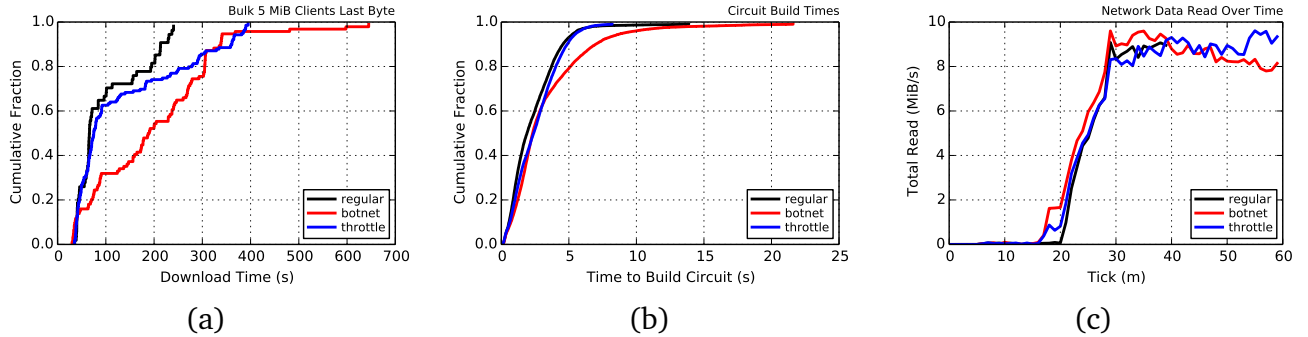


Figure 4: Results of guard throttling: 20 relays, 200 clients, 500 bots. (a) 5MiB download times, (b) Circuit build times, (c) Total bytes read

4.2 Can we throttle at the entry guard?

A more direct approach would be to simply have guard nodes rate-limit the number of `EXTEND` cells they will process on a given connection. If the entry guard won't process the `EXTEND` cell needed to build a circuit, the hidden server's OP can't flood the network with onion-skins. Notice that this measure won't prevent *bots* from flooding the network with circuit requests; it simply makes the network ineffective from the botmaster's standpoint and thus, hopefully, encourages botmasters to find some other C&C channel that causes less stress on the Tor network.

Effective circuit throttling at the guard node faces a number of challenges, however. Biryukov *et al* [3] found that the most popular hidden services see over 1000 requests per hour; if we assume that these hidden services won't modify Tor's default behavior, then guard nodes need to allow each client to extend over 300 circuits per hour; but since there are currently over 1200 relays with the guard flag, a single C&C server run by an adaptive botmaster could build 360 000 circuits per hour at this rate. We could decrease the cap and try to make it easier for busy hidden servers to increase their guard count, but this significantly increases the chance that a hidden server chooses a compromised guard and can be deanonymized.

One possibility would be to use *assigned guards*. In this approach, ordinary clients would pick guards as usual, and guards would enforce a low rate-limit r_{default} on circuit extensions, for example 30 circuits per hour.¹ OPs that need to build circuits at a higher rate r_{server} – say, 2000 per hour – could follow a cryptographic protocol that would result in a verifiable token that assigns a deterministic, but unpredictable, guard node for the OP when running on a given IP address. These OPs could then show this token to the assigned guard and receive a level of service sufficient for a busy hidden server, but not for a flood of circuit extensions. An example of this type of protocol appears as Protocol 3 (section 3.3) in the *BRAIDS* design by Jansen *et al.* [9]. The rates r_{default} and r_{server} could appear in the network consensus, to allow adjustments for the volume of traffic in the network. Figure 4 shows the result of simulating this strategy with $r_{\text{default}} = 10$ and $r_{\text{server}} = 2000$ using the *shadow* simulator [8]; despite nearly identical bandwidth usage, the throttled simulation has performance characteristics similar to the simulation with no botnet.

¹Naturally, finding the right number to use for this default rate is also an interesting research challenge: a very low rate-limit could prevent bots from flooding the network but might also disrupt legitimate hidden service clients

An additional technical challenge associated with guard throttling is the need to enforce the use of entry guards when building circuits. If the C&C server joins the network as a relay, `CREATE` cells coming from the hidden service would be indistinguishable from `CREATE` cells coming from other circuits running through the relay, effectively circumventing the rate limit. In principle this could be detected by a distributed monitoring protocol, but designing secure protocols of this type that avoid adversarial manipulation has proven to be a difficult challenge.

4.3 Can we throttle by Tor version?

Another strategy that may be useful against a nonadaptive botmaster is to throttle (perhaps temporarily) by Tor version. In this strategy, nodes running old versions of Tor could be either severely throttled or completely ignored by nodes with the latest software. Since relays are typically kept up-to-date and most clients access Tor using the Tor Browser Bundle (which can inform clients on start-up when new versions are recommended or, in the future, might incorporate auto-update functionality) the impact on typical users would be minimal, but a bot that packages the Tor software may not be able to respond to updates. The technical challenges here involve designing reliable methods to generate new software versions, (sufficiently) secure protocols to check that a client is running the software, and analyzing the use cases that might affect clients' ability to update frequently and the impact of throttling in these cases.

5 Client-side circuit-building adjustments

Although an adaptive botnet could always modify the Tor client code, the regular user base still represents a large fraction of the load on the network, so another set of potential solutions would be to focus on modifying the behavior of ordinary clients to reduce the circuit-building load.

5.1 Can we reuse failed partial circuits?

Part of the problem caused by the heavy circuit-building load is that when a circuit times out, the entire circuit is destroyed. This means that for every failed `CREATE`, at least three new `CREATE` cells will be added to the network's load. If we model the entire Tor network as having probability p of having a `CREATE` cell timeout, then the expected number of `CREATE` cells needed to successfully build a circuit will be the X_0 satisfying the linear system:

$$\begin{aligned} X_0 &= pX_0 + (1-p)X_1 && +1 \\ X_1 &= pX_0 && + (1-p)X_2 && +1 \\ X_2 &= pX_0 && && +1, \end{aligned}$$

where X_i is the expected number of cells to complete a partial circuit with i hops. This gives us $X_0 = \frac{p^2 - 3p + 3}{(1-p)^3}$.

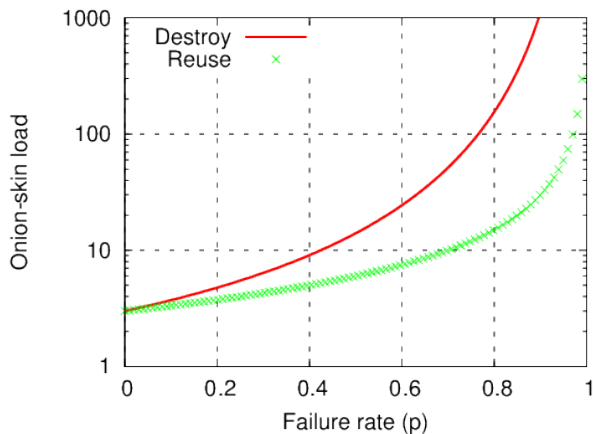


Figure 5: Expected onion-skin load per circuit created, for failure rate p

Conceptually, we can reduce this load by re-using a partially-built circuit, e.g. when a timeout occurs, we truncate the circuit and attempt to extend from the current endpoint. In this case, the expected number of CREATE cells needed to build a circuit will be simply $X'_0 = \frac{3}{1-p}$. Figure 5 shows plots of both functions. We can see that for high enough failure rates, this change causes a substantial reduction in load for the network. Figure 3 shows typical failure rates for a stable (TAP) and release candidate (ntor) roughly one month after the beginning of the botnet event; we can see that at the observed failure rates ranging from 10%-25%, reusing partial circuits would reduce the load on the network by 10-30%.

Of course, this model ignores the fact that failure probabilities are neither static nor uniform across the entire Tor network, and the fact that many nodes use “create fast” cells to exchange a first-hop key without using Diffie-Hellman key exchange. Reducing the load introduced by failures will also reduce the rate of circuit failures overall, but since CPU capacities vary widely across the Tor network (and load balancing is by the essentially uncorrelated bandwidth of nodes) the size of the actual effect due to this change is difficult to predict. Further evaluation will be needed. Additionally, this change would also somewhat increase the power of selective denial of service attacks [5], although such attacks typically only become noticeably effective in situations where we would already consider Tor to be compromised.

5.2 Can we build circuits less often?

One reason for the relatively mild impact of the Mevade incident is that, while hidden services initially build more circuits than ordinary downloads, Tor is configured to aggressively avoid repeating this process through the use of “circuit dirtiness” defaults. When a Tor client builds an ordinary circuit, it is marked as “clean” until it is first used for traffic, at which point it becomes “dirty.” After a configurable amount of time (by default, 10 minutes) a “dirty” circuit cannot be used for new traffic and will be closed after existing streams close. In contrast, the “dirty timer” for a hidden service circuit restarts every time it is used, so that as long as a hidden service is visited once every 10 minutes, there is no need to build a new circuit.

Thus, another set of technical approaches to dealing with circuit stress would be to investigate ways to reduce the number of circuits that Tor clients build:

- Network consensus documents could include a “recommended max dirtiness” parameter that would adjust the lifetime of ordinary circuits, and a separate, potentially longer “max dirtiness” parameter for hidden services. The technical challenges here are: first, balancing the tradeoffs between decreased anonymity for individual users (the longer a circuit is used, the more chance that some user activities will use the same circuit and be linked by an adversary) versus allowing more users on the network; and second, reliably setting this parameter in an automated fashion.
- Similarly, Tor clients currently build extra circuits for internal purposes like descriptor fetching and timeout testing, that could be disabled by an appropriately set consensus parameter.
- Finally, while “ordinary” circuits are built preemptively for use when needed, some tasks such as hidden service connections, connecting to “rare” exit ports, and others currently build “on-demand” circuits. Finding ways to avoid constructing new circuits for these tasks, and analyzing the impact on anonymity and security, could reduce the amount of circuit building and also the perceived performance impact of a botnet.

6 Can we isolate Hidden Service circuits?

Another approach to protect the regular users of the Tor network from resource depletion by a hidden-service botnet would be to isolate hidden service onion-skin processing from ordinary processing. By introducing a mechanism that allows relays to recognize that an `EXTEND` or `CREATE` cell is likely to carry hidden service traffic, we could provide a means to protect the rest of the system from the effects of this traffic, by scheduling priority or simple isolation.

An example of how this might work in practice is to introduce new `NOHS-EXTEND`/`NOHS-CREATE` cell types with the rule that a circuit that is created with an `NOHS-CREATE` cell will silently drop a normal `EXTEND` cell, or any of the cell types associated with hidden services. If relays also silently drop `NOHS-EXTEND` cells on circuits created with ordinary `CREATE` cells, then `NOHS-CREATE` circuits are guaranteed not to carry hidden service traffic. Updated clients would then create all circuits with `NOHS-CREATE` unless connecting to a hidden service. When a sufficient number of clients and relays update their Tor version, a consensus flag could be used to signal relays to begin isolating processing of ordinary `CREATE` cells. For example, these cells might only be processed in the last 20ms of each 100ms period, leaving 80% of processing capacity available for regular traffic. The flag could be triggered when hidden service circuits exceed a significant fraction of all circuits in the network.²

This solution protects the network and typical users from a massive botnet hidden service, but would, unfortunately, intensify the effect on users of legitimate hidden services in time

²Detecting this condition in a privacy-preserving manner represents another technical challenge requiring further research.

periods when an attack was detected. As with guard throttling, the intended effect would thus be to encourage botmasters to develop C&C channels that do not stress the Tor hidden service ecosystem, while providing stronger protection against botnet clients flooding the network.

One privacy concern related to this approach is that as the network upgrades to versions of Tor supporting NOHS-CREATE, identification of hidden-service traffic approaches deterministic certainty. By contrast, current hidden service circuits follow traffic patterns that allow them to be identified with high statistical confidence [4] only. Because (excluding botnet traffic) the base rates of hidden service traffic compared to all other traffic are low, this will also decrease the privacy of hidden service users. One potential mitigation mechanism would be to have clients only use NOHS-CREATE when the consensus flag for hidden service isolation is activated, which would indicate that hidden service clients would already have a large anonymity set.

7 Conclusions

As of this writing, further evaluation is needed before recommending any approach. In the medium-term, throttling Tor versions prior to 0.2.4.17-rc may provide the most immediate relief. In the long term, partial circuit reuse seems safe, pending complete analysis of the impact on anonymity; further evaluation is needed to determine the effectiveness of other measures.

8 Acknowledgements

Roger Dingledine, Ian Goldberg, Rob Jansen, Mike Perry, and Max Schuchard all contributed significantly to the ideas, analysis, and discussion in this tech report.

References

- [1] Adam Back et al. Hashcash-a denial of service counter-measure, 2002.
- [2] Marco Valerio Barbera, Vasileios P. Kemerlis, Vasilis Pappas, and Angelos Keromytis. CellFlood: Attacking Tor onion routers on the cheap. In *Proceedings of ESORICS 2013*, September 2013.
- [3] Alex Biryukov, Ivan Pustogarov, and Ralf-Philipp Weinmann. Content and popularity analysis of tor hidden services. *arXiv [cs.CR]*, August 2013.
- [4] Alex Biryukov, Ivan Pustogarov, and Ralf-Philipp Weinmann. Trawling for tor hidden services: Detection, measurement, deanonymization. In *Proceedings of the 2013 IEEE Symposium on Security and Privacy*, May 2013.
- [5] Nikita Borisov, George Danezis, Prateek Mittal, and Parisa Tabriz. Denial of service or denial of security? How attacks on reliability can compromise anonymity. In *Proceedings of CCS 2007*, October 2007.

- [6] Xiang Cai, Xincheng Zhang, Brijesh Joshi, and Rob Johnson. Touching from a distance: Website fingerprinting attacks and defenses. In *Proceedings of the 19th ACM conference on Computer and Communications Security (CCS 2012)*, October 2012.
- [7] Ian Goldberg, Douglas Stebila, and Berkant Ustaoglu. Anonymity and one-way authentication in key exchange protocols. *Designs, Codes and Cryptography*, 67(2):245–269, May 2013.
- [8] Rob Jansen and Nicholas Hopper. Shadow: Running Tor in a Box for Accurate and Efficient Experimentation. In *Proceedings of the Network and Distributed System Security Symposium - NDSS'12*. Internet Society, February 2012.
- [9] Rob Jansen, Nicholas Hopper, and Yongdae Kim. Recruiting new Tor relays with BRAIDS. In Angelos D. Keromytis and Vitaly Shmatikov, editors, *Proceedings of the 2010 ACM Conference on Computer and Communications Security (CCS 2010)*. ACM, October 2010.
- [10] Rob Jansen, Aaron Johnson, and Paul Syverson. LIRA: Lightweight Incentivized Routing for Anonymity. In *Proceedings of the Network and Distributed System Security Symposium - NDSS'13*. Internet Society, February 2013.
- [11] Douglas W Jones. Chain voting. In *Workshop on Developing an Analysis of Threats to Voting Systems, National Institute of Standards and Technology*, 2005.
- [12] Prateek Mittal, Ahmed Khurshid, Joshua Juen, Matthew Caesar, and Nikita 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.
- [13] W. Brad Moore, Chris Wacek, and Micah Sherr. Exploring the potential benefits of expanded rate limiting in tor: Slow and steady wins the race with tortoise. In *Proceedings of 2011 Annual Computer Security Applications Conference (ACSAC'11), Orlando, FL, USA*, December 2011.
- [14] Tsuen-Wan “Johnny” Ngan, Roger Dingledine, and Dan S. Wallach. Building Incentives into Tor. In Radu Sion, editor, *Proceedings of Financial Cryptography (FC '10)*, January 2010.
- [15] Andriy Panchenko, Lukas Niessen, Andreas Zinnen, and Thomas Engel. Website fingerprinting in onion routing based anonymization networks. In *Proceedings of the Workshop on Privacy in the Electronic Society (WPES 2011)*. ACM, October 2011.
- [16] Michael K Reiter, XiaoFeng Wang, and Matthew Wright. Building reliable mix networks with fair exchange. In *Applied Cryptography and Network Security*, pages 378–392. Springer, 2005.
- [17] Tao Wang and Ian Goldberg. Improved website fingerprinting on tor. In *Proceedings of the Workshop on Privacy in the Electronic Society (WPES 2013)*. ACM, November 2013.

- [18] Charles Wright, Scott Coull, and Fabian Monrose. Traffic morphing: An efficient defense against statistical traffic analysis. In *Proceedings of the Network and Distributed Security Symposium - NDSS '09*. IEEE, February 2009.