Poster: Mitigating OnionBots

Amirali Sanatinia, Guevara Noubir
Northeastern University, Boston, USA
{amirali,noubir}@ccs.neu.edu

***Abstract*—Over the last decade botnets have become a serious security threat. They have evaded mitigation and take overs by adopting an increasing sophisticated strategies. At the same time the rise and success of privacy infrastructures, has opened new possibilities of abuse by malicious users. Tor is a prominent example of such infrastructure, which allows users to hide their activities and location from government agencies and corporations. Furthermore, it also offers anonymity for servers through hidden services. Recent statistics about hidden services clearly indicates changes in their popularity and use. For instance, the number of hidden services has abruptly doubled in the last year (Figure 1), which clearly indicates the presence of some coordinated massive use. We envision a next generation of cryptographic, resilient, stealthy botnets, OnionBots, that subverts privacy infrastructures for cyber attacks, by completely decoupling their operation from the infected host IP address. Furthermore, they rely on disturbed self-healing network formation that is simple to implement, yet achieves a low diameter and a low degree, and is robust to partitioning attacks. As a result, the current detection and mitigation strategies would be inadequate against them. We devise a mitigation mechanism that uses OnionBots' very own capabilities to neutralize them. In light of the potential of such botnets, we believe that the research community should proactively develop detection and mitigation methods to thwart OnionBots, potentially making adjustments to privacy infrastructure. The preliminary results of this work have been presented [1]**

I. OnionBot: A Cryptographic P2P Botnet

OnionBots form a peer-to-peer, self-healing network that maintains a low degree and a low diameter with other bots to relay messages. The already existing peer-to-peer networks are generic in terms of their operations. Therefore, their design and resiliency is based on different assumptions and requirements. We propose a Dynamic Distributed Self Repairing (DDSR) graph construction that is simple, stealthy and resilient.

The DDSR construct is inspired by the knowledge of Neighbors-of-Neighbor. Where each node has the knowledge about its immediate neighbors. Consider graph G with n nodes (V), where each node $u_i \in V$, $0 \leq i < n$, is connected to a set of nodes. The neighbors of u_i, are denoted as $N(u_i)$. Furthermore, u_i has the knowledge of nodes that are connected to $N(u_i)$. Meaning that each node also knows the identity of its neighbor's neighbors. In the context of our work the identity is the `.onion` address. Having this information enables the botnet to repair its graph formation and maintain its connectivity in a distributed setting. When a node u_i is deleted, each pair of its neighbors u_j, u_k will form an edge (u_j, u_k) if $(u_j, u_k) \notin E$, where E is the set of existing edges. The aforementioned basic DDSR graph does not deal with the growth in the connectivity degree of each node, denoted by $d(u)$; after multiple deletions the degree of some nodes can increase significantly. Such increase is not desirable for the resiliency and the stealthy operation of the botnet. Therefore we introduce the concept of pruning in node deletion to address this challenge. When a node is removed, each neighboring node of the deleted node (u_i), deletes the highest degree node from its peer list. If there is more than one such candidate, it randomly selects one among those for deletion, until its degree is in the desired range. Figure 2 depicts the impact of pruning on the degree and closeness centrality of the botnet's connectivity network.

In the proposed OnionBot, nodes forget the `.onion`

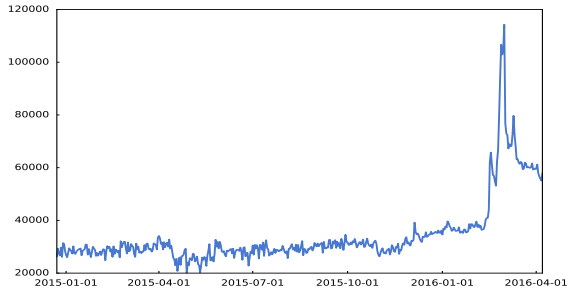

Fig. 1: The increase in the number of Hidden Services, based on the statistics collected by Tor.

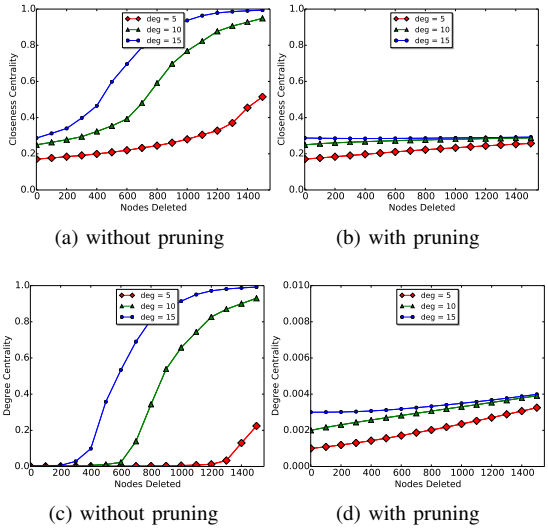

(a) without pruning (b) with pruning

(c) without pruning (d) with pruning

Fig. 2: The average closeness centrality, and degree centrality of nodes in a *k-regular* graph, ($k = 5, 10, 15$) with 5000 nodes after 30% node deletions, with and without pruning.

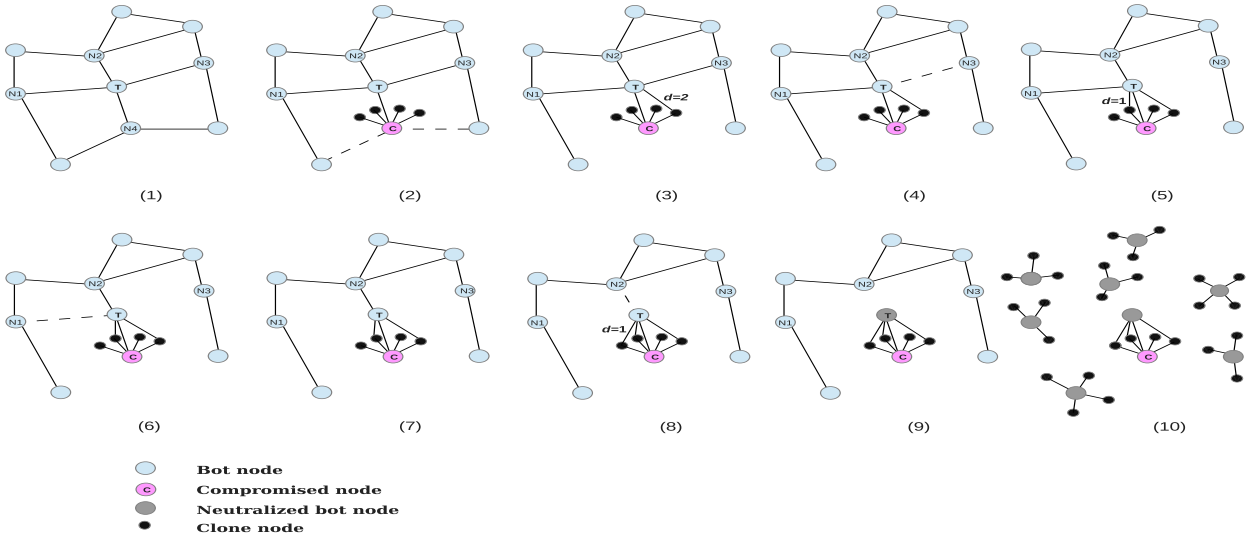

Fig. 3: SOAP: node T is under attack by the compromised node C and its clones. In each step one of the clones initiates the peering process with the T, until it is contained. After several iterations, the network is partitioned, and the botnet is neutralized.

address of the pruned nodes. Additionally, to avoid discovery, mapping and further blocking, each bot can periodically change his `.onion` address and announce the new address to his current peer list. The new `.onion` address is generated based on a secret key and time. This periodic change is possible because of the decoupling between IP address and the bots, which is provided by Tor.

II. Mitigating OnionBots

Many of the current detection and mitigation mechanisms are IP-based, and rely on the network traffic patterns or DNS queries to distinguish legitimate traffic from malicious traffic. However, current solutions do not work with OnionBots, since the Tor traffic is encrypted, non IP-based, and there are no conventional DNS queries. Malicious traffic detection mechanisms in Tor [2] are the first step in the mitigation. However, we need to adapt our detection and mitigation methods to address the evasion mechanism of OnionBots. We devised a mitigation mechanism that uses OnionBots' very own capabilities (e.g., the decoupling of IP address and the host) against them. Figure 3 depicts the soaping attack in different steps. Node T is the target of the soaping attack, nodes N_i, are its neighboring bot nodes, and nodes C are the adversary, and his clones, which are represented with small black circles. In step 1, the botnet is operating normally, and none of T's neighbors are compromised. In step 2, one of its peers, N_4, is compromised. Then, N_4 (now depicted as C), makes a set of clones (the small black circles). In step 3, a subset of C's clones, start the peering process with T, and declare their degree to be a small random number, which changes to avoid detection (e.g., d=2). Doing so increases the chances of being accepted as a new peer, and replacing an existing peer of T. In step 4, T forgets about one of its neighbors with the highest degree, N_3, and peers with one the clones. The clones repeat this process until T has no more benign neighbors (steps 5-8). As a result, T is surrounded by clones and is contained (step 9).

III. Conclusion and Future Work

Privacy infrastructures such as Tor had a tremendous impact on society, protecting users anonymity and rights to access information in the face of censorship. It also opened the door to abuse and illegal activities, such as ransomware, and a marketplace for drugs and contraband [3], [4], [5]. In this work we envisioned OnionBots, and investigated the potential of subverting privacy infrastructures for cyber attacks. We presented the design of a robust and stealthy botnet that lives symbiotically within these infrastructures to evade detection, measurement, scale estimation and observation. Additionally, OnionBots rely on a resilient self-healing network formation that is simple to implement, yet it has desirable features such as low diameter and low degree. We developed soaping, a novel mitigation attack that neutralizes the OnionBots. However, there are still many challenges that need to be preemptively addressed by the security community, such as the byzantine behavior of OnionBots. We hope that this work ignites new ideas to proactively design mitigations against the new generations of crypto-based botnets.

References

[1] A. Sanatinia and G. Noubir, "Onionbots: Subverting privacy infrastructure for cyber attacks," in *The Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, 2015.

[2] Z. Ling, J. Luo, K. Wu, W. Yu, and X. Fu, "Torward: Discovery, blocking, and traceback of malicious traffic over tor," in *IEEE Transactions on Information Forensics and Security*, 2015.

[3] A. Biryukov, I. Pustogarov, and R.-P. Weinmann, "Content and popularity analysis of tor hidden services," *arXiv preprint arXiv:1308.6768*, 2013.

[4] N. Christin, "Traveling the silk road: A measurement analysis of a large anonymous online marketplace," in *Proceedings of the International Conference on World Wide Web*, 2013.

[5] M. Thomas and A. Mohaisen, "Measuring the leakage of onion at the root: A measurement of tor's. onion pseudo-tld in the global domain name system," in *Proceedings of the Workshop on Privacy in the Electronic Society*, 2014.