

Towards Incentivizing ISPs To Mitigate Botnets

Qasim Lone, Giovane C. M. Moura, and Michel Van Eeten

Economics of Cybersecurity Group
Faculty of Technology, Policy and Management
Delft University of Technology
Delft, The Netherlands
{q.b.lone, g.c.moreiramoura, m.j.g.vaneeten}@tudelft.nl

Abstract. ISPs form a centralized point to control botnet infections. However, they do not have enough incentives to invest in mitigation of botnets. In this paper, we propose an approach based on comparative metrics to incentivize ISPs to mitigate botnets. This research is still in its initial phase and will contribute to a Ph.D. thesis after four years.

1 Introduction

A botnet is network of compromised machines, controlled by a botmaster, which is used to carry out attacks [1]. Some of the common attacks botnets in which partake include spam, phishing, distributed denial-of-service (DDoS), credential theft, and click fraud. These attacks incur significant financial losses; for instance, it is estimated that spam causes losses of US\$ 20 billion [2] yearly, only in the United States.

Previous studies have shown that most of the malicious hosts are concentrated in a small number of Internet Service Providers (ISP). Van Eeten *et al.* [3] found that 50 ISPs account for around half of all spamming IP addresses worldwide. In another study, Moura [4] found that 20 Autonomous Systems (AS), out of 42,201, were responsible for 50% of all spamming IP addresses. Similar trends were also found in [5–7].

Taking these observations into account, ISPs would then form a centralized control point and, this concentration in a small number of ISPs would make it easier for them to mitigate botnets. There are a number of steps ISPs can take to reduce infected machines in their networks, including quarantining, providing links to antivirus software, and notifying customers about infections in their computers. However, there is little evidence that ISPs are taking many concrete actions. For example, in [3], Van Eeten *et al.* found that only 10% of infected customers were notified by their Dutch ISPs.

2 Research Problem

Such a low rate of ISPs actions can be due to several reasons, such as, content filtering may violate user’s privacy of the customers according to the legislation of

some of the countries [8]. However, one of the prime reasons is *lack of incentives* for ISPs to invest in mitigation of botnets [3]: if the market for Internet access is characterized by price competition, ISPs would be strongly discouraged to invest more in botnet mitigation than their competitors, *i.e.*, they would be disincentivised to contact and quarantine more infected customers than their competitors.

We do not really know to what extent individual providers actually and effectively fight botnets. This information asymmetry impedes the functioning of markets and may even result in market failure. It weakens the incentives to invest in mitigation, because users and other stakeholders cannot tell good performing providers from bad ones. In order to improve incentive structure of ISPs, analytical models based on game theory can be utilized to explore and evaluate relative security performance of ISPs. Various studies [9–11] based on empirical data suggest that metrics can be an effective way to measure cyber security performance. In this research our focus is to develop comparative metrics to evaluate ISPs efforts to mitigate botnets. The proposed research problem leads to the following research questions.

1. What kind of network measurement data is required to statistically account for botnet population in the networks of ISPs ?
2. How to turn the measurements into comparative relative metrics for ISPs performance in botnet mitigation ?
3. How can these metrics contribute to evaluate and incentivizing botnet mitigation by ISPs ?

3 Approach

Research question 1, from previous section, focuses on types of measurement data we can use to statistically estimate botnet population in ISP. We will obtain access to data which is collected by various collaborators of the project. There are two types of measurement data for botnets: data which is collected *outside* or *inside* botnets.

The first type of data is obtained by observing direct attacks from infected machines, for example, machines taking part in spam or participating in DDoS attack. It is collected using various approaches, including, honeypots, spam traps, intrusion detection system, sinkholes. This type of data helps capture wide range of botnets. However, captured data might have false positives and false negatives, due to limitation in detection capabilities of these systems. We already have access to DShiled and Spam trap from this category of datasets.

In the second type of measurement, data is obtained by taking over command and control centers (C&C) of botnets. The advantage of this approach is that we have accurate data but the downside is that measurements are only limited to a single botnet, such data is not representative of the total population. We have access to ZeroAccess, Conficker and Zeus botnet data sets, which were collected by taking over the respective botnet.




Fig. 1. Relationship between ISPs, botnet and home users

Hence, a detailed study will be performed on available data sources keeping in view advantages and limitations of these sources, so that we can statistically estimate the presence of botnet in the networks of ISPs.

The goal of research question 2 is to extract botnet metrics from network measurement data collected in research question 1. The developed metrics need to be consistent, normalized (for *e.g.* by number of customers per ISP), easily understandable by customers, and validated [12], *i.e.* to prove that they in fact capture the behavior they are supposed to capture.

To create such metrics, there are many challenges to overcome for example, consider measuring botnet presence in the networks of ISPs. To illustrate this, consider Figure 1. In this Figure, we see that a subscriber of ISP A is using a home router (with DHCP and NAT) to connect three laptops to the Internet. Laptop 1 has two malware instances running, while laptop 3 has one and laptop 2 has none. There are three bots which are operating from two different laptops and are hiding behind a single public routable IP address.

This exemplifies how complex it is to count botnet presence in ISP networks, and how IP addresses do not correspond to the number of botneted computers [13]. To show how the number of IP addresses may significantly differ from the actual number of hosts, we have analyzed the variation on the number of IP addresses of 1,064 RIPE Atlas probes [14]¹, over a 1 year period. As can be seen, there is a significant variation among the probes and, on average, each probe had 24 IP addresses (1:24). In another study, Stone-Gross *et al.* [15] hijacked the Torpig botnet for 10 days, and found that on average, each bot had (1:7) IP addresses, varying significantly according to ISP and country.

Some of the major challenges in developing these metrics include, bot counting, partial view, false positives/ negatives, and relative potency of botnets. Therefore, in research question 2, we will carry out a detailed literature review and various types of networks measurements to develop these metrics.

¹ Atlas probes are small hardware devices distributed all over the world and used to measure Internet connectivity and reachability, developed and maintained by the *Réseaux IP Européens* Network Coordination Centre (RIPE NCC).

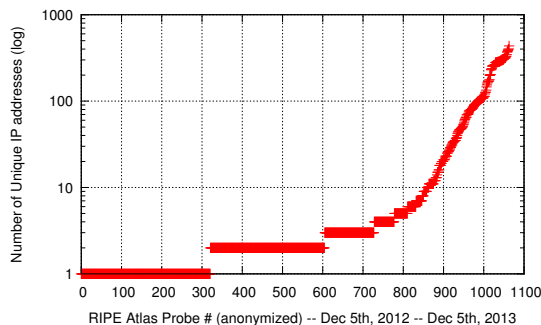


Fig. 2. Number of unique IP addresses per RIPE probe

In research question 3, we will investigate how the developed comparative metrics can be utilized to maximize economic incentives of ISPs. Tang *et al.* [9] found a total of 16% reduction in spam after spam rankings were published on a website. Similarly, there are yearly/quarterly reports published by various security companies [16–19] on security measurements. However, these studies are usually too limited to certain type of infection, do not rank ISPs based on performance or are not transparent on how these rankings were developed. Hence our goal in research question 3, will to not only to publish these rankings frequently, but to also make them accessible and understandable for majority of Internet consumers.

4 Final Considerations

As discussed in Section 1, malicious hosts are concentrated in small number of ISPs, which makes it easier to mitigate botnets. However, ISPs have limited incentives to invest in botnet mitigation. The effectiveness of mitigation measures cannot be established without accurate and reliable reputation metrics [9]. Without such metrics, there is only anecdotal evidence that cannot be reliably interpreted. Additionally, we can also evaluate effectiveness of mitigation strategies relative to each other. This Ph.D. research aims at designing comparative metrics to incentivize ISPs to take countermeasures for botnet mitigation. The goals of this work should be achieved within a period of four years, as part of Ph.D. thesis.

Acknowledgments: This work was partly funded by the Advanced Cyber Defense Centre (*ACDC*) project (#325188), which is supported by the European Commission under its ICT Policy Support Programme as part of the Competitiveness and Innovation Framework (CIP-PSP).

References

1. G. Gu, R. Perdisci, J. Zhang, W. Lee *et al.*, “Botminer: Clustering analysis of network traffic for protocol-and structure-independent botnet detection.” in *USENIX Security Symposium*, 2008, pp. 139–154.
2. J. M. Rao and D. H. Reiley, “The economics of spam,” *The Journal of Economic Perspectives*, vol. 26, no. 3, pp. 87–110, 2012.
3. M. van Eeten, J. M. Bauer, H. Asghari, S. Tabatabaie, and D. Rand, “The Role of Internet Service Providers in Botnet Mitigation: An Empirical Analysis Based on Spam Data,” in *WEIS 2010: Ninth Workshop on the Economics of Information Security*, 2010.
4. G. C. M. Moura, “Internet Bad Neighborhoods,” Ph.D. dissertation, University of Twente, Enschede, The Netherlands, March 2013. [Online]. Available: <http://dx.doi.org/10.3990/1.9789036534604>
5. M. P. Collins, T. J. Shimeall, S. Faber, J. Janies, R. Weaver, M. De Shon, and J. Kadane, “Using Uncleanliness to Predict Future Botnet Addresses,” in *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, ser. IMC ’07. New York, NY, USA: ACM, 2007, pp. 93–104.
6. A. Ramachandran and N. Feamster, “Understanding the Network-level Behavior of Spammers,” in *Proceedings of the 2006 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, ser. SIGCOMM ’06. New York, NY, USA: ACM, 2006, pp. 291–302.
7. W. van Wanrooij and A. Pras, “Filtering Spam from Bad Neighborhoods,” *International Journal of Network Management*, vol. 20, no. 6, pp. 433–444, November 2010.
8. G. Huston, “Opinion: The ISP – The Uncommon Carrier,” *The Internet Protocol Journal*, vol. 5, no. 3, pp. 23–27, 2002. [Online]. Available: http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_5-3/uncommon_carrier.html
9. Q. Tang, L. Linden, J. Quarterman, and A. Whinston, “Improving internet security through social information and social comparison: A field quasi-experiment,” *WEIS 2013: Twelfth Workshop on the Economics of Information Security*, 2013.
10. W. H. Baker, L. P. Rees, and P. S. Tippet, “Necessary measures: Metric-driven information security risk assessment and decision making,” *Commun. ACM*, vol. 50, no. 10, pp. 101–106, Oct. 2007. [Online]. Available: <http://doi.acm.org/10.1145/1290958.1290969>
11. L. P. Rees, J. K. Deane, T. R. Rakes, and W. H. Baker, “Decision support for cybersecurity risk planning,” *Decision Support Systems*, vol. 51, no. 3, pp. 493 – 505, 2011. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0167923611000728>
12. C. Kaner, S. Member, and W. P. Bond, “Software engineering metrics: What do they measure and how do we know?” in *METRICS 2004, IEEE CS. Press*, 2004.
13. M. Fabian and M. Terzis, “My botnet is bigger than yours (maybe, better than yours): Why size estimates remain challenging,” in *Proceedings of the 1st USENIX Workshop on Hot Topics in Understanding Botnets*, Cambridge, USA, 2007.
14. RIPE_NCC, “RIPE Atlas.” [Online]. Available: <https://atlas.ripe.net>
15. B. Stone-Gross, M. Cova, L. Cavallaro, B. Gilbert, M. Szydlowski, R. Kemmerer, C. Kruegel, and G. Vigna, “Your botnet is my botnet: analysis of a botnet takeover,” in *Proceedings of the 16th ACM conference on Computer and communications security*. ACM, 2009, pp. 635–647.

16. Microsoft, "Microsoft Security Intelligence Report ." [Online]. Available: <http://www.microsoft.com/security/sir/default.aspx>
17. McAfee, "McAfee Threats Report." [Online]. Available: <http://www.mcafee.com/us/resources/reports/rp-quarterly-threat-q1-2013.pdf>
18. TrendMicro, "Security Research and Threat Analysis." [Online]. Available: <http://www.trendmicro.com/us/security-intelligence/research-and-analysis/index.html>
19. European Union Agency for Network and Information Security, "Cybersecurity cooperation - Defending the digital frontline ." [Online]. Available: http://www.enisa.europa.eu/activities/Resilience-and-CIIP/critical-applications/botnets/botnets-measurement-detection-disinfection-and-defence/at_download/fullReport