

Characterizing Botnets-as-a-Service

Wentao Chang¹ An Wang¹ Aziz Mohaisen² Songqing Chen¹

¹Department of Computer Science, George Mason University
{wchang7, awang10, sqchen}@gmu.edu

²VeriSign Labs
amohaisen@verisign.com

Categories and Subject Descriptors

C.2.0 [Computer Communication Networks]: General—*data communications, security and protection.*

General Terms

Measurement, Security

Keywords

Botnet, measurement, collaborations

1. INTRODUCTION

Botnets are one of today's most challenging cybersecurity threats, and promise to remain a serious threat for many years to come. Bots today are not limited to sophisticated machines, such as servers and personal computers: recent DDoS attacks were reportedly utilizing fridges [3], and other massive scanning activities were using embedded devices, including IP monitoring cameras and security doors [5].

Driven by economical profit, botnets are arising in what has been coined as "Botnet-as-a-Service (BaaS)" [2]. Many of today's botnets are designed and developed to be loaned easily to third parties. Reportedly botnet controllers can make a large amount of money by loaning the service in the mature underground market [4]. Understanding such phenomena through analysis has been the goal of the research community for a while to develop effective defense mechanisms and to guide disinfections.

As the arms race between the malware developers and defenders is endless, it is essential to continuously track and understand the latest strategies of attackers in manipulating botnets for attacks. A timely understanding can provide important insights to guide the building of effective defenses. Therefore, we set to investigate modern botnets from the service perspective, focusing on its elasticity and stability.

Data analyzed in our study is obtained from the monitoring and attribution unit in a private security company that is located in the United States, with partnerships of traffic sharing across the globe. Malware samples used in launching various attacks are reverse engineered and labeled to a known malware family using best practices. A honeypot is then created to emulate the operation of the reverse-engineered malware sample that belongs to a given botnet and to enumerate all bots across the globe participating in that particular botnet. Traces of traffic associated with various botnets are then collected at various anchor points on the Internet, via the

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).

SIGCOMM'14, August 17–22, 2014, Chicago, IL, USA.

ACM 978-1-4503-2836-4/14/08.

<http://dx.doi.org/10.1145/2619239.2631464>.

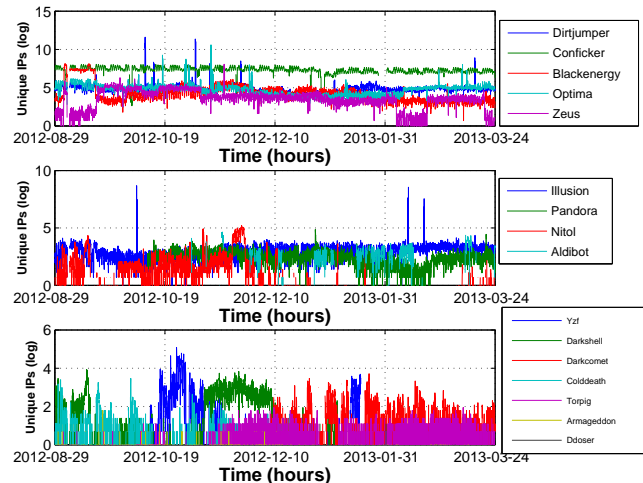


Figure 1: Botnet stability of three groups (based on botnet size)

cooperation of more than 300 ISPs, and analyzed to attribute and characterize attacks. The monitors of the company track temporal activities of 23 different known botnet families in the wild, and generate hourly log dumps from 08/29/2012 to 03/24/2013, a total of 207 days.

Our preliminary study reveals several interesting new trends of botnet management:

- Large botnets often maintain a dynamic stability: keep a stable number of live bots but also keep rotating individual bots.
- From the perspective of attack magnitude, large botnets are more elastic than small ones.
- There is a clear rising trend for botnets to collaborate on campaigns, concurrently or in turn.

2. BaaS CHARACTERIZATION

We briefly present some preliminary results of our BaaS analysis from the following three perspectives: stability, elasticity, and collaboration.

2.1 Botnet stability

Similar to the metrics used in [1], we define the botnet size as the total number of unique IPs that were once recruited in their lifetime by the specific botnet. We find that the botnet size of different families varies significantly. Based on the botnet size, we can classify them into 3 different groups from top to down: large, medium and small (Figure 1).

Figure 1 shows the simultaneous live bots along time for these three groups. Note the *y*-axis is in log scale. To some extent, the simultaneous live bots distribution along time can indicate the stability of the corresponding botnet family. As shown in the figure,




Figure 2: Botnet Elasticity

besides the difference on the absolute numbers, the results indicate that the botnets in *large* group have maintained a relatively stable army of active bots with a few spikes, while there are more fluctuations in botnets in the *medium* and *small* groups. While this is kind of expected, it also raises a question on whether such stability in large botnets is due to the same sets of bots being active or more sophisticated strategies (e.g., bot rotation) have been used by botnets. Our further investigation (omitted due to space limit) indicates such stability is not static, but dynamic. That is, the live bots are being strategically rotated with a short online duration to maintain such dynamic stability.

2.2 Botnet elasticity

As Figure 1 shows, the number of live bots will surge within a very limited time frame, usually a couple of days, when the botnet is instructed to engage in certain attacks. In our observation, a well managed botnet could rapidly recruit new bots in the order of hundreds of thousands. This is another important perspective when we evaluate the potential of a botnet’s attacking power, particularly as today botnets are moving towards BaaS. Therefore, we further define a metric *elasticity* of a botnet, which is the ratio of the maximum over the minimum number of simultaneously live bots.

Figure 2 shows the elasticity result for all families—sorted by size in descending order. As indicated by the figure, overall, the elasticity decreases with the decreasing botnet size. An exception is Conficker. As Figure 1 shows, compared to other large botnets, Conficker maintains a very stable army of bots. From the service perspective, the larger the elasticity value, the more capable a botnet is upon a demand for attacks. From the defense perspective, the more elastic a botnet, the harder to shut it down.

2.3 Service collaborations

For each family, we have a list of botnet identifiers derived from different malware signatures we detected during the 7 months, and we cross-compare the set of active bots from different botnet identifiers over time.

Some collaborative attacks are concurrent. Some different botnet identifiers within the same family have extremely high concurrent usage over the same set of bots. We speculate the attacker employs multiple botnets to launch the same attack. Figure 3 shows an example, demonstrating the collaborations of two botnets within the Blackenergy family. In this figure, the *x*-axis represents daily timestamps when collaboration happened, the *y*-axis on the left represents the index of subnet involved in the collaborations, and the *y*-axis on the right represents the the count of bots. For the scatter plot, each dot with different color shows which botnet the

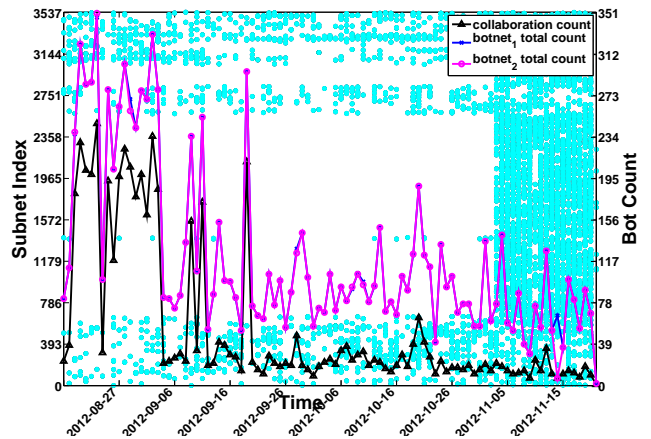


Figure 3: Intra-family collaboration for Blackenergy

subnet belongs to at the time of collaborations, and cyan color indicates that they are employed concurrently by both botnets. Three curves for the total number of bots from each botnet and the number of collaborating bots are also plotted. From the collaboration curve and the total count curve, we can clearly see that activities of both botnets are well synchronized. Also, a dedicated group of bots that belong to both botnets are responsible for the surge events.

Some botnets take turns in collaborative attacks. We notice that some different botnet identifiers own a large number of the same bots, but those bots are rarely used concurrently. Instead, they were solely used by botnet 1 at a time, and later majority of those bots are transferred to botnet 2. We suspect the same set of bots are leveraged to participate in different campaigns. This temporal pattern suggests that different botnets could be controlled by the same botmaster. They are essentially the same botnet, with some difference in their code base that resulted in a different signature.

3. DISCUSSION

Botnets have been widely used for various Internet attacks. Driven by profit, “Botnet-as-a-Service (BaaS)” is on the rise. In this work, we set to examine botnets from the perspective of services, focusing on the service stability, elasticity, and collaborations. Our results indicate that to remain active and profitable, BaaS is very versatile and adaptive by constantly and continuously adopting new techniques, to evade from being detected and to better serve their underground customers. We are actively conducting in-depth analysis and we seek to offer more insights to the defense community in a timely manner.

This work is partially supported by NSF under grants CNS-0746649 and CNS-1117300.

4. REFERENCES

- [1] M. Abu Rajab, J. Zarfoss, F. Monrose, and A. Terzis. A multifaceted approach to understanding the botnet phenomenon. In *IMC*, 2006.
- [2] P. McDougall. Microsoft: Kelihos ring sold ‘botnet-as-a-service’. <http://ubm.io/MtCSr7>, September 2011.
- [3] M. Starr. Fridge caught sending spam emails in botnet attack. <http://bit.ly/1j5Jac1>, Jan 2014.
- [4] M. Vicario. Four ways cybercriminals profit from botnets. <http://bit.ly/1e1SIiP>, Nov 2010.
- [5] Wikipedia. Carna botnet. <http://bit.ly/1slx1E6>, 2014.