

# POSTER: How Distributed Are Today's DDoS Attacks?

An Wang<sup>1</sup> Wentao Chang<sup>1</sup> Aziz Mohaisen<sup>2</sup> Songqing Chen<sup>1</sup>  
<sup>1</sup>Department of Computer Science, George Mason University  
{awang10, wchang7, sqchen}@gmu.edu  
<sup>2</sup>VeriSign Labs  
amohaisen@verisign.com

## ABSTRACT

Today botnets are responsible for most of the DDoS attacks on the Internet. Understanding the characteristics of such DDoS attacks is critical to develop effective DDoS mitigation schemes. In this poster, we present some preliminary findings, mainly concerning the distribution of the attackers, of today's DDoS attacks. Our investigation is based on 50,704 different Internet DDoS attacks collected within a seven-month period for activities across the globe. These attacks were launched by 674 botnet generations from 23 different botnet families with a total of 9026 victim IPs belonging to 1074 organizations that are collectively located in 186 countries. We find that different from the traditional widely distributed intuition, most of these DDoS attacks are not widely distributed as the attackers are mostly from the same region, i.e., highly regionalized. We also find that different botnet families have strong target preferences in the same area as well. These findings refresh our understanding on the modern DDoS attacks.

## Categories and Subject Descriptors

C.2 [Computer Communication Networks]: General—*data communications, security and protection*

## General Terms

Measurement, Security

## Keywords

Botnet, DDoS attacks, attack characterization

## 1. INTRODUCTION

Driven by the underlying profit motive, various botnets on the Internet are constantly being leveraged for Internet Distributed Denial of Services (DDoS) attacks. To understand the fundamentals of DDoS attacks and defend against them, enormous efforts have been continuously made from both academia and industry [1, 3, 4]. The ever-improving defenses lift the bar of DDoS attacks, and make the attacking strategies more and more sophisticated as well.

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
CCS '14, November 3–7, 2014, Scottsdale, Arizona, USA  
Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

Therefore, a timely and accurate understanding of latest DDoS attack strategies is the key to enhance the existing defenses. But most of the existing measurements are based on the indirect traffic measures, such as backscatters, or traffic collected locally, such as an ISP or by infiltration into a botnet. A recent and large scale view of today's Internet DDoS attacks is demanded.

In this preliminary study, we gain access to a DDoS attack workload collected around the globe. Our dataset is provided by the monitoring and attribution unit in a DDoS mitigation company in the United States, with partnerships with various major ISPs. The workload we obtained ranges from August 28, 2012 to March 24, 2013, a total of 209 days (about seven months of valid and marked attack logs). In this seven-month period, a total of 50,704 different DDoS attacks were observed, which were launched by 674 different botnets coming from 23 different botnet families.

Figure 1 shows the distribution of the DDoS attacks along time. They are launched by different botnet families. In this figure, the  $x$ -axis represents the timestamp when DDoS attacks happened and the  $y$ -axis represents the aggregated number of DDoS attacks at the same time over the period of 28 weeks. As the figure shows, there is no obvious patterns concerning time. This might indicate that most DDoS attacks happening today are profit driven. As we can observe from the figure, sometimes there are several hundred of simultaneous DDoS attacks. For example, on October 27th, 2012, there were 261 DDoS attacks happening at the same time.

Our initial study on this workload mainly concerns the distribution of these DDoS attacks, i.e., how distributed are these attacks. Through careful investigation, we find that different from our intuition, most of these DDoS attacks are not widely distributed, but highly regionalized. Furthermore, the attacks of different botnet families often have target preferences as well. These findings update our DDoS understanding and could be leveraged to develop new mitigation schemes.

## 2. DDOS GEOLOCATION ANALYSIS

In this section, we present some preliminary results of our study.

### 2.1 Attacker Geolocation Affinity

We first look into the geographical distributions of these DDoS attacks from attacker's perspective. For this purpose, we obtained the mapping database of IP addresses to corresponding organizations/countries from third-party resources. Then for these attacks, we extract the IP addresses and map them into the corresponding organization in real-time.

Figure 2 shows the CDF of country coverage of each single DDoS attack launched by each of the top-10 botnet family. From this figure, we can see that for more than 80% of the attacks, there are less than 7 countries which indicate that most DDoS attacks




Figure 1: Attack distribution

are not widely distributed. And we can also tell that Optima has the broadest coverage among all these families. We plot the similar figure for organization coverage, shown in Figure 3. These two figures show very similar trends. This indicates that for a single country, there are usually only a few organizations involved.




Figure 2: CDF of country coverage for each family




Figure 3: CDF of organization coverage for each family

We use Dirtjumper as an example to further illustrate some details. Figure 4 shows one of the DDoS attacks launched by Dirtjumper towards a Russian website starting at 17:12:03 on November 9th, 2012. To make it clearer, we use log scale on the numbers of bots involved. For this attack, most bots were located in US and Israel. 3242 bots involved in this attack were from Israel. Figure 5, tells more details about the affinity at the organization level.

In Figure 5, the size of the circles represents the number of attacks launched from that specific organization. This figure aggregates all the attacks launched in October 2012 by Dirtjumper. Even though there are some small dots scattered randomly, generally speaking, most attackers were located in US, South Africa, Europe area and Southeast Asia. After further looking into the most popular organizations involved, we found that these organizations are

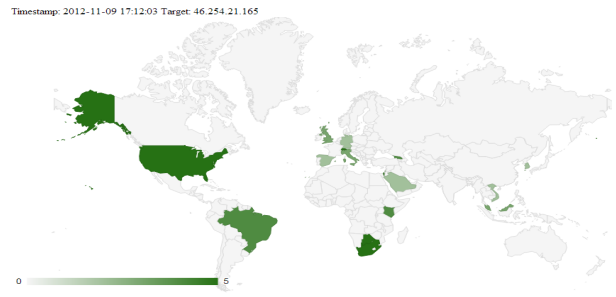


Figure 4: Dirtjumper attacker preference (country-level)

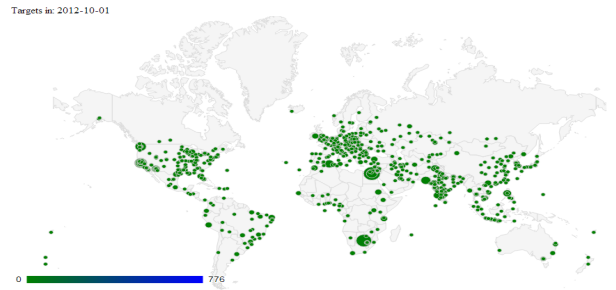


Figure 5: Dirtjumper attacker preference (organization-level)

almost all large-scale ISPs, which possess the capability to inundate the network.

## 2.2 Target Geolocation Affinity

After finding that these attacks are originated from attackers in same or close regions, we conduct a similar analysis on the targets of these DDoS attacks. Interestingly, we found a similar affinity pattern of these targets. In our dataset, there are 9026 unique targets were attacked only once, the most popular target was attacked 940 times over the same period. After looking up the most popular target IP address, we found that this IP address belongs to the domain of HostGator, which is a Houston-based web hosting service, indicating that the real target could be an organization hosted by this service. These results indicate that most botnet families have some target preference.

Consider one of the most active families, Dirtjumper, as an example. Figure 6 shows the target country distribution in October, 2012. In this figure, the shades of color represent the popularity of

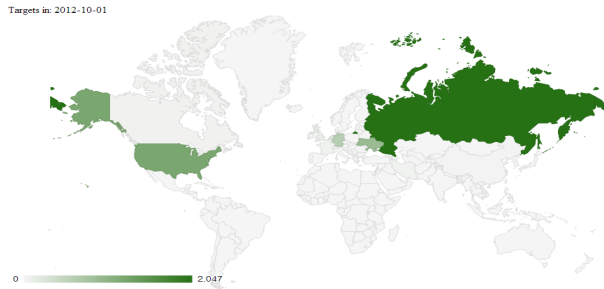


Figure 6: Dirtjumper target preference (country-level)

different target areas. The deeper the color, the more attacks happened in that area. From this figure, we can see clearly that most attacks were targeted US and Russian areas. While this figure shows the country-level analysis, Figure 7 shows the organization-level analysis.

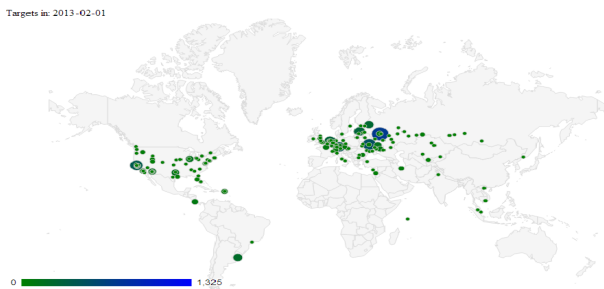


Figure 7: Dirtjumper target preference (organization-level)

In Figure 7, the size of the circles on the map again represents the number of attacks happened towards a specific target. From this figure, we could see clearly that among all the targets there are some hotspots. Our further analysis indicates that most attacks were aimed towards web hosting services, large-scale cloud providers and data centers, Internet domain registrars and backbone autonomous system. They normally possess massive network resources and play a critical function in the operation of other Internet services.




Figure 8: Pandora bots shift (country-level)

Above are static analyses of geolocation affinity. Now we will briefly display some preliminary results of dynamic analysis. Here we use Pandora as an example since it is active enough to show clearly the shift pattern. Figure 8 shows the result of our dynamic analysis. In this figure, the  $x$ -axis represents the whole 28 weeks timespan and the  $y$ -axis represents all the countries involved in the attacks. The size of the markers in the figure represents the number of bots involved and located in the corresponding countries. This figure further confirms the geolocation affinity since most of the

shifts happened within the same countries. Note here we aggregate the shifts of all the attacks happened at the same time. We will do more detailed dynamic analyses by inspecting each single DDoS attack in the future.

### 3. DISCUSSION

DDoS attacks are the most popular large scale attacks frequently launched on the Internet via botnets. While most existing studies have mainly focused on designing various defense schemes [2], the measurement and analysis of large scale Internet DDoS attacks are not very common. Our preliminary study covers the botnet geolocation affinity analyses. All these findings from our analysis could be applied to practical defense designs. For example, the target affinity analysis reveals that more than half of the targets were attacked more than once. This clearly indicates that the current defense mechanisms in practice, if any, are insufficient, and the response from the defense side is slow. Furthermore, most botnet families have the long-term targets, which hints us to develop specific botnet based defenses. This could be further improved by combining the dynamics of botnet behaviors.

Our investigation of these DDoS attacks reveals a trend of not massively distributed DDoS attacks. This behavior change will make network anomaly detection more difficult especially for those detection mechanisms that depend on the location or distribution information. Our country and organization level affinity analysis may shed some light on the potential solutions. For example, different countries from different regions can prioritize their efforts in dealing with attacks from some different botnet families. Organizations can also better provision their limited defense resources to maximize the protection surface and capabilities.

We will continue our analysis and track how each botnet family deploys the attack forces along time and how they shift the bots involved. By utilizing machine learning techniques, we plan to build botnets dynamics models for each family to predict their behavior. Also, from our study, the geolocation affinity has become so popular among botnet families. This behavior change should be related to the behavior of botmasters. To further validate this, we will investigate more about how botmasters control their bots in our future work.

### 4. ACKNOWLEDGEMENT

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

### 5. REFERENCES

- [1] M. Casado, T. Garfinkel, W. Cui, V. Paxson, and S. Savage. Opportunistic measurement: Extracting insight from spurious traffic. In *Proc. 4th ACM Workshop on Hot Topics in Networks (Hotnets-IV)*, 2005.
- [2] C. Douligeris and A. Mitrokotsa. DoS attacks and defense mechanisms: classification and state-of-the-art. *Computer Networks*, 2004.
- [3] S. Jin and D. Yeung. A covariance analysis model for ddos attack detection. *IEEE International Conference on Communications*, 2004.
- [4] Z. M. Mao, V. Sekar, O. Spatscheck, J. van der Merwe, and R. Vasudevan. Analyzing Large DDoS Attacks using Multiple Data Sources. In *Proceedings of ACM SIGCOMM Workshop on Large-Scale Attack Defense*, 2006.