Profiling DRDoS Attacks with Data Analytics Pipeline

Laure Berti-Équille, Yury Zhauniarovich

► To cite this version:

Laure Berti-Équille, Yury Zhauniarovich. Profiling DRDoS Attacks with Data Analytics Pipeline. the 2017 ACM, Nov 2017, Singapore, Singapore. 10.1145/3132847.3133155 . hal-01829013

HAL Id: hal-01829013

<https://hal.science/hal-01829013>

Submitted on 3 Jul 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Profiling DRDoS Attacks with Data Analytics Pipeline

Laure Berti-Equille

LIF, CNRS, Aix Marseille University, Polytech
Marseille, France

laure.berth@lif.univ-mrs.fr

Yury Zhauniarovich

Qatar Computing Research Institute, HBKU
Doha, Qatar

yzhauniarovich@qf.org.qa

ABSTRACT

A large amount of Distributed Reflective Denial-of-Service (DRDoS) attacks are launched every day, and our understanding of the *modus operandi* of their perpetrators is yet very limited as we are submerged with so Big Data to analyze and do not have reliable and complete ways to validate our findings. In this paper, we propose a first analytic pipeline that enables us to cluster and characterize attack campaigns into several main profiles that exhibit similarities. These similarities are due to common technical properties of the underlying infrastructures used to launch these attacks. Although we do not have access to the ground truth and we do not know how many perpetrators are acting behind the scene, we can group their attacks based on relevant commonalities with cluster ensembling to estimate their number and capture their profiles over time. Specifically, our results show that we can repeatably identify and group together common profiles of attacks while considering domain expert's constraint in the cluster ensembles. From the obtained consensus clusters, we can generate comprehensive rules that characterize past campaigns and that can be used for classifying the next ones despite the evolving nature of the attacks. Such rules can be further used to filter out garbage traffic in Internet Service Provider networks.

CCS CONCEPTS

• **Information systems** → **Clustering**; *Data analytics*;
• **Networks** → **Denial-of-service attacks**; • **Theory of computation** → **Unsupervised learning and clustering**; • **Computing methodologies** → *Ensemble methods*;

1 INTRODUCTION

Over the last few years, we are witnessing the massive recrudescence of DDoS (Distributed Denial of Service) attacks affecting every online service: SSH, email, Web, gaming, etc. [2, 16] with system outages that can have tremendous repercussions on the business and Internet users' everyday-life¹. Amongst the various types of DDoS nuisances [8], UDP-based

Distributed Reflective Denial of Service (DRDoS) attacks are among the most impactful due to their unique properties. First, they preserve the attacker's anonymity through IP address spoofing. Therefore, it is tough to identify the attackers and block their services. Second, these attacks abuse some UDP-based network protocols that send considerably larger response comparing to the size of the request. Several studies have shown that bandwidth amplification of UDP-based DRDoS attacks can multiply the traffic up to a factor of 500 [5, 14]. The combination of these two properties allows attackers to launch massive attacks while preserving the anonymity of the attacking infrastructure [12]. Such properties provoke interest in DRDoS attacks from researchers and analysts for finding answers to many questions raised by these attacks: Are there patterns in the attacks, and commonalities in the attacks' characteristics that can be discovered? What are the softwares that are used to generate the spoofed packets? How many machines are there in the attacker net? To fight effectively these types of DDoS attacks, we need to collect evidences from the data left behind by the attackers and identify their profiles as it is done for any other criminal activities. In line with recent studies [1, 11, 13] at improving our understanding of the DRDoS threat, this paper precisely attempts to answer some of the aforementioned questions by analyzing traffic data collected from an instance of the amplification honeypot we have adapted and deployed.

Our main contributions in this paper are the following: (i) we developed an amplifier honeypot to collect fine-grained data about hundreds-of-millions-packets scale attacks happening every day; (ii) we cluster the attacks with clustering ensemble, explore, and discover patterns and commonalities of the attack campaigns; and (iii) we built an analytic pipeline for exploration and clustering of attack profiles. Our aim is to reveal common technical characteristics of the infrastructures responsible for typical attacks. We believe that an in-depth understanding of the patterns of DRDoS attacks can help to characterize the attackers' *modus operandi* providing the tools and knowledge to combat with this type of DDoS attacks.

2 DATA COLLECTION

In order to analyze UDP-based DRDoS attack hyperplane, Kramer et al. [11] developed a honeypot, called AMPOT. This type of honeypot pretends to run services known to be vulnerable to amplification attacks, such as DNS or NTP, so that when they are scanned and exploited by real attackers, researchers can collect useful information to better understand the intrinsic characteristics of amplification attacks and the targets of the attackers. Unlike traditional honeypots, amplification honeypots can only be used to analyze data related to the victims, not the attackers.

¹<https://www.akamai.com/us/en/multimedia/documents/state-of-the-internet/akamai-q2-2016-state-of-the-internet-security-report.pdf>

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.

CIKM'17, November 6–10, 2017, Singapore, Singapore

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-4918-5/17/11...\$15.00

<https://doi.org/10.1145/3132847.3133155>

To collect the data for our analysis, we deployed a modified version of AMPPOT on a cloud provider, labeled as a Honeypot in Figure 1. We modified the honeypot in several aspects. The original implementation had several traffic limiting strategies to reduce the amount of incoming data. So, we first disabled all limitation mechanisms such that all packets coming to the honeypot could be recorded. Second, we added to our implementation the code that extracts additional information from the incoming packets. In particular, for every incoming packet, we extract ‘‘Do not Fragment’’ (DF) flag. We compute the UDP checksum value, and compare it with the data from the UDP header. If the value does not coincide (and the UDP header value is not equal to 0), we report the packet as **corrupted**. Third, our modification enables the storage of all data characterizing all incoming packets, avoiding throttling used in the original implementation. Similar to the original AMPPOT, we split the data on a daily basis resetting the initial state of the honeypot every day. At the time of dataset collection for this paper, our honeypot has been active since May 2016.

3 FEATURE DESCRIPTION

For every raw incoming packet, we extract primary data features. Then, we group the packets related to the same attack originating from one machine together (according to our understanding). Further, for every such group, we also extract a set of group-level features described hereafter.

Primary Data Features. Primary data features describe an attack from a fine-grained raw packet perspective. These are the features extracted from every packet such as the victim IP address, source and destination ports, TTL (time-to-live), UDP checksum, ID number of IP header, DF flag, and a timestamp (with microsecond precision), when the packet is received by our AMPPOT.

Group-level Features. Group-level features are aggregated values that are computed from packets that are grouped together when they belong to the **same** attack originating from **one** single attacking machine. These groups have the same values for the following fields: **Source IP Address**, **Destination Port Number**, **Internet Header Length**, **UDP Length**, **Don’t Fragment (DF) flag**, **Time-To-Live (TTL)**, and **Corrupted**. Such grouping allows us to reduce the amount of data as input of our analytical pipeline from hundreds of millions of packets down to thousands of groups. Then, for every group we extract relevant group-level features, such as, for example, the attack durations (in ms), the frequency of changes in the source port numbers, the entropy of the port number usage, etc. These features contain the traces that may shed some light on how the infrastructure for launching attacks is used. Other distributional characteristics such as skewness, kurtosis, min, and max values of each feature have been computed to better characterize the data distribution in each group.

Data Reduction and Normalization. Furthermore, we used data normalization for group-level features and applied feature and record selection to avoid multi-collinearity between features (i.e., when one or more features are strongly correlated and can introduce bias in the analysis). In particular, for each pair of features exhibiting strong linear

correlation (above .75), we selected the most meaningful feature from the experts’ point-of-view. Additionally, we filtered out the groups containing less than 10 packets as we do not consider them as potential attacks but rather scanning attempts.

4 ANALYTIC PIPELINE

After these crucial steps of data collection, feature selection, and feature engineering, we developed a unique analytic pipeline that combines constraint-based clustering ensemble and rule discovery for profiling DRDoS attacks and discovering interesting patterns in the attack campaigns. First, we give an overview of the pipeline as it is illustrated in Figure 1 (from left to right). As described earlier, DRDoS attack data are collected by our deployed AMPPOT instance, stored and queried in Elasticsearch. Our analysis has been applied at various levels of time granularity considering all packets received either every 5mn interval, every hour, or every day. In Stage 1 of the pipeline, for each time granularity level, primary features and group-level features are analyzed. Interestingly, the number of groups with common values of a subset of primary features can give us a bound of the maximum number of attacks with similar modus operandi for a given time window. We observed that this bound was fairly stable between 2 to 5 thousands similar attack profiles per day over the period of our study. In Stage 2 of the pipeline, we apply various clustering techniques to the group-level data for capturing complementary aspects and data clustering properties with diverse density-based and distance-based clustering techniques. We tested various methods:

- kMeans [7] that requires the specification of k , the number of clusters;
- HDBSCAN [3] that requires the specification of the minimal cluster size;
- Self-Organizing Map (SOM) [10] that requires the specification of the size of the grid for data reduction and transfer to a neural network;
- Hierarchical clustering (HCLUST) [6] that captures the clusters’ embedding structure, such as hierarchies of clusters;
- EM (expectation maximisation)-based clustering that assigns a probability distribution to each packet which indicates the probability of it belonging to each of the clusters. EM can decide how many clusters to create by cross validation.

Since no single clustering algorithm is optimal, we selected five representative clustering approaches, tuned their parameters with multiple hypothesis testing, and compared their results under three dimensions: (1) traditional clustering quality metrics (e.g., silhouette coefficient, SSE – Sum of Squared Error, and cluster entropy); (2) overlapping between the clusters returned by the different methods; and (3) lineage over time to check whether the characteristics of particular clusters of a given day are repeatedly preserved in the clusters of the next days (using confusion matrices and Jaccard index). The pipeline is modular and can seamlessly integrate other clustering or ensembling techniques that can be selected with respect to a set of labeled data when available.

Figure 1: Analytic Pipeline for Profiling DRDoS Data from Amplifier Honeypot.

Date	Total Nb of distinct				Nb of Groups	Nb of Clusters (single method)					Nb of Clusters (Ensemble)		
	Packets	Victims	TTLs	Src. Ports		kMeans	HDBSCAN	SOM	HCLUST	EM	Co-Assoc.	Voting	Consensus
10/12/2016	180,161,186	2,199	188	12	5035	15	912	13	15	10	10	34	13
11/12/2016	136,137,749	2,065	183	12	5302	15	941	13	14	10	10	41	12
12/12/2016	195,212,682	1,764	186	11	5223	15	969	12	14	10	10	37	13
13/12/2016	198,782,496	1,013	185	12	2763	15	493	10	13	9	9	26	11
14/12/2016	151,586,854	1,046	183	12	2521	15	455	10	12	8	8	25	11

Table 1: 5-Days Sample of AMPOT dataset with clustering results.

In Stage 3, clustering ensembles are built using two ensembling techniques: (1) in the first ensembling technique, co-association matrices are computed based on multiple data partitions obtained from Stage 2, then a similarity-based clustering algorithm based on single-link and normalized cut is applied to the co-association matrices in order to obtain the final partition of the data; (2) in the second ensembling technique, the Hungarian method is applied to solve the assignment problem of the multiple cluster labels in order to re-label the data. Finally, a voting process is applied to select the final consensus cluster for the labeling. In both techniques, instance-based constraints from the experts are used to guide the ensembling process in the form of “must-link” and “cannot-link” instances of the consensus clustering results and they encode the domain-specific knowledge of DRDoS attacks similar to previous work of [4]. Finally, in Stage 4, M5 and decision tree-based rules are discovered from each consensus cluster. These rules are submitted to the domain experts and the most relevant rules can be included as new constraints for Stage 3.

The main objective of this pipeline was twofold: first, to check whether the results obtained from various clustering techniques were consistent to characterize attack profiles for a given time window and over time; and second, whether these results could be corroborated, combined and further explored, for instance, with rules discovered from the consensual clusters as typical profiles of the attacks. From the technical challenge point-of-view, the evaluation of the effectiveness of the method that combines constraint-based clustering ensemble and rule discovery for DRDoS attack profiling is very challenging due to the lack of ground truth and the difficulty to set up a small scaled controlled experiment to show how well this method performs on labeled data. Although we could pay for launching our own DRDoS attacks, there is no guarantee that they will be fully captured

Figure 2: Consensus Clusters (k=13).

by our honeypot and the lack of coverage and traceability will impair the evaluation anyway.

Due to space limitation, details describing how we select the optimal parameter settings for each clustering method are omitted. For illustrating the results, Table 1 only presents the results for each clustering method and their ensembles for 5 representative days. Figure 2 is a SOM-like representation of the consensus clustering result obtained for one day, Dec. the 10th, 2016 where the weight vectors across the map are representative of the distribution of the packets and variables listed in the legend. Figure 3 gives an example of the decision rule we can obtain from the consensus cluster# 6 that is

```

Rule: 10 — On Dec. 10, 2016
IF      src_port_min_norm <= 0.509
      src_port_max_usage_time_norm > .701
      attack_duration_norm >= 20,102,452.3 ms
      attack_duration_norm <= 21,261,873.6 ms
THEN   Cluster# = 6   for [143/73.475%]

```

Figure 3: Example of decision rule from Cluster# 6.

satisfied for 73.475% of the groups that belong to this cluster. The rule indicates that for a particular range of the attack duration in ms and a certain bound on the maximal usage time of the source ports (all being normalized), 143 groups have the same attack profile in cluster 6. The next step for domain experts is to filter and link the information provided by such rules to technical characteristics of the attackers' machines. In summary for the whole period, the average number of packets is $182,376,195.4 \pm 27,499,207.4$ with $4,168.8 \pm 1399.7$ groups and 54.2 ± 17.4 M5 rules generated per day.

5 RELATED WORK

Our work focuses on data analytics applied to one of the prevalent threats of cybercrime commoditization, namely the amplified DDoS attacks where low-cost DDoS services (aka booters) can be purchased by anyone to launch Gbps-scale attacks to exhaust the bandwidth of a victim whose IP address has been spoofed. With a modest subscription fee, abusive users are empowered to harass, block, extort, and intimidate their competitors (e.g., in business or online gaming). To meet the increasing demand for DDoS attacks, such services have scaled up their infrastructure as demonstrated by recent studies exposing various technical and business aspects of DDoS-as-a-service market such as their attack infrastructure, payment ecosystem [9, 15], or characterization of their victims [13]. DDoS botnets and their attack monitoring have been an extensively studied topic [2, 18], followed by amplification attacks [14]. The use of honeypots as "baits" mimicking services that are vulnerable to amplification DDoS attacks has been recently advocated by [11] to be a very useful mean for gaining attack intelligence regardless of the particular protocols vulnerable to amplification and/or the specifics of a booter service used for launching the attacks. Close to this work, we take advantage of the deployment of honeypots and we extensively analyze the data collected from amplification attacks and extract typical profiles and fine-grained characteristics. This level of detail in our data analysis reveals interesting patterns for fingerprinting the attacks, which have not been investigated earlier. Clustering ensemble methods have been designed for combining partitionial and hierarchical clustering methods, but they generally do not take advantage of the domain-related constraints. Constraining the clustering ensembles is an innovative aspect of our work where multiple clustering techniques are combined and domain expert's constraints are used to optimize the clustering results, advancing the state-of-the art on this topic [17].

6 CONCLUSIONS

In this paper, we presented the results of our analysis of DDoS amplification honeypot data from actual network traffic with hundreds of millions packets that have disrupted more than

ten thousands of IP addresses every day within the last 6 months. In particular, we make the following contributions. First, we deployed a modified instance of AMPOT to collect and analyze fine-grained data in order to characterize the profiles of the attacks received by our honeypot. Second, we present an analytic pipeline to group, cluster, and characterize the attacks combining multiple clustering techniques. Third, from the consensus clusters we have obtained, we can discover *a posteriori* M5 and decision tree-based rules explaining attack clusters' characteristics that we will further enable the domain experts to infer technical characteristics of the attackers' infrastructure. Future work will include the development of near real-time detection and profiling.

REFERENCES

- [1] AUPETIT, M., ZHAUNAROVICH, Y., VASILADIS, G., DACIER, M., AND BOSHMAF, Y. Visualization of Actionable Knowledge to Mitigate DRDoS Attacks. In *Proceedings of the 2016 IEEE Symposium on Visualization for Cyber Security (VizSec)* (Oct 2016), pp. 1–8.
- [2] BÜSCHER, A., AND HOLZ, T. Tracking ddos attacks: Insights into the business of disrupting the web. In *Presented as part of the 5th USENIX Workshop on Large-Scale Exploits and Emergent Threats* (2012).
- [3] CAMPELLO, R. J. G. B., MOULAVI, D., AND SANDER, J. *Density-Based Clustering Based on Hierarchical Density Estimates*. 2013, pp. 160–172.
- [4] CHAHDI, H., GROZAVU, N., MOUGENOT, I., BERTI-EQUILLE, L., AND BENNANI, Y. On the use of ontology as a priori knowledge into constrained clustering. In *DSAA* (2016), pp. 632–641.
- [5] CZYZ, J., KALLITSIS, M., GHARAIBEH, M., PAPADOPOULOS, C., BAILEY, M., AND KARIR, M. Taming the 800 Pound Gorilla: The Rise and Decline of NTP DDoS Attacks. In *Proc. of the Internet Measurement Conf. (IMC)* (2014), pp. 435–448.
- [6] HARTIGAN, J. A. *Clustering Algorithms*, 99th ed. John Wiley & Sons, Inc., 1975.
- [7] KANUNGO, T., MOUNT, D. M., NETANYAHU, N. S., PIATKO, C. D., SILVERMAN, R., AND WU, A. Y. An efficient k-means clustering algorithm: Analysis and implementation. *IEEE Trans. Pattern Anal. Mach. Intell.* 24, 7 (2002), 881–892.
- [8] KARAMI, M., AND MCCOY, D. Understanding the emerging threat of ddos-as-a-service. In *Proc. of 6th USENIX Workshop on Large-Scale Exploits and Emergent Threats* (2013).
- [9] KARAMI, M., PARK, Y., AND MCCOY, D. Stress Testing the Booters: Understanding and Undermining the Business of DDoS Services. In *Proc. of the 25th Int. Conf. on World Wide Web, WWW* (2016), pp. 1033–1043.
- [10] KOHONEN, T., SCHROEDER, M. R., AND HUANG, T. S., Eds. *Self-Organizing Maps*, 3rd ed. Springer-Verlag New York, Inc., 2001.
- [11] KRÄMER, L., KRUPP, J., MAKITA, D., NISHIZOE, T., KOIDE, T., YOSHIOKA, K., AND ROSSOW, C. *AmpPot: Monitoring and defending against amplification DDoS attacks*, vol. 9404. 2015, pp. 615–636.
- [12] KÜHRER, M., HUPPERICH, T., BUSHART, J., ROSSOW, C., AND HOLZ, T. Going Wild: Large-Scale Classification of Open DNS Resolvers. In *Proc. of the Internet Measurement Conf. (IMC)* (2015).
- [13] NOROOZIAN, A., KORCZYŃSKI, M., GAÑAN, C. H., MAKITA, D., YOSHIOKA, K., AND VAN EETEN, M. *Who Gets the Boot? Analyzing Victimization by DDoS-as-a-Service*. 2016, pp. 368–389.
- [14] ROSSOW, C. Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In *Proc. of the 2014 Network and Distributed System Security Symp. (NDSS)* (2014).
- [15] SANTANNA, J. J., VAN RIJSWIJK-DEIJ, R., SPEROTTO, A., HOFSTEDE, R., WIERBOSCH, M., GRANVILLE, L. Z., AND PRAS, A. Booters - an analysis of ddos-as-a-service attacks. In *IFIP/IEEE Int. Symp. on Integrated Network Management (IM)* (2015).
- [16] SPECHT, S. M. Distributed denial of service: taxonomies of attacks, tools and countermeasures. In *Proc. of the Int. Workshop on Security in Parallel and Distributed Systems* (2004), pp. 543–550.
- [17] VEGA-PONS, S., AND RUIZ-SHULCLOPER, J. A survey of clustering ensemble algorithms. *IJPRAI*, 25, 3 (2011), 337–372.
- [18] WELZEL, A., ROSSOW, C., AND BOS, H. On Measuring the Impact of DDoS Botnets. In *Proc. of the 7th European Workshop on System Security (EuroSec)* (2014), pp. 3:1–3:6.