Poster: DDoSGrid: a Platform for the Post-mortem Analysis and Visualization of DDoS Attacks

Muriel Franco[1], Jan von der Assen[1], Luc Boillat[1], Christian Killer[1],
Bruno Rodrigues[1], Eder Scheid[1], Lisandro Granville[2], Burkhard Stiller[1]
[1]*Communication Systems Group CSG, Department of Informatics IfI, University of Zürich UZH*
Binzmühlestrasse 14, CH—8050 Zürich, Switzerland
[2]*Computer Networks Group, Institute of Informatics, Federal University of Rio Grande do Sul UFRGS*
Av. Bento Gonçalves, 9500, Porto Alegre, Brazil
E-mail: [franco, killer, rodrigues, scheid, stiller]@ifi.uzh.ch[1],
[jan.vonderassen, lucthierrynicolas.boillat]@uzh.ch[1], granville@inf.ufrgs.br[2]

***Abstract*—Distributed Denial-of-Service (DDoS) attacks remain one of the top reasons for business disruption and financial losses. Although mitigation solutions are available on the market, there is still a need for approaches that help network operators understand attack characteristics and behaviors, resulting in better planning of companies' cybersecurity strategies.**

This paper introduces *DDoSGrid*, a platform for the analysis and visualization of DDoS attacks. DDoSGrid implements an extensible set of miners to extract, process, and analyze information from network traces (*i.e.*, PCAP files) to provide insightful visualizations for a better understanding and in-depth analysis of DDoS attacks in different scenarios. A case study was performed using an HTTP flood attack scenario to evaluate the feasibility of the approach. DDoSGrid enables real-world DDoS scenarios' analysis, providing an intuitive interface integrated with extensible insightful visualizations and data miners.

I. Introduction

The amount of cyberattacks continues to rise in many sectors with different motivations, such as cyberwarfare, politics, and business competition [8]. Distributed Denial-of-Service (DDoS) attacks remain one of the most dangerous threats to companies and service providers worldwide. For example, DDoS attacks are responsible for a high number of occurrences impacting service downtime and performance degradation [5]. Furthermore, the growing number of unsecured Internet-of-Things (IoT) devices ease the spreading of botnets being able to launch massive attacks on service providers [1].

Therefore, besides the mitigation of imminent DDoS attacks, it is also important to consider approaches that help to understand their characteristics and plan of protection against future threats. These approaches include tools for analyzing network traffic [2], extraction of information for pattern recognition [7], and visualization of traffic behavior [9]. Although there are well-known and established tools for the analysis and visualization of security data placed in the market (*e.g.*, Elastic Stack (ELK) and Splunk), there are still open challenges [2] and opportunities for new tools that simplifies the understanding of cybersecurity without the need of dedicated staff and cybersecurity skills. Also, these kind of tools can have an important role in cybersecurity research and education activities.

Thus, this paper develops *DDoSGrid*, an dedicated open-source platform [4] for the post-mortem analysis and visualization of DDoS attacks. *DDoSGrid* addresses the lack of integrated approaches for processing, analyzing, and visualizing records of DDoS attacks. *DDoSGrid* implements an extensible set of miners to extract, process, and analyze information from network traces (*i.e.*, Packet Capture files, PCAP) to provide insightful visualizations for a better understanding and especially an in-depth analysis of different types of DDoS attacks (*e.g.*, SYN flood, Ping-of-Death, and HTTP flood). According to the complexity of a cyberattack, both miners and visualizations developed are extensible to address different scenarios and requirements. For evaluation purposes a case study on a DDoS attack was performed to demonstrate the key features of *DDoSGrid* and provide measurable evidence of the solution's feasibility and its benefits.

The remainder of this paper is organized as follows. While Section II introduces *DDoSGrid*, its architectural components, and key implementation details, Section III provides insights into the case study. Finally, Section IV summarizes the paper and points to further work needed.

II. *DDoSGrid* Overview

DDoSGrid consists of *(a)* miners able to decode PCAP files and to extract features from different protocols (*e.g.*, Ethernet, IP, TCP, or HTTP), *(b)* a Web-based interface that allows users to interact with *DDoSGrid*, and *(c)* visualizations and statistics to enable users to reach detailed insights regarding the dataset under investigation. Thus compared to related work, *DDoSGrid* simplifies the analysis of complex and large log files to provide a complete analysis of a post-mortem DDoS attack. The extensibility of *DDoSGrid* allows users to apply features already provided, but also to implement their customized miners (*i.e.*, components responsible for the extraction of information) and visualizations. Therefore, *DDoSGrid* has potential applications not only for real-world security analysis (*e.g.*, attack identification and planning), but

Annex to ISBN 978-3-903176-39-3 © 2021 IFIP

also in addition for cybersecurity-related research and teaching activities due to its open-source nature.

Figure 1 depicts *DDoSGrid*'s architecture. A user accesses the Web-based interface to analyze a dataset available (*i.e.*, a PCAP file) or uploads a new one. The *Data Manager* is in charge of handling user requests to store and access data related to a cyberattack. When the user uploads a new dataset, it is sent to the *Data Layer* for data extraction and processing of all relevant features of the cyberattack. In turn, an optimized data structure for the *Visualization Module* is build to enable different visualizations according to dedicated user interactions. The communication between the *Data Layer* and the *User Layer* is done though the *Communication API* (Application Programming Interface). An *Integration API* is available, which allows external solutions to request information and reuse monitors, thus, the integration of *DDoSGrid* with other state-of-the-art solutions, such as systems to recommend and offer protections against cyberattacks [3], was reached.

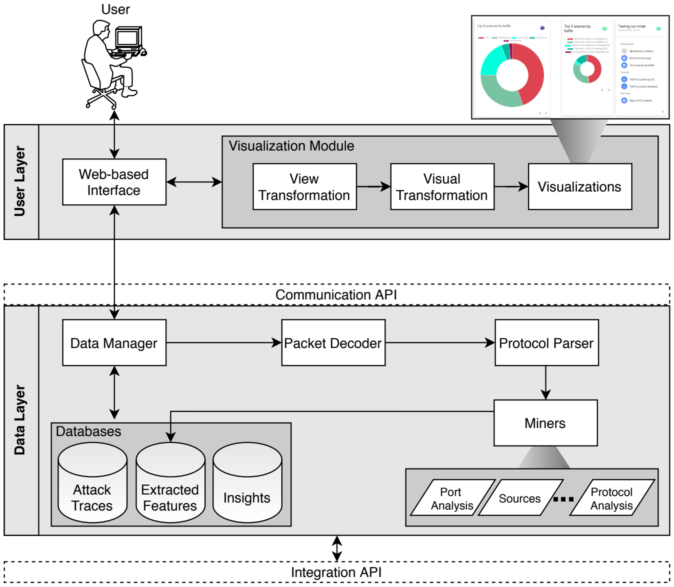

Fig. 1: *DDoSGrid*'s Architecture

Table I lists examples of miners implemented by *DDoSGrid*. All these miners can be applied separately or in combination to extract meaningful information about the traffic available in a PCAP file. This information can be combined to generate different reports and visualizations. *E.g.,* the combination of the TCP States Analyzer, Device Analyzer, and Port Analyzer can be used to identify Command and Control traffic between IoT devices of a botnet. This traffic usually shows activity on port 23 (*i.e.*, Telnet), with most of the packets being SYN packets, and occasionally keep-alive packets (*e.g.*, PSH and ACK) can be observed [1].

A running prototype and its source code is available [4], including a demo-box to be used in production. In order to support the implementation of different miners, a list of protocol parsers was developed for the *DDoSGrid*, such as Ethernet, ARP, IPv4, and TCP/UDP. An integration with the pilot for the European DDoS Clearing House [6] enables an information sharing between users, thus, allowing for the import/export between *DDoSGrid* and the DDoSDB, a platform, part of the DDoS Clearing House, for helping victims of DDoS attacks, the academic community, and the security community to share and get access to basic and enriched information of DDoS attacks [6]. Thus, *DDoSGrid* and the DDoSDB can communicate via APIs provided by both solutions.

TABLE I: Examples of Miners Implemented by *DDoSGrid*

Miner	Target Data	Outcome
Metrics Analyzer	Attack duration, number of packets, IPs and ports	Overview of metrics associated to a cyberattack log file
IP Protocol Analyzer	IPv4, IPv6 packets	Analysis of packets according to IP protocol versions being used
Port Analyzer	UDP and TCP ports	Overview of the most used UDP/TCP ports by number of segments
Top Source Hosts Extractor	Source address	Overview of the hosts sending more traffic and requests
TCP States Analyzer	TCP flags	Analysis of the frequency of TCP flags in the packets, such as ACK, SYN, and FIN
Device Analyzer	HTTP User Agent	Identifies which type of device is being used for the request
Browser and OS Analyzer	HTTP User Agent	Identifies the browser and operating system being used for the request
HTTP Analyzer	HTTP Verbs and End-points	Analysis of the most used HTTP requests (*i.e.*, GET and POST) and end-points accessed via the HTTP protocol

In order to simplify the analysis of cyberattacks and to support the information sharing between users, *DDoSGrid* implements a set of components to ease the communication with other platforms. Currently, the prototype is fully integrated with DDoSDB. For that, a local instance of DDoSDB was deployed together with *DDoSGrid*. New features were added for users *(a)* to visualize datasets while using DDoSDB and *(b)* to import or export log files in both directions. Thus, *DDoSGrid* calls the components *dissector* and *converters* of DDoSDB to generate fingerprints of DDoS attacks and translates its fingerprints to mitigation rules using these converters implemented by DDosDB. Thus, both platforms are fully integrated from the technical perspective, making calls for all components available in a unified architecture.

III. Case Study

Assuming a user with access to a set of network log files from his/her business (*e.g.*, a department store with an e-commerce channel), such a log file contains traffic from a time window within which communication partners experienced problems, for instance service downtime, and customers complaint regarding the services' performance. Thus, this user wants to analyze the problem to define a strategy and, in turn, to avoid this kind of problem impacting their business and its partners in the future. For that, the user accesses *DDoSGrid* and uploads that log file containing the traffic of his/her server maintaining the e-commerce platform under threat.

Initially, the user imports a log file (*i.e.*, in a PCAP format) with 4 minutes of unnatural traffic identified by the *DDoSGrid*. Figure 2 presents the dashboard of the *DDoSGrid* for this dataset, containing a summary of metrics extracted by *DDoSGrid*'s miners. As visible, many packets were received

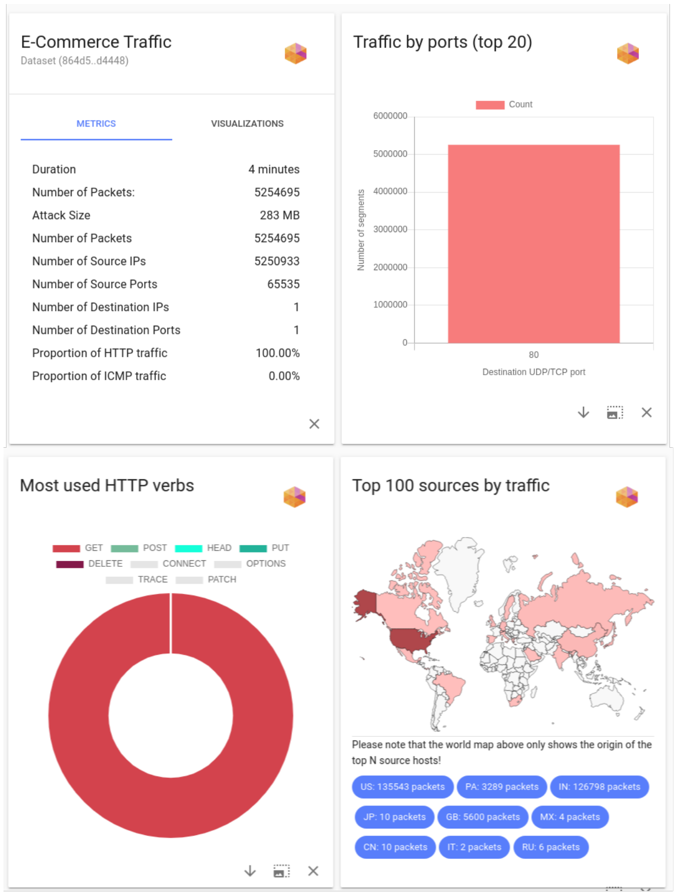

E-Commerce Traffic

Dataset (864d5..d4448)

Duration	4 minutes
Number of Packets:	5254695
Attack Size	283 MB
Number of Packets	5254695
Number of Source IPs	5250933
Number of Source Ports	65535
Number of Destination IPs	1
Number of Destination Ports	1
Proportion of HTTP traffic	100.00%
Proportion of ICMP traffic	0.00%

Traffic by ports (top 20)

Most used HTTP verbs

Top 100 sources by traffic

Please note that the world map above only shows the origin of the top N source hosts!

Fig. 2: HTTP Flood Attack on the E-Commerce Platform — Scenario *(a)*

during the last 4 minutes, all of them with the same destination IP (*i.e.*, the E-commerce server). As highlighted by the bar and pie chart, 100% of the traffic is HTTP GETs on Port 80. Besides that, the user verifies that most of these requests come from North America, which is not common for the business, since most of his/her business customers are during regular days from Europe. Based on that analysis, supported by additional visualizations and information, the user realizes that an HTTP Flood, one of the most common DDoS attacks, targets his/her system. Thus, further measures (*e.g.*, a better cybersecurity planning and an acquisition of protections schemes) should be taken into account to avoid further impacts of this kind of attack in the future.

This sample scenario shows possible measurable aspects that can be covered by *DDoSGrid*, highlighting different features, visualizations, and miners available. Therefore, *DDoSGrid* can *(a)* help during the attack investigation or *(b)* provide a complementary view of malicious traffic, which, together with additional cybersecurity tools, is needed during different steps related to the planning and investments in cybersecurity.

IV. Summary and Future Work

This paper introduced *DDoSGrid*, an open-source approach for the analysis and visualization of DDoS attacks. *DDoSGrid* stands out as a dedicated planning and support tool to help network operators and decision-makers to gain insights about possible occurrences, impacts, and behaviors of different types of DDoS attacks (*e.g.*, SYN flood, Ping-of-Death, and HTTP flood). For that, *DDoSGrid* implements a set of components to process information from log files (*e.g.*, PCAP files) and presents such information in a structured, interactive, and user-friendly manner.

Next steps in this context includes *(a)* the support of real-time traffic by integrating *DDoSGrid* with different monitors and traffic files (*e.g.*, Netflow records and IoT sensors), *(b)* the implementation of additional features based on Machine Learning (ML) supervised techniques (*e.g.*, K-nearest and Random Forest) to identify and classify types of cyberattacks in given traffic automatically, and *(c)* the support of a different number cyberattacks, thus, allowing users to analyze malicious traffic and behaviors of others threats, such as malware families, brute-force, and botnets. Although evidence as above already shows that *DDoSGrid* can scale for large datasets, an in-depth analysis of *DDoSGrid*'s performance will be performed to validate these findings together with an overall usability analysis with real-world cybersecurity operators.

Acknowledgements

This paper was supported partially by *(a)* the University of Zürich UZH, Switzerland and *(b)* the European Union's Horizon 2020 Research and Innovation Program under Grant Agreement No. 830927, the CONCORDIA project.

References

[1] A. Kumar, T. Joon Lim, "Early Detection of Mirai-Like IoT Bots in Large-Scale Networks through Sub-sampled Packet Traffic Analysis," in *Advances in Information and Communication*, K. Arai and B. Rahul, Ed. Springer International Publishing, 2020, pp. 847–867.

[2] M. Cinque, D. Cotroneo, and A. Pecchia, "Challenges and Directions in Security Information and Event Management (SIEM)," in *2018 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW 2018)*, Memphis, USA, 2018, pp. 95–99.

[3] M. Franco, B. Rodrigues, and B. Stiller, "MENTOR: The Design and Evaluation of a Protection Services Recommender System," in *15th International Conference on Network and Service Management (CNSM 2019).* Halifax, Canada: IEEE, October 2019, pp. 1–7.

[4] J. von der Assen, L. Boillat, M. Franco, "DDoSGrid - Source Code," 2021, https://www.csg.uzh.ch/ddosgrid/, Last visit April 2021.

[5] R. Maciel, J. Araujo, J. Dantas, C. Melo, E. Guedes, and P. Maciel, "Impact of a DDoS Attack on Computer Systems: an Approach based on an Attack Tree Model," in *IEEE International Systems Conference (SysCon 2018)*, Vancouver, Canada, 2018, pp. 1–8.

[6] R. Poortinga, J. Ceron, J. Santanna, C. Hesselman, "European DDoS Clearing House Pilot," 2019, https://github.com/orgs/ddos-clearing-house, Last visit April 2021.

[7] A. A. Sallam, M. N. Kabir, Y. M. Alginahi, A. Jamal, and T. K. Esmeel, "IDS for Improving DDoS Attack Recognition Based on Attack Profiles and Network Traffic Features," in *16th IEEE International Colloquium on Signal Processing Its Applications (CSPA 2020)*, Langkawi, Malaysia, 2020, pp. 255–260.

[8] Verizon, "Data Breach Investigations," March 2020, https://enterprise.verizon.com/resources/reports/2020-data-breach-investigations-report.pdf, Last visit April 2021.

[9] C. Wu, S. Sheng, and X. Dong, "Research on Visualization Systems for DDoS Attack Detection," in *IEEE International Conference on Systems, Man, and Cybernetics (SMC 2018)*, Miyazaki, Japan, 2018, pp. 2986–2991.