

Playing in the Sandbox: A Step Towards Sound DDoS Research Through High-Fidelity Evaluation

Lumin Shi, Samuel Mergendahl, Devkishen Sisodia, Jun Li

ABSTRACT

While volumetric distributed denial-of-service (DDoS) attacks evolve into stealthier and more disruptive threats, real-world network operators often ignore the over two decades of DDoS defense research and still rely on basic defense solutions that cannot properly defend against these advanced attacks. One likely explanation for this contradiction lies in the lack of *sound* empirical evaluation of a DDoS defense solution; prior to the deployment of a DDoS defense system, a network operator must understand the impact of the defense system specifically in their network. Without such knowledge, the network operator may fear poor efficacy of the defense due to known issues such as an increased false positive rate from domain shift or negative effects on legitimate traffic from coarse-grained mitigation techniques. In fact, many of the most cited academic solutions for DDoS defense often lack this crucial insight. In order to provide network operators assurance of defense efficacy in their network, we propose a DDoS emulation platform that can evaluate state-of-the-art DDoS detection and mitigation solutions in various real-world scenarios. Our platform emulates the real-world Internet topology, fine-grained application traffic from actual applications, and a user-friendly interface for network operators/researchers to implement different attacks/defenses. Moreover, we demonstrate the usefulness of our DDoS emulation platform through a comprehensive study of existing DDoS defense solutions.

1. INTRODUCTION

Advanced distributed denial-of-service (DDoS) attacks, such as the CrossFire attack [1] and CICADAS [2], seriously challenge the efficacy of typical, basic DDoS defense strategies deployed by network operators. A basic DDoS defense consists of two main entities: 1) simple threshold-based DDoS detection/classification systems [3], and 2) coarsely-grained mitigation solutions [4]. Unfortunately, these advanced attacks exacerbate the common concerns of basic DDoS defense systems. Namely, a threshold-based detection solution suffers from high false positive rates and a coarsely-grained mitigation solution filters legitimate traffic.

Despite years of DDoS defense research that outlines these concerns, at the time of this writing, network operators still rely on the basic DDoS defense strategies that cannot sufficiently prevent advanced attacks. Without assurance of the efficacy of a defense in a specific target network, many network operators may not risk the deployment of cutting edge research defenses. In particular, we surveyed well-received DDoS detection papers, and found that even highly-cited

detection solutions frequently lack the thorough evaluations of realistic network settings that could provide such an assurance (e.g., a defense may not evaluate its efficacy with traffic flows generated by heterogeneous applications). What is worse, advanced DDoS attacks, which may not be visible to the targeted edge networks, require access to telemetry data from transit networks for detection. Unfortunately, not all researchers can access such information. Similarly, as a detection system will eventually generate false positives. In order to evaluate the efficacy of a DDoS mitigation system, a network operator must understand the often complex consequences of mitigation under false positives.

We propose a DDoS emulation platform, which we refer to as a DDoS sandbox, as a step towards sound empirical evaluation for DDoS research. At a high level, the sandbox provides 1) packet-level mimicry of real networks (e.g., fine-grained traffic flow rates from real IP addresses), 2) a mininet-based emulation environment that creates inferred network topologies, 3) a set of sandbox interfaces that allow users to manipulate networks and end-host behavior (e.g., network routing, end-host applications), and 4) a set of DDoS attack and defense implementations based on the highly-cited DDoS research papers.

2. A HIGH-FIDELITY DDOS SANDBOX

In this project, we hope to bridge the gap between DDoS defense research and real-world deployment. In particular, the DDoS sandbox is designed to facilitate network operators to 1) evaluate off-the-shelf solutions in an emulated environment that mimics their networks and 2) provide them with more confidence when deploying DDoS defense solutions. We also believe DDoS and network research communities can benefit from this project by evaluating their projects inside of the sandbox. Our proposed sandbox for DDoS experiments consists of three components: *Topology Builder*, *Traffic Mimicker* and *DDoS Repositories*, as shown in Figure 1.

2.1 Topology Builder

Since network operators are most likely not willing to deploy defense solutions in their networks without thorough testing, and because most researchers do not have access to a real network to evaluate their research, we believe the best way forward is to emulate a realistic network environment that network operators and researchers can use to study defense solutions. As the foundation of our sandbox, the Topology Builder consists of three steps toward building a realistic network environment. First, it processes the traffic telemetry data of a network (e.g., a campus network) to

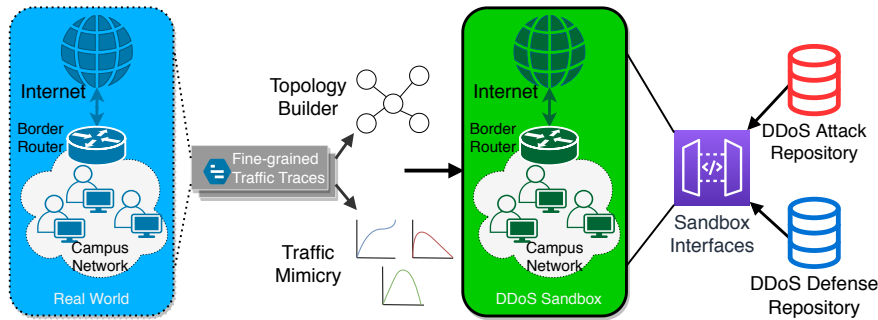


Figure 1: DDoS Sandbox Workflow

collect the source and destination IP addresses which will be emulated as the end-hosts in the sandbox’s network topology. The traffic telemetry data is provided as input by network operators to be processed by the Topology Builder. The telemetry data can be represented as summarized flow information (e.g., source/destination IP/port, protocol, number of packets/bytes transferred, etc.), or packet-level information (e.g., pcap files). Second, the Topology Builder then uses the unique source and destination IPs to build an inferred Internet AS-level topology of the network using the AS relationships dataset provided by CAIDA [5]. Such a topology or blueprint is especially important when it comes evaluate defense solutions deployed at different Internet locations. Finally, the Topology Builder uses the above topology/blueprint to create a software-defined network in a many-core system. Such a network serves the base for facilitating packet-level mimicry of a network and the launch of DDoS attacks and defenses. Also at this step, we provide a reference routing implementation for the network to enable end to end communications. Previous emulation projects, such as Mininet [6], have build a solid foundation for us to build such a network in a reasonable time frame.

2.2 Traffic Mimicker

The Traffic Mimicker re-creates the network traffic of a network in the sandbox, and it provides realistic background traffic of the network for evaluating DDoS defense solutions. We use the collected telemetry data of a network as mentioned in Sec. 2.1 and assign each host in the sandbox its observed/captured flows from the original network. Then, at the evaluation time, each host establishes flow connections to their destination hosts, and faithfully follow the packet sending rate as observed from the original network.

Since we are creating an emulation framework, we should not limiting ourselves to layer 4 flows. The sandbox will allow network operators to run real applications on each host, as long as a sufficient amount of telemetry data is captured from real networks. With that said, as our initial goal, we will create a reference Traffic Mimicker that utilizes the real layer 4 stack in Linux to generate packet-level flows.

2.3 DDoS Repositories

As our main contribution to sound empirical DDoS research, we create a repository for different DDoS attack and defense implementations. In this research, we first try to collect DDoS attack/defense implementations from their original authors. If we cannot obtain the implementation from its author(s), we will implement the solution while making it publicly accessible. Of course, due to the time

constraint, we cannot implement all the DDoS attacks and defenses. Instead, we only implement the attacks/defenses in well-received DDoS papers (e.g., the CrossFire attack [1]).

3. CONCLUSION

There is a significant gap between existing DDoS defense research and its real-world adoption by the network community. This gap led us to propose a high-fidelity DDoS sandbox. With this sandbox, we hope to help network operators evaluate existing DDoS attacks and defenses in emulated networks that faithfully mimic their real networks. To the best of our knowledge, our proposed system is the first attempt to shrink the gap between DDoS research and its real-world deployment. Our vision is to build a DDoS evaluation platform that enables realistic empirical evaluations of DDoS research. We invite DDoS researchers to evaluate their attacks and defenses with such a sandbox as a benchmarking tool. Lastly, we acknowledge that the engineering challenges are not trivial, and there are open questions yet to be answered (e.g., What are the hardware requirements needed to emulate an edge network with high-fidelity? To what extent and at what granularity can one mimic the observed traffic of a network?). Nonetheless, we believe such a system must exist for sound DDoS research.

4. REFERENCES

- [1] Min Suk Kang, Soo Bum Lee, and Virgil D. Gligor. The Crossfire Attack. In *2013 IEEE Symposium on Security and Privacy*, 2013.
- [2] Yu-Ming Ke, Chih-Wei Chen, Hsu-Chun Hsiao, Adrian Perrig, and Vyas Sekar. CICADAS: Congesting the Internet with Coordinated and Decentralized Pulsating Attacks. In *Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security*, 2016.
- [3] FastNetMon. Testimonials. <https://fastnetmon.com/client-testimonials>, 2020.
- [4] Warren Kumari and Danny McPherson. Remote Triggered Black Hole Filtering with Unicast Reverse Path Forwarding (uRPF). <https://tools.ietf.org/html/rfc5635>, 2009.
- [5] CAIDA. The CAIDA AS relationships dataset. <http://www.caida.org/data/as-relationships>, 2017.
- [6] Bob Lantz, Brandon Heller, and Nick McKeown. A network in a laptop: Rapid prototyping for software-defined networks. In *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks, Hotnets-IX*, 2010.