Stratum Filtering: Cloud-based Detection of Attack Sources

Amir Herzberg(1), Haya Shulman(2), Michael Waidner(2)
(1)Bar Ilan University, (2)Fraunhofer Institute for Secure Information Technology SIT and
Center for Research in Security and Privacy, Darmstadt CRISP

Abstract

Denial of Service (DoS) attacks pose a critical threat to the stability and availability of the Internet. In Distributed DoS (DDoS) attacks multiple attacking agents cooperate in an attempt to cause excessive load in order to disconnect a victim. The frequency and volume of DoS attacks continue to break records, reaching 400Gb/s. Although many defenses were proposed, very few are adopted, due to low effectiveness, high costs and the changes required to integrate them into the existing infrastructure.

To improve resilience against DDoS attacks the service providers move their operations to cloud platforms. Unfortunately, even if the cloud applies filtering, rate limiting and deep packet inspection, the attacker can subvert those defenses by distributing the attack among multiple attacking IP addresses and aiming the flood at the victim.

In this talk we focus on DDoS attacks which disrupt the availability of a service by depleting the bandwidth or the resources of an operating system or application on the server side. Such attackers typically employ a botnet to generate large traffic volumes. A botnet consists of bots (compromised computers) located in different parts of the Internet. The bots, depending on their privileges on the victim host, send multiple packets either from spoofed or using their real IP addresses.

We utilize the cloud platform to implement Stratum Filtering, a novel mechanism aimed at protecting the availability and resilience of the web servers hosted on clouds. Our mechanism is easy to integrate into the cloud platform and does not require changes to the existing infrastructure nor the protected servers. Stratum Filtering facilitates the large IP address blocks allocated to the clouds, distributed availability zones and the support of service migration within the cloud platforms. These advantages offered by clouds enable us to restrict the attacker to a naive strategy where the best possible attack is to simply flood the *entire* IP address block allocated to the cloud. However, such an attack requires huge volume of traffic exposing malicious sources. In addition, controlling and coordinating a large number of bots that would suffice for disconnecting a cloud is not trivial to accomplish.

Stratum Filtering is comprised of three layers, such that each successive layer applies filtering targeted at blocking a different type of attack traffic on network, transport or application layers.

The filtering uses the difference in behavior of legitimate clients vs bots, to identify and filter traffic arriving from non-standard clients. To characterize and model the legitimate behavior of web clients we perform large scale Internet measurements using a distributed ad-network, and collect data across multiple networks and geographical areas. The client connections to the protected webservers via Stratum Filtering consist of three steps, as follows:

1. Anti-spoofers map to proxy: A client makes a DNS request and is mapped to one of multiple proxies, chosen pseudorandomly as function of client's IP address. As a result, the spoofers cannot receive the response, hence this step foils spoofer-only attacks; responses are also not sent to black-listed IP addresses (detected in previous interactions).

2. Redirect to server: A client sends the HTTP request to the proxy, which redirects non-blacklisted clients to pseudorandomly assigned server IP and port, with cookie for validation at next step. This light-weight process facilitates effective filtering and blacklisting to protect the ("heavy") processing step (next), and supports load-balancing, scalability and separation of authenticated vs. unauthenticated clients. Since this step is "light", it is hard to attack or to disrupt already-connected clients.

3. Server filtering: allow only valid server IP, port pairs by validating and processing requests to the web server. Validation is performed after the TCP handshake has completed correctly, hence unlikely to be spoofed, and failure of validation is a likely indication that the IP address is controlled by the adversary, resulting in its blacklisting. The validation checks include: (1) cookie (from proxy), (2) match between the server IP and port, and the client's IP address; (3) match between the client's IP address and the TTL value; (4) the number of connections (SYN packets) and of out-of-order packets does not exceed threshold.

We show that Stratum Filtering does not harm the performance of legitimate clients, while effectively filtering and dropping traffic elicited by the malicious hosts to the protected web servers.

Keywords: Denial of service attacks; defenses against DDoS; cloud platforms; IP spoofing

Short Bio

Michael Waidner is the Director of the Fraunhofer Institute for Secure Information Technology (Fraunhofer SIT) in Darmstadt and Chair Professor for Security in IT at the Technische Universität Darmstadt. He is also the Speaker of the Center for Research in Security and Privacy (CRISP), which is funded by the German federal and state governments, and he is a Director of the Fraunhofer Project Center for Cybersecurity at the Hebrew University of Jerusalem in Israel (Fraunhofer SIT/IL)

CCSW'16, October 28, 2016, Vienna, Austria.
ACM ISBN 978-1-4503-4572-9/16/10.
DOI: http://dx.doi.org/10.1145/2996429.2996440