ITS: A DDoS Mitigating Architecture

Hikmat Farhat
Notre Dame University
Lebanon

Abstract- We propose a DDoS mitigation architecture that protects legitimate traffic from the large volume of malicious packets during a DDoS bandwidth attack. The system keeps a legitimacy list and gives higher priority to those packets that are on the list. The legitimacy list is kept up to date by keeping only the entries that complete the TCP three-way handshake and thus defeats IP spoofing. Entries in the list contain the IP address and the path signature of active TCP connections. A packet obtains high priority if its path signature strongly correlates with the corresponding path signature stored in the legitimacy list. We show that the scheme is efficient when deployed incrementally by using priority queuing at perimeter routers. An autonomous system (AS) can immediately benefit from our proposed system when deployed even if other ASs do not deploy it.

I. INTRODUCTION

While the open nature of the Internet is key to its success, it is also the main reason to its vulnerabilities. One of the most serious abuses of the Internet is the Denial of Service Attack (DoS) and its distributed version, the Distributed Denial of Service attacks (DDoS). Some of the DoS attacks received wide publicity, but, DoS attacks are in fact much more frequent [10]. Due to the relative simplicity of initiating them, bandwidth DoS attacks, in particular, have proven a difficult nut to crack. Today many services and critical infrastructure depend on the Internet, therefore protecting them from DoS attacks has become a crucial objective.

Protecting systems from bandwidth DDoS attacks has been an elusive goal. Detecting that a system or a network is under attack is in itself a difficult task. Even when we know that an attack is under way it is still a challenge to distinguish legitimate packets from malicious ones. While some DoS attacks use *spoofed* source IP addresses others use the legitimate address of compromised hosts, called *zombies*, to launch the attack. Because many attacks use IP spoofing, the determination of the source of the attack is an important step forward, albeit not enough. Toward this end many traceback schemes have been proposed [23,16,12,15]. A good number of DoS mitigation methods use a two-step approach. The first step consists of a learning phase (e.g. reconstruction of the attack paths) is needed to identify attack traffic signatures. In the second step filtering rules obtained from the "learning phase" are used to drop malicious traffic. The efficiency of the approach depends on the efficiency of the "learning" phase. Furthermore, traceback techniques usually take long time when the number of attackers is large and therefore cannot be used as a real-time response to attacks. In addition, the filters installed after the learning phase often block legitimate traffic as well leading to *collateral damage* which in itself is a form of DoS.

In a previous work [5] we have proposed the Implicit Token Scheme (ITS) to mitigate IP *spoofing*. The proposed method, while promising, had two shortcomings: it requires the estimate of the number of hops a packet has traveled, and could not be deployed incrementally, which is a serious shortcoming. We fix the above mentioned shortcomings in this paper by using a priority queuing at the perimeter routers of an ISP with packets having "better" signature matches given higher priority. This is made possible by the way the path signature is recorded. The same method obviates the need for the estimation of the number of hops.

The rest of the paper is organized as follows. Previous works on the DDoS problem are introduced in Section II. In Section III we present the objectives and assumptions of a DDoS mitigating architecture. Details of our proposed scheme are discussed in Section IV. The performed simulations and their results are shown in Section V. We conclude and provide pointer for future research in Section VI.

II. RELATED WORK

Most of the existing DDoS defense schemes are reactive in nature. The defense system becomes active when an attack is detected. Ideally, one would like to halt the attack and also determine the attacking host(s). One method, IP traceback, has focused on allowing the victim to trace the origin of malicious packets, which are usually spoofed [12,15,16,4,14]. The traceback methods require routers to stamp, with a certain probability, a mark in the IP header. When enough such packets have been collected the victim starts the process of reconstructing the path(s) that the attack packets have followed. While traceback schemes are important in finding the location of the attackers, they suffer from two shortcomings. First, the cost of the reconstruction algorithm becomes prohibitive when the number of attacking hosts is large. Second, most traceback approaches do not specify how (Ref [16] is a notable exception) to mitigate the DDoS attack once path reconstruction is completed. Many approaches have been proposed to solve the above-mentioned problems. Yaar et al. [23] have proposed a traceback algorithm that scale to thousand of attackers. Sung and Xu [16] use the concept of sub-channels to preferentially filter attack packets using the reconstructed paths.

Other approaches attempted to defend against DDoS attacks by filtering out IP packets with spoofed source addresses. One of the earliest such methods was the work by Ferguson et al. [6].

T. Sobh (ed.), Innovations and Advanced Techniques in Computer and Information Sciences and Engineering, 537–541.

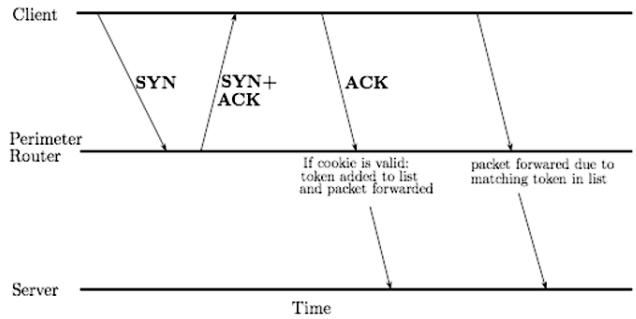

Fig. 1. Typical session of a legitimate client. Note that the SYN+ACK is received only by non-spoofed clients.

Their method requires the installation of ingress filtering at every ISP. Even if ingress filtering is universally deployed, an unrealistic assumption, IP addresses in the local network can still be spoofed. Another approach to ingress filtering is the SAVE protocol proposed by Li. et al. [9]. The distributed packet filtering method proposed by Park and Lee [11] discards spoofed IP packets using a route-based detection method even if only 20% of autonomous systems install such filters. Unfortunately their method requires the cooperation of thousands of autonomous systems. Jin [7] proposed a scheme where packets with spoofed addresses are identified by their hop count. The idea is to build a table in "peace time" that maps the source address of clients to the number of hops they need to reach the victim. During a DoS attack each packet is compared to the corresponding entry in the table. In practice, routes change every frequently, which makes the table entries obsolete very quickly.

The original idea for deterministic path identification is due to Yaar et al. [21]. They use the path identification which is a deterministic, as opposed to the probabilistic one used by most IP traceback methods, mark stamped by the intermediate routers on every packet as a way to distinguish malicious from legitimate users. Even if one assumes that the malicious signatures can be clearly identified the number of malicious and legitimate users having the same signature grows as the number of attackers grows, which quickly leads to self-inflicted DoS.

Anderson et al. [1] introduced the concept of capabilities whereby a sender first obtains a permission to send to the receiver from a Request-To-Send server (RTS) whose addresses are advertised by BGP as a community attribute. Yaar et al. [22] extended the idea where the sender obtains the permission explicitly from the receiver via a handshake protocol. Both method require routers to compute per packet hash functions and are vulnerable to attackers colluding with hosts co-located with the victim. Furthermore, the capabilities approach is vulnerable to the denial of capabilities attack [24]. A similar approach was proposed by Xu [20] to sustain the availability of Web servers and it uses HTTP redirect requests to prevent spoofed packets from reaching the victim.

III. Assumptions and Design Objectives

Approaches based on building filters for malicious traffic will inevitably drop legitimate traffic along with illegitimate traffic, as the two cannot be totally distinguished from each other. ITS instead uses the opposite approach: it drops any traffic that has not proven itself to be genuine.

In building the Implicit Token Scheme we were guided by the following design objectives:

1. Require minimal or no changes to Internet protocols. Particular attention should be given to the feasibility of deployment.
2. The implemented method should achieve zero false positives otherwise it might lead to self-inflicted DoS.
3. The algorithms implemented by the routers for the common case should be simple and fast. This object is important for two reasons
   - Any change requiring significant per-packet computation by routers is unlikely to be accepted.
   - Intensive computation leads to slower router performance and might even lead to a DoS if the extra computation is slower than a table look up.

IV. The ITS Architecture

In the ITS architecture an ISP provides filtering service for its customers at its Point of Presence (POP). The filtering is done by maintaining a legitimacy list composed of *tokens* for all active TCP connections. An entry in the list is added only when a client completes the TCP handshake thus making sure that the entry is legitimate and up to date. Every packet carries a token, composed of the source IP address and path signature. The path signature is build by having intermediate routers stamp their mark in the identification field in the IP header. When a packet arrives at the perimeter router, its token is compared to tokens in the legitimacy list. The result of this comparison decides the fate of the packet.