Computer Communications 46 (2014) 10–21

Contents lists available at ScienceDirect

Computer Communications

journal homepage: www.elsevier.com/locate/comcom

A moving target DDoS defense mechanism

Huangxin Wang *, Quan Jia, Dan Fleck, Walter Powell, Fei Li, Angelos Stavrou

Department of Computer Science, George Mason University, Fairfax, VA 22030, USA

ARTICLE INFO

Article history:
Available online 20 March 2014

Keywords:
DDoS
Moving target defense
Secret proxy
Insider
Shuffling

ABSTRACT

In this paper, we introduce a moving target defense mechanism that defends authenticated clients against Internet service DDoS attacks. Our mechanism employs a group of dynamic, hidden proxies to relay traffic between authenticated clients and servers. By continuously replacing attacked proxies with backup proxies and reassigning (*shuffling*) the attacked clients onto the new proxies, innocent clients are segregated from malicious insiders through a series of shuffles. To accelerate the process of insider segregation, we designed an efficient greedy algorithm which is proven to have near optimal empirical performance. In addition, the insider quarantine capability of this greedy algorithm is studied and quantified to enable defenders to estimate the resource required to defend against DDoS attacks and meet defined QoS levels under various attack scenarios. Simulations were then performed which confirmed the theoretical results and showed that our mechanism is effective in mitigating the effects of a DDoS attack. The simulations also demonstrated that the overhead introduced by the shuffling procedure is low.

1. Introduction

Distributed denial-of-service (DDoS) attacks are a rapidly growing problem which poses an immense threat to the Internet. Arbor Networks reported a significant increase in the prevalence of large-scale distributed denial-of-service (DDoS) attacks in recent years [1]. In 2010, the largest reported bandwidth achieved by a flood-based DDoS attack reached 100 Gbps. Even as the bandwidth of attacks has increased, the cost of performing a DDoS attack has turned out to be surprisingly low. A Trend Micro white paper [2] reported that the price for a 1-week DDoS attack could be as low as $150 on the Russian underground market.

A number of mechanisms have been proposed in the past to prevent or mitigate the impact of DDoS attacks. Filtering-based approaches [3–5] use ubiquitously deployed filters that block unwanted traffic sent to the protected nodes. Capability-based defense mechanisms [6–9] endeavor to constrain resource usage by senders to beneath a threshold defined by the defended system. Secure overlay solutions [10–15] interpose a network of proxy nodes that redirect packets between clients and the protected nodes and are designed to absorb and filter out attack traffic. All these mechanisms are effective to varying degrees; but these static defense mechanisms either rely on the global deployment of additional functionalities on Internet routers or require large, robust, virtual networks designed to withstand the ever-larger attacks. Due to the large investment required and the vulnerability to sophisticated attacks such as sweeping [11] and adaptive flooding attacks [12], the development of novel, effective, efficient, and low cost defense mechanisms continues to be a high priority, but elusive goal.

Motivated by the aforementioned elusive goal, we propose *MOTAG* [16], a MOving Target defense mechanism AGainst Internet DDoS attacks. This dynamic DDoS defense mechanism implements a scheme of moving proxy nodes to protect centralized online services. In particular, *MOTAG* offers DDoS resilience for authenticated clients of security sensitive services such as online banking and e-commerce. *MOTAG* employs a layer of secret moving proxy nodes to relay communications between clients and the protected application servers.

The proxy nodes in *MOTAG* have two important characteristics. First, the proxy nodes are "secret" in that their IP addresses are concealed from the general public and are exclusively known only to legitimate clients and only after successful authentication. In order to avoid unnecessary information leakage, each authenticated client is provided with the IP address of only a single proxy node at any given time. Existing Proof-of-Work (PoW) schemes [17–20] are employed to protect the client authentication channel. Second, the proxy nodes are "moving". As soon as an active proxy node is attacked, it is replaced by a set of alternate proxy nodes instantiated at a different IP address; and the clients associated with the attacked proxy node are migrated to alternative proxy

* Corresponding author. Tel.: +1 7039019768.
E-mail addresses: hwang14@gmu.edu (H. Wang), qjia@gmu.edu (Q. Jia), dfleck@gmu.edu (D. Fleck), wpowell@gmu.edu (W. Powell), fli4@gmu.edu (F. Li), astavrou@gmu.edu (A. Stavrou).

http://dx.doi.org/10.1016/j.comcom.2014.03.009
0140-3664/© 2014 Elsevier B.V. All rights reserved.

node(s). We show that this migration to “secret” proxy nodes not only enables the *MOTAG* mechanism to mitigate brute-force DDoS attacks, but also provides a means to discover and isolate malicious insiders designed to divulge the location of the secret proxy nodes to external attackers. The malicious insiders are isolated via a shuffling process that reassigns and migrates clients through sequential sets of instantiated of proxy nodes. This paper presents the algorithms developed to (1) accurately estimate the number of insiders and (2) to dynamically determine client-to-proxy assignment that will “save” the largest number of legitimate clients after each shuffle.

Unlike previously proposed DDoS defense mechanisms, *MOTAG* does not rely on global adoption on Internet routers or collaboration across different ISPs to function. Also *MOTAG* neither depends on resource-abundant overlay networks to out-muscle high bandwidth attacks nor uses filters to provide fault tolerance. Instead, we take advantage of our proxy nodes’ secrecy and mobility to fend off powerful DDoS attacks including sweeping and adaptive flooding attacks. Employing the *MOTAG* DDoS defense mechanism requires lower deployment costs while offering substantial defensive agility which results in effective and cost-efficient DDoS protection.

This paper is an extension of the work we presented in [16]. The main contributions of this paper are:

- A discussion of the reduction of the computational complexity of greedy algorithm from $O(N \cdot N_i)$ to $O(1)$ where N is the number of total clients and N_i is the number of insiders.
- A theoretical and empirical analysis of the insider quarantine capability of the greedy algorithm.
- A discussion of a special DDoS attack case for which a simple and elegant shuffling mechanism is designed to segregate innocent clients from insiders.

2. Threat model and assumptions

Instead of targeting open and general-purpose web services, we focused on protecting security sensitive online services against *network flooding attacks*. We assume that legitimate clients of the protected services are pre-authorized and their identities can be authenticated before they are served. We assume the availability of a cloud environment with sufficient computing power and bandwidth to instantiate numerous backup proxy nodes. Since only a small group of proxy nodes are active at any time, a cloud environment in which customers are charged only for running instances would be ideal to avoid extensive operational costs. We further assume that although powerful attackers with a high aggregate bandwidth are capable of simultaneously overwhelming many stand-alone machines on the Internet, attackers cannot saturate the well-provisioned Internet backbone links of ISPs, data centers, and cloud service providers.

We also assume that attackers, in case of uncertainty, can first perform reconnaissance attacks (e.g., IP and port scanning) to pinpoint targets for the subsequent flooding attacks. With knowledge of the *MOTAG* mechanism, attackers could attempt to flood the authentication channel through which the legitimate clients are admitted. Successful attacks against the authentication server are considered unlikely because it employs proof-of-work (PoW) schemes to prevent both computational and network flooding attack. *MOTAG* takes advantage of PoW schemes that are designed to prevent computational attacks. In a cloud environment, the ability of lightweight PoW schemes to quickly reject non-authentication requests makes it resistant to flooding attacks. Since it is significantly harder for attackers to pass strong authentication by brute force and reach the proxy nodes as legitimate clients, some attackers will attempt to uncover the network locations of proxy nodes and may plant ”insiders” by compromising legitimate clients or eavesdropping on legitimate clients’ network connections. However, the number of such insiders in a protected system is assumed to be limited.

3. MOTAG architecture

The proposed *MOTAG* mechanism employs a group of dynamic proxy nodes that relay traffic between servers and authenticated clients and that provide a moving target defense mechanism that mitigates Internet service DDoS attacks. The IP address of the proxy nodes are hidden from clients (and potential attackers), and each client can only see the IP address of the proxy node to which he is randomly assigned. Therefore, insiders will only be able to attack the proxy node(s) to which they are assigned, and the innocent clients who are impacted by the attack will be only those who share the proxy node(s) with the insiders. In order to separate affected innocent clients from insiders, *MOTAG* instantiates proxy nodes and performs client-to-proxy reassignment under the guidance of the shuffling algorithm discussed in Sections 4 and 5. In the following sections, we first give an overview of *MOTAG*, and then introduce the main components in greater detail.

3.1. MOTAG overview

Fig. 1 shows the overall architecture of *MOTAG* which consists of four inter-connected components: the authentication server, the proxy nodes, the filter ring, and the application server. The application server provides the online services (e.g., banking or e-commerce services) that we want to protect and make accessible to authenticated clients. The IP address of the application server is concealed from all clients and all traffic is relayed to the application server via the proxy nodes. The filter ring, similar to what was described in [12], is comprised of a number of high speed routers placed around the application server which allows inbound traffic only from valid proxy nodes. The proxy nodes are a group of dynamic and distributed cloud instantiations that relay communications between clients and the application server. The authentication server is responsible for authenticating clients, assigning legitimate ones to individual proxy nodes, and coordinating the shuffling of clients.

MOTAG allows a client to access the application server via a proxy node only if the client can be successfully authenticated. One simple solution is to associate the application domain name with the IP address of the authentication server during DNS registration. Each successfully authenticated client is then randomly

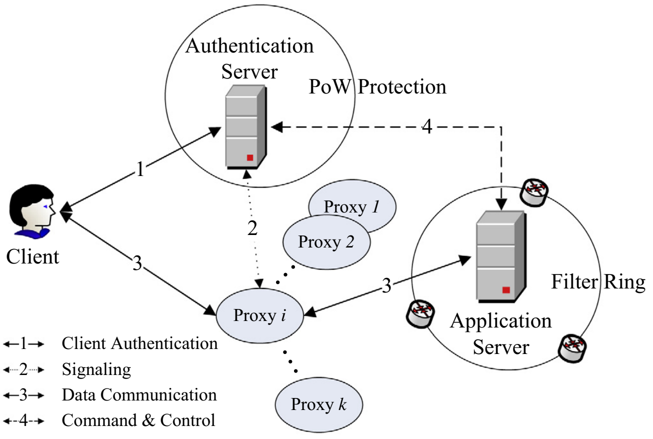

Fig. 1. Overview of the MOTAG Architecture.