Cloud-Base Defense against DRDoS Attacks

H. Fujinoki

Southern Illinois University Edwardsville, Edwardsville, IL. USA

***Abstract*--In this paper, we present our new approach for protecting production servers from distributed reflector denial of service attacks, while the volume of their attacking traffic has been exponentially increasing. We apply a cloud to effectively sieving out the reflectors of the attacking traffic. We assessed the convergence delay for sieving using simulations. Our simulation results indicated that the proposed protection method promising, showing the approach is scalable in both attacking traffic volume and attacker ratio.**

I. INTRODUCTION

Most of the known DDoS (Distributed Denial of Service) attacks these days use intermediate host computers as bots or reflectors these days [1, 2]. The DDoS attacks that especially utilize reflectors are called "Distributed Reflector Denial of Service (DRDoS) attacks" [3, 4, 5]. This paper presents a solution that minimizes and mitigates the damages from DRDoS attacks using cloud-base infrastructure, as well as the results of our performance assessments.

II. PROPOSED CLOUD-BASE DEFENSE

The proposed solution mitigates DRDoS attacks using cloud environments. It consists of four different types of servers: production servers, temporary mirror servers, shadow production servers, and the scrum master, which are set up in a cloud environment. Its procedure consists of three phases of attack detection, sieving out DRDoS reflectors, and evacuations of legitimate traffic sources to shadow production servers, as described below:

Phase 1: each production server constantly monitors the incoming traffic load by monitoring pps (packet per second) and bps (bits per second). When the incoming traffic load reaches the threshold of the caution level, a production server sends the list of the network addresses (i.e., IP addresses) of the currently active traffic sources (including both DRDoS reflectors and legitimate sources) to the scrum master. The network address of the scrum master is known only to production servers, temporally mirror servers, and shadow production servers.

Phase 2: When the incoming traffic load exceeds the limit of "caution level", the production server notifies its "warning" status to the scrum master. On the arrival of a warning notification from a production server, the scrum master initiates the process of sieving out DRDoS reflectors. First, the scrum master activates two temporary miller servers in a cloud, each of which is running at a different host (each of which is assigned a different IP address). Then, the scrum master randomly splits the currently active traffic sources to the two new mirror servers. The scrum master notifies the production server the network addresses of the two temporary mirror servers. The production server then notifies each of the active traffic sources the address of one of the two new temporary mirror servers. After the active traffic sources are split to the two temporary mirror servers, then each mirror server monitors the incoming traffic load (pps and bps), just like the production server does.

Phase 3: Each temporary mirror server recursively executes Phase 1 and 2. This process eventually allows the scrum master to identify the network address of DRDoS reflectors as shown in Figure 1. After the network addresses of DRDoS reflectors are identified, the scrum master announces each of the legitimate traffic sources the network address of a shadow production server. Figure 2 visualizes the overview of the entire defense system.

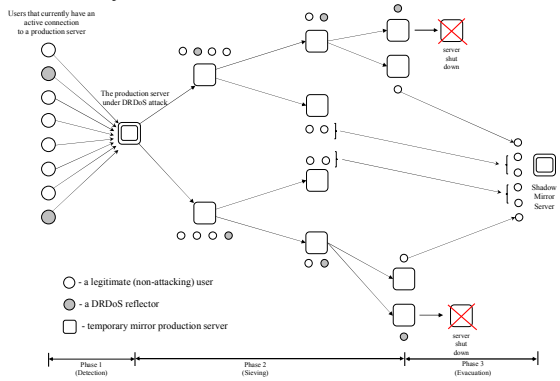

Figure 1 - Visualization of DRDoS reflector sieving at a production server using a cloud environment

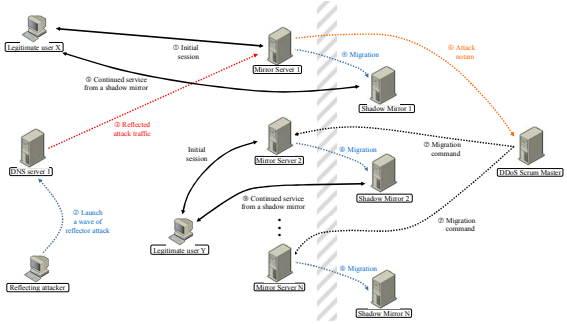

Figure 2 - overview of the proposed DRDoS defense system

III. PERFORMANCE EVALUATION

We assessed the expected performance of the proposed solution using simulations. The metrics we focused on are: (i) the number of concurrent temporary mirror servers necessary to complete the sieving process and (ii) the elapsed time since the start of a round of DRDoS attack to a production server until all the DRDoS reflectors are successfully identified.

Figure 3 shows the results when DRDoS attacking traffic from 128 distinct DRDoS reflectors is submitted to 16 production servers while the 16 production servers deal with

896 legitimate traffic sources. Our simulations assumed 5 seconds to detect incoming DRDoS attack after a wave of attack reaches a server. The overhead time for splitting active traffic sources to two temporary mirror servers is ignored in our experiments this time. Figure 4 shows the results when only the number of the initial production servers was reduced from 16 to 4 (i.e., there are 1024 traffic sources to the four production servers). In Figure 3 and 4, the majority of the DRDoS reflectors are successfully identifies (and eliminated) approximately in 600 seconds, with 6.4 and 10.9 temporary mirror servers are used for each production server, respectively.

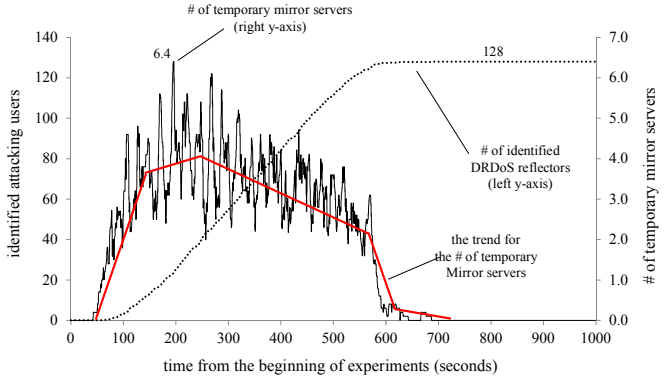

Figure 3 - convergence delay and number of the temporary mirror servers used

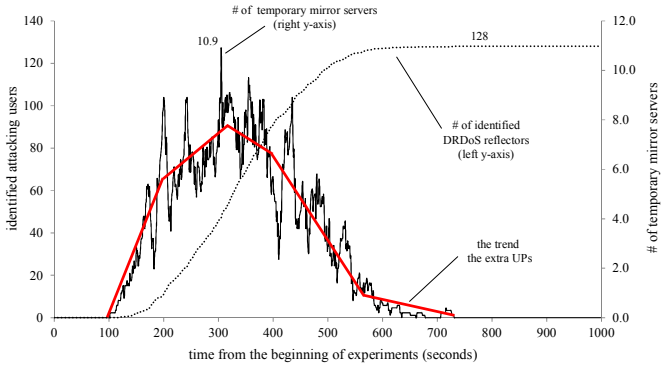

Figure 4 – convergence delay when four initial production servers under attack

Figure 5 shows the convergence delay when only the number of traffic sources are changed to 512, 1024, 2048, 4096, and 8192. Figure 6 shows the convergence delay when the ratio of attacking:legitimate traffic sources is 1/16, 1/4, and 1/2 (instead of 1/8).

In Figure 5 and 6, the convergence delay takes a shape of parabola curve. The observed results indicate that the proposed cloud-based method is capable of identifying DRDoS reflectors and that it is scalable in both the total incoming traffic load and attacker ratio.

IV. CONCLUSIONS AND FUTURE WORK

In this paper, we presented a new approach for protecting network servers from DRDoS attacks. Our simulation results indicate promising results, showing the effectiveness in the solution. While a single-point protection is known to be ineffective to DRDoS attacks [6, 7, 8], sieving out DRDoS reflectors using a cloud environment can be one of the main streams for server protections from DRDoS threats. We are currently testing the proposed cloud-base protection for more different scenarios using experiments in a prototype testbed.

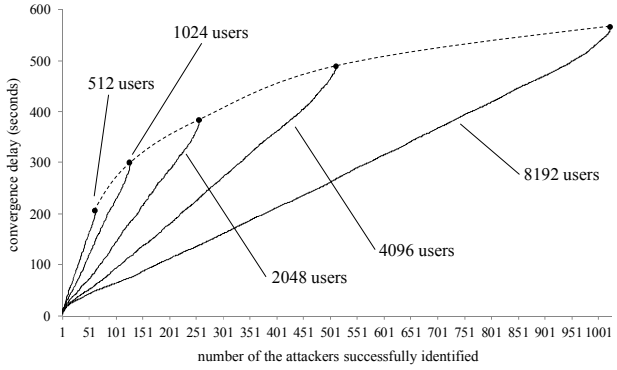

Figure 5 - convergence delay for different number of users

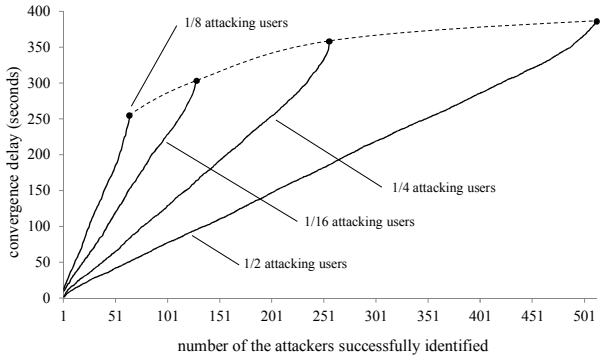

Figure 6 – convergence delay for different attack/legitimate connection ratio

REFERENCE

[1] Q. Yan, F. R. Yu, Q. Gong, J. Li, "Software-Defined Networking and Distributed Denial of Service Attacks in Cloud Computing Environments: A Survey, Some Research Issues, and Challenges," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 1, pp. 602-622, 2016.

[2] J. Czyz , M. Kallitsis, M. Gharaibeh, C. Papadopoulos, M. Bailey, and M. Karir, "Taming the 800 Pound Gorilla: The Rise and Decline of NTP DDoS Attacks," *Proceedings of the Conference on Internet Measurement*, pp. 435-448, 2014.

[3] V. Paxson, "An Analysis of Using Reflectors for Distributed Denial-of-Service Attacks," SIGCOMM Computer Communication Review, vol. 31, no. 3, pp. 38–47, 2001.

[4] C. Fachkha, E. Bou-Harb, and M. Debbabi, "Inferring Distributed Reflection Denial of Service Attacks from Darknet," *Computer Communications*, vol. 62, no. 15, pp. 59-71, 2015.

[5] M. Kuhrer, T. Hupperich, C. Rossow, and T. Holz, "Exit from Hell? Reducing the Impact of Amplification DDoS Attacks," the Proceedings of the USENIX Security Conference, pp. 111-125, 2014.

[6] A. D. Olaniyi, R. Christoph, S. Adesina Simon, A. Adio Taofeek, and B. S. Badmus, "Resolving DRDoS Attack in Cloud Database Service Using Common Source IP and Incremental Replacement Strategy," *Proceedings of SAI Intelligent Systems Conference*, pp. 725-737, 2016.

[7] Q. Yan and F. R. Yu, "Distributed Denial of Service Attacks in Software-Defined Networking with Cloud Computing," *IEEE Communications Magazine*, vol. 53, no. 4, pp. 52-59, 2015.

[8] R. Rijswijk-Deij, A. Sperotto, and A. Pras, "DNSSEC and its Potential for DDoS Attacks: a Comprehensive Measurement Study," *Proceedings of the Internet Measurement Conference*, pp. 449-460, 2014.