Discovering Attack Propagation Patterns in Honeypots

Ariel Bar, Bracha Shapira, Lior Rokach and Moshe Unger

Department of Information Systems Engineering and
Telekom Innovation Laboratories
Ben-Gurion University of the Negev Beer-Sheva, Israel
{arielba, bshapira, liorrk, mosheun}@bgu.ac.il

During the last decade, both the magnitude and sophistication of cyber-attacks are substantially increasing. Botnets, worms, malware, viruses and denial of service attempts are required to be recorded, investigated and analyzed in order to cope with these growing threats. Honeypots are computer resources that are used to detect and deflect attacks on a protected system, while their value lies in being probed, attacked or compromised. Honeypot frameworks (a set of honeypots which are monitored together) are designed to attract attackers by presenting false or misleading information that an attacker will want to gain, attack or control. By definition, honeypots are supposed to have zero false alarms, since no legitimate access to them is expected. Analyzing the captured data from honeypots may reveal new attack patterns or emerging threats, and aid security administrators or designers to support decision making, and improve defense mechanisms.

In this work we present a novel approach for detecting attack propagation patterns within a honeypot framework. A propagation pattern among two honeypots is assumed to occur in situations where the same attacker attacks several honeypots in a sequential order [1]. Examples of such situations might occur as a result of a scanning activity [2, 3] or from the propagation of worms [4, 5]. Modeling these kinds of phenomena reveals information about attack trends which span over multiple honeypots, rather than traditional analysis which focuses on the individual honeypots separately.

The suggested method aims at finding these kinds of sequential patterns, and in addition tries to expose interaction patterns between the different honeypots in the framework. To the best of our knowledge, the suggested model is the first that is capable of addressing the following analysis questions:

- What is the probability distribution of the next attacked honeypot in a session?
- Which honeypots are attacked together?
- Which attack propagation patterns are new and evolving ?

- Which honeypots contribute the most to the attack propagation trends?
- Are there any contextual differences between attack propagations? (E.g. Do attack propagation patterns vary by the source of the attacking country?).

In practice, we suggest a modeling process which contains two phases: First, we train a Markov Chains model [6] in order to discover sequential attack patterns in the data. Second, we use the graph structure of the trained Markov Chains model and apply algorithms from the Complex Networks [7] research area in order to discover various interactions between the honeypots. The resulting model includes three components: a) a Markov Chains model which holds the sequential attack propagation probabilistic distributions, b) Segmentation of all honeypots into communities, and c) Ranking scores for each honeypot, describing the relative importance level of each of them in the context of enabling attack propagation.

In order to evaluate the proposed method, we analyzed a massive real-world dataset of attack records on a honeypot framework. The analyzed dataset was collected via the T-Pot Multi-Honeypot Platform – a globally deployed honeypot framework which contains multiple well-established honeypot technologies. T-Pot emulate various types of network protocols and services, including, HTTP, SSH, SMB, FTP, MSSQL and VOIP. Overall, we analyzed over 160 million attacks on the framework between July 2014 and August 2015. The attacks targeted 267 distinct honeypots deployed around the globe. The attacks originated from more than 900,000 distinct IP addresses that span over thousands of different ISPs in more than 200 countries around the world.

Results indicate that the analyzed honeypot framework can be partitioned into several segments as presented in figure 1. Each node represents a honeypot, while the edges between the nodes represent a propagation pattern between the matching honeypots. Overall, we were able to detect six well defined honeypot communities. Each community was assigned with a matching label and color. Finally, the size of each node is proportional to the "Betweenness Centrality"[8] measure which quantifies the number of times a node acts as a bridge on the shortest path between two other nodes in the graph.

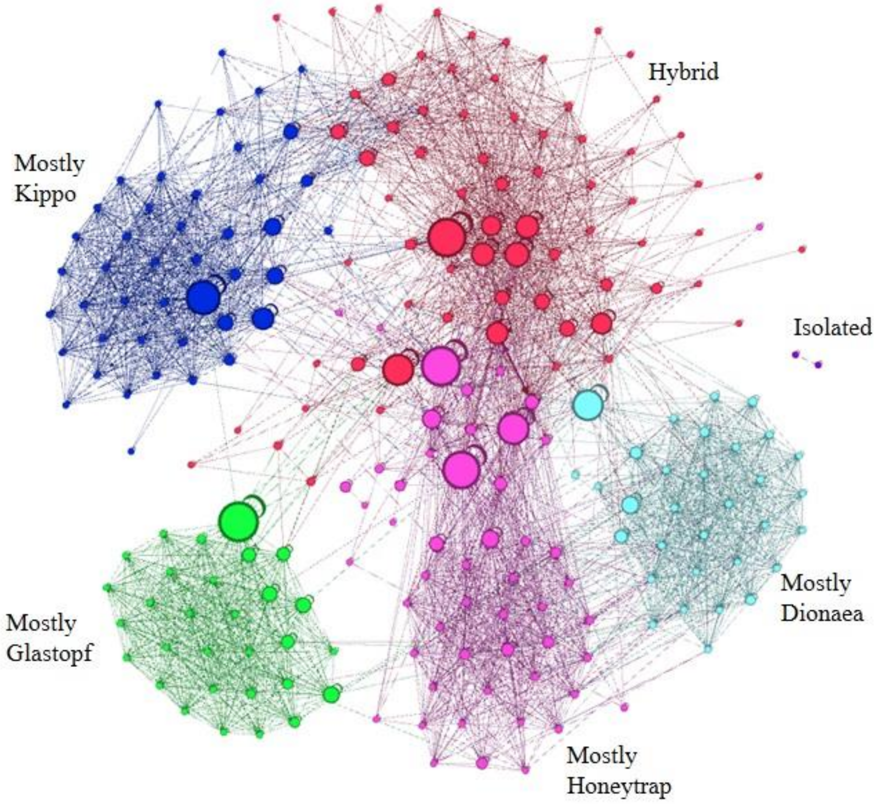

Figure 1: Honeypot Propagation Graph

We also observed that attack propagations within each segment are very common, while propagations between different segments are rarer. We believe that propagations between segments are more interesting since each segment empirically had its own characteristics. Hence, these kinds of attacks are expected to be more sophisticated. Moreover, we were able to detect the top honeypots which contribute the most to attack propagation. Thus, security administrators can replicate their configuration in order to increase the magnitude of attack propagation in the system and by doing so, exposing more information about the attackers. Finally, our analysis discovered that attack propagation patterns may vary according to the origin of the attacking country.

References:

[1] M. Kaaniche, Y. Deswarte, E. Alata, M. Dacier, and V. Nicomette, “Empirical analysis and statistical modeling of attack processes based on honeypots,” in Workshop on Empirical Evaluation of Dependability and Security (WEEDS), Philadelphia, USA, June 2006, pp. 119–124.

[2] Goseva-Popstojanova, Katerina, et al. "Characterization and classification of malicious Web traffic." Computers & Security 42 (2014): 92-115.

[3] Goseva-Popstojanova, Katerina, et al. "Quantification of attackers activities on servers running Web 2.0 applications." Network Computing and Applications (NCA), 2010 9th IEEE International Symposium on. IEEE, 2010.

[4] Newsome, James, Brad Karp, and Dawn Song. "Polygraph: Automatically generating signatures for polymorphic worms." Security and Privacy, 2005 IEEE Symposium on. IEEE, 2005.

[5] Su, Ming-Yang. "Internet worms identification through serial episodes mining." Electrical Engineering/Electronics Computer Telecommunications and Information Technology (ECTI-CON), 2010 International Conference on. IEEE, 2010.

[6] Grinstead, C. M., & Snell, J. L. (2012). Introduction to probability. American Mathematical Soc.

[7] Boccaletti, Stefano, et al. "Complex networks: Structure and dynamics." Physics reports 424.4 (2006): 175-308.

[8] Brandes, Ulrik. "A faster algorithm for betweenness centrality*." Journal of mathematical sociology 25.2 (2001): 163-177.