Towards Reconstructing Multi-Step Cyber Attacks in Modern Cloud Environments with Tripwires

Mario Kahlhofer
mario.kahlhofer@dynatrace.com
Dynatrace Research
Linz, Austria

Michael Hölzl
michael.hoelzl@dynatrace.com
Dynatrace Research
Linz, Austria

Andreas Berger
andreas.berger@dynatrace.com
Dynatrace Research
Linz, Austria

ABSTRACT

Rapidly-changing cloud environments that consist of heavily interconnected components are difficult to secure. Existing solutions often try to correlate many weak indicators to identify and reconstruct multi-step cyber attacks. The lack of a true, causal link between most of these indicators still leaves administrators with a lot of false-positives to browse through. We argue that cyber deception can improve the precision of attack detection systems, if used in a structured, and automatic way, i.e., in the form of so-called tripwires that ultimately span an attack graph, which assists attack reconstruction algorithms. This paper proposes an idea for a framework that combines cyber deception, automatic tripwire injection and attack graphs, which eventually enables us to reconstruct multi-step cyber attacks in modern cloud environments.

CCS CONCEPTS

• **Security and privacy → Intrusion/anomaly detection and malware mitigation**; *Web application security*; Network security.

KEYWORDS

multi-step cyber attacks, intrusion detection, cyber deception, cyber kill-chain, attack graphs, honeypots, honeytokens, tripwires

1 INTRODUCTION

Both academia and industry continue to work on systems that detect cyber attacks. Promising results have been achieved for identifying vulnerabilities, malware, and malicious behavior, but, those systems still struggle to distinguish between mere anomalies and security-relevant incidents [4, 9]. Recent work focuses on correlating security alarms from various sources, to detect advanced threats that consist of multiple steps and evolve over long time periods [6]. Especially in the domain of enterprise networks, such attacks are popularly known as Advanced Persistent Threats (APTs) [1] and are

often modelled by a cyber kill-chain [2]. Since cloud threats cannot be represented well with this model we instead use the general term multi-step cyber attack.

Although these approaches yield more interpretable alarms, they still have troubles filtering out false-positives when applied in the real world [9]. Furthermore, in the reconstruction of multi-step attacks, common approaches correlate IP addresses, alarm types, or time windows [6], however, no approach fundamentally knows the true, causal link between alarms.

Securing cloud environments that consist of heavily interconnected components is particularly difficult. Adversaries profit from the broadened attack surface and increased number of attack vectors. Cyber deception, e.g., deploying honeypots, which are purposefully vulnerable entities that should attract adversaries, counteracts their unfair advantage by slowing them down [8]. However, honeypots are often configured and deployed manually, making them scale poorly with rapidly-changing cloud environments.

We focus on three aspects to progress on these open issues:

(1) Cyber Deception. To improve the precision of detection methods, we argue that it is important to have strong indicators of compromise (IoCs), instead of correlating many weak indicators. To achieve this, we introduce tripwires, which describe a deceptive scenario, and combine honeypots and honeytokens [10].

(2) Automatic Tripwire Injection. To address the complex nature of cloud environments, tripwires are automatically, and strategically injected into existing applications. Our framework manages their deployment and reacts to changes in the cloud environment.

(3) Attack Graph Reconstruction. A tripwire consists of multiple, related deceptive components, e.g., a private SSH key that acts as the lure, and the SSH server that is the decoy. Thereby, tripwires naturally form an attack graph [7], e.g., if the decoy is accessed, we know that the adversary must have gotten the keys first, which we deliberately placed on a different entity to trace the attack path. Incoming alarms can then be projected onto the attack graph to reconstruct multi-step attacks.

To this end, we ask the following research question: *"Are automatically injected tripwires suitable to reconstruct multi-step cyber attacks in modern cloud environments?"*. This paper proposes an idea for a framework that can answer this question.

2 FRAMEWORK

Our framework is designed for cloud environments, i.e., applications that are primarily exposed on the public internet. Applications have in common that they often use libraries for common use cases such as web services or database communication. We automatically inject tripwires into those libraries with the 7-component framework depicted in Figure 1. These components interact as follows:

Deploy Modules (DMs). Process hooks can identify used libraries and even inject code into applications at runtime. Thereby, we can automatically inject deceptive components into applications, e.g., additional HTTP endpoints on web applications. Likewise, a process that runs on every host can deploy honeypots, or, create and modify files. We call the component that injects lures or decoys, and optionally sets up an alarm system, a DM. The associated alarm system watches access attempts to the deployed lures or decoys.

DM Registry (DMR). After scanning the cloud environment for injectable processes and hosts, many such DMs are instantiated. Each of them registers itself in the DMR.

Tripwire Pool (TP). A tripwire is a manually-created definition of a deceptive scenario. Many instances of these are stored inside a TP. Take the AWS S3 bucket honeytoken[1] as a concrete example for the abstract model in Figure 2. There, we need at least two DMs: One creates an empty bucket (the target[2]) and monitors accesses attempts to it, while one or more other DMs deploy the access token lure. Various DMs allow to inject access token strings (e.g., into files, HTTP headers, cookies, REST endpoints).

Deployment Controller (DC). The DC first queries all available DMs from the DMR and then strategically decides which DMs to use. Its goal is to cover the entire cloud environment with deceptive components and react to re-deployments of applications.

Attack Graph (AG). The DC also populates the AG to keep track of the deployment of deceptive components and their relationships, i.e., which lure is required to access which decoy.

Alarm Store (AS). Every DM notifies the AS upon new alarms that were caught by an alarm system. While the technical implementation might be different, we imagine that the AS stores alarms in a condensed form, which makes querying the AS more efficient.

Attack Reconstruction (AR). The reconstruction algorithm takes the AS and the AG as its input, and is tasked to reconstruct attacks. Given this data, we intend to use common backward and forward tracking algorithms [3, 5] or variations of them to reconstruct multi-step cyber attacks.

3 CONCLUSION

We proposed an idea for a framework that strategically deploys deceptive components in modern cloud environments. We described both the general framework, as well as tripwires, which are a definition of managed, adaptive, automatically injected, and strategically placed deceptive components, combining lures and decoys.

We envision that automatically injected deceptive components that form an attack graph offer a way to find true, causal links between alarms, thus improving the reconstruction of multi-step cyber attacks. In future work, we plan to implement such a system and evaluate various attack reconstruction algorithms.

ACKNOWLEDGMENTS

---

[1] For an implementation, see https://bitbucket.org/asecurityteam/spacecrab

[2] The terms lure and decoy in Figure 2 solely mean an abstract concept of them. They do not represent a target. Eventually, a DM either creates a concrete instance, which then becomes a target, or, a DM injects something into an existing target (e.g., applications, files). The bucket tripwire specifies that it requires an empty bucket decoy to work. A DM with access to the AWS API will ultimately create that concrete decoy.
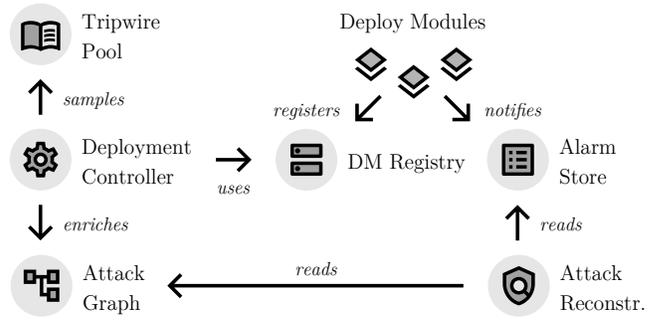
Figure 1: The framework describes the life cycle of tripwires in cloud environments, from deployment, alarm and attack graph storage, to attack reconstruction.
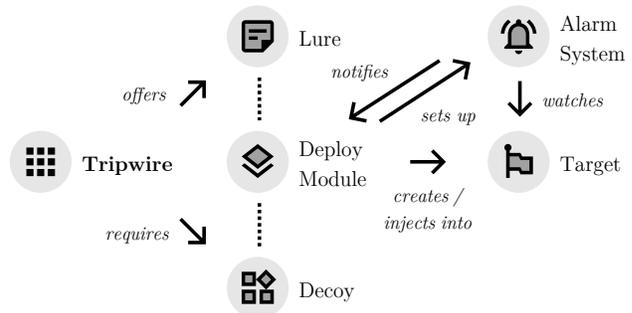
Figure 2: A tripwire describes the relation between lures, decoys, their deployment on some target via a deploy module, and its associated alarm system.

REFERENCES

[1] Ping Chen, Lieven Desmet, and Christophe Huygens. 2014. A Study on Advanced Persistent Threats. In *Communications and Multimedia Security (CMS '14) (LNCS, Vol. 8735)*. Springer, Berlin, Heidelberg, 63–72. https://doi.org/10/gg33gx

[2] Eric Hutchins, Michael Cloppert, and Rohan Amin. 2011. Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains. In *6th International Conference on Information Warfare and Security (ICIW '11)*. ACPI, Washington, DC, 113–125.

[3] Samuel T. King and Peter M. Chen. 2003. Backtracking Intrusions. In *Proceedings of the Nineteenth ACM Symposium on Operating Systems Principles (SOSP '03)*. ACM, Bolton Landing, NY, 223–236. https://doi.org/10/b47kcm

[4] Ram Shankar Siva Kumar, Andrew Wicker, and Matt Swann. 2017. Practical Machine Learning for Cloud Intrusion Detection: Challenges and the Way Forward. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security (AISec '17)*. ACM, Dallas, Texas, 81–90. https://doi.org/10/ggkqcp

[5] Yushan Liu, Mu Zhang, Ding Li, Kangkook Jee, Zhichun Li, Zhenyu Wu, Junghwan Rhee, and Prateek Mittal. 2018. Towards a Timely Causality Analysis for Enterprise Security. In *Proceedings 2018 Network and Distributed System Security Symposium (NDSS '18)*. Internet Society, San Diego, CA. https://doi.org/10/ggk6gj

[6] Julio Navarro, Aline Deruyver, and Pierre Parrend. 2018. A Systematic Survey on Multi-Step Attack Detection. *Computers & Security* 76 (July 2018), 214–249. https://doi.org/10/gdv95j

[7] Cynthia Phillips and Laura Painton Swiler. 1998. A Graph-Based System for Network-Vulnerability Analysis. In *Proc. of the 1998 Workshop on New Security Paradigms (NSPW '98)*. ACM, New York, NY, 71–79. https://doi.org/10/cxxqrd

[8] Neil C. Rowe and Julian Rrushi. 2016. *Introduction to Cyberdeception*. Springer International Publishing, Cham. https://doi.org/10/d65q

[9] Robin Sommer and Vern Paxson. 2010. Outside the Closed World: On Using Machine Learning for Network Intrusion Detection. In *2010 IEEE Symposium on Security and Privacy (S&P '10)*. IEEE Computer Society, Oakland, CA, 305–316. https://doi.org/10/cgp43q

[10] Lance Spitzner. 2003. Honeytokens: The Other Honeypot. https://www.symantec.com/connect/articles/honeytokens-other-honeypot