

Poster: Effectiveness of Moving Target Defense Techniques to Disrupt Attacks in the Cloud

Salman Manzoor
Lancaster University
United Kingdom
s.manzoor1@lancaster.ac.uk

Antonios Gouglidis
Lancaster University
United Kingdom
a.gouglidis@lancaster.ac.uk

Matthew Bradbury
Lancaster University
United Kingdom
m.s.bradbury@lancaster.ac.uk

Neeraj Suri
Lancaster University
United Kingdom
neeraj.suri@lancaster.ac.uk

ABSTRACT

Moving Target Defense (MTD) can eliminate the asymmetric advantage that attackers have in terms of time to explore a static system by changing a system's configuration dynamically to reduce the efficacy of reconnaissance and increase uncertainty and complexity for attackers. To this extent, a variety of MTDs have been proposed for specific aspects of a system. However, deploying MTDs at different layers/components of the Cloud and assessing their effects on the overall security gains for the entire system is still challenging since the Cloud is a complex system entailing physical and virtual resources, and there exists a multitude of attack surfaces that an attacker can target. Thus, we explore the combination of MTDs, and their deployment at different components (belonging to various operational layers) to maximize the security gains offered by the MTDs. We also propose a quantification mechanism to evaluate the effectiveness of the MTDs against the attacks in the Cloud.

CCS CONCEPTS

• Security and privacy → Information flow control; Distributed systems security.

1 INTRODUCTION

Current IT systems operate in a relatively static configuration which gives attackers the advantage of time, as attackers can plan, perform reconnaissance, and execute attacks without time constraints. Conventional security measures rely on patching individual vulnerabilities, which can be cumbersome, time-consuming, and risk introducing configuration errors in the system. Moreover, it is difficult (and likely impossible) for system defenders to eliminate all vulnerabilities in a system; this provides the attackers a window of opportunity to compromise the system. Consequently, Moving Target Defense (MTD) [8] techniques are advocated as a *proactive* approach to improve a system's security. The basic premise behind MTD is that introducing increased uncertainty and complexity for the attackers reduces the likelihood of them successfully exploiting the system. For example, a dynamic run-time environment could change the execution environment presented to the application [2], while MTD techniques focusing on the network layer dynamically modifies the network characteristics (e.g., IP/MAC addresses) [4] to change the attack surfaces presented to the adversary.

The previously mentioned MTDs are typically applied to individual aspects of a system (IP addresses, OS replicas, etc.). Their effectiveness is measured through their capability to mitigate attacks targeting these components. However, the application of the MTDs considering a holistic view of the Cloud is limited [1] due to the complexity of the Cloud environment. The complexity stems from the coupling of physical and virtual resources, which can be

instantiated, migrated, and decommissioned to provide elasticity to users. Furthermore, many attack surfaces exist across different layers of the Cloud [3]. For example, an attacker can utilize VMs to implement malware in the first stage of the attack [5]. During the second stage, an attacker can utilize the underlying hypervisor's vulnerabilities to target services connected with the hypervisor [7]. Consequently, deploying an individual MTD, e.g., OS diversity at the VM level or changing the hypervisor, might not be sufficient to mitigate such attacks. Therefore, a combination of MTDs at different layers is desired to maximize the security gains of the MTD techniques.

To facilitate the application and to evaluate the effectiveness of MTDs across different layers of the Cloud, we develop a framework to quantify the effectiveness of MTDs deployed across the operational stack of the Cloud. The basis for the framework is the development of a model to determine the operations involved in the Cloud. Consequently, this model forms the basis for tracking the subsequent steps at an attacker's disposal in a multi-stage attack, i.e., if the triggering source of an attack is a hypervisor, the successive options available to an attacker are the services that interact with the hypervisor. Accordingly, enumerating the services is critical to understanding attack paths and evaluating the MTD's placements along the attack path. Overall, the contributions can be summarized as follows:

- Evaluating the effectiveness of aggregated MTD techniques across different layers of the Cloud to maximize security gains, reduce the attack window or increase the attacker's cost.
- Assessing the placement of MTD application in the Cloud by considering potential actions of an attacker.
- Quantification method to determine the effectiveness of the proposed MTD approach in disrupting both single and multi-stage attacks in the Cloud.

2 MTD FRAMEWORK

The proposed MTD framework, depicted in Figure 1, comprises of three layers. The first is the set of available MTD techniques such as IP shuffling, OS diversity, etc. The second layer is the threats that MTDs *can impact*. To attack a system, an attacker must know its weaknesses and how to exploit them. Therefore, threats have primarily two main blocks. Preconditions encompass the knowledge about the weaknesses of a system, and exploitation is the process of using the weaknesses to compromise the system. It might be noted here that an MTD can affect any or both of these blocks. Finally, the third layer is the system under investigation. In our case, we show a subset of the Cloud operations since our goal is to determine the effectiveness of MTDs (or a combination of MTDs) considering a

holistic view of the system. To achieve this, we need to enumerate possible paths that an attacker can take in multi-stage attacks.




Figure 1: Proposed MTD framework

We define both the Cloud and the threats layer, presented in Figure 1, in the following.

Definition 2.1 (Cloud). A Cloud is modeled as a directed graph representing interconnections among the services/components, i.e., $G = (V, E)$, where G represents the Cloud system, V is a finite set of services and $E \subseteq V \times V$ is the set of ordered connecting edges between the services such that a communication link exists between $v_i \in V$ and $v_j \in V$. An attacker might use the latter to propagate further into the system.

Definition 2.2 (Threat). A Threat is modeled as a tuple $T = (V, E, I)$, where: V is a set of vulnerabilities that an attacker can exploit to target a service, E is the process of exploiting the vulnerability, and I is the impact of successfully exploiting the vulnerability in the system. E can be further defined in the form of a tuple (Rec, Ex) such that $Rec : \mathbb{N} \rightarrow A$ denotes the sequence of actions that an attacker can take during the reconnaissance process, and $Ex : \mathbb{N} \rightarrow A$ signifies the corresponding sequence of actions to exploit the vulnerability.

Having defined the Cloud system and the threats, the third layer enumerates the choices of the MTDs. We can now define an overall MTD framework as a 6-tuple $(M, T, S, R_{mt}, R_{ms}, R_{ts})$ where

- M is a finite set of MTDs.
- T is a finite set of threats and is defined in Definition 2.2.
- S is a finite set of services in the Cloud and is defined in Definition 2.1.
- $R_{mt} \subseteq M \times T$ is the relationship between MTDs and the threats. It determines the corresponding impact of MTDs on threats.
- $R_{ms} \subseteq M \times S$ defines the relationship between MTDs and services in the Cloud. It signifies the applicability of an MTD

on the service. For instance, dynamic application data is only applicable to the application running on the VM.

- $R_{ts} \subseteq T \times S$ defines the relationship between threats and the targeted services. It identifies the threats that could potentially target a specific service.

Consequently, the selection and placement of MTD techniques in the Cloud can be formulated as an optimization problem. The objective is to maximize the effectiveness of MTDs and to minimize the cost associated with their deployment. Furthermore, we define constraints associated with the system, attacker, and defender. Thus the optimization problem can be defined as:

$$\begin{aligned} & \text{maximize} && E(M_m, T_t, S_s) \\ & \text{minimize} && C(M_m, S_s) \\ & \text{subject to} && C1, C2, \text{ and } C3 \end{aligned} \quad (1)$$

- C1 (System's Constraint) Proper termination of the system, i.e., VM run state, is eventually reached.
- C2 (Attacker's Constraint) A threat with maximum likelihood is used by an attacker at any given time.
- C3 (Defender's Constraint) An MTD can only be used if it can target a threat and is deployable on the targeted service. If a single MTD is insufficient to thwart an attack, multiple MTDs can be deployed.

We define the relationship among different layers of the MTD framework by the following:

Threat to service relationship: A service can be compromised by exploiting vulnerabilities that target the service. However, if there does not exist a vulnerability that can target a service then the service cannot be compromised, i.e., $(t, s) \notin R_{ts} \implies \Pr(X_t^s = t) = 0$.

Satisfying a threat's precondition: To compromise a service, a vulnerability has to be exploited. However, preconditions might need to be satisfied to exploit the specific vulnerability, which can be further broken down into the fulfillment of each individual precondition. Pre is the set of all preconditions, and MSP is the minimum set of preconditions that must be satisfied to exploit the vulnerability.

$$\begin{aligned} SP(s, t) &= \{ p \mid p \in \mathcal{P}(Pre) \wedge SAT(s, t, p) \} \\ MSP(s, t) &= \arg \min_{p \in SP(s, t)} |p| \end{aligned} \quad (2)$$

Effectiveness of MTDs: We have defined the relationship between threats and the respective services and the minimum set of preconditions necessary to exploit a vulnerability. Now, we define the relationship between MTDs, threats, and services. Initially, we start with evaluating the effectiveness of a single MTD technique against a threat targeting a particular service. We plan to relax this assumption in future work by incorporating many-to-many relationships amongst the layers. We assume that in case an MTD is not deployed, an attacker can successfully compromise the service, i.e., the probability of the attack success is 1. We can calculate the effectiveness of a single instance of MTD on a threat targeting a service as:

$$E(M_m, T_t, S_s) = (T_t, S_s) \in R_{ts} \wedge (M_m, T_t) \in R_{mt} \wedge (M_m, S_s) \in R_{ms} \quad (3)$$

Equation (3) determines the effectiveness of an MTD technique against a threat to a service. We assume that an MTD is either effective or ineffective against a threat; however, in the future, we will relax this assumption to assign different values to each MTD technique against a threat and evaluate the effectiveness of multiple MTDs targeting a threat. In other words, Equation (3) also dictates that an MTD has not been deployed, and therefore, the attacker’s success is one as no mitigation strategy has been deployed.

Cost function of a single MTD on a service: Cost function maps the cost associated with deploying an MTD on a service. This cost can be further divided into the operational cost of deploying an MTD as well overhead caused by an MTD.

$$C(M_m, S_s) = Deployment(M_m, S_s) + Overhead(M_m, S_s) \quad (4)$$

Solving an optimization problem with conflicting objective functions is challenging. The goal of the proposed work is to find the optimal MTDs placement across the Cloud that maximize their effectiveness in mitigating threats (or offering higher security gains) while minimizing the cost associated with the MTDs deployment.

3 SECURITY ANALYSIS

An example of the proposed quantification framework is shown in Figure 2 which shows all the layers of the framework. The threats are extracted from the national vulnerability databases [6] and the respective probabilities of exploiting each vulnerability can be assigned from the database. For example, CVE-2011-1751 affects the networking service of the Cloud and the corresponding MTD technique (dynamic networks) can be applied to mitigate the threat. Similarly, dynamic platforms can be applied to either the hypervisor level or at the VM level to mitigate CVE-2008-7096 and CVE-2010-2938 respectively. Therefore, these MTDs can be deployed at the respective services to protect them from being vulnerable. Currently, we focus on each MTD affecting the service, however, in the future, we plan to relax this assumption by including multiple MTDs impacting each service and threat. Furthermore, it is evident that the relationship between the layers is many to many. For example, a threat can impact multiple services, and an MTD can be applied to a number of services. Therefore, we conclude that finding an optimal solution is challenging considering these relationships.

4 CONCLUSION

We have presented the problem of evaluating the effectiveness of MTDs and their deployment across the operational stack of the Cloud. The approach enables evaluating the security gains on the overall system instead of limiting the effectiveness of MTDs to individual components. The security analysis presented encompasses different services involved in launching a VM, threats targeting these services, and the respective MTDs. Initially, we target finding the optimal placement of MTDs for single-stage attacks and plan to extend the security analysis to multi-stage attacks.

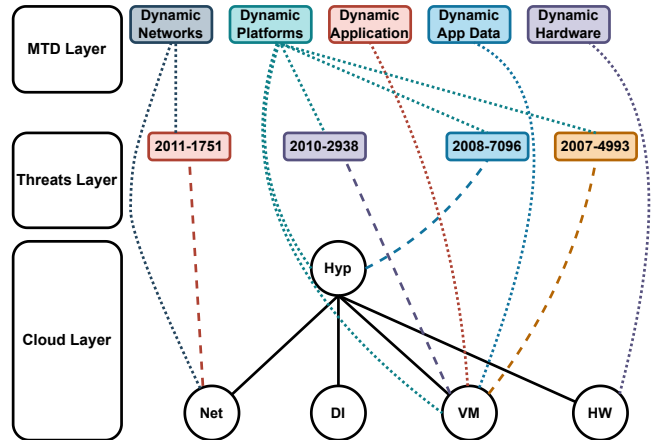


Figure 2: Single stage attack scenario

ACKNOWLEDGMENTS

This research was supported, in part, by EC H2020 CONCORDIA GA No. 830927 and by the UKRI Trustworthy Autonomous Systems Node in Security [EPSRC grant EP/V026763/1].

REFERENCES

- [1] Hooman Alavizadeh, Julian Jang-Jaccard, and Dong Seong Kim. 2018. Evaluation for Combination of Shuffle and Diversity on Moving Target Defense Strategy for Cloud Computing. In *International Conference On Trust, Security And Privacy In Computing And Communications*. IEEE, New York, NY, USA, 573–578. <https://doi.org/10.1109/TrustCom/BigDataSE.2018.00087>
- [2] Ping Chen, Jun Xu, Zhiqiang Lin, Dongyan Xu, Bing Mao, and Peng Liu. 2015. A Practical Approach for Adaptive Data Structure Layout Randomization. In *ESORICS*. Springer, 69–89. https://doi.org/10.1007/978-3-319-24174-6_4
- [3] Nils Gruschka and Meiko Jensen. 2010. Attack Surfaces: A Taxonomy for Attacks on Cloud Services. In *Proceedings of the International Conference on Cloud Computing*. IEEE, Miami, FL, USA, 276–279. <https://doi.org/10.1109/CLOUD.2010.23>
- [4] Haadi Jafarian, Ehab Al-Shaer, and Qi Duan. 2015. An Effective Address Mutation Approach for Disrupting Reconnaissance Attacks. *IEEE Transactions on Information Forensics and Security* 10, 12 (2015), 2562–2577. <https://doi.org/10.1109/TIFS.2015.2467358>
- [5] Samuel King and Peter M. Chen. 2006. SubVirt: Implementing malware with virtual machines. In *IEEE Symposium on Security and Privacy*. IEEE, 14–327. <https://doi.org/10.1109/SP.2006.38>
- [6] NIST. n.d. National Vulnerability Database. Retrieved 2022-06-01 from <https://nvd.nist.gov/>
- [7] Diego Perez-Botero, Jakub Szefer, and Ruby Lee. 2013. Characterizing Hypervisor Vulnerabilities in Cloud Computing Servers. In *Proceedings of the ACM International Workshop on Security in Cloud Computing*. ACM, 3–10. <https://doi.org/10.1145/2484402.2484406>
- [8] Rui Zhuang, Scott A DeLoach, and Xinming Ou. 2014. Towards a Theory of Moving Target Defense. In *Proceedings of the First ACM Workshop on Moving Target Defense*. 31–40. <https://doi.org/10.1145/2663474.2663479>