

Don't Wait to be Breached!

Creating Asymmetric Uncertainty of Cloud Applications via Moving Target Defenses

Kennedy A. Torkura

Christoph Meinel

Nane Kratzke

Hasso Plattner Institute
University of Potsdam, Germany
Email: kennedy.torkura@hpi.de

Hasso Plattner Institute
University of Potsdam, Germany
christoph.meinel@hpi.de

Lübeck University of Applied Sciences
Lübeck, Germany
nane.kratzke@th-luebeck.de

Abstract—Cloud applications expose besides service endpoints also potential or actual vulnerabilities. Therefore, cloud security engineering efforts focus on hardening the fortress walls but seldom assume that attacks may be successful. At least against zero-day exploits, this approach is often toothless. Other than most security approaches and comparable to biological systems we accept that defensive “walls” can be breached at several layers. Instead of hardening the “fortress” walls we propose to make use of an (additional) active and adaptive defense system to attack potential intruders - an immune system that is inspired by the concept of a moving target defense. This “immune system” works on two layers. On the infrastructure layer, virtual machines are continuously regenerated (cell regeneration) to wipe out even undetected intruders. On the application level, the vertical and horizontal attack surface is continuously modified to circumvent successful replays of formerly scripted attacks. Our evaluations with two common cloud-native reference applications in popular cloud service infrastructures (Amazon Web Services, Google Compute Engine, Azure and OpenStack) show that it is technically possible to limit the time of attackers acting undetected down to minutes. Further, more than 98% of an attack surface can be changed automatically and minimized, which makes it hard for intruders to replay formerly successful scripted attacks. So, even if intruders get a foothold in the system, it is hard for them to maintain it. Therefore, our proposals are robust and dynamically change due in response to security threats similar to biological immune systems.

Keywords—zero-day; exploit; moving target defense; microservice; cloud-native; application; security; asymmetric

I. INTRODUCTION

This paper extends ideas presented in [1] to improve cloud application security in the context of unknown zero-day exploits and reports on ongoing research in this field. Cloud computing enables a variety of innovative IT-enabled businesses and service models. Several research studies and programs focus on responsibly developing systems to ensure the security and privacy of users. But compliance with standards, audits, and checklists, does not automatically equal

security [2] and there is a fundamental issue remaining. Zero-day vulnerabilities are computer-software vulnerabilities that are unknown to those who would be interested in mitigating the vulnerability (including the entity responsible for operating a cloud application). Until a vulnerability is mitigated, hackers can exploit it to adversely affect computer programs, data, additional computers or a network. For zero-day exploits, the probability that vulnerabilities are patched is zero, so the exploit should always succeed. Therefore, zero-day attacks are a severe threat, and we have to draw a scary conclusion: **In principle, attackers can establish footholds in our systems whenever they want.**

This contribution deals with the question how to build “unfair” cloud systems that permanently jangle attackers nerves. We present the latest results from our ongoing research that applies Moving Target Defense (MTD) principles on cloud runtime environment and cloud application layer.

Recent research [3], [4] made successfully use of elastic container platforms (see Table I) and their “designed for failure” capabilities to realize transferability of cloud-native applications at runtime. By transferability, the conducted research means that a cloud-native application can be moved from one IaaS provider infrastructure to another without any downtime. These platforms are more and more used as distributed and elastic runtime environments for cloud-native applications [5] and can be understood as a kind of cloud infrastructure unifying middleware [6]. It should be possible to make use of the same features to immunize cloud applications simply by moving an application within the same provider infrastructure. To move anything from A to A makes no sense at first glance. However, let us be paranoid and aware that with some probability and at a given time, an attacker will be successful and compromise at least one virtual machine [7]. A transfer from A to A would be an effective countermeasure – because the intruder immediately loses any hijacked machine that is moved. To understand that, the reader must know that our approach does not effectively move a machine, it regenerates it. To move a machine means to launch a compensating machine unknown to the intruder and to terminate the former (hijacked) machine. Whenever an application is moved its virtual machines are regenerated. Moreover, this would effectively eliminate undetected hijacked machines.

However, attackers can run automated attacks against regenerated machines that will incorporate the same set of vulnerabilities. Therefore, this extended paper shows how we further can improve the regenerating security measure by

TABLE I. Some popular open source elastic platforms

Platform	Contributors	URL
Kubernetes	Cloud Native Found.	http://kubernetes.io
Swarm	Docker	https://docker.io
Mesos	Apache	http://mesos.apache.org/
Nomad	Hashicorp	https://nomadproject.io/

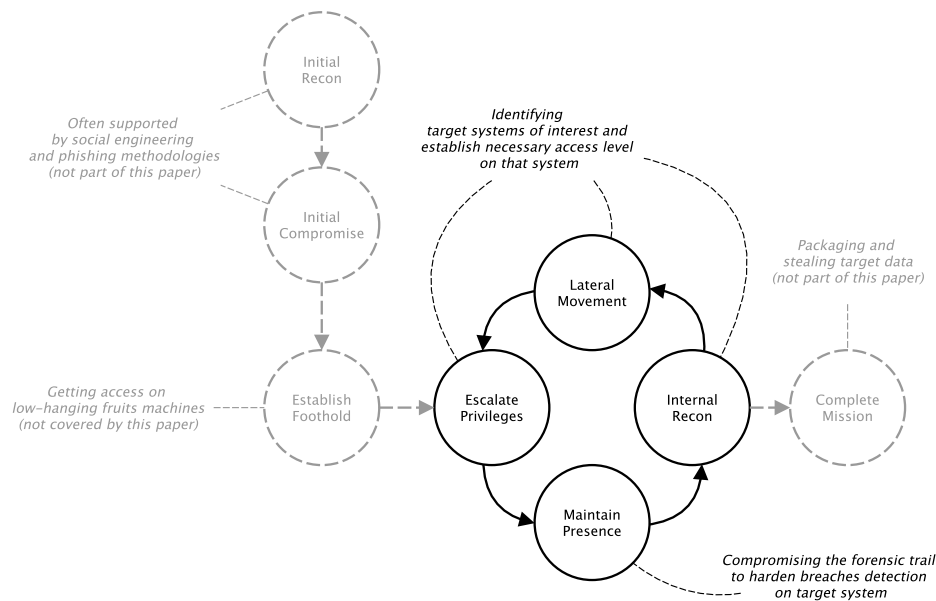


Figure 1. The cyber attack life cycle model. Adapted from the cyber attack lifecycle used by the M-Trends reports, see Table II.

employing MTD at the *application layer* [8] to change the attack surface of the application itself to let even automated and formerly successful attack scripts fail (at least partly). Primarily, this is achieved by diversifying the application in a way that its containerized components are dynamically transformed at *runtime*. The two abstraction layers that compose microservice applications (application layer and the container image layers) are dynamically changed by changing the programming languages of the applications, and consequently, the container images are built to conform to the requirements of the corresponding applications. This combined approach is enforced at runtime to transform the attack surface of cloud-native applications, thereby reducing the possibility of successful attacks.

The remaining of this paper is **outlined** as follows: Section II presents a cyber-attack lifecycle model to show where our approach intends to break the continuous workflow of security breaches. Section III presents an approach on how MTD can be applied on cloud runtime environment (infrastructure) level to regenerate the "*infrastructure cells*" of a system continuously, leveraging the inherent "*designed-for-failure*" capabilities of modern container platforms like Kubernetes, Swarm, or Mesos. This continuing regeneration will wipe out even undetected attackers in a system. However, attackers might recognize that they periodically loose foothold in a hijacked system and might try to automate their breaches. To overcome this, Section IV will present how even the attack surface of an application can be continuously changed and therefore extends our ideas shown in [1]. We have to consider that our approach has some limitations. We will discuss these limitations in Section V and present corresponding related work in Section VI. We conclude our findings in Section VII.

II. CYBER ATTACK REFERENCE MODEL

Figure 1 shows the cyber attack life cycle model, which is used by the M-Trends Report [9] to present current devel-

opments in cyber attacks over the years. According to this model, an attacker passes through different stages to complete a cyber attack mission. It starts with initial reconnaissance and compromising of access means. Social engineering methodologies [10] and phishing attacks [11] very often supports these steps. Intruders aim to establish a foothold near the target. All these steps are not covered by this paper, because technical solutions are not able to harden the weakest point in security – the human being. The following steps of this model are more important for this paper. According to the life cycle model, the attacker's goal is to escalate privileges to get access to the target system. Because this leaves trails on the system, which could reveal a security breach, the attacker is motivated to erase this forensic trail. According to security reports, attackers increasingly employ counter-forensic measures to hide their presence and impair investigations. These reports refer to techniques used to clear event logs and securely delete arbitrary files. The technique is simple, but the intruders' knowledge of forensic artifacts demonstrate increased sophistication, as well as their intent to persist in the environment. With a barely detectable foothold, the internal reconnaissance of the victim's network is carried out to allow lateral movement to target systems. This process is a complex and lengthy process and may even take weeks. So, infiltrated machines and application components have worth for attackers and tend to be used for as long as possible. Table II shows the average period an intruder remains on a victim system undetected. So, basically there is the requirement, that **(1) an undetected attacker should lose access to compromised nodes of a system as fast as possible**. Furthermore, there is the requirement that it **(2) must be hard for an attacker to regain foothold in a system by automating successful attacks**. However, how?

Section III will deal with the **(1) requirement** showing that it is possible to regenerate possibly compromised infrastructure continuously even to get rid of undetected attackers. Section

TABLE II. Undetected days on victim systems reported by M-Trends. External and internal discovery data is reported since 2015. No data could be found for 2011.

Year	External notification	Internal discovery	Median
2010	-	-	416
2011	-	-	?
2012	-	-	243
2013	-	-	229
2014	-	-	205
2015	320	56	146
2016	107	80	99

IV will deal with the (2) requirement and demonstrate that it is possible to change attack surfaces of applications in a way that successful attacks cannot be repeated 1:1.

III. MOVING TARGET DEFENSE MECHANISMS ON THE CONTAINER RUNTIME ENVIRONMENT LEVEL

Our recent research dealt [12] mainly with vendor lock-in and the question how to design cloud-native applications that are transferable between different cloud service providers. One aspect that can be learned from this is that there is no common understanding of what a cloud-native application is. A kind of software that is “intentionally designed for the cloud” is an often heard but empty phrase. However, noteworthy similarities exist between various viewpoints on *cloud-native applications* (CNA) [5]. A common approach is to define maturity levels in order to categorize different kinds of cloud applications (see Table III). [13] proposed the IDEAL model for CNAs. A CNA should strive for an **isolated state**, is **distributed**, provides **elasticity** in a horizontal scaling way, and should be operated on **automated deployment machinery**. Finally, its components should be **loosely coupled**.

Balalaie et.al [14] stressed that these properties are addressed by cloud-specific architecture and infrastructure approaches like **Microservices** [15], **API-based collaboration**, adaption of **cloud-focused patterns** [13], and **self-service elastic platforms** that are used to deploy and operate these microservices via self-contained deployment units (containers). Table I lists some of these platforms that provide additional operational capabilities on top of IaaS infrastructures like automated and on-demand scaling of application instances, application health management, dynamic routing and load balancing as well as aggregation of logs and metrics [5].

A. Regenerating cloud application runtime environments continuously

If the reader understands and accepts the commonality that cloud-native applications are operated (more and more often) on elastic – often container-based – platforms, it is an obvious idea to delegate the responsibility to immunize cloud applications to these platforms. Recent research showed that the operation of these elastic container platforms and the design of applications running on top of them should be handled as two different engineering problems. This point of view often solves several issues in modern cloud-native application engineering [4]. Also, that is not just true for the transferability problem but might be an option to tackle zero-day exploits. These kinds of platforms could be an essential part of the immune system of modern cloud-native applications.

Furthermore, **self-service elastic platforms** are really “bulletproofed” [17]. *Apache Mesos* [18] has been successfully operated for years by companies like Twitter or Netflix to consolidate hundreds of thousands of compute nodes. Elastic container platforms are **designed for failure** and provide self-healing capabilities via auto-placement, auto-restart, auto-replication, and auto-scaling features. They will identify lost containers (for whatever reasons, e.g., process failure or node unavailability) and will restart containers and place them on remaining nodes. These features are necessary to operate large-scale distributed systems resiliently. However, the same features can be used intentionally to **purge “compromised nodes”**.

In [3], a software prototype that provides the control process shown in Figure 2 and Figure 3 was presented. This process relies on an *intended state* ρ and a *current state* σ of a container cluster. If the intended state differs from the current state ($\rho \neq \sigma$), necessary adaption actions are deduced (creation and attachment/detachment of nodes, creation and termination of security groups) and processed by an execution pipeline fully automatically (see Figure 3) to reach the *intended state* ρ . With this kind of control process, a cluster can be simply resized by changing the intended amount of nodes in the cluster. If the cluster is shrinking and nodes have to be terminated, affected containers of running applications will be rescheduled to other available nodes.

The downside of this approach is that this will only work for Level 2 (cloud resilient) or Level 3 (cloud-native) applications, (see Table III) which by design, can tolerate dependent service failures (due to node failures and container rescheduling). However, for that kind of Level 2 or Level 3 application, we can use the same control process to regenerate nodes of the container cluster. The reader shall consider a cluster with $\sigma = N$ nodes. If we want to regenerate one node, we change the intended state to $\rho = N + 1$ nodes, which will add one new node to the cluster ($\sigma' = N + 1$). Moreover, in a second step, we will decrease the predetermined size of the cluster to $\rho' = N$ again, which affects that one node of the cluster is terminated ($\sigma'' = N$). So, a node is regenerated simply by adding one node and deleting one node. We could even regenerate the complete cluster by changing the cluster size in the following way: $\sigma = N \mapsto \sigma' = 2N \mapsto \sigma'' = N$. However, this would consume much more resources because the cluster would double its size for a limited amount of time.

TABLE III. Cloud Application Maturity Model, adapted from OPEN DATA CENTER ALLIANCE Best Practices [16]

Level	Maturity	Criteria
3	Cloud native	- Transferable across infrastructure providers at runtime and without interruption of service. - Automatically scale out/in based on stimuli.
2	Cloud resilient	- State is isolated in a minimum of services. - Unaffected by dependent service failures. - Infrastructure agnostic.
1	Cloud friendly	- Composed of loosely coupled services. - Services are discoverable by name. - Components are designed to cloud patterns. - Compute and storage are separated.
0	Cloud ready	- Operated on virtualized infrastructure. - Instantiatable from image or script.

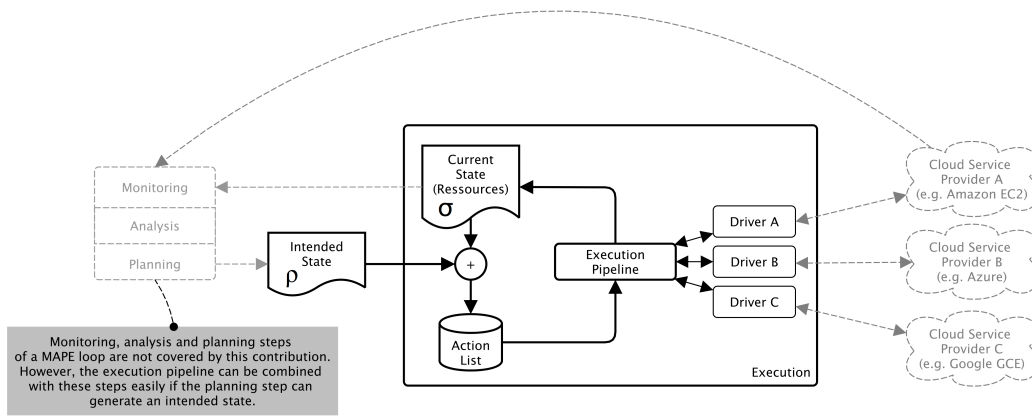


Figure 2. The control theory inspired execution control loop compares the intended state ρ of an elastic container platform with the current state σ and derives necessary scaling actions. These actions are processed by the execution pipeline explained in Figure 3. So, platforms can be operated elastically in a set of synchronized IaaS infrastructures. Explained in details by [3].

A more resource efficient way would be to regenerate the cluster in N steps: $\sigma = N \mapsto \sigma' = N + 1 \mapsto \sigma'' = N \mapsto \dots \mapsto \sigma^{2N-1} = N + 1 \mapsto \sigma^{2N} = N$. The reader is referred to [4] for more details, especially if the reader is interested in the multi-cloud capabilities that are not covered by this paper due to page limitations.

Whenever such regeneration is triggered, all – even undetected – hijacked machines would be terminated and replaced by other machines, but the applications would be unaffected. For an attacker, this means losing their foothold in the system entirely. Imagine this would be done once a day or even more frequently?

B. Evaluation

The execution pipeline presented in Figure 3 was evaluated by operating and transferring two elastic platforms (*Swarm Mode of Docker 17.06* and *Kubernetes 1.7*). The platforms operated a reference “sock-shop” application being one of the most complete reference applications for microservices architecture research [19]. Table IV lists the machine types that show a high similarity across different providers [20].

The evaluation of [4] demonstrated that most time is spent on the IaaS level (creation and termination of nodes and security groups) and not on the elastic platform level (joining,

draining nodes). The measured differences on infrastructures provided by different providers are shown in Figure 4. For the current use case, the reader can ignore the times to create and delete a security group (because that is a one time action). However, there will be many node creations and terminations. According to our execution pipeline shown in Figure 3, a node creation ($\sigma = N \mapsto \sigma' = N + 1$) involves the durations to **create a node** (request of the virtual machine including all installation and configuration steps), to **adjust security groups** the cluster is operated in and to **join the new node** into the cluster. The shutdown of a node ($\sigma = N \mapsto \sigma' = N - 1$) involves the **termination of the node** (this includes the platform draining and deregistering of the node and the request to terminate the virtual machine) and the necessary **adjustment of the security group**. So, for a complete regeneration of a node ($\sigma = N \mapsto \sigma' = N + 1 \mapsto \sigma'' = N$) we have to add these runtimes. Table V lists these values per infrastructure.

Even on the “slowest” infrastructure, a node can be regenerated in about 10 minutes. In other words, one can regenerate six nodes every hour or up to 144 nodes a day or a cluster of 432 nodes every 72h (which is the reporting time requested by the EU General Data Protection Regulation). If the reader compares a 72h regeneration time of a more than 400 node cluster (most systems are not so large) with the median value of 99 days that attackers were present on a victim system in 2016 (see Table II) the benefit of the proposed approach should become apparent.

TABLE IV. Used machine types and regions for evaluation

Provider	Region	Master type	Worker type
AWS	eu-west-1	m4.xlarge	m4.large
GCE	eu-west-1	n1-standard-4	n1-standard-2
Azure	europewest	Standard_A3	Standard_A2
OS	own datacenter	m1.large	m1.medium

TABLE V. Durations to regenerate a node (median values)

Provider	Creation	Secgroup	Joining	Term.	Total
AWS	70 s	1 s	7 s	2 s	81 s
GCE	100 s	8 s	9 s	50 s	175 s
Azure	380 s	17 s	7 s	180 s	600 s
OS	110 s	2 s	7 s	5 s	126 s

IV. MOVING TARGET DEFENSE MECHANISMS ON THE MICROSERVICE ARCHITECTURE LEVEL

MTD techniques introduce methods for improving the security of protected assets by applying *security-by-diversity* tactics and *security diversification* concepts. While most MTD techniques do not have formal requirements for diversifying, i.e., when, how and why to diversify, we employ a *cyber risk-based* technique as the primary diversification decision making factor on the application level [8]. Our motivation for this is to overcome the high number of vulnerability infection among container images as shown by several recent researchers[21], [22]. Therefore, our MTD techniques are designed to improve this *state of insecurity* by reducing the *window of vulnerability*

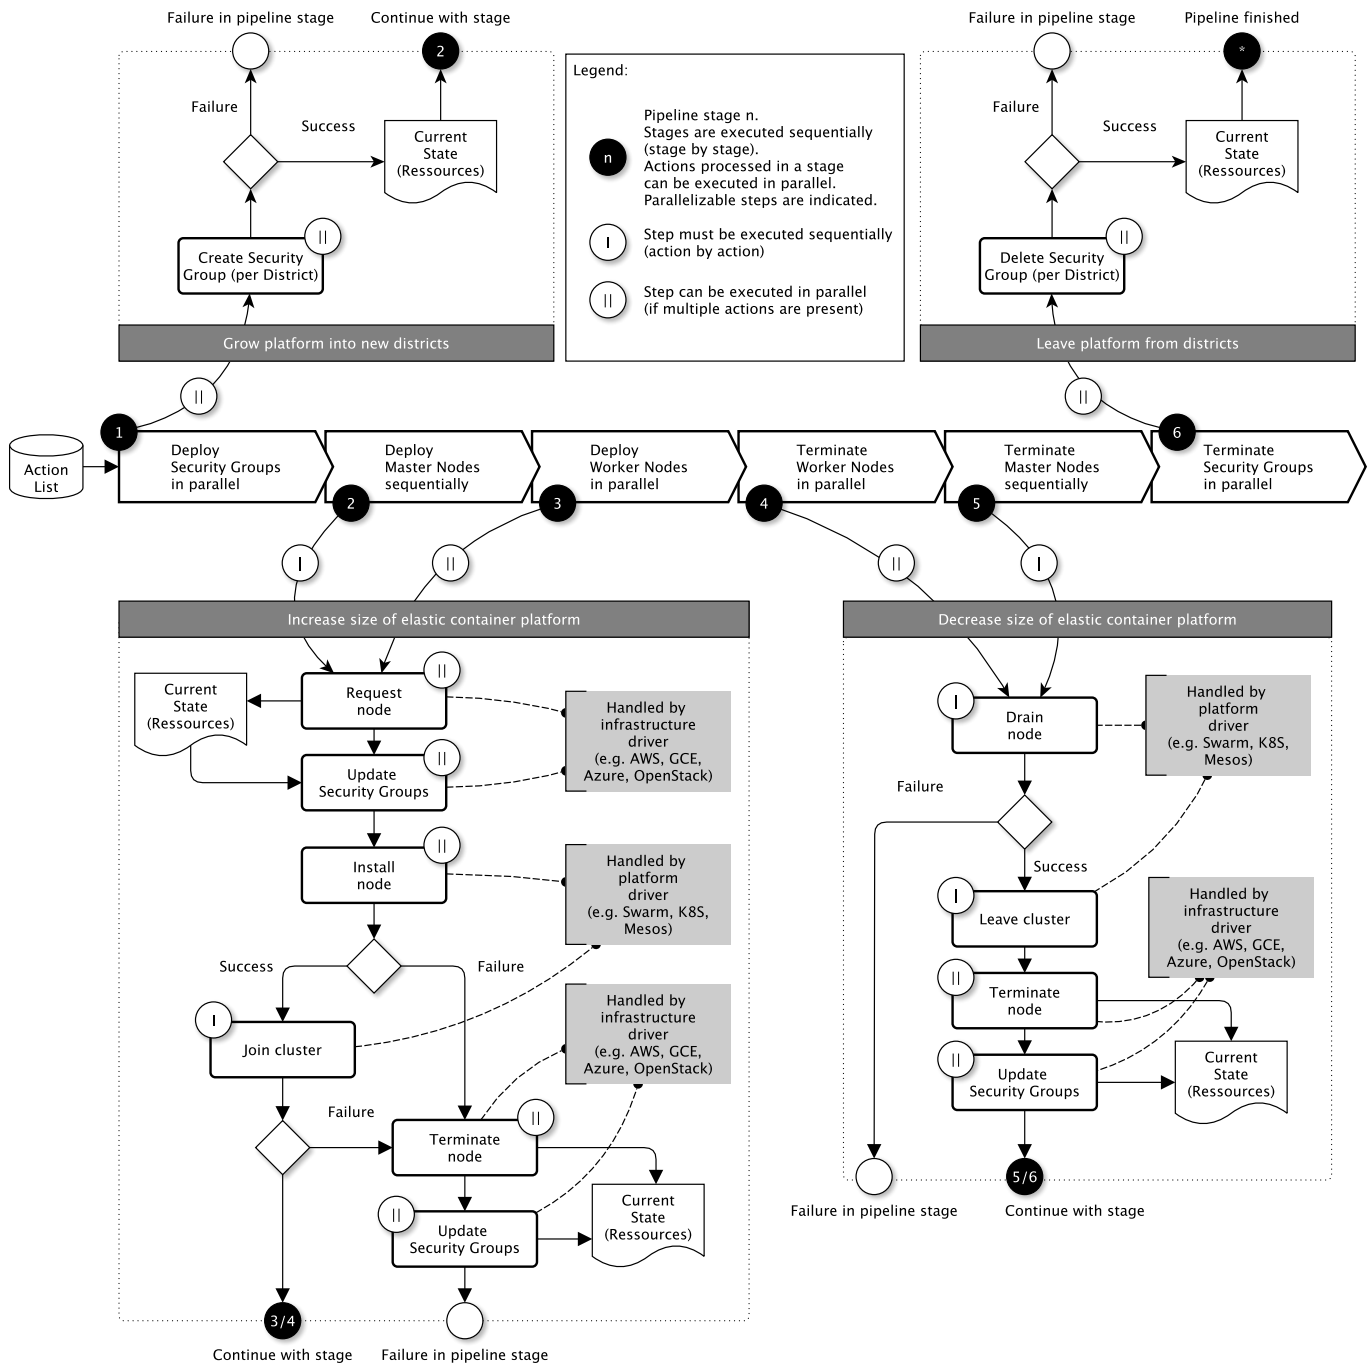


Figure 3. The execution pipeline processes necessary actions to transfer the current state σ into the intended state ρ . See [4] for more details.

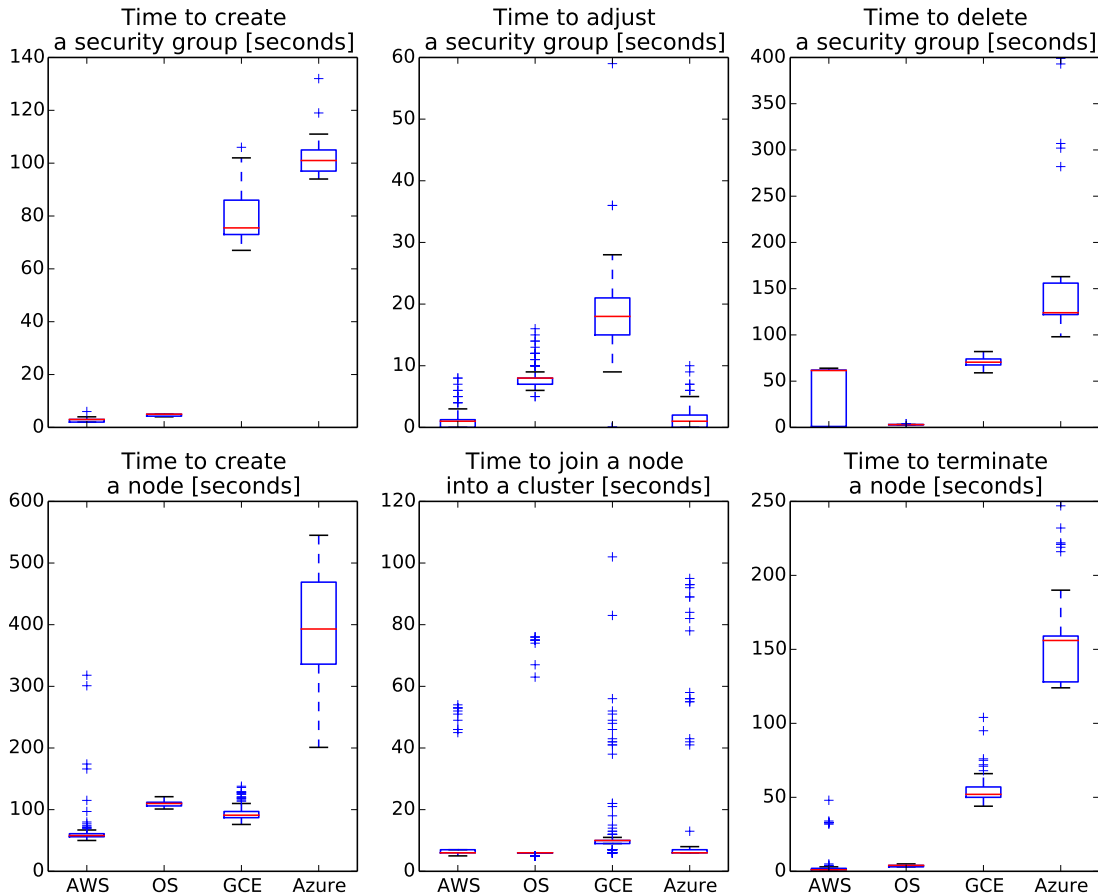


Figure 4. Infrastructure specific runtimes of IaaS operations see [4].

exposure via diversification and commensurate attack surface randomization.

A. Cyber Risk Analysis for Microservice Diversification

Larsen et al. [23] assert that a common challenge when employing diversification strategies is deciding on *when*, *how* and *where* to diversify. We present a cyber risk procedure to support decision making or satisfy the afore-mentioned requirements. We leverage security metrics to design a cyber risk-based mechanism, and security metrics are useful tools for risk assessment. These metrics are computed by deriving security risks per microservice and after that employing *vulnerability prioritization* such that diversification is a function of microservice risk assessment, i.e., microservices are diversified in order risk severity. We introduce the notion of **Diversification Index** - D_i as an expression of the depth of diversification to be implemented. D_i defines if microservices are to be globally or selectively diversified e.g., diversifying 2 out of 4 microservices can be expressed as 2:4. We formally define D_i as:

$$D_i = \frac{m_d}{m} \quad (1)$$

where,

- m_d = number of microservices to be diversified,
- m = total number of microservices in the application.

In order to decide on the value of D_i , it is important to analyze the security state of the microservice application. There are several techniques for arriving at what may be defined as the *security state* of a microservice application. Traditionally, several security metrics have been defined to evaluate the security state of computing environments, therefore we employ security risk techniques, following two approaches:

1) *Risk Analysis Using CVSS*: The Common Vulnerability Scoring System CVSS [24] is a widely adopted vulnerability metrics standard. It provides vulnerability *base scores*, which express the severity of damage the referred vulnerability might impact upon a system if exploited. In order to derive the microservice security state (Security Risk - SR), base scores of all the vulnerabilities detected can be summed and averaged as expressed below:

$$SR = \frac{1}{N} \sum_{i=1}^N V_i \quad (2)$$

where SR is the Security Risk, V_i is the CVSS base score of vulnerability i , and N is the total number of vulnerabilities detected in microservice m . However, averaging vulnerabilities to obtain a single metric to signify a system's security state is not optimal. Derived values are not sufficiently representative of other factors such as the public availability of exploits. Therefore, we employ another scoring technique called *shrink-*

age estimator, an approach, which has been popularly used for online rating systems, e.g., IMDB. The shrinkage estimator considers the average rating and the number of votes. Hence, it provides a more precise value for SR, than mere averaging (Equation (2)). Therefore, leveraging the shrinkage estimator, we can derive a more precise *SR* as follows:

$$SR = \frac{v}{v+a}R + \frac{a}{v+a}C \quad (3)$$

where,

v = the total number of vulnerabilities detected in a microservices,

a = minimum number of vulnerabilities to be detected in a microservice assessment before it added in the risk analysis,

C = the mean severity score of vulnerabilities detected in a microservice

R = the average severity score of all vulnerabilities infecting a microservice-based application

The Pearson's correlation coefficient is derived to determine the dependence relationship between the microservices.

2) *Risk Analysis Using OWASP Risk Rating Methodology*: The risk assessment method described in the previous subsection is limited to vulnerabilities contained in the Common Vulnerability Enumeration (CVE) dictionary. CVE is a public dictionary for publishing known vulnerabilities. These vulnerabilities are analyzed and assigned vulnerability security metrics using the CVSS. However, the CVE contains only a handful of web application vulnerabilities. Thus, we need to derive another risk assessment methodology for application layer vulnerabilities. This additional step is necessary since microservices are essentially web/REST-based applications. We opt for the OWASP Risk Rating Methodology (ORRM), which is specifically designed for web applications [25]. This methodology is based on two core risk components: *Likelihood* and *Impact* formally expressed as:

$$Risk = Likelihood * Impact \quad (4)$$

In order to derive these metrics, risk assessors are required to consider the threat vector, attacks to be used and the impacts of successful attacks.

B. Dissecting Microservice Attack Surfaces

An important aspect of our *security-by-diversity* tactics is to manipulate microservice attack surfaces against possible attackers through random architectural transformations. Therefore, the attack surfaces are altered by randomizing the entry and exit points, which are commonly used for identifying attack surfaces [27], [28]. A detailed understanding of these attack surfaces is imperative. Therefore, we categorize microservice attack surface into: *horizontal* and *vertical* attack surfaces and thereafter employ *vulnerability correlation* to identify vulnerability similarities.

1) *Horizontal Vulnerability Correlation*: The objective of correlating vulnerabilities horizontally is to analyze the relationship of vulnerabilities along the horizontal attack surface, i.e., the parts of the applications users directly interact with. Figure 5 illustrates the multi-layered attack surface of the PetClinic application [26]. The application layer horizontal attack surface consists of the interactions and exit/entry points




Figure 5. Typical Microservice Attack Surfaces illustrated with the PetClinic Application [26]

from the API gateway to the Vets, Visits and Customer services application layers. Requests and responses are transversed along this layer, providing attack opportunities for attackers. The vulnerability correlation process is similar to *security event correlation* techniques [29], though rather than clustering similar attributes, e.g., malicious IP addresses, we focus on Common Weakness Enumeration (CWE) Ids. The CWE is a standardized classification system for application weaknesses [30]. For example, CWE 89 categorizes all vulnerabilities related to *Improper Neutralization of Special Elements used in an SQL Command (SQL Injection)* [31] and can be mapped to several CVEs, e.g., *CVE-2016-6652*[32], a SQL injection vulnerability in Spring Data JPA. If this vulnerability exists in all PetClinic's microservices, an attacker could easily conduct a correlated attack (*Attack Paths 2, 4, 5, and 6* of Figure 5) resulting to correlated failures and eventual application failure since each microservice works ultimately to the successful functioning of the PetClinic application.

2) *Vertical Vulnerability Correlation*: The vertical correlation technique is similar to the horizontal correlation. However, the interactions across application-image layers are analyzed. This analysis, therefore, employs *security-by-design* tactics across the vertical attack surface. Attack Path 1 illustrates the exploitation of vulnerability across the vertical attack surface, and the attacker initiated an attack against the API Gateway of the PetClinic application, from the application layer to the image layer. From there, another attack is launched to the Customers service application layer, across the image layer and finally, the database is compromised. The same attack can be repeated against the other microservices if affected by the vulnerabilities. Hence we need to express such casual relationships in vulnerability correlation matrices.

Correlated vulnerabilities can be represented with correlation matrices, more specifically referred to as *microservices vulnerability correlation matrix*. Therefore, we are influenced by [33] to define the microservices vulnerability correlation matrix as a *mapping of vulnerabilities to microservice instances in a microservice-based application*. The *microservices vulnerability correlation matrix* presents a view of vulnerabilities that concurrently affect multiple microservices. An

$$\begin{matrix}
 & V_1 & V_2 & \dots & V_n \\
 M_1 & \begin{bmatrix} 1 & 1 & \dots & \dots \\ 1 & 0 & \dots & \dots \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \dots & \dots \end{bmatrix} \\
 M_2 & \\
 \vdots & \\
 M_n &
 \end{matrix}$$

Figure 6. Microservice Vulnerability Correlation Matrix

TABLE VI. Vulnerabilities Detected in PetClinic App-Layer

CWE-ID	API-GATEWAY	CUSTOMERS-SERVICE	VETS-SERVICE	VISITS-SERVICE
CWE-16	31	4	2	2
CWE-524	48	17	6	11
CWE-79	0	3	0	1
CWE-425	0	0	20	0
CWE-200	14	6	0	0
CWE-22	0	1	0	0
CWE-933	1	0	0	0
TOTAL	94	31	28	14

example of the microservice correlation matrix is Figure 6, where the microservices M_1 and M_2 will have a correlated failure under an attack that exploits vulnerability V_1 since they share the same vulnerability. However, an attack that exploits V_2 can only affect M_1 , while M_2 remains unaffected.

C. Evaluation

The PetClinic application was used for our evaluation. PetClinic is a Spring application that demonstrates cloud-native application capabilities. It has been used for demonstration purposes in several industry and academic scenarios [19]. However, we were forced to modify the original PetClinic by adding OpenAPI support. Two experiments have been conducted: (1) Security risk comparison to verify the efficiency of our security-by-diversity tactics (2) Attack surface analysis to evaluate the improvement in the horizontal and vertical attack surfaces.

In order to perform Security Risk analysis, we leveraged the Cloud Aware Vulnerability Assessment System (CAVAS) [34]. The vulnerability scanners integrated into CAVAS (Anchore and OWASP ZAP), are used for launching vulnerability scans against PetClinic images and microservice instances respectively. The detected vulnerabilities were persisted in the Security Reports and CMDB. First, the diversification index

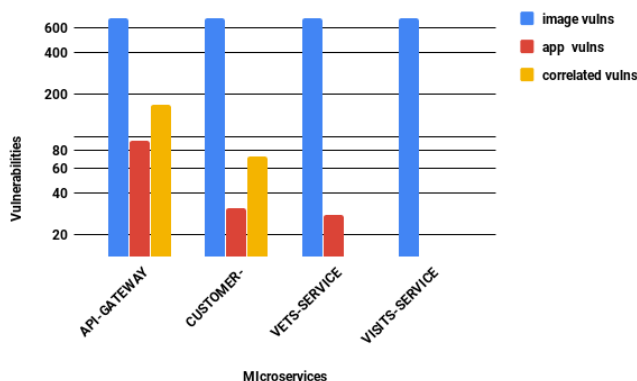


Figure 7. Homogeneous PetClinic Vulnerabilities

TABLE VII. Risk Scores By CWE

CWE-ID	OWASP T10 Risk Category	Risk Score
CWE-16	A6 - Security Misconfiguration	6.0
CWE-524	Not Listed	3.0
CWE-79	A6 - Security Misconfiguration	6.0
CWE-425	Not Listed	3.0
CWE-200	A3 - Sensitive Data Exposure	7.0
CWE-22	A5 - Broken Access Control	6.0
CWE-933	Not Listed	3.0

is derived by computing risks per PetClinic microservices to obtain the *Security Risk - SR*. Hence, we inspect the results for the image vulnerability scan and notice that the vulnerabilities are too similar (Figure 7). Therefore, *SR* will be too similar for meaningful vulnerability prioritization. Since the prioritization step is imperative for ranking microservices in order of risk severity, we compute *SR* using the ORRM (Section IV-A2). The application layer scan results are retrieved from the database and analyzed. Scores are assigned to the detected vulnerabilities based on the risk scores for OWASP Top-10 2017 web vulnerabilities. This is a reasonable approach given OWASP uses ORRM for deriving the Top-10 web application vulnerability scores. Also, this affords objective assignment of scores [35], which are publicly verifiable. Table VI is the distribution of detected vulnerabilities, while a subset of the mapping between CWE-Ids and OWASP Top-10 is on Table VII. From Table VII, it is obvious that the API-Gateway has the most severe risks followed by the Customer, Vets, and Visits microservices. Therefore, we apply diversification based on this result using a *diversification index of 3:4*, i.e., three out of four microservices. The diversified PetClinic is *retested* and the results are shown in Figure 8. We observe that the diversified PetClinic application layer vulnerabilities are reduced with about 53.3%. However, the image vulnerabilities increased especially for the Customer and Vets service, which are transformed to NodeJS and Ruby respectively. Importantly, the microservices are no longer homogeneous, and the possibilities for correlated attacks have been eliminated. Also, the vulnerabilities in the API Gateway's image are drastically reduced from 696 to 6, while the application layer vulnerabilities reduced from 94 to 24. The reduction is due to reduced code base size, a distinct characteristic of Python programming model. The API Gateway is the most important microservice since it presents the most vulnerable and sensitive attack surface of the application, therefore consider the security of PetClinic improved, our results mean that out of 94 opportunities for attacking the API Gateway, only 24 were left.

D. Attack Surface Analysis

Here we analyze the attack surfaces of the homogeneous and diversified PetClinic versions. We consider direct and indirect attack surfaces, i.e., vulnerabilities that directly/ indirectly lead to attacks respectively. From the vulnerability scan reports, each detected vulnerability is counted as an attack surface *unit* (*attack opportunities concept* [36], [37]). Figure 9 compares the horizontal app layer attack surface for both PetClinic apps. Notice a reduced attack surface in the diversified version, showing better security. Essentially, the attackability of PetClinic has been reduced. However, the results for the vertical attack surface are different. This attack surface portrays attacks transversing the app-image layer




Figure 8. Diversified PetClinic Vulnerabilities

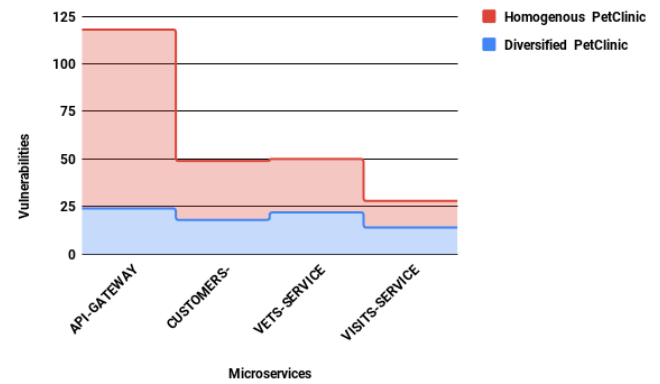


Figure 9. Horizontal Attack Surface Analysis

(Figure 5). While there are fewer correlated vulnerabilities in the diversified API-Gateway, correlated vulnerabilities in the Customers and Vets Services have increased. This increment is due to the corresponding increase of image vulnerabilities. However, the attackability due to homogeneity is reduced. We want to emphasize that intruders would observe this approach as permanently changing attack surfaces increasing dramatically the effort to breach the system.

V. CRITICAL DISCUSSION

The idea presented in Section III of an immune system like approach to remove undetected intruders in virtual machines seems to a lot of experts intriguing. Nevertheless, according to the state of the art, this is currently not done. There might be reasons for that and open questions the reader should consider.

It is often remarked that the proposal can be compared with the approach to restart periodically virtual machines that have memory leak issues and has apparently nothing to do with security concerns, and could be applied to traditional (non-cloud) systems as well. So, the approach may have even a broader focus than presented (which is not a bad thing).

Another question is how to detect “infected” nodes? The presented approach selects nodes simply at random and will hit every node at some time. The same could be done using a round-robin approach, but a round-robin strategy would be better predictable for an attacker. However, both strategies will create a lot of additional regenerations, and that leaves room for improvements. It seems obvious to search for solutions like presented by [38], [39] to provide some “intelligence” for the identification of “suspicious” nodes. Such a kind of intelligence would limit regenerations to likely “infected” nodes. In all cases, it is essential for anomaly detection approaches to secure the forensic trail [40], [41].

Furthermore, to regenerate nodes periodically or even randomly is likely nontrivial in practice and depends on the state management requirements for the affected nodes. Therefore, this paper proposes the approach only as a promising solution for Level 2 or 3 cloud applications (see Table III) that are operated on elastic container platforms. These kinds of applications have desirable state management characteristics. However, this is a limitation to applications following the microservice architecture approach.

One could be further concerned about exploits that are adaptable to bio-inspired systems. Stealthy resident worms dating back to the old PC era would be an example. This concern might be especially valid for the often encountered case of not entirely stateless services when data-as-code dependencies or code-injection vulnerabilities exist. Furthermore, attackers could shift their focus to the platform itself in order to disable the regeneration mechanism as a first step. On the other hand, this could be easily detected – but there could exist more sophisticated attacks. In order to efficiently employ this strategy, efficient *real-time* security monitoring is required. This could be achieved via two major approaches, the first requires log aggregation and analysis using either machine learning practices or other anomaly detection techniques. Otherwise, it is also possible to deploy run-time security monitoring agents in the cloud as recommended by the NIST Application Container Security Guide [42]. For example, Falco [43] is an open-source behavioral activity monitor for detecting anomalous activities in containers. When deployed, it can trigger the regeneration of new cells when malicious activities are detected, however the system has to automate the management of traffic (requests and responses) in a manner that evades service disruption.

The “immunization” results on the infrastructure level (see Section III) are impressive but should be combined with secure coding practices in development pipelines, i.e., coordinated with continuous security assessments. We presented how to automate security in CNA development environments [34]. In these cases, detected web vulnerabilities, e.g., *X-Content-Type-Options Header Missing*, can be resolved by appending appropriate *headers*, as described and advised in CAVAS reports. Furthermore, image vulnerabilities can be reduced by using more secure container images. For example, *Alpine Linux* images can replace *Ubuntu* images as base images due to smaller footprint, which equals smaller attack surfaces [44].

Of course, this is in line with the current trend of moving security of microservices *leftwards*, i.e., integrating security in the development pipelines [45]. Similarly, it is imperative to implement continuous security monitoring techniques for deployed containers to detect vulnerabilities discovered after the deployment of containers in production environments. This approach provides an efficient possibility of patching vulnerable images and redeploying the appropriate containers. However, it adds an overhead for the development pipelines

since it requires additional deployment, management and involvement of security teams/architects.

Our MTD approach presented in Section IV leverages automatic code generation techniques on the application level via Swagger CodeGen library. We discovered that over 150 companies/projects use Swagger CodeGen in production [46], hence the library is mature and capable of transforming large microservice applications. Nevertheless, in this work a basic application has been used to introduce the concepts, more complex applications will be tested in the future. There are however several obstacles to the realization of the proposed MTD, our techniques can be applied only to OpenAPI compatible microservices. This implies that the target microservices have to be compatible with the OpenAPI standard, this is achievable via several programming frameworks, and also a positive addition since OpenAPI seems to be the leading API exchange standard. Also, Swagger Codegen currently supports about 30 programming languages/frameworks and this might be a limitation in terms of possible combinations (entropy), although more languages can be added via customization. There might be a need for manual efforts to check if the transformation output is functionally compatible especially for complex applications. A possible challenge for real deployments might be the need to have development teams that are proficient in multiple programming languages.

An event-based technique might interestingly enhance our MTD technique by detecting attacks and triggering commensurate diversification. Conventionally, Web Application Firewalls (WAF) are deployed in front of web applications to detect and stop malicious traffic (which might also indicate an ongoing attack). Hence WAF can be deployed at the API Gateway and configured with *attack thresholds*. Once a threshold is breached, the WAF would trigger the diversification of the entire microservice application or only the endangered microservices. A scheduled diversification routine might support this methodology. These techniques can comfortably be applied across cloud platforms using orchestration technologies, e.g., Kubernetes.

VI. RELATED WORK

To the best of the authors' knowledge, there are currently no approaches making intentional use of virtual machine regeneration for security purposes neither on the infrastructure nor on the application level. However, the proposed approach is derived from multi-cloud scenarios and their increased requirements on security. Moreover, several promising approaches are dealing with multi-cloud scenarios. So, all of them could show equal opportunities. However, often, these approaches come along with much inherent complexity. A container-based approach seems to handle this kind of complexity better. There are some good survey papers on this [47], [48], [49], [50].

MTD via software diversity was first introduced by Forest et al. [51], since then the concept has been applied at different abstraction levels. Baudry et al. [52] introduced *sosiefication*, a diversification method, which transforms software programs by generating corresponding replicas through statement deletion, addition or replacement operators. These variants still exhibit the same functionality but are computationally diverse. Williams et al. [53] presented *Genesis*, a VM-based dynamic diversification system. Genesis employed the *Strata* VM to distribute software components such that every version became

unique, hence difficult to attack. A detailed comparison of automated diversification techniques was presented in [23]. The authors have not found a prior work that applied MTD concepts to microservices.

VII. CONCLUSION

There is still no such thing as an impenetrable system. Once attackers successfully breach a system, there is little to prevent them from doing arbitrary harm but we can reduce the available time for the intruder to do this. Moreover, we can make it harder to replay a successful attack. The presented approach evolved mainly from transferability research questions for cloud-native applications. Therefore, it is limited to microservice-based application architectures but provides some unusual characteristics for thinking about security in general.

Basically we proposed an "immune system" inspired approach to tackle zero-day exploits. The founding cells are continuously regenerated. The primary intent is to reduce the time for an attacker acting undetected massively. Therefore, this paper proposed to regenerate virtual machines (the cells of an IT-system) with a much higher frequency than usual to purge even undetected intruders. Evaluations on infrastructures provided by AWS, GCE, Azure, and OpenStack showed that a virtual machine could be regenerated between two minutes (AWS) and 10 minutes (Azure). The reader should compare these times with recent cybersecurity reports. In 2016 an attacker was undetected on a victim system for about 100 days. The presented approach means for intruders that their undetected time on victim systems is not measured in months or days any-more, it would be measured in minutes.

However, regenerated virtual machines will incorporate the same set of application vulnerabilities. So, a reasonable approach for intruders would be to script their attacks and rerun it merely. Although they might lose their foothold within minutes in a system, they can regain it automatically within seconds. Therefore, we propose to alter the attack surface of applications by randomizing the entry and exit points, which are commonly used for identifying attack surfaces [27], [28]. Based on horizontal and vertical microservice attack surfaces we demonstrated how to employ a *vulnerability correlation* to identify vulnerability similarities on the application layer and how to adapt the attack surface accordingly. This attack surface modification would let even automated and formerly successful attack scripts fail (at least partly). We propose and demonstrate the feasibility to diversify the application via dynamic transformations of its containerized components at *runtime*. In our presented use cases, we could show that it is possible to change the attack surface of a reference application incorporating over 600 container image vulnerabilities and approximately 80 application vulnerabilities to a surface with no image vulnerabilities and only 24 application vulnerabilities anymore. That is a reduction of almost 98%. What is more, the surface of the application can be changed continuously resulting that scripted attacks fail with each surface change. That is a nightmare from an intruders point of view.

The critical discussion in Section V showed that there is a need for additional evaluation and room for more in-depth research on both levels: continuously infrastructure regeneration and application surface modifying. However, several reviewers remarked independently that the basic idea is so "intriguing", that it should be considered more consequentially.

ACKNOWLEDGMENT

This research is partly funded by the Cloud TRANSIT project (13FH021PX4, German Federal Ministry of Education and Research). The authors would like to thank Bob Duncan from the University of Aberdeen for his inspiring thoughts on cloud security challenges.

REFERENCES

- [1] N. Kratzke, "About an Immune System Understanding for Cloud-native Applications - Biology Inspired Thoughts to Immunize the Cloud Forensic Trail," in Proc. of the 9th Int. Conf. on Cloud Computing, GRIDS, and Virtualization (CLOUD COMPUTING 2018, Barcelona, Spain), 2018.
- [2] B. Duncan and M. Whittington, "Compliance with standards, assurance and audit: does this equal security?" in Proc. 7th Int. Conf. Secur. Inf. Networks - SIN '14. Glasgow: ACM, 2014, pp. 77–84, [Accessed 10 February 2019]. [Online]. Available: <http://dl.acm.org/citation.cfm?doi=2659651.2659711>
- [3] N. Kratzke, "Smuggling Multi-Cloud Support into Cloud-native Applications using Elastic Container Platforms," in Proc. of the 7th Int. Conf. on Cloud Computing and Services Science (CLOSER 2017), 2017.
- [4] —, "About the complexity to transfer cloud applications at runtime and how container platforms can contribute?" in Cloud Computing and Service Sciences: 7th International Conference, CLOSER 2017, Revised Selected Papers, Communications in Computer and Information Science (CCIS). Springer International Publishing, 2018, to be published.
- [5] N. Kratzke and P.-C. Quint, "Understanding Cloud-native Applications after 10 Years of Cloud Computing - A Systematic Mapping Study," Journal of Systems and Software, vol. 126, no. April, 2017.
- [6] N. Kratzke and R. Peinl, "ClouNS - a Cloud-Native Application Reference Model for Enterprise Architects," in 2016 IEEE 20th Int. Enterprise Distributed Object Computing Workshop (EDOCW), September 2016.
- [7] L. Bilge and T. Dumitras, "Before we knew it: an empirical study of zero-day attacks in the real world," in ACM Conference on Computer and Communications Security, 2012.
- [8] K. A. Torkura, M. I. Sukmana, and A. V. Kayem, "A cyber risk based moving target defense mechanism for microservice architectures," in 2018 IEEE Intl Conf on Parallel & Distributed Processing with Applications, Ubiquitous Computing & Communications, Big Data & Cloud Computing, Social Computing & Networking, Sustainable Computing & Communications (ISPA/IUCC/BDCloud/SocialCom/SustainCom). IEEE, 2018, pp. 932–939.
- [9] FireEye, "M-trends 2019 report," [Accessed 07 November 2017]. [Online]. Available: <https://www.fireeye.com/current-threats/annual-threat-report.html>
- [10] K. Krombholz, H. Hobel, M. Huber, and E. Weippl, "Advanced social engineering attacks," Journal of Information Security and Applications, vol. 22, 2015.
- [11] S. Gupta, A. Singhal, and A. Kapoor, "A literature survey on social engineering attacks: Phishing attack," 2016 International Conference on Computing, Communication and Automation (ICCCA), 2016, pp. 537–540.
- [12] N. Kratzke and P.-C. Quint, "Technical Report of the Project Cloud-TRANSIT - Transfer Cloud-native Applications at Runtime," October 2018, technical report.
- [13] C. Fehling, F. Leymann, R. Retter, W. Schupeck, and P. Arbitter, Cloud Computing Patterns: Fundamentals to Design, Build, and Manage Cloud Applications. Springer Publishing Company, Incorporated, 2014.
- [14] A. Balalaie, A. Heydarnoori, and P. Jamshidi, "Migrating to Cloud-Native Architectures Using Microservices: An Experience Report," in 1st Int. Workshop on Cloud Adoption and Migration (CloudWay), Taormina, Italy, 2015.
- [15] S. Newman, Building Microservices. O'Reilly Media, Incorporated, 2015.
- [16] S. Ashtikar, C. Barker, B. Clem, P. Fichadia, V. Krupin, K. Louie, G. Malhotra, D. Nielsen, N. Simpson, and C. Spence, "Open Data Center Alliance Best Practices: Architecting Cloud-Aware Applications Rev. 1.0," 2014, [Accessed 10 February 2019]. [Online]. Available: https://www.opendatacenteralliance.org/docs/architecting_cloud_aware_applications.pdf
- [17] M. Stine, Migrating to Cloud-Native Application Architectures. O'Reilly, 2015.
- [18] B. Hindman, A. Konwinski, M. Zaharia, A. Ghodsi, A. D. Joseph, R. H. Katz, S. Shenker, and I. Stoica, "Mesos: A Platform for Fine-Grained Resource Sharing in the Data Center." in 8th USENIX Conf. on Networked systems design and implementation (NSDI'11), vol. 11, 2011.
- [19] C. M. Aderaldo, N. C. Mendonça, C. Pahl, and P. Jamshidi, "Benchmark requirements for microservices architecture research," in Proc. of the 1st Int. Workshop on Establishing the Community-Wide Infrastructure for Architecture-Based Software Engineering, ser. ECASE '17. Piscataway, NJ, USA: IEEE Press, 2017.
- [20] N. Kratzke and P.-C. Quint, "About Automatic Benchmarking of IaaS Cloud Service Providers for a World of Container Clusters," Journal of Cloud Computing Research, vol. 1, no. 1, 2015.
- [21] J. Gummaraju, T. Desikan, and Y. Turner, "Over 30% of official images in docker hub contain high priority security vulnerabilities," BanyanOps, Tech. Rep., 2015.
- [22] R. Shu, X. Gu, and W. Enck, "A study of security vulnerabilities on docker hub," in Proceedings of the Seventh ACM on Conference on Data and Application Security and Privacy, 2017.
- [23] P. Larsen, S. Brunthaler, L. Davi, A.-R. Sadeghi, and M. Franz, "Automated software diversity," Synthesis Lectures on Information Security, Privacy, & Trust, vol. 10, no. 2, 2015, pp. 1–88.
- [24] P. Mell, K. Scarfone, and S. Romanosky, "Common vulnerability scoring system," IEEE Security & Privacy, 2006.
- [25] OWASP, "Owasp risk rating methodology," [Accessed 08 January 2019]. [Online]. Available: https://www.owasp.org/index.php/OWASP_Risk_Rating_Methodology
- [26] Pivotal, "Distributed version of spring petclinic built with spring cloud," [Accessed 08 January 2019]. [Online]. Available: <https://github.com/spring-petclinic/spring-petclinic-microservices>
- [27] A. Younis, Y. K. Malaiya, and I. Ray, "Assessing vulnerability exploitability risk using software properties," Software Quality Journal, 2016.
- [28] P. K. Manadhata, Y. Karabulut, and J. M. Wing, "Report: Measuring the attack surfaces of enterprise software." ESSoS, vol. 9, 2009, pp. 91–100.
- [29] M. Ficco, "Security event correlation approach for cloud computing," International Journal of High Performance Computing and Networking 1, 2013.
- [30] S. Christey, J. Kenderdine, J. Mazella, and B. Miles, "Common weakness enumeration," Mitre Corporation, 2013.
- [31] M. Corporation., "Cwe-89: Improper neutralization of special elements used in an sql command ('sql injection')," [Accessed 16 May 2019]. [Online]. Available: <https://cwe.mitre.org/data/definitions/89.html>
- [32] NIST, "Cve-2016-6652 details," [Accessed 16 May 2019]. [Online]. Available: <https://nvd.nist.gov/vuln/detail/CVE-2016-6652>
- [33] P.-Y. Chen, G. Kataria, and R. Krishnan, "Correlated failures, diversification, and information security risk management," MIS quarterly, 2011, pp. 397–422.
- [34] K. A. Torkura, M. I. Sukmana, and C. Meinel, "Cavas: Neutralizing application and container security vulnerabilities in the cloud native era (to appear)," in 14th EAI International Conference on Security and Privacy in Communication Networks. Springer, 2018.
- [35] OWASP, "Top 10-2017 details about risk factors," 2017, [Accessed 07 January 2019]. [Online]. Available: https://www.owasp.org/index.php/Top_10-2017_Details_About_Risk_Factors
- [36] M. Howard, J. Pincus, and J. M. Wing, "Measuring relative attack surfaces," in Computer security in the 21st century. Springer, 2005, pp. 109–137.
- [37] OWASP, "Attack surface analysis cheat sheet," [Accessed 08 January 2019]. [Online]. Available: https://www.owasp.org/index.php/Attack_Surface_Analysis_Cheat_Sheet
- [38] Q. Fu, J.-G. Lou, Y. Wang, and J. Li, "Execution Anomaly Detection in Distributed Systems through Unstructured Log Analysis," in 2009 Ninth IEEE Int. Conf. on Data Mining, 2009.

- [39] M. Wurzenberger, F. Skopik, R. Fiedler, and W. Kastner, "Applying High-Performance Bioinformatics Tools for Outlier Detection in Log Data," in CYBCONF, 2017.
- [40] B. Duncan and M. Whittington, "Cloud cyber-security: Empowering the audit trail," *Int. J. Adv. Secur.*, vol. 9, no. 3 & 4, 2016, pp. 169–183.
- [41] —, "Creating an Immutable Database for Secure Cloud Audit Trail and System Logging," in *Cloud Comput. 2017 8th Int. Conf. Cloud Comput. GRIDs, Virtualization*. Athens, Greece: IARIA, ISBN: 978-1-61208-529-6, 2016, pp. 54–59.
- [42] M. Souppaya, J. Morello, and K. Scarfone, "Application container security guide," 2017. [Online]. Available: <https://doi.org/10.6028/NIST.SP.800-190>
- [43] FalcoSecurity, "Falco: Container native runtime security," [Accessed 08 January 2019]. [Online]. Available: <https://github.com/falcosecurity/falco>
- [44] H. Gantikow, C. Reich, M. Knahl, and N. Clarke, "Providing security in container-based hpc runtime environments," in *International Conference on High Performance Computing*. Springer, 2016.
- [45] H. Myrbakken and R. Colomo-Palacios, "Devsecops: a multivocal literature review," in *International Conference on Software Process Improvement and Capability Determination*. Springer, 2017, pp. 17–29.
- [46] SmartBear, "Swagger codegen repository," [Accessed 08 January 2019]. [Online]. Available: <https://github.com/swagger-api/swagger-codegen>
- [47] A. Barker, B. Varghese, and L. Thai, "Cloud Services Brokerage: A Survey and Research Roadmap," in *2015 IEEE 8th International Conference on Cloud Computing*. IEEE, June 2015.
- [48] D. Petcu and A. V. Vasilakos, "Portability in clouds: approaches and research opportunities," *Scalable Computing: Practice and Experience*, vol. 15, no. 3, October 2014.
- [49] A. N. Toosi, R. N. Calheiros, and R. Buyya, "Interconnected Cloud Computing Environments," *ACM Computing Surveys*, vol. 47, no. 1, May 2014.
- [50] N. Grozev and R. Buyya, "Inter-Cloud architectures and application brokering: taxonomy and survey," *Software: Practice and Experience*, vol. 44, no. 3, March 2014.
- [51] S. Forrest, A. Somayaji, and D. H. Ackley, "Building diverse computer systems," in *Operating Systems, 1997., The Sixth Workshop on Hot Topics in*. IEEE, 1997, pp. 67–72.
- [52] B. Baudry, S. Allier, and M. Monperrus, "Tailored source code transformations to synthesize computationally diverse program variants," in *Proceedings of the 2014 International Symposium on Software Testing and Analysis*. ACM, 2014.
- [53] D. Williams, W. Hu, J. W. Davidson, J. D. Hiser, J. C. Knight, and A. Nguyen-Tuong, "Security through diversity: Leveraging virtual machine technology," *IEEE Security & Privacy*, 2009.