


RESEARCH

Open Access



Classifying malware attacks in IaaS cloud environments

Noëlle Rakotondravony^{1*} , Benjamin Taubmann¹, Waseem Mandarawi¹, Eva Weishäupl², Peng Xu³, Bojan Kolosnjaji³, Mykolai Protsenko⁴, Hermann de Meer¹ and Hans P. Reiser¹

Abstract

In the last few years, research has been motivated to provide a categorization and classification of security concerns accompanying the growing adaptation of Infrastructure as a Service (IaaS) clouds. Studies have been motivated by the risks, threats and vulnerabilities imposed by the components within the environment and have provided general classifications of related attacks, as well as the respective detection and mitigation mechanisms. Virtual Machine Introspection (VMI) has been proven to be an effective tool for malware detection and analysis in virtualized environments. In this paper, we classify attacks in IaaS cloud that can be investigated using VMI-based mechanisms. This infers a special focus on attacks that directly involve Virtual Machines (VMs) deployed in an IaaS cloud. Our classification methodology takes into consideration the source, target, and direction of the attacks. As each actor in a cloud environment can be both source and target of attacks, the classification provides any cloud actor the necessary knowledge of the different attacks by which it can threaten or be threatened, and consequently deploy adapted VMI-based monitoring architectures. To highlight the relevance of attacks, we provide a statistical analysis of the reported vulnerabilities exploited by the classified attacks and their financial impact on actual business processes.

Keywords: IaaS, Malware, VM, Classification

Introduction

The cloud computing market continues to grow with spendings on public IaaS clouds having reached 38 billion U.S. dollars in 2016 [1]. Virtualization technology is the key enabler for such computing infrastructure services. Despite all advances in IT security in the past three decades, recent statistics also indicate the growth of malware activities, with a record number of over 140 million new malware samples detected in 2015 [2].

Over the years, many publications such as [3] have presented comprehensive analysis of security threats, vulnerabilities, example incidents, and countermeasures in IaaS cloud. While such reports provide a good overview on the wide range of potential problems, few publications specifically focus in-depth on the problem of malware in the context of IaaS environments and virtualization technology.

Besides, VMI is a set of techniques that allow for the inspection of VMs from outside the guest OS and the analysis of the running programs inside of it. In the VMI approach, the security monitoring software is isolated from the monitored guest VMs [4]. This isolation ensures both stealthiness and higher integrity of the diagnosis, which encourages practitioners to bring VMI capabilities into IaaS Cloud [5].

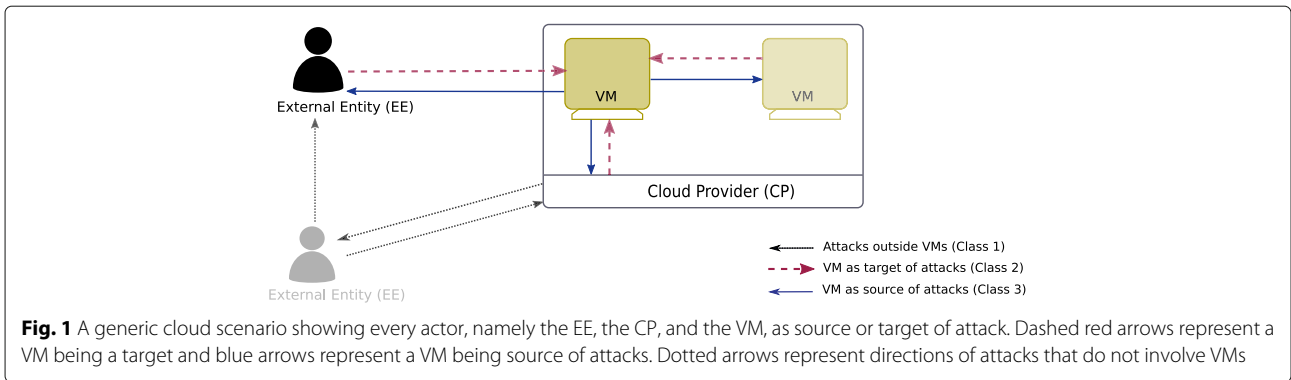
In this paper, we classify malware attacks in IaaS cloud taking into consideration their origin and target among the different actors in IaaS environments (see Fig. 1). We give an overview of attacks by which each actor can be threatened or with which it could harm other entities in the environment. To be able to deploy effective VMI-based mechanisms for analysis, monitoring or detection, in the IaaS cloud, it is necessary to have a knowledge of existing virtualization-related attacks that can be addressed using VMI. Therefore, we put a particular focus on attacks which directly involve VMs.

Overall, our study provides insight into the threats caused by malware against IaaS environments. Our results

*Correspondence: nr@sec.uni-passau.de

¹University of Passau, Passau, Germany

Full list of author information is available at the end of the article



can be used for targeting future research to develop and enhance Security-as-a-Service offerings in public cloud environments, and to raise awareness to any party willing to get involved in a cloud scenario and use VMI-based approaches as its security mechanisms against the different existing attacks and threats. The paper is structured as follows. Section “Related work” discusses related work on classifying security issues, vulnerabilities, or attacks in virtualized environments. We present our classification methods in “Threat model” section and detail each class in “Attacks outside of VM”, “VM as target of attacks” and “VM as attack source” section. “Evaluation” section presents our findings regarding the practical relevance of the identified categories and the analysis of the focus of related work. “Conclusion” section summarizes our work.

Related work

In literature, several surveys and classifications of cloud computing security issues have been presented. They include discussions from experts on the most significant threats in industrial practices [6] and from academic researchers on the risks and threats related to privacy, confidentiality, integrity, availability and accountability that the nature of the cloud service models poses [7, 8].

Diogo et al. [9] discuss the service delivery models (IaaS, Software as a Service (SaaS), Platform as a Service (PaaS)) of the cloud and the deployment models (public, private, hybrid, community, virtual private cloud). They also discuss the features that cloud computing offers but which may introduce security vulnerabilities. These are the virtualization capabilities, the multi-tenancy nature of the cloud, the trust and the standardization. Alongside the state-of-the-art of security issues in the cloud, the authors also propose a taxonomy of attacks, threats and vulnerabilities of relevant cloud components: software, storage and computing, virtualization, Internet and services, network, access, trust, and compliance and legality.

Huang et al. [10] survey the security issues in public IaaS clouds. The authors present a comparison between the perspectives of industry and academia, and demonstrate

how industry deals with threats in IaaS cloud and how researchers consider solutions related to IaaS computing and storage. The authors classify attacks based on two relevant security threats which they have identified: threats from a malicious cloud service provider, and from the clients.

Modi et al. [11] outline several vulnerabilities that are relevant to cloud environments, such as those concerning virtualization, Internet protocols, or unauthorized access to the management interface. They also discuss the potential effects of their exploitation. Furthermore, the threats are discussed according to affected cloud services and possible solutions. The corresponding attack types which jeopardize the confidentiality, integrity, and availability of cloud resources, are listed together with their effects and mitigation solutions. Finally, the discussed security issues are classified with respect to the different levels in the cloud infrastructure that they affect.

Khalil et al. [12] identify 28 security issues arising in cloud computing environments and group those into five categories: security standards, network, access control, cloud infrastructure, and data. Correspondingly, nine known attack groups have been defined, such as Denial-of-Service (DoS), cloud malware injection, and cross-VM side channels. The countermeasures to the presented attacks have been discussed and evaluated. Additionally, the authors provide an overview of the previous research in the field of cloud computing security.

Ardagna et al. [13] present a wide overview and classification of 306 scientific publications addressing vulnerabilities, threats, attacks in clouds and corresponding solutions, classified in different levels: application, tenant-on-tenant, and provider-on-tenant/tenant-on-provider. Their classification considered security properties: confidentiality, integrity, availability, authenticity, and privacy. Beyond the discussion of security issues, the authors also define security assurance as a broader term embracing methodologies for collecting and validating evidences that support security properties. The security assurance techniques considered in surveyed publications have

been classified into testing, monitoring, certification, audit/compliance, and Service Level Agreement (SLA).

Each work that deals with the study of malware and attacks in the cloud also provides recommendations for future research topics including and the design of security solutions, that are mostly based on traditional deployments such as in-guest Intrusion Detection System (IDS), in which the security monitoring agent runs on the component to be monitored. One disadvantage of in-guest deployment is that if the host operating system is tampered (e.g. with rootkit), the monitoring software becomes vulnerable and, therefore, it might not reliably return correct diagnosis.

Garfinkel et al. [4] introduce the VMI-based intrusion detection architecture in which the monitoring agent is pulled outside the monitored VM. The VMI IDS analyzes machine states and events through the Virtual Machine Monitor (VMM) interface. It leverages the isolation, inspection, and interposition properties of the VMM. Isolation protects the VMI IDS from being tampered even if the monitored VM is corrupted. Inspection and interposition properties give the VMI IDS the capabilities to inspect all states of hosted VMs and notify the user about any change or suspicious behavior. Advantages of VMI IDS are better integrity of the diagnosis compared to in-guest mechanisms and stealthiness of the monitoring system as the introspection can be performed without the introspected VM being aware of it.

The Garfinkel et al. work has been followed by various attempts to bring VMI to the IaaS cloud environments. Baek et al. [14] have developed CloudVMI which aims at bringing VMI functionality to cloud users. CloudVMI virtualizes the VMI interface by wrapping the Remote Procedure Calls (RPCs) of the introspection library LibVMI¹. Zach et al. [15] have extended the Application Programming Interface (API) of the cloud management system functions to analyze running VMs. LiveCloudInspector provides users functionalities such as taking main memory snapshots and executing volatility scripts to inspect their VMs.

In summary, the related works analyze the different possible malware attacks that might target an IaaS cloud and provide the different mitigation methodologies that can be adopted. In our work, we classify malware attacks in IaaS clouds based on the attack direction, i.e. the origin or target of the attacks. Moreover, we take into consideration the promising advantages of VMI-based security monitoring mechanisms. Therefore, we focus on attacks which directly involve VMs as a source or target. This direction-based classification allows different actors in a cloud scenario to assess the different malware attacks by which they could be threatened or might threaten other actors in the cloud, and consequently design the adequate VMI-based detection and mitigation mechanisms.

Threat model

We define a generic IaaS *cloud scenario* as the set of the following three distinct actors:

- *Cloud provider (CP)*: hosts the hardware and infrastructure, to which it has a physical access. Provider's tasks include the VMs deployment and system maintenance.
- *External entity (EE)*: can be a customer or non-customer². A cloud customer or tenant rents, uses, and manages the VMs allocated by the cloud provider according to the contracted terms of a service agreement.
- *Virtual machine (VM)*: is instantiated at the provider side upon request from a registered external entity. The conditions of VMs deployment follow a contracted agreement between a cloud provider and a customer. The VMs host the customer's services, and applications.

Several reasons put the VMs under the focus of our analysis and classification:

- In the IaaS cloud model, the known reasons behind the adoption and migration to the cloud by the users include mainly the unlimited and on-demand computational capabilities of VMs.
- Moreover, VMs host user's application ranging from daily calculation software to critical services. This makes the VMs critical components in this service model. The security of VMs has been addressed by many research works in virtualization, by leveraging traditional monitoring mechanisms at network, application, system and hardware level.
- The use of VMI-based mechanisms for security monitoring of VMs is an emerging and promising approach that has motivated research works in finding adequate heuristics that best fit users' needs and bringing its capabilities as a service in IaaS Cloud. With leveraged access control, the security monitoring of VMs can benefit from the advantages of the VMI-based techniques such as isolation and stealthiness of the monitoring agent [5].

Figure 1 sketches our classification and the interaction between entities in our classification. Our analysis aims at providing the necessary knowledge about attacks and threats surrounding each actor in a generic cloud scenario. To support future research in designing adequate VMI-based architectures, attacks that can be investigated using VMI-based mechanisms, therefore directly involving VMs, are specially covered. As Fig. 1 shows, the class of attacks between an external entity and the cloud provider (dashed line) does not involve VMs.

Attack classes

We distinguish the following classes of attacks:

- Class 1: attacks in which the malicious activities take place outside VMs.
- Class 2: attacks targeting a VM. These are attacks that let the attacker gain control over a VM or bring a malicious VM into the environment.
- Class 3: attacks originating from a VM, carried out by malware running in a VM.

We detail each class of attacks in the following sections and provide the relevant categories and taxonomy. In a general consideration, attacks from and to *external entity* apply to both client and non-client of the cloud, unless mentioned when distinction is necessary. Each analysis is summarized in a table with the attacks' source in the vertical dimension and target in the horizontal dimension.

Attacks outside of VM

In this category, we include malicious activities that neither originate from a VM nor aim at infecting a VM with malware. A unique characteristic of attacks in this category is the impossibility to detect them with mechanisms that externally monitor VMs.

EE as source of attacks

In this subsection, we consider different attacks originating from malicious external entities and targeting cloud users or the cloud provider. We distinguish between malicious cloud customers and non-customers. In fact, certain attacks can be only realized when the attacker has access to services or resources only available to registered clients.

From EE cloud customers. By overusing its allocated resources, a customer can degrade the service quality of other clients in the same cloud sector and eventually prevent others from accessing their critical resources [16]. A reported case describes a malicious customer

creating forwarding loops inside one Content Delivery Network(CDN) or across multiple CDNs [17]. Such forwarding loops caused one request to be processed repeatedly or even indefinitely, resulting in undesired resource consumption and potential DoS attacks.

From EE customers or non-customers. The Cloud Security Alliance (CSA) has listed the 12 cloud computing top threats of 2016 [18] to support cloud customers and providers in building their defenses methodologies. According to the report, traditional access and management of the computation infrastructures through dedicated APIs, or software and the usage of registered accounts are attractive for attackers. Users' accounts are targets of phishing, fraud, and credential theft attacks. Additionally, software vulnerabilities and API misconfiguration can also be exploited by attackers. These give enough power to control a user's virtual infrastructure, lead further activities at the user's expense, or expose cloud user's data at risk [16, 19]. Social engineering also need to be properly handled as it helps attackers collect valuable information from users or trigger them into manipulating their cloud services by providing a false alarming information.

Attacks targeting the cloud provider do not directly (or not necessarily) harm the VMs, but first tamper with the functionality of the cloud management system. Design flaws in the management stack can help an external attacker to infer information about the virtualized environment, execute unauthorized arbitrary commands at the cloud management, or perform a DoS attack, for instance when exploiting vulnerabilities such as described in CVE-2015-1842³ Fig. 2.

CP as source of attacks

In this subsection, we consider attacks originating from the cloud provider and targeting cloud users or the cloud provider.

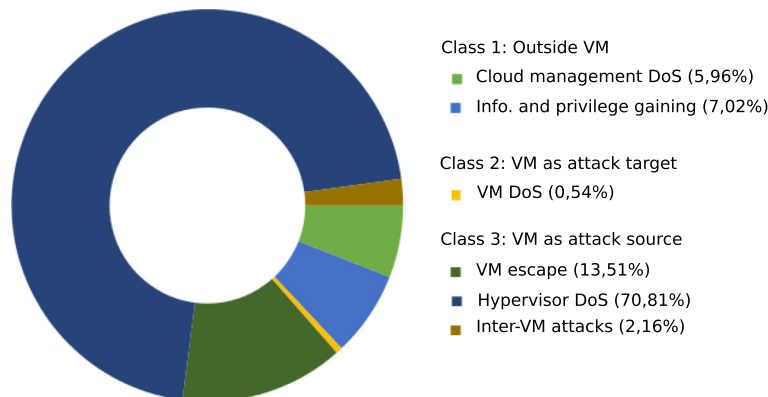


Fig. 2 CVE entries for the three classes of attacks

The cloud customers and provider can also be target of malicious activities or misbehavior from certain employees also called *malicious insiders*. The The Computer Emergency Response Team (CERT) defines a malicious insider as "a (former) employee, contractor or business partner who has or had authorized access to an organization's network system or data and intentionally exceeded or misused that access in a manner that negatively affected the confidentiality, integrity, or availability of the organization's information or information systems" [20]. According to [16], in 2014, 8% of reported cases of data breaches, data theft and illegitimate modification in public clouds were the results of insider theft conducted by malicious administrators who could directly access to customer's data. Even if the cloud services feature data encryption, malicious insiders still can access the private keys when they are stored on provider's storage components [20].

Additional harm by malicious insider includes tarnishing the company's reputation [18] that is considerably dependent on its public image, the clients' satisfaction on the delivered services, costs and security aspects [21].

A summary of the class of attacks that take place outside the VMs is given in Table 1. An external entity can be attacked by both customer and external adversary. Attacks that necessarily require the attackers to be a registered cloud customers are in **bold font**.

VM as target of attacks

In this category, we discuss malware attacks that target VMs in a cloud environment. The malicious activities attempt to corrupt a VM (bringing malware in a VM as an objective of the attack) or morph it to behave maliciously towards its owner or environment (attack serving mainly

as an intermediate stage to a more complex attack). In both cases, attacks in this class may originate from both external entities or from the cloud provider.

From external entities

Compromised repositories. A repository is a service that allows users and provider to store the VM templates or images. These are prepackaged software containing the configurations files that are used to create VMs. Securing the VM repository is the responsibility of the provider while securing the VMs is that of the customer [22]. Injecting malware into a VM image repository is a way of bringing malicious VMs into the cloud. Infected VMs are intentionally uploaded to the repository by malicious users who upload images containing malware, or unintentionally by non-malicious users who upload misconfigured templates. In general, only registered customers have the rights to upload VM templates.

Repositories containing infected images are also called *bad repositories* [23]. They pose challenging issues to cloud users and providers. In fact, the sharing of images becomes a threat as anyone who retrieves them will initialize compromised VMs, introducing vulnerabilities to his infrastructure and eventually to the cloud environment.

Co-location. In co-location attacks, an attacker aims at locating the host on which a target VM is deployed and placing his VM on the same host as the victim VM. Co-locations usually constitutes a prerequisite for other attacks such as cross-VM side-channel attacks. In general, an attacker might:

- exploit the weaknesses in VM placement algorithms and the lack of location privacy in cloud environments to gain and verify co-location. Using network-based methods, internal IP addresses assigned to instances can be mapped to availability zones and instance types, as described by the case-study for Amazon's Elastic Compute Cloud (EC2) [24].
- profile the service provided by the target, such as en-/decryption. By observing cache usage patterns, it becomes possible to verify co-location with a certainty of 50% in the worst case and up to 90% in the best case [25].
- access a metered rack Power Distribution Unit (PDU) [26], a common hardware in data centers to monitor the power consumption. Using Simple Network Management Protocol (SNMP) to access the PDU, the attacker can introduce varying load to the target VM to identify the server on which the target resides.

Attacks on migrated VMs. One of the main features of IaaS clouds is the live migration of VMs in which a running VM is moved to another server with the least possible

Table 1 Class 1: Attacks taking place outside of VMs, entries in bold require the adversary to be a registered cloud customer

Source	Target		
	External Entity	VM	Cloud Provider
External adversary	Exploit of access management fault	-	Exploit of cloud management design flaws
	Exploit of misconfigured API		Arbitrary code execution
	Social engineering		DoS attacks
	DoS attacks		
	Resource overuse		
	Forwarding loop in CDN(s)		
Virtual Machine	-	-	-
Cloud Provider	Data breach	-	Malicious insider
	Data theft		
	Data modification		

interruption. Live migration improves the flexibility of the virtual environment and allows the provider to keep the VM running with the required performance when the original host is overloaded or has to be isolated for maintenance or because of an error or attack. Live migration can also be utilized for improving the operating costs (e.g. energy) by consolidating VM. The security challenges in the cloud are more serious when migration is used, in particular if migration is performed between different widely distributed data centers [27]. Two main threats are imposed by live migration: the exploitation of the migration itself (see Section “Exploiting live migration”), and attacks on the customer VMs during migration. In the second threat, the migrated VMs might face different attacks such as man-in-the-middle, DoS and stack over-flow [28]. These attacks can be either active attacks that change the migrated data or passive attacks that perform an eavesdropping on the VM to extract sensitive data such as passwords [29]. Other migration data such as kernel memory, application state, and keys might also be sniffed or tampered if transmitted without encryption, thus compromising the integrity and confidentiality of the VM data [28].

Attacks on hosted applications. These attacks are considered in our classification because VMs might also host services that are publicly accessible. However, such attacks are not exclusive to cloud. Known attacks are for instance: SQL injection, buffer overflow, flooding and browser attacks. An attacker’s goal is to tamper with the application hinder the VM to properly deliver its services.

From cloud provider

This category of attacks from the cloud provider targeting VMs is referred to as malicious insider. In our analysis, we limit the objective of the malicious insider to owning the infrastructure or gaining customer’s confidential information contained in the VMs. However, even if owning the cloud infrastructure is not the final objective, having control on it is a necessary step to have enough capabilities to lead further damaging attacks.

Potential attack scenarios are:

- Analyzing the memory dump of targeted VM to find clear text password using a hypervisor specific command or extract user private key using tools like *rsakeyfind* [30].
- Accessing the victim’s disk partition in order to extract confidential information from its active logical volumes [30].
- Performing cold boot attacks or tampering the hardware to which the insider has access [31].
- Exploiting flaws in the integrity-protection mechanisms of the hypervisor to divert a victim VM

to a machine under his control using the basic relocation functionality [30, 31].

- Using VMI-based inspection techniques to illegally extract any information from a target VM [32].

Hyperjacking A hyperjacking attack tries to take control of the hypervisor that manages the virtual environment. This is possible when attacker injects an hypervisor beneath the original one, or has direct access to the original hypervisor. Hijacking an Operating System (OS) using a hypervisor is associated to the emergence of VM-based rootkit (VMBR) with *Subvirt*, *Vitriol* and *Blue Pill* as illustrating proof-of-concepts [33]. When the target is the host OS, hyperjacking becomes a serious attack as once the hypervisor is owned, attacker can take full control of the environment, and use any guest OS as a staging ground to attack other guests.

A direct consequence of VMBR and the hyperjacking attack on VMs is the complete control of the machines. The malicious hypervisor has the required privileges to subvert the behavior of a VM, hence causing a partial compromise (when it leaks sensitive information about the VM) or a total compromise (when it executes arbitrary code on the host with a VMM process privilege or serves as a staging ground to attack other VMs) [34].

VM theft. In this attack, an attacker copies a VM over the network or to a portable storage media in order to mount and run it elsewhere [35]. However, an attacker might also manipulate the control panel of VMM that manages live migration, in order to initiate unauthorized migration of guest VM to his own cloud infrastructure [36].

Table 2 summarizes the class of attacks by which VMs are targeted in IaaS cloud environment. For attacks which originate from other VMs we refer the reader to Table 3.

Table 2 Class 2: Attacks targeting VMs

Source	Target		
	External Entity	Virtual machine	Cloud Provider
External Entity	-	Bad repositories Attacks on apps. Co-location Attacks on migrated VMs (man-in-the-middle)	-
Virtual Machine	-	(see Table 3)	-
Cloud Provider	-	Password & key theft Disk partition access VMI-based attacks Hyperjacking VM theft	-

Table 3 Class 3: Attacks originating from VMs

Source	Target		
	External Entity	Virtual machine	Cloud Provider
External Entity	-	-	-
Virtual Machine	Botnet	Net. spoofing/sniffing	Attacks on hypervisor
	Spaming	DoS attacks	VM escape
	Resource overuse	Side channel	Hypercall attack
	Anti-VMI attack	Exploiting live migration (DoS)	DoS attacks
	Split-personality		Exploiting live migration (DoS)
Cloud Provider	-	-	-

VM as attack source

In this category, we describe the behavior of malware within a VM. In the same way as on a physical host, malware can run in a cloud VM and perform attacks against external entities outside the cloud, such as running a botnet that attacks third parties by DoS attacks, sending spam, trying to exploit other external vulnerabilities. Similarly, the Malware could just excessively use the resources and cause large cost for the owner of the VM. Such attacks are not exclusive to cloud environments. In our analysis, we distinguish three possibilities of malware behavior:

- Malware that attack other (co-located) VMs (inter-VM attacks),
- Malware that attack the hosting hypervisor in the cloud infrastructure.
- Malware that hinder legitimate introspection actions on user's VMs.

Inter-VM Attacks

Inter-VM attacks are launched from one VM to another co-residing VM through shared memory, network connections, and other shared resources without compromising the hypervisor layer. Such attacks are potentially damaging as once a VM in a cluster is compromised, other VMs become more vulnerable [37]. We highlight the following attacks:

Virtual network sniffing/spoofing. In the virtualized environment, the bridge mode, by which VM share a virtual hub to communicate, allows the guests to sniff each others' traffic. The route mode, in which a virtual switch connects each VM, allows malicious VM(s) to perform an Address Resolution Protocol (ARP) cache poisoning to redirect packets to or from other guests [38].

Denial-of-Service attacks. Inter-VM communication mechanisms targeting high-efficiency by design might sacrifice certain level of security and security requirement,

hence exposing the hosts to inter-Distributed Denial-of-Service (DDoS) attacks that exploit weak authentication mechanisms in the inter-VM communication [37]. Similarly to DoS attacks, *shared-resource consumption* attacks hinder other customers' VMs and hinder the VM from delivering hosted services.

Cross-VM side channel attacks. The isolation between co-residing VMs belonging to different customers is a critical security problem in IaaS. When an attacker succeeds to place his VM in the same host as the victim VM, for example through co-location attack (see Sect. 4), a weak implementation of the isolation mechanisms allows a malicious VM to perform sophisticated side-channel attacks by monitoring the computing resources of the host. To the best of our knowledge, most of the researches about cross-VM side-channel attacks focus on memory- and cache- based side-channels, since memory and cache are shared between VMs on the same host and hold easily accessible run-time data of VMs. The most known attacks are:

- **Prime and Probe attack.** Prime and probe attack operates in three steps. In the *Prime* phase, the attacker fills a portion of the cache with his data. In the *Trigger* phase, the attacker waits for the victim to access the memory. In the *Probe* phase, the attacker measures the time of memory access. If it is higher than a certain threshold, the accessed memory page is not cached anymore. Therefore, a memory page associated with a certain cache line was accessed by the victim (or other VMs). Such attacks have been shown to recover a full 2048 bit Rivest-Shamir-Adleman (RSA) key in Amazon [39].
- **Flush/Reload Attack.** Memory deduplication is a memory utilization optimization technique introduced by Bugnion et al. [40]. It seeks for memory pages with the same content and merges them to one physical memory page. Merged memory pages are tagged as read-only to prevent one owner from modifying them. If the page needs to be

modified, its modified version is stored in another memory page called Copy-On-Write (CoW). In a cross-VM scenario, by adding a detection step for running encryptions, a full AES key could be recovered by observing only 30,000 encryptions for the fully asynchronous attack and 10,000 for the semi-asynchronous attack [41].

Exploiting live migration. An example of such attacks is forcing the cloud management system to create many migrations, leading to a DoS attack on the VMs and the involved hosts. This can be done for example by varying the resource usages of a malicious VM to trigger live migration [42]. Another example is using a fake migration to inject a malicious VM into a host to perform VM-escape (see Section “VM-escape”) or side-channel attacks against the host and other VMs respectively.

Attacks on the Hypervisor

In virtualized environments, hypervisors are not safe from malicious VMs: 71.2% of all Xen and 65.8% of all KVM vulnerabilities could be exploited by a guest VM. The study in [43] has demonstrated how Xen’s memory is hijacked from the Dom0 domain. A CPU testing method has uncovered 117 bugs on the Linux KVM hypervisor, which introduced security vulnerabilities to the VMs, and flaws in Intel virtualization technology that cause differences in the observable behavior of code running on virtual and bare-metal servers [44]. In the following, we describe attacks targeting the hypervisor.

VM-escape. VM escape is an exploitation by which a malware running in a VM bypasses the isolation between the host and VMs and interacts directly with the hypervisor. This provides the attacker a root privilege, access to the host OS, and possibly full control over the environment.

For example, *Virtunoid* [45] is a functional guest-to-host privilege escalation attack against KVM that exploits the CVE-2011-1751⁴ bug. Another example of exploitable vulnerability for VM escape attack is described by the CVE-2008-0923 vulnerability which allows guests in some VMware products to write arbitrary file in the guest OS.

Hypercall attacks. Hypercall attacks consist of an intrusion by a malicious guest VM to the VM using well-defined hypercall interfaces and exploiting vulnerabilities in a VMM’s hypercall handler [46]. Such attacks might lead to an alteration of the VMM’s functionalities or “host crash”, upon the execution of a malicious code with VMM privileges. Recent exploitable vulnerabilities of the Xen hypervisor, enabling hypercall attacks, are detailed in the CVE-2015-7812, CVE-2015-4163, CVE-2015-4164 and CVE-2015-2752 records.

Denial-of-service (DoS) attack. In DoS attacks against the hypervisor, attackers aim at utilizing high amount of resources that is enough to considerably degrade the performance of the environment, by leveraging the hypervisor’s misconfiguration or design flaws [47]. One example by Zhou et al. [48] consists of manipulating the hypervisor’s scheduling mechanisms by a malicious guest to obtain up to 98% of total CPU cycles.

Anti-VMI Attacks

“VMI is a powerful technique for determining the specific aspects of guest VM execution from outside” [49]. In IaaS environments, VMI allows a user to run intrusion detection or necessary monitoring actions to keep his VMs secure and running properly. The monitoring tools analyze the kernel data structure of the monitored VM. Hence, a malware injected into a vulnerable guest kernel running in a VM might subvert an introspection tool and its analysis, if it succeeds to effectively manipulate the kernel data structure by adding or removing certain field of the data structure, changing its semantic, or both simultaneously [49].

Another method of anti-VMI attacks is **split personality malware**, in which a malware analyzes the environment in which it is running, and when detecting a virtualized environment, it runs only harmless code and behaves in a benign manner escaping detection [50].

Table 3 summarizes the class of malware attacks which originate from a VM and target the different actors of an IaaS cloud.

Evaluation

In this section, we consider several aspects of the classified malware attacks that involve VMs. Our objective is to help a (future) user of IaaS to have a general understanding of relevant security aspects and purpose of the attacks threatening his infrastructure, so that adequate VMI-based mitigation mechanisms can be designed.

First, we summarize the different characteristics of the classified attacks in terms of attack complexity, security impact and proposed defense measurements in the literature. Then, we analyze the statistics about the virtualization vulnerabilities exploited by the attacks, reported in publicly disclosed databases, and highlight their evolution in the course of time. Finally, we present the economic impact that attacks had and might have on business processes.

Attacks characteristics

Our classification focuses on attacks which directly involve VMs as source and target. Such attacks can be mitigated using VMI-based mechanisms. In Table 4, we summarize the VM attacks by presenting their relevant characteristics which are described as follows:

Table 4 Characterization of the different attacks originating from VMs

Attacks	Impact	Detectability	Countermeasures	Complex.	Ref.
Network-based	C/A	E: Net. monitoring	Secure channel (encryption)	Low	[38, 55]
VM-DoS	A	E: Net. monitoring	QoS management	Low	[37]
Cache side-channels	C	D: heuristics andcode, RTSC	Static & dynamic page coloring	High	[56]
Mem. deduplication side-Channels	C	D: heuristics andcode, RTSC	Fuzzy timers	High	[40, 57]
Co-residency detection	none	D: Attack specific	Unresolved problem	Low	[24]
VM Escape	C/A/I	Unresolved problem	Patching, software engineering formal verification	Low/Med.	[45, 46]
Hypervisor-DoS	A	E: Availability monitoring	Good isolation	Low	[58–60]
Hypervisor info. Gain	C	D/unresolved	Good isolation	Low	[43]
Anti-VMI Attacks	none	Unresolved problem	Difficult: Attack specific	Medium	[49, 50]

- **Impact:** if upon success of the attacks, an attacker consequently gains information to which he should not have access, executes malicious codes or modifies the legitimate functionalities of the system, or hinders the victim to deliver expected services, then the attacks have a direct impact on confidentiality (C), integrity (I), or availability (A), respectively.
- **Detectability:** defines the difficulty to detect the attacks, ranging from easy (E) to difficult (D), and the deployed methods to detect the attacks.
- **Countermeasures:** enumerates existing techniques to mitigate the attacks.
- **Complexity:** defines the difficulties faced by an attacker when performing the attacks ranging from low, medium to high.
- **References:** enumerates referenced works that detail the characterized attacks.

Table 4 describes how the VM attacks are addressed in literature. Major countermeasures deal with attacks by tackling the vulnerabilities in the design or construction of virtualization components such as hypervisors. In Section “Attacks and Vulnerability Reports”, we present a statistical analysis of the vulnerabilities in most popular hypervisors.

Attacks and Vulnerability Reports

Hypervisor, the key enabler of IaaS cloud virtualization, can also be a single point of failure [37]. Most damaging threats to a cloud platform take benefits of hypervisor vulnerabilities. In the following, we analyze the vulnerabilities of most popular hypervisors in data centers (Xen, KVM) [43] and OpenStack virtualization core components from the perspective of publicly disclosed vulnerabilities. Our directive is as follows:

- To eliminate duplication, we focus on CVE details (<http://www.cvedetails.com/>) in which a unique identifier is assigned to each vulnerability report.

- We considered entries that describe the vulnerabilities exploited by the attacks in our classification.

We consulted a total of 185 reports (Xen (159), KVM (8), OpenStack (18)) associated to our classes of attacks, listed in Section “Attack classes”. Among the reports, 12,98% relate to vulnerabilities that allow attacks in *class 1* (attacks outside VMs), 0,54% for the attacks in *class 2* (VM as attack target) and 86,5% for attacks in *class 3* (VM as attack source), see Fig. 2.

Most vulnerabilities in hypervisors are exploited to achieve hypervisor DoS attacks, VM-escape and attacks targeting the cloud management. Attacks that target VMs from outside the environment are not proper to cloud infrastructures while inter-VM attacks are not yet sufficiently studied [37].

From the considered CVEs, Fig. 3 shows the proportion of exploited vulnerabilities for DoS attacks over the last six years.

The higher percentage are observed for vulnerabilities arising from the system design which are the most exploited by attackers to realize DoS attacks against the hypervisor. Other attacks that exploit hardware issues or hypercall-handling flaws are harder to perform as they require certain administrator privileges or additional staging steps.

Economic aspect

IaaS cloud computing offers the possibility of cost saving and optimized and efficient computing, following provider’s policies and contracted agreements in which the user only pays for resources that are consumed by his deployed VMs[51]. However, the security threats by attacks involving the VMs can lead to financial harm and impacts for both the user and provider. In IaaS clouds, DoS and DDoS constitute the most common attacks and “largest threat”[51]. According to a survey [52], these attacks cost to small-to-medium sized organizations an




Fig. 3 Vulnerabilities exploited for hypervisor DoS

average of \$52,000 per incident, while \$444,000 for larger enterprises. An additional cost to the cloud provider is the penalties negotiated with users, which are due whenever a certain Quality-of-Service (QoS) constraint specified in the SLAs is not fulfilled [51].

For the cloud customer, Economic Denial of Sustainability (EDoS) is a type of DoS attacks proper to cloud environments [53]. It usually takes the form of Fraudulent Resource Consumption (FRC) where the attacker's goal is to cause considerable financial loss to its victim. For example, the victim VM continuously sends requests that consume bandwidth and causes consequent billings to its owner [9].

Both the cloud provider and the cloud user should implement cost effective solutions that reduce the risk of security breaches. Cloud providers invest in security measurements including prevention software (e.g., antivirus software) and hardware, and IT security employees to prevent security threats. In case of malware attack, the financial damages to the cloud provider include the cost of working hours for analyzing, repairing and disinfecting the systems and the losses in productivity and revenue [54]. Furthermore, the long-term damage in the provider's reputation needs to be considered: if the breach is publicly announced it can be a deterrent for future customers.

Conclusion

In this paper, we presented a classification of malware attacks in IaaS cloud environments. We defined the cloud scenario in which external entities, the VMs and the cloud provider can be both source and target of attacks. Our approach takes into consideration the origin and target of attacks to distinguish three different classes: attacks that take place outside a VM, attacks targeting VMs and attacks originating from VMs. A common characteristic of attacks in the second and third class is that they can be addressed using adequately deployed VMI-based techniques. Therefore, we added a focus on attacks that directly involve VMs as both source and target of attacks. Our classification supports practitioners at early stage of

the design of VMI-based mitigation mechanisms by identifying relevant attacks which threaten their VM or by which it can harm co-located VMs.

A statistical analysis of CVE reports on popular virtualization products highlighted how most vulnerabilities allow attackers to exploit flaws in the product design, especially to achieve DoS attacks which, from economic perspective, remain the most damaging attack and most expensive regarding financial loss for the victim or cloud provider.

Endnotes

¹ <http://libvmi.com/>

² Unless mentioned as malicious adversary, we do not make a distinction of the external entity as cloud customer or non-customer

³ <http://www.cvedetails.com/cve/CVE-2015-1842/?q=CVE-2015-1842>

⁴ All CVE details are found at <http://www.cvedetails.com/>

Acknowledgements

The research leading to these results was supported by the "Bavarian State Ministry of Education, Science and the Arts" as part of the FORSEC research association.

Authors' contributions

NR coordinated the whole paper and collected results for the evaluation. NR and HR elaborated the classification in Sections "VM as target of attacks", "Attacks outside of VM" and "Threat model" described attacks that take place outside of VM. BT provided the material on VMI based attacks. WM and HM discussed attacks concerning colocation and migration, as well as sidechannel attacks. EW described the economical aspects of the attacks. PX and BK contributed to the CDN attacker model. MP discussed in detail related work concerning threats to cloud computing. All authors contributed to the definition of the threat model, and discussion on related work. All authors read and approved the final manuscript.

Competing interests

The authors declare that they have no competing interests.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Author details

¹University of Passau, Passau, Germany. ²University of Regensburg, Regensburg, Germany. ³Technische Universität München, München, Germany. ⁴Fraunhofer AISEC, Garching bei München, Germany.

Received: 3 March 2017 Accepted: 22 November 2017

Published online: 08 December 2017

References

- Statista (2016) Global public cloud infrastructure hardware/software spending 2015-2026, by segment (fee based). Online at <http://www.statista.com/statistics/507952/worldwide-public-cloud-infrastructure-hardware-and-software-spending-by-segment/>. Accessed 1 Dec 2017
- AV-Test Institute (2016) Statistics – new malware. Online at <https://www.av-test.org/en/statistics/malware/>. Accessed 1 Dec 2017
- Hashizume K, Rosado DG, Fernández-Medina E, Fernandez EB (2013) An analysis of security issues for cloud computing. *J Internet Serv Appl* 4(1):1–13
- Garfinkel T, Rosenblum M (2003) A Virtual Machine Introspection Based Architecture for Intrusion Detection. In: *Proceedings Network and Distributed Systems Security Symposium*. The Internet Society, Reston. pp 191–206
- Taubmann B, Rakotondravony N, Reiser HP (2016) Cloudphylactor: Harnessing mandatory access control for virtual machine introspection in cloud data centers. In: *International Conference on Trust, Security and Privacy in Computing and Communications, 2016. 15th Annual. IEEE Computer Security, Los Alamitos*
- Top Threats Working Group (2013) The notorious nine: cloud computing top threats in 2013. *Cloud Secur Alliance*. Online at https://downloads.cloudsecurityalliance.org/initiatives/top_threats/The_Notorious_Nine_Cloud_Computing_Top_Threats_in_2013.pdf. Accessed 1 Dec 2017
- Xiao Z, Xiao Y (2013) Security and privacy in cloud computing. *IEEE Commun Surv Tutor* 15(2):843–859
- Subashini S, Kavitha V (2011) A survey on security issues in service delivery models of cloud computing. *J Netw Comput Appl* 34(1):1–11
- Fernandes DA, Soares LF, Gomes JV, Freire MM, Inácio PR (2014) Security issues in cloud environments: a survey. *Int J Inf Secur* 13(2):113–170
- Huang W, Ganjali A, Kim BH, Oh S, Lie D (2015) The state of public infrastructure-as-a-service cloud security. *ACM Comput Surv (CSUR)* 47(4):1:68–68:31
- Modi C, Patel D, Borisaniya B, Patel A, Rajarajan M (2013) A survey on security issues and solutions at different layers of cloud computing. *J Supercomput* 63(2):561–592
- Khalil IM, Khreishah A, Azeem M (2014) Cloud computing security: A survey. *Computers* 3(1):1–35
- Ardagna CA, Asal R, Damiani E, Vu QH (2015) From security to assurance in the cloud: A survey. *ACM Comput Surv* 48(1):2:1–2:50
- Baek Hw, Srivastava A, Merwe Jvd (2014) Cloudvmi: Virtual machine introspection as a cloud service. In: *2014 IEEE International Conference on Cloud Engineering*. IEEE Computer Society, Washington. pp 153–158
- Zach J, Reiser HP (2015) Livecloudinspector: Towards integrated IaaS forensics in the cloud. In: *IFIP International Conference on Distributed Applications and Interoperable Systems*. Springer, New York. pp 207–220
- Wueest C, Barcena MB, O'Brien L (2015) White paper: Mistakes in the IaaS cloud could put your data at risk. Technical report, Symantec
- Jianjun Chen HD, Zheng X, Liang J, Jiang J, Li K, Wan T, Paxson V (2016) Forwarding loop attacks in content delivery networks (CDN). In: *NDSS2016*. Citeseer. The Internet Society, Reston
- Top Threats Working Group (2016) The Treacherous 12: Cloud Computing Top Threats in 2016. *Cloud Security Alliance*. online at https://downloads.cloudsecurityalliance.org/assets/research/topthreats/Treacherous12_CloudComputing_TopThreats.pdf. Accessed 1 Dec 2017
- Somorovsky J, Heiderich M, Jensen M, Schwenk J, Gruschka N, Lo Iacono L (2011) All your clouds are belong to us: security analysis of cloud management interfaces. In: *Proceedings of the 3rd ACM workshop on Cloud computing security workshop*. ACM, New York. pp 3–14
- Cappelli DM, Trzeciak RF (2009) Best practices for mitigating insider threat: Lessons learned from 250 cases. In: *RSA Conference*. Springer, New York
- Luna J, Ghani H, Vateva T, Suri N (2012) Quantitative assessment of cloud security level agreements: A case study. In: *Proc. of Security and Cryptography*. pp 64–73
- Al Morsy M, Grundy J, Müller I (2010) An analysis of the cloud computing security problem. In: *Proceedings of APSEC 2010 Cloud Workshop*, Sydney, Australia, 30th Nov, x
- Liao X, Alrwais S, Yuan K, Xing L, Wang X, Hao S, Beyah R (2016) Lurking malice in the cloud: Understanding and detecting cloud repository as a malicious service. In: *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, New York. pp 1541–1552
- Ristenpart T, Tromer E, Shacham H, Savage S (2009) Hey, you, get off of my cloud: Exploring information leakage in third-party compute clouds. In: *Proceedings of the 16th ACM Conference on Computer and Communications Security, CCS '09*. ACM, New York. pp 199–212
- Inci MS, Gulmezoglu B, Eisenbarth T, Sunar B (2016) Co-location detection on the cloud. *Cryptology ePrint Archive*, Report 2016/284
- Hlavacs H, Treutner T, Gelas J-P, Lefevre L, Orgerie A-C (2011) Energy consumption side-channel attack at virtual machines in a cloud. In: *Dependable, Autonomic and Secure Computing (DASC), 2011 IEEE Ninth International Conference on*. IEEE Computer Society, Washington. pp 605–612
- Bin Sulaiman N, Masuda H (2014) Evaluation of a secure live migration of virtual machines using ipsec implementation. In: *Advanced Applied Informatics (IIAIAI), 2014 IIAI 3rd International Conference on*. IEEE Computer Society, Los Alamitos. pp 687–693
- Aiash M, Mapp G, Gemikonakli O (2014) Secure live virtual machines migration: Issues and solutions. In: *Advanced Information Networking and Applications Workshops (WAINA), 2014 28th International Conference on*. IEEE Computer Society, Los Alamitos. pp 160–165
- Anala M, Shetty J, Shobha G (2013) A framework for secure live migration of virtual machines. In: *Advances in Computing, Communications and Informatics (ICACCI), 2013 International Conference on*. IEEE, Piscataway. pp 243–248
- Rocha F, Correia M (2011) Lucy in the sky without diamonds: Stealing confidential data in the cloud. In: *2011 IEEE/IFIP 41st International Conference on Dependable Systems and Networks Workshops (DSN-W)*. IEEE, Los Alamitos. pp 129–134
- Santos N, Gummadi KP, Rodrigues R (2009) Towards trusted cloud computing. *HotCloud* 1:9
- Jain B, Baig MB, Zhang D, Porter DE, Sion R (2014) SoK: Introspections on trust and the semantic gap. In: *2014 IEEE Symposium on Security and Privacy*. IEEE Computer Society, Los Alamitos. pp 605–620
- Carbone M, Zamboni D, Lee W (2008) Taming virtualization. *IEEE Secur Priv* 6(1):65–67
- Ormandy T An Empirical Study into the Security Exposure to Hosts of Hostile Virtualized. Online at <http://taviso.decsystem.org/virtsec.pdf>. Accessed 1 Dec 2017
- Garfinkel T, Rosenblum M (2005) When virtual is harder than real: Security challenges in virtual machine based computing environments. In: *HotOS*. USENIX Association, Berkeley
- Oberheide J, Cooke E, Jahanian F (2008) Empirical exploitation of live virtual machine migration. In: *Proc. of BlackHat DC convention*. The Pennsylvania State University, Citeseer
- Zhang S (2013) Deep-diving into an easily-overlooked threat: Inter-vm attacks Whitepaper, provided by Kansas State University, TechRepublic/US2012
- Wu H, Ding Y, Winer C, Yao L (2010) Network security for virtual machine in cloud computing. In: *Computer Sciences and Convergence Information Technology (ICCIT), 2010 5th International Conference on*. IEEE, Los Alamitos. pp 18–21
- Inci MS, Gulmezoglu B, Irazoqui G, Eisenbarth T, Sunar B (2015) Seriously, get off my cloud! cross-VM RSA key recovery in a public cloud. Technical report, IACR Cryptology ePrint Archive
- Bugnion E, Devine S, Govil K, Rosenblum M (1997) Disco: Running commodity operating systems on scalable multiprocessors. *ACM Trans Comput Syst* 15(4):412–447
- Gülmezoğlu B, İnci MS, Irazoqui G, Eisenbarth T, Sunar B (2015) A faster and more realistic flush+reload attack on AES. In: *Constructive Side-Channel Analysis and Secure Design: 6th International Workshop, COSADE*. Springer, New York. pp 111–126

42. Yeh JR, Hsiao HC, Pang AC (2016) Migrant attack: A multi-resource dos attack on cloud virtual machine migration schemes. In: 2016 11th Asia Joint Conference on Information Security (AsiaJClS). IEEE Computer Society, Los Alamitos. pp 92–99
43. Perez-Botero D, Szefer J, Lee RB (2013) Characterizing hypervisor vulnerabilities in cloud computing servers. In: Proceedings of the 2013 international workshop on Security in cloud computing. ACM, New York. pp 3–10
44. Amit N, Tsafrir D, Schuster A, Ayoub A, Shlomo E (2015) Virtual CPU validation. In: Proceedings of the 25th Symposium on Operating Systems Principles, SOSP '15:311–327
45. Elhage N Virtunoid: A KVM Guest → Host privilege escalation exploit. Black Hat USA. Online at http://www.hakim.ws/BHUS2011/materials/Elhage/BH_US_11_Elhage_Virtunoid_WP.pdf. Accessed 1 Dec 2017
46. Milenkoski A, Payne BD, Antunes N, Vieira M, Kounev S (2013) HInjector: injecting hypercall attacks for evaluating VMI-based intrusion detection systems. In: Poster Reception at the 2013 Annual Computer Security Applications Conference (ACSAC 2013). ACM, New York
47. Booth G, Soknacki A, Somayaji A (2013) Cloud security: Attacks and current defenses. In: 8th Annual Symposium on Information Assurance (ASIA'13). NYSCSC, New York. pp 56–62
48. Zhou F, Goel M, Desnoyers P, Sundaram R (2013) Scheduler vulnerabilities and coordinated attacks in cloud computing. *J Comput Secur* 21(4):533–559
49. Bahram S, Jiang X, Wang Z, Grace M, Li J, Srinivasan D, Rhee J, Xu D (2010) DKSM: subverting virtual machine introspection for fun and profit. In: 29th IEEE Symposium on Reliable Distributed Systems. IEEE Computer Society, Los Alamitos. pp 82–91
50. Vishnani K, Pais AR, Mohandas R (2011) Detecting & defeating split personality malware. In: The Fifth International Conference on Emerging Security Information, Systems and Technologies. IEEE Computer Security, Los Alamitos
51. Irimie B-C, Petcu D (2015) Scalable and fault tolerant monitoring of security parameters in the cloud. In: 2015 17th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC). IEEE, Los Alamitos. pp 289–295
52. Kaspersky (2014) Global it security risks survey 2014 distributed denial of service (ddos) attacks. <https://media.kaspersky.com/en/B2B-International-2014-Survey-DDoS-Summary-Report.pdf>. Accessed 1 Dec 2017
53. Bhingarkar AS, Shah BD (2015) A survey: Securing cloud infrastructure against EDoS attack. In: Proceedings of the International Conference on Grid Computing and Applications (GCA), The Steering Committee of The World Congress in Computer Science, Computer Engineering and Applied Computing (WorldComp). CSREA Press, Athens. pp 16–22
54. Masud MM, Al-Khateeb TM, Hamlen KW, Gao J, Khan L, Han J, Thuraisingham B (2011) Cloud-based malware detection for evolving data streams. *ACM Trans Manag Inf Syst (TMIS)* 2(3):16:1–16:27
55. Miao R, Potharaju R, Yu M, Jain N (2015) The dark menace: Characterizing network-based attacks in the cloud. In: Proceedings of the 2015 ACM Conference on Internet Measurement Conference. ACM, New York. pp 169–182
56. Irazoqui G, Inci M, Eisenbarth T, Sunar B (2014) Fine grain cross-vm attacks on xen and vmware. In: Big Data and Cloud Computing (BdCloud) 2014 IEEE Fourth International Conference on. IEEE Computer Society, Los Alamitos. pp 737–744
57. Vattikonda BC, Das S, Shacham H (2011) Eliminating fine grained timers in xen. In: Proceedings of the 3rd ACM Workshop on Cloud Computing Security Workshop, CCSW '11. ACM, New York. pp 41–46
58. Wojtczuk R Subverting the Xen hypervisor. Black Hat USA. Online at <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.167.5640&rep=rep1&type=pdf>. Accessed 1 Dec 2017
59. Rutkowska J, Wojtczuk R (2008) Preventing and detecting Xen hypervisor subversions. Blackhat Brief USA. Online at <http://invisiblethingslab.com/resources/bh08/part2full.pdf>. Accessed 1 Dec 2017
60. Rutkowska J, Tereshkin A (2008) Bluepilling the Xen hypervisor. Black Hat USA. Online at <http://invisiblethingslab.com/resources/bh08/part3.pdf>. Accessed 1 Dec 2017

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► springeropen.com
