A Systematic Review of Vulnerabilities in Hypervisors and Their Detection

Full Paper

Alan Litchfield
School of Engineering, Computer and Mathematical Sciences
Auckland University of Technology
Alan.litchfield@aut.ac.nz

Abid Shahzad
Service and Cloud Computing Research Lab
Auckland University of Technology
Abid.Shahzad@aut.ac.nz

Abstract

The paper presents a systematic review of risk assessment processes to provide an overview of the risks to cloud computing and identify future research directions. This paper also provides an analysis of sophisticated threats to hypervisors and highlights vulnerabilities and exploits. Virtualization is a core feature of Cloud Computing and it is often a target for attackers. The hypervisor, which provides the virtualization layer, if compromised, can result in loss or damage to critical assets owned by Cloud Service Providers and their customers. The exploitation of hypervisor vulnerabilities provide opportunities for an attacker to launch sophisticated attacks such as Cross-VM Side Channel, Denial of Service, and Hypervisor Escape. The rate of adoption of cloud services is reflected in the lack of security controls against such sophisticated attacks and the resulting lack of trust, therefore we argue that risk assessment for hypervisors' is significant for Cloud Service Providers.

Keywords

Hypervisor, vulnerabilities, risk assessment, vulnerability exploits, vulnerability assessment, zero-day threat

Introduction

The objective of this paper is to highlight threats and vulnerabilities that permit successful attacks on common hypervisors, to find previously unknown vulnerabilities in a hypervisor, and apply measures to mitigate sophisticated threats to improve security controls. The paper also assesses means of vulnerability and risk assessment. The paper summarizes a systematic review process and notes key findings from the study.

Virtualization allows Cloud Service Providers (CSPs) to create a multi-tenanted architecture with the result that more revenue may be generated from the more effective use of infrastructure, by creating and providing virtual infrastructure services (Shoaib & Das, 2014). Virtualization involves the emulation of hardware and/or software upon which other software and applications can run. It allows the sharing of hardware resources between different tenants by creating Virtual Machines (VMs) (Chhabra & Dixit, 2015) that emulate physical server systems. The virtualization management system, the hypervisor, ensures that VMs running on the shared hardware are isolated from each other. By providing software layer between hardware and the operating system of the VM, the hypervisor ensures that VMs exist as separate entities (You, Peng, Liu, & Xue, 2012). The hypervisor controls the flow of instructions between the operating system of the guest VM and the physical hardware (Ayala, Vega, & Vargas-Lombardo, 2013), including peripheral items such as CPU core, cache memory, main memory, hard drives and network interface cards.

There are different kinds of hypervisor that provide varying levels of service and functionality. Native or bare-metal virtualization hypervisors, such as XenServer, Microsoft Hyper-V, and VMware ESX/ESXi, run on the underlying hardware. Such hypervisors have direct access to the hardware of the physical server, resulting in high levels of efficiency and performance. Consequently, guest VMs running through these hypervisors are unaware that they are running in a virtual environment. Hosted virtualization

hypervisors, for example VMware Workstation/Player, Redhat Kernel-based VM (KVM), Microsoft Virtual PC, and Oracle VirtualBox (Pek, Buttyán, & Bencsáth, 2013), run on top of the host or existing Operating System (OS). The host OS provides communication between the hypervisor and hardware. Such hypervisors are not considered to be as efficient as bare metal hypervisors.

Systematic Review Search Results

Kitchenham, Brereton, Budgen, Turner, Bailey, & Linkman (2009) describe a process for undertaking a systematic literature review in the field of software engineering. This process has been applied to address research questions in this study, to analyze the results, to identify gaps in the existing body of research, and to propose work that needs to be undertaken to address the apparent gap. For data collection, initial search questions are developed along with the literature search scope and inclusion/exclusion criteria. Once the data is collected and relevant material extracted, an analysis of the data is undertaken using NVivo software. After reviewing critical assets, threats to hypervisors, and vulnerability assessment, a thematic analysis of risk assessment frameworks is undertaken. Six themes and two research questions including sub questions are derived from the analysis of existing risk assessment frameworks.

Research questions

SLRQ1: How can the assessment of unknown vulnerabilities help to mitigate zero-day threats to hypervisors?

SLRQ2: How can the risk of the exploitation of unknown vulnerabilities' be quantified?

Scope

SLRQ1: Peer reviewed journal or conference papers, book chapters, and theses published since 2005. One exception is an article from 2000 that describes basic information and concepts about vulnerability assessment processes.

SLRQ2: Peer reviewed journal or conference papers, book chapters, and theses published since 2010.

Keywords

"vulnerability assessment", "unknown vulnerabilities assessment", "hypervisors vulnerability assessment", "zero-day threat assessment", "risk assessment" + "risk assessment for cloud computing"

Sources

Science Direct, Elsevier, Springer, IEEE, and ACM online databases.

Inclusion/exclusion criteria

The papers are included if they are written in English and incorporate either or both of two notions:

Vulnerability Assessment: Unknown vulnerability discovery process, vulnerability discovery techniques, vulnerability discovery modeling, predicting unknown vulnerabilities and vulnerability assessment process.

Risk Assessment: Security risks in cloud computing, risk assessment, risk assessment methods and risk management. They are excluded if they do not incorporate one or more of those.

55 papers were found, to which inclusion and exclusion criteria are applied. From this, 39 papers are selected. 16 research papers are excluded because they did not meet the inclusion criteria. 17 papers selected for review are related to LRQ1 and 22 papers to LRQ2. 23 papers are from conferences, 11 from journals, three are book chapters, and two dissertations.

The year range indicates that the number of papers published each year is relatively consistent (of between two and five each year). Major topics are around vulnerability discovery, risk assessment, and risk management. Other topics include vulnerability analysis (two papers), vulnerability assessment (two papers), quantitative impact and risk assessment (one paper), risk assessment model or method (two papers), and software tools (two papers).

For the purpose of the systematic literature review, the major topics are analyzed in relation to one primary notion, risk assessment. By using NVivo, a thematic analysis of risk assessment frameworks is undertaken and this produced six themes: risk assessment-as-a-service, quantitative and qualitative risk assessment, risk assessment based on graph analysis, hierarchical assessment of risk, security matrix assessment, and risk mitigation process.

Thematic analysis

In this section, the thematic analysis is summarized. Each theme is presented and briefly described, with references to the papers.

Theme 1: Risk Assessment-as-a-Service

Risk assessment-as-a-service, as an online or real time solution, may be an approach suitable for CSPs, but any implementation needs to avoid bottlenecks. For example, with the idea of offering countermeasures as a service for CC, Hussain & Abdulsalam (2011) present a theoretical implementation of Security-as-a-service (SECaaS) that can be extended through access control, auditing, and instruction detection. To assess trust for a specific CSP, Kaliski Jr and Pauley (2010) propose a risk assessment-as-a-service model for CSPs and customers. The model is applied in real-time and covers applications, tenants, and the whole cloud infrastructure. Similarly, Alhomidi and Reed (2013) present a security risk analysis-as-a-service (SRAaaS) framework that focuses on IaaS, networks, and offline storage.

Theme 2: Quantitative and Qualitative Risk Assessment

For an efficient and optimized risk evaluation process, both qualitative and quantitative risk analysis methods should be considered to provide a more thorough analysis. The literature presents both approaches, for example, to assess security of cloud based system based on six key categories, Saripalli and Walters (2010) present a quantitative impact and risk assessment framework. The categories are confidentiality, integrity, availability, trust, mutual audit ability and usability. A difficulty in determining probability is establishing a suitable basis. In this case risk is defined as a product of probability of a threat event and its impact (consequence) by establishing a security objective for each threat. Similar work is presented by Fitó & Guitart (2014) but with the basis of impact on business objectives. Both cases focus on business objectives of CSPs and not cloud customers. Combining qualitative and quantitative measures, Fitó, Macías Lloret & Guitart Fernández (2010) seek to prioritize and categorize asset risk.

Theme 3: Risk Assessment based on Graph Analysis

By analyzing attack vectors and possibilities, graphs can be used in risk assessment. This approach merges quantitative and qualitative methods. For example, Leitold and Hadarics (2012) present a mathematical assessment model (a directed graph and a matrix to discover risk) for threats that considers the communication of risks for separate entities and calculates risk for the target infrastructure. Tanimoto, Hiramoto, Iwashita, Sato, and Kanai (2011) use decision tree analysis for quantitative analysis and a risk matrix for qualitative analysis as well. Risk is determined by extracting risk factors from the users' point of view using a Risk Breakdown Structure method. The risk matrix identifies risk avoidance, mitigation, acceptance, and transference.

Theme 4: Hierarchical Assessment of Risk

One normally seeks to place risk into some kind of scale or hierarchy, for example Peiyu and Dong (2011) present a security risk assessment model that has three layers: a decomposition layer (the problem is formulated); in layer two, eight factors are used to assess risk; layer three has thirty nine factors specific to local conditions. Risk is still determined through input from domain experts. Using a different approach, Zhang, Sun, and Zhai (2012) construct an indicator system that includes principles/indicators, and sub-indicators are introduced for assessment. For example, the risk to an overall cloud environment, risk to cloud storage, and so on can be used as first indicator, then risk to Operating System, applications, or loss of availability can be used as secondary indicators for assessment. A third approach that provides a risk

	VMWare	Xen	KVM	Hyper-V
1999	1			
2000	1			
2001	1			
2002	1			
2003	4			
2004	4			
2005	8			
2006	6			
2007	25			
2008	31	2		
2009	20	2	5	
2010	24		5	1
2011	18			1
2012	34	35	5	
2013	18	43	7	1
2014	15	43	10	
2015	15	41	4	2
Total	226	166	36	5

Table 1: CVE vulnerabilities by year

management framework to support to CSPs, is described by Zhang, Wuwong, Li, and Zhang (2010) and Albakri, Shanmugam, Samy, Idris, and Ahmed (2014).

Alruwaili & Gulliver (2014) identify limitations to Zhang et al. (2010) because it doesn't cover key elements such as a cloud risk control matrix, threats, vulnerabilities, prevention and detection, risk control strategies, secure service level agreement parameters, and compliance and monitoring processes. Moreover, Xie, Peng, Zhao, Chen, Wang, and Huo (2012) raise concerns that the framework is similar to traditional security management approaches and lacks a risk communication feature, which is important for both CSP and customer to have an effective risk management for cloud environment. Therefore, making the framework ineffective for assessing risks to the cloud.

Theme 5: Security Matrix Assessment

Chandran and Angepat (2010) provide a two-factor trust matrix for risk assessment, data cost and provider history. Whereas, Zhao (2012) use existing organizational IT policies and standards, and security management. This framework highlights the necessity to start by identifying assets before evaluating risk.

Theme 6: Risk Mitigation Process

Cloud customers often own data and services, rather than the CSP. To include customer requirements in all assessments may not be feasible, but cloud customers must be considered because of the effect on their business. The framework provided by Zhang et al. (2010) undertakes risk mitigation for CSPs. The seven-process framework clarifies roles, responsibilities and accountabilities for each step and addresses the risk assessment and treatment phases. However, some components of context establishment and asset

identification, which are the primary steps to be performed, are not considered (Damenu & Balakrishna, 2015). Tanimoto et al. (2011) classify risk as risk transference, mitigation, acceptance, and avoidance.

Discussion and Conclusion

Security Vulnerabilities in Hypervisors

In this section, to further understand the nature of risk, a detailed analysis of hypervisor vulnerabilities is provided. According to Perez-Botero, Szefer, & Lee. (2013), four hypervisors provide up to 93% of the commercial market, VMware, Hyper-V, XEN, and KVM. A survey to determine the market share for hypervisors (Gallagher & White, 2012) polled around 4,000 users over two months and shows 58% of users use VMware, 16% use Hyper-V, and 13% use either KVM or XEN Server. Moreover, for the next 12 months 56% users prefer VMware, 17% prefer Hyper-V, 14% prefer KVM, and 13% prefer XEN Server. These hypervisors feature in posted vulnerabilities.

In order to highlight the security concerns of common hypervisors, a reputable vulnerability database, Common Vulnerabilities and Exposures (CVE) Details, is investigated. Common types of vulnerabilities are Denial of Service (DoS), stack overflow, memory corruption, directory traversal, and information gain. The results below are as of December 2015. The vulnerabilities are illustrated by software package.

The data provided in Table 1 is what has been made available. While it is possible this data is not complete, an initial scan shows a consistent number of vulnerabilities found in VMWare until the mid-2000's. The sudden upsurge may be explained by an increased interest in the product and CC in general rather than any specific or inherent weakness in the software design. This view is formed on the number of vulnerabilities found in the other three software packages. That is, prior to the mid-2000's there were no reported vulnerabilities in Xen, KVM, or Hyper-V. Additionally, Hyper-V has a much lower number of reported vulnerabilities than the other three packages. While this is not a definitive explanation, it is probable that the software owner (Microsoft) would prefer not to have vulnerabilities known publicly, to protect its reputation and reduce the chances of mass attacks against its OS's that use Hyper-V internally. This is a view supported by Nature (2010) and indeed Ghose, Smith, & Telang (2006) find that not publishing known vulnerabilities and publishing unpatched vulnerabilities (those that have not yet been exploited) attract fewer attacks than the publication of known and already patched vulnerabilities.

For the sake of completeness for this paper, the terms used in Table 2 are defined as follows:
Denial of Service (DoS): An attack that is intended to consume so many resources that a server or network is unavailable for legitimate users. In a DoS attack a target server or network are flooded with a large number of formatted requests (Deshmukh et al. 2015).**Code Execution:** Some software systems allow the uploading of file extensions to a web server. This results in code execution if a file is executable, such as *.asp, *.php, or *.shtml (Christey and Martin 2007).
Stack Overflow: A successful overflow occurs when a function copies data into a buffer without doing bounds checking (Christiansen 2007).
Memory Corruption: Errors in code allows an attacker to launch a memory corruption attack such as, buffer overflow, heap corruption, and format string. A successful memory corruption exploit often forces a program to crash (Chen, Xu, Nakka, Kalbarczyk, and Iyer 2005).
Cross-Site Scripting (XSS): As an attack vector, an attacker uses program input in web based applications without input validation to steal web browser cookies and credentials. It is a form of injection attack in which malicious scripts are sent to an end user. This results in the theft of cookies, user tokens, or other information being stored by a victim's web browser (Gupta and Sharma 2012). Scripts may also be able to change the content on a page, for example by redirecting a click through destination.
Directory Traversal: The purpose of the attack is to gain unauthorized access to the server file system. Typically, the attack involves the exploitation of weak sanitization or security validation of user supplied input names by including "../" or variants to access files. These vulnerabilities are not only restricted to local attack vectors and can be launched from to access the files outside of an authorized directory (Xu et al. 2006).
HTTP response splitting: Otherwise known as CRLF injection, HTTP a response splitting attack allows the insertion of a carriage return and line feed command in an HTTP header. This in turn allows an attacker to take control of the HTML page header and page content. It also leverages other attacks such as, cross-site scripting and web cache poisoning (Kshirsagar, Kumar, and Purohit 2015).

	VMWare	Xen	KVM	Hyper-V	Total
DoS	66	131	30	5	232
Code Execution	48	12	2	3	65
Stack overflow	30	28	4	1	63
Memory Corruption	8	10	2	1	21
Cross-Site Scripting	13				13
Directory Traversal	11				11
HTTP response splitting	1				1
Bypass something	5		3	1	9
Gain information	17	16	2		35
Gain privileges	54	24	7		85
Cross Site Request Forgery	3				3
Total	256	221	50	11	

Table 2: Vulnerabilities by type

Cross Site Request Forgery (CSRF): Allows an attacker to instruct the victim's browser to perform unwanted actions on a trusted website using a link or any other content. A successful attack forces a user to take an unwanted action (Lin et al. 2009).
Bypass Something: Bypassing something such as authentication allows an attacker to gain the same privileges as legitimate users of a system. In an attack scenario an attacker can bypass input validation processes and send invalid input to a server. For example, a client-side cross-site script is deployed to override input validation (Tajpour, Ibrahim, and Sharifi 2012).
Gain Information: The system configuration or software vulnerability is exploited to expose system or network access information. While exposure may not directly result in the compromise of a system, it may be an important part of a successful attack later.
Gain Privileges: Gaining high-level privilege without going through an authentication process. After exploiting a vulnerability of the system, an attacker that has gained unauthorized access to the system may execute operations like an authenticated user. Access may be limited if users have been granted the least privilege available (Provos, Friedl, and Honeyman 2003).

The pattern of vulnerabilities in the four packages in Table 2 shows a preponderance of type DoS, gain privilege, code execution, and stack overflow. By far, the weakest point appears to be immunity to DoS attacks. However, this may also be a result of an increased awareness of this type of attack and thus, the factors leading to a successful attack are better known. It is also probable that as a commonly known attack vector, there are a large number of published exploits and scripts available for attackers to draw from.

Table 1 shows 226 VMWare vulnerabilities were identified. In Table 2, with 66 DoS vulnerabilities reported, VMware is has been vulnerable to these types of attacks. However, that is not necessarily the case now if most or all of the possible DoS vulnerabilities have been found. Similarly, gain privilege (54), execute code (48), and stack overflow (30) vulnerabilities raise concerns about the security of the VMware hypervisor. If the assumption is true that there exists a relationship between the size and complexity of a software package and the number of potential vulnerabilities, then given the large size of the VMWare suite, it may be that there still exist unknown vulnerabilities that have yet to be found.

Compared with VMWare, the 166 XEN hypervisor vulnerabilities from 2008 to 2015 appear much lower (VMWare still had 175 in the same period). However, the largest number of Xen vulnerabilities, DoS, indicates a potential weakness in this area. This may also reflect the type of hypervisor (bare metal) whereas hypervisors such as KVM would benefit from the underlying OS providing a degree of protection

through resource and memory management. Notable security concerns for Xen are stack overflow (28) and gain privilege (24) vulnerabilities are.

KVM has 30 DoS vulnerabilities reported. As with VMWare and Xen, the upsurge in this type of vulnerability may be to do with a greater awareness of the type of attack and how to leverage vulnerabilities for successful exploitation. On the other hand, other types of vulnerability present in VMWare and Xen, such as cross-site scripting CRLF injection, are controlled/prevented by the underlying OS in the case of KVM and Hyper-V. That is, we would not expect to find these types of vulnerability or exploit.

Only 5 vulnerabilities have been reported for Hyper-V. As previously discussed, companies are reluctant to publish information about vulnerabilities until after they have been patched. Those that are published include DoS (5) and execute code (3). It is possible that there exist other unpublished vulnerabilities that have either been patched or not.

Vulnerabilities in common hypervisors may result in the compromise of confidentiality, integrity and availability of the critical cloud assets if they are exploited. The DoS vulnerability is the most commonly reported vulnerability. However, the evidence does not show what of these led to an exploit and observing the trend, it appears that there is a link between the number of vulnerabilities found and knowledge of the type of vulnerability in the population. We would hypothesize that as the knowledge of a type of vulnerability grows (the triggers, what code, software behaviors and so on), then more of that type are found in software. In this sense, knowledge and subsequent of vulnerability types are the consequence of a cognitive bias. Additionally, this trend also seems to indicate that there exists a perceptual fluency (Bornstein and D'Agostino 1994) toward the discovery of vulnerabilities and that this represents an instance of the mere exposure effect (Reber, Winkielman, and Schwarz 1998). That is, that people express a preference towards something merely because they are exposed to it and that the more they are exposed, the greater the effect.

Security Challenges

The dynamic nature of the CC environment makes it a very complex environment and critical assets present a potential target for attackers, undertaking a process for the assessment of existing security measures with a view to improving security is made more difficult. Protecting the hypervisor from attack is challenging as the number of zero-day vulnerability exploits increases. Furthermore, compared to a conventional client server data center environment, the deployment of security applications and solutions for the virtual environment is complicated. Thus, assessment of the cloud virtual environment for risk and recommendations for mitigation techniques is a challenging research problem. To convince cloud customers to adopt CC, free of security concerns, CSPs need an efficient and secure platform. In the study, the question of how to assess the risk posed by vulnerabilities is addressed.

Vulnerability assessment and penetration testing is a proactive approach to assess a hypervisor for vulnerabilities that may otherwise be exploited. Vulnerability assessment can be performed in two ways (Liu et al. 2012), vulnerability analysis and discovery. To find the root cause and results of an exploit, vulnerability analysis considers known vulnerabilities that have been exploited. The lessons learned may be applied to prevent a repeat of the event. Vulnerability discovery, on the other hand, seeks to find unknown vulnerabilities that may exist in software but are not yet discovered. Finding vulnerabilities and applying patches before they are exploited lessens the need for expensive remediation. Therefore, vulnerability discovery is preferred but given the complex nature of hypervisor code, it is unlikely that all vulnerabilities will be found. The influence of cognitive bias appears in both cases where the searcher may be led by their perceptual fluency to seek particular vulnerabilities. Additionally, once a vulnerability has been exploited, it is difficult to attribute any level of potential threat or harm.

The more effective the vulnerability assessment process, the more effective can the risk to a hypervisor be assessed. Thus, the vulnerability assessment process leads to a risk assessment process. In general, to cause harm or damage to an asset and decrease its value, a flaw or weakness in an asset can be exploited. According to Blank (2011), risk assessment is the combination of three key aspects: risk modeling, risk assessment approach, and risk analysis approach. The risk assessment process allows an organization to identify threats to assets, operations and other services, zero-day unknown vulnerabilities, the harm a

threat can cause to assets by exploiting a vulnerability, and the likelihood that harm may occur. A risk assessment approach can be quantitative, qualitative or a combination of the two. These approaches, defined by the modeling of risk, provide a range of values that determine risk. The modeling of risk factors and how they are combined, identified and evaluated allow for the assessment of overall risk to virtual assets.

The determination of what constitutes risk factors specific to the cloud environment is made difficult when there are competing claims for what should or should not be included. For example, Alberts, Behrens, Pethia, and Wilson (1999) claim that due to the dynamic capabilities of the hypervisor, such as multi-tenancy, on-demand-self-service, rapid elasticity, and the wide range of network access opportunities, conventional security controls and solutions are ineffective. This claim is open to challenge given that many of the same risk factors exist in cloud and non-cloud based environments. Furthermore, cloud offers service models and each requires layers of security. The security concerns also increase when services are moved from private to public cloud services (Theoharidou, Tsalis, and Gritzalis 2013). So, risk assessment is a challenge when it comes to the hypervisor-based environment. Thus, the need for an effective cloud security risk assessment process and framework that realizes an effective security plan and provides greater surety is apparent.

Conducting Risk Assessment

The objective of risk assessment is to identify the risks to virtual assets managed by a hypervisor. An optimized risk process needs to be determined so that security risks to virtual assets can be prioritized. To ensure the security of cloud customers' assets, to determine risk level, and make decisions later on, a process needs to provide some essential elements: Based on the characteristics of an attacker such as capability and motivation to exploit a vulnerability, the process reveals threats; estimate the probability that an exploit will occur; to determine the severity of a threat and assess the level of exposure; the impact that an exploit will have on virtual assets that belong to cloud customers; the likelihood that the threat will manifest; and, limitations of the risk assessment process.

As an example, the Xen hypervisor is used to assess risk. Xen is a powerful, open source hypervisor used by large CSPs and favored by open source communities. Amazon Elastic Compute Cloud (EC2) and Slicehost launched their cloud services in 2006 and both selected the Xen hypervisor. Also, Rackspace Cloud, Linode and other large CSPs use Xen to provide IaaS. As a core hypervisor for large enterprises, the security of Xen is critical. Like other hypervisors, Xen is also prone to attacks and the number of exploits reported to vulnerability databases there may be good reason for concern. Thus, an optimized risk assessment process for assessing risks to the Xen hypervisor is desirable as a platform for the development of a framework. The framework would then need to be generalizable across other hypervisors.

Summary

From the analysis of the themes, there appears to be no concise method or framework for analyzing and assessing security risks for hypervisors, thus it is possible there is a lack of consistent risk assessment process and mitigation. Furthermore, immature security controls result in lower cloud service adoption. Moreover, most of the solutions view risk from a broad perspective, instead of targeting key technologies such as virtualization or the hypervisor, which provides the base for cloud infrastructure services. Therefore, there exists a need for a risk assessment platform for the virtualization layer, specifically targeting common hypervisors, to improve the security of cloud environment and to mitigate sophisticated threats such as DoS, Cross-VM side channel and Hypervisor Escape. Such an assessment framework needs to enable CSPs to assess their virtual infrastructure and demonstrate how security risks are managed and mitigated. Such a framework ought to enable the cloud customers to review and accept the low risk percentage and define security requirements accordingly.

It is accepted that the vulnerabilities are genrally fixed, so the issue or the authors in this instance is to use this data in order to how it is that future vulnerbailities may become zero day exploits. The data provides the opportunity to find risk factors in code that lead to exploitation. The systematic review assists in understanding risk assessment models and frameworks proposed for assessing the risks to hypervisors. However, to improve and optimize the risk assessment process and specifically, risk assessment of virtualization technology, the review leads to further questions: How can the risk of sophisticated threats

to hypervisors be assessed? Based on the capabilities and motivation of the attacker to a vulnerability, how can we identify the threats and determine impact level? Can vulnerability severity and threat level be used to assess the risks to critical virtual assets provided by the underlying hypervisor? How can risks to hypervisors be mitigated? What optimized security controls may be developed? How can the application of security controls be shown to reduce risk?

As part of the ongoing research, further analysis of the environment is being undertaken and identification of measurable factors, domain experts, users of assessment tools (practitioners and auditors), and probabilistic and statistical methods are being collected. It is clear that an industry led approach is required and so partners are being canvassed. Given concerns with gaining access to proprietary systems' detailed information, initial efforts are targeted at the open source Xen hypervisor.

References

Albakri, S. H., Shanmugam, B., Samy, G. N., Idris, N. B. & Ahmed, A. 2014. "Security risk assessment framework for cloud computing environments." Security and Communication Networks, (7:11), pp 2114–2124.

Alberts, C. J., Behrens, S. G., Pethia, R. D., Wilson, W. R. 1999. "Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE) Framework," DTIC Document.

Alhomidi, M. A. & Reed, M. J. 2013. "Security risk analysis as a service." In 8th international conference for Internet technology and secured transactions. pp. 156–161.

Ayala, I. D. C. L., Vega, M., Vargas-Lombardo, M. 2013. "Emerging Threats, Risk and Attacks in Distributed Systems: Cloud computing," in Innovations and Advances in Computer, Information, Systems Sciences, and Engineering, New York: Springer, pp. 37–51.

Blank, R. M. 2011. Guide for Conducting Risk Assessments, NJ: CiteSeerX.

Bornstein, R. F., D'Agostino, P.R. 1994. "The Attribution and Discounting of Perceptual Fluency: Preliminary Tests of a Perceptual Fluency/Attributional Model of the Mere Exposure Effect". *Social Cognition*. (12: 2), pp. 103–128. doi:10.1521/soco.1994.12.2.103.

Chandran, S & Angepat, M. 2010. "Cloud Computing: Analyzing the risks involved in cloud computing environments". Proceedings of Natural Sciences and Engineering. Pp2-4.

Chen, S., Xu, J., Nakka, N., Kalbarczyk, Z. Iyer, R. K. 2005. "Defeating Memory Corruption Attacks via Pointer Taintedness Detection", International Conference on Dependable Systems and Networks, Yokohama, Japan, pp. 378-387.

Chhabra, S., Dixit, V. 2015. "Cloud Computing: State of the Art and Security Issues," ACM SIGSOFT Software Engineering Notes (40:2), pp. 1–11.

Christey, S. and Martin, R. A. 2007. "Vulnerability Type Distributions in CVE," Mitre Report, May, pp. 1-38.

Christiansen, M. 2007. "Stack Based Overflows: Detect & Exploit," SANS Institute.

Damenu, T. K. & Balakrishna, C. 2015. "Cloud security risk management: A critical review." In 9th international conference on next generation mobile applications, services and technologies, 2015 pp. 370–375.

Deshmukh, R. V. & Devadkar, K. K. 2015. "Understanding DDoS Attack & its Effect in Cloud Environment," Procedia Computer Science (49), pp. 202–210.

Fitó J O, Guitart J. 2014 "Business-driven Management of Infrastructure-Level Risks in Cloud Providers," Future Generation Computer Systems, (32), p. 41–53.

Fitó J O, Lloret M M, Fernández, J G. 2010 "Toward Business-driven Risk Management for Cloud Computing," International Conference on Network and Services Management, Niagara Falls, p. 238–241.

Gallagher, E., and White, L. 2012 "Nexenta Releases Server Hypervisor Market Share Survey Results," https://nexenta.com/company/media/press-releases/nexenta-releases-server-hypervisor-market-share-survey-results.

Ghose, A., Smith, M., Telang. R. 2006 "Does Information Security attack frequency increase with Vulnerability disclosure? - An Empirical Analysis." Information Systems Frontier. Springer. Doi:10.1007/s10796-006-9012-5

Gupta, S. and Sharma, L. 2012. "Exploitation of Cross-site Scripting (XSS) Vulnerability on Real World Web Applications and its Defense," International Journal of Computer Applications (60:14), New York.

Hussain, M. & Abdulsalam, H. 2011. "Secaas: security as a service for cloud-based applications." In Proceedings of The Second Kuwait Conference on e-Services and e-Systems. p. 8.

Kaliski Jr, B. S. & Pauley, W. 2010. "Toward risk assessment as a service in cloud environments." Hotcloud.

Kitchenham, B., Brereton, O. P., Budgen, D., Turner, M., Bailey, J. & Linkman, S. 2009. "Systematic literature reviews in software engineering–a systematic literature review." Information and software technology, (51:1), pp 7–15.

Kshirsagar, D., Kumar, S., and Purohit, L. "Exploring Usage Of Ontology For HTTP Response Splitting Attack." *Proceedings On 2015 1St International Conference On Next Generation Computing Technologies, NGCT 2015* Proceedings on 2015 1st International Conference on Next Generation Computing Technologies, NGCT 2015 (2016): 437-440.

Leitold, F. & Hadarics, K. 2012. "Measuring security risk in the cloud-enabled enterprise." In 2012 7th international conference on malicious and unwanted software.

Lin, X., Zavarsky, P., Ruhl, R., Lindskog, D. 2009. "Threat Modeling for CSRF Attacks," International Conference on Computational Science and Engineering, Vancouver, Canada, pp. 486-491.

Liu, B., Shi, L., Cai, Z., and Li, M. 2012. "Software Vulnerability Discovery Techniques: A Survey," in proceedings of the Fourth International Conference on Multimedia Information Networking and Security (MINES), Nanjing, CN, pp. 152-156.

Nature. 2010. "Security ethics - Manufacturers of computer systems should welcome researchers' efforts to find flaws." Editorial in Nature. Nature 463, 136 (14 January 2010) | doi:10.1038/463136a.

Peiyu, L., & Dong, L. 2011. "The new risk assessment model for information system in cloud computing environment". Procedia Engineering. (15:1). Pp3200-3204.

Pek, G., Buttyán, L., & Bencsáth, B. 2013 "A survey of security issues in hardware virtualization". ACM Computing Surveys (CSUR). (45:3).

Perez-Botero, D., Szefer, J., Lee, R. B. 2013. "Characterizing Hypervisor Vulnerabilities in Cloud Computing Servers," in proceedings of the 2013 International Workshop on Security in Cloud Computing, Hangzhou, CN, pp. 3-10.

Provos, N. Friedl, M. Honeyman, P. 2003. "Preventing Privilege Escalation," USENIX Security (3).

Reber, R., Winkielman, P., & Schwarz, N. 1998. "Effects of Perceptual Fluency on Affective Judgments". *Psychological Science*. (9:1): 45–48. doi:10.1111/1467-9280.00008

Saripalli, P. & Walters, B. 2010. "QUIRC: A quantitative impact and risk assessment framework for cloud security." In 2010 IEEE 3rd international conference on cloud computing. pp. 280–288.

Shoaib, Y., Das, O. 2014. "Pouring Cloud Virtualization Security Inside Out," arXiv preprint arXiv:1411.3771, November, pp. 1-13.

Tajpour, A., Ibrahim, S., Sharifi, M. 2012. "Web Application Security by SQL Injection Detectiontools," IJCSI International Journal of Computer Science Issues (9:2), pp. 332–339.

Tanimoto, S., Hiramoto, M., Iwashita, M., Sato, H. & Kanai, A. 2011. "Risk management on the security problem in cloud computing. In first ACIS/JNU International Conference on Computers, Networks, Systems and Industrial Engineering (cnsi), 2011 pp. 147–152.

Theoharidou, M., Tsalis, N., and Gritzalis, D. 2013. "In cloud we trust: Risk-Assessment-as-a Service," in proceedings of the IFIP International Conference on Trust Management, Malaga, ES, pp. 100-110.

Xie, F., Peng, Y., Zhao, W., Chen, D., Wang, X. & Huo, X. 2012. "A risk management framework for cloud computing." In 2012 IEEE 2nd international conference on cloud computing and intelligence systems (1:1) pp. 476–480.

Xu, W., Bhatkar, S., Sekar, R. 2006. "Taint-Enhanced Policy Enforcement: A Practical Approach to Defeat a Wide Range of Attacks," 15th Usenix Security Symposium, Stony Brook, NY, pp. 121-136.

You, P., Peng, Y., Liu, W., Xue, S. 2012. "Security Issues and Solutions in Cloud Computing," in proceedings of the 32nd International Conference on Distributed Computing Systems Workshops (ICDCSW), Macau, CN, pp. 573–577.

Zhang, J., Sun, D. & Zhai, D. 2012. "A research on the indicator system of cloud computing security risk assessment." In 2012 international conference on quality, reliability, risk, maintenance, and safety engineering.

Zhang, X., Wuwong, N., Li, H. & Zhang, X. 2010. "Information security risk management framework for the cloud computing environments." In 2010 IEEE 10th international conference on Computer and information technology, pp. 1328–1334.

Zhao, G. 2012. "Holistic framework of security management for cloud service providers". IEEE 10th International Conference on Industrial Informatics. pp 852-856.