Vulnerability and Risk Assessment of XEN Hypervisor

Completed Research

Alan Litchfield
Service and Cloud Computing Research Lab
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

A vulnerability prediction and risk assessment process for the Xen hypervisor that predicts the number of vulnerabilities and levels of risk a specific software version provides is presented. The hypervisor is a key component of virtualisation and is thus a target of attackers. When such critical infrastructure is compromised, then the assets of service consumers are consequently at risk. The benefit of a risk analysis process is that it provides surety for Cloud services consumers (making the Cloud Computing option more attractive) and assists Systems Administrators in decision making about software choices and version upgrades. The process has been tested on three popular open source, infrastructure level software packages. In each case, the level of predictive accuracy is excellent to good. The study combines quantitative and qualitative methods to predict vulnerabilities and determine risk levels.

Keywords

Vulnerability and risk assessment, Xen hypervisor, threat actors, user classification, threat level identification.

Introduction

In this paper, a vulnerability and risk assessment process for hypervisors is presented. Working from the understanding that virtualisation is the basis of Infrastructure as a Service (IaaS) (Shoaib & Das, 2014) and that Virtual Machines (VMs) emulate the physical server system (Chhabra & Dixit, 2015), IaaS allows the sharing of hardware resources between tenants. Thus, the hypervisor enables the creation and management of VMs. However, hypervisors provide a point of weakness for attackers to gain unauthorised access to cloud assets. A successful exploitation of the hypervisor can allow an attacker to take control of all its guest VMs and also the data stored on those VMs (Shahzad & Litchfield, 2015). A fallacy exists, that the hypervisor is assumed to be secure and robust, but just like any other software package, they contain vulnerabilities. Exploitation of vulnerabilities provide opportunities to launch further attacks to cause harm or damage to critical virtual assets (Kortchinsky, 2009; Wojtczuk, 2008; Elhage, 2011; Rutkowska & Wojtczuk, 2008). Also, hypervisors are vulnerable to a range of threats that raise further hypervisor security concerns (Brohi, Bamiah, Brohi, & Kamran, 2012; Reuben, 2007; Bazargan, Yeun, & Zemerly, 2012; Sabahi, 2011).

Vulnerability exploitations in common hypervisors may result in the compromise of confidentiality, integrity and availability of the critical cloud assets. For this study, we hypothesise that as the knowledge of a type of vulnerability grows (the triggers, what code, software behaviours, and so on), then more of that type are found in software. In this sense, knowledge and subsequent exploitation of vulnerability types are the consequence of cognitive bias (that is, we tend to find or see that which is in our consciousness). Since Xen hypervisor is a large open-source software system, some unknown vulnerabilities are likely to exist and due the effect of cognitive bias, Xen becomes a suitable candidate for vulnerability examination.

Further, unknown vulnerability exploitations via sophisticated threats pose a serious risk to cloud assets. The risk assessment of cloud infrastructure is very challenging for the customers because Cloud Service Providers (CSPs) typically do not share location information, infrastructure details, security policies, and other details with customers. So assessing and managing risk in a virtual cloud environment is a challenging problem (Cayirci, 2015). However, to make informed security decisions, customers desire to know the risk and severity levels. A review of current literature alleviates some of those concerns, but the research tends to present risk assessment from a broad perspective instead of focusing on the risk assessment of specific infrastructure elements such as the hypervisor. It is this reason that we address this specific issue in this study.

To make the process relevant to cloud service consumers, we have conducted the assessment from the IaaS consumer perspective. The assessment process addresses unknown vulnerabilities and determines their impact ratings. In this sense, the process may help IaaS consumers identify threats and determine their likelihood levels then vulnerability impact ratings and threat likelihood levels are mapped to determine risk and severity levels. Ultimately, in deciding to adopt Xen-based infrastructure services, risks can be prioritised against severity level acceptability.

Prior research

Hypervisors manage virtual infrastructure and allow CSPs to offer IaaS (Shoaib & Das, 2014). However, like other software packages, hypervisors contain vulnerabilities such as insufficient Authentication, Authorization and Accounting (AAA) controls, lack of resource isolation, allows network probing, lack of prevention of side-channel communications, and eavesdropping of communication between VM and the host OS (Catteddu & Hogben, 2009). Exploitation of these vulnerabilities results in the loss of confidentially, data integrity and availability, access to IaaS resources, and may destroy the Cloud environment (Dawoud, Takouna, & Meinel, 2010).

While risks vary between CSPs, depending on the type of hypervisor, security controls, security procedures, and risk management methods in place. Vulnerability exploitation scenarios raise concerns about the security of hypervisors (Wang, Liu, & Liu, 2012). The exploitation of vulnerabilities from sophisticated threats lead to risks such as loss of business reputation due to co-tenant activities, isolation failure, malicious insider behaviour, interception of data in transit, data leakage, malicious network probes or scans, and privilege escalation (Catteddu & Hogben, 2009).

What the literature highlights is that instead of targeting the hypervisor, it takes a broad approach to the risk to Cloud Computing (CC) (Hussain & Abdulsalam, 2011; Saripalli & Walters, 2010; Fitó & Guitart, 2014; Leitold & Hadarics, 2012; Tanimoto, Hiramoto, Iwashita, Sato, & Kanai, 2011; Litchfield & Shahzad, 2017). That is, the literature considers risks to investment and general security risks.

From these issues, it is our position that vulnerability and risk assessment for hypervisors and accountability of CSP security controls and procedures are significant. We believe that the determination of acceptable risk severity levels would encourage IaaS customers to adopt cloud services. We consider that risk is related to vulnerabilities and threats and the likelihood of threat occurrence (Zhang, Wuwong, Li, & Zhang, 2010).

Vulnerability & Risk Assessment Process

Since existing research primarily focuses on risk assessment of CC from the CSP perspective, the involvement of customers in the risk assessment process is very limited. In most cases, the customer needs to trust the risk assessment and mitigation processes, security controls and procedures, dictated by the CSP. The vulnerability and risk assessment process we present (**Error! Reference source not found.**) would benefit IaaS customers to independently assess risk and make informed decisions. We have taken a probabilistic approach, Time Series Winter's Method (TSWM), to risk assessment in order to present what may occur, and if it does, how severe.

In the first step, TSWM is used to predict unknown Xen vulnerabilities that may appear in the near future (Roumani et al., 2015). TSWM is useful when there is no fundamental knowledge about the factors that can affect the data used as variables for the predictions. The time series seeks patterns in variables from past movements and uses those patterns to predict future movements in data. During the risk assessment, CVSS is used to score vulnerabilities and produce impact ratings (Mell, Scarfone, & Romanosky, 2006, 2007) as Critical, High, Medium, and Low. They are then used to score three Xen vulnerability exploitation scenarios. However, since the ratings do not cover all the scenarios, the European Union Agency for Network and Information Security Agency (ENISA) risk framework (Catteddu & Hogben, 2009) has been adopted to realise a complete risk assessment process. Note that to score and determine the impact ratings of Apache HTTP and Squid Proxy servers, CVSS is used instead of ENISA.

Then to determine the likelihood of threats to Xen, a structured analysis approach using attack trees and graphs are used (Noel, Jajodia, Wang, & Singhal, 2010; Hutle, Hansch, & Fitzgerald, 2015). An attack graph is a combination of attack trees and typically focuses on the actions of an attacker and how it interacts with the target system whereas attack trees typically analyse the consequences of a successful attack. An attack graph uses a model analysis technique to identify all possible actions and Attack Vectors (AVs) that an attacker might employ to exploit Xen hypervisor. Capability and motivation characteristics of two Threat Actors (TAs) are used to identify initial threat levels. The threat levels are assigned to the source nodes of

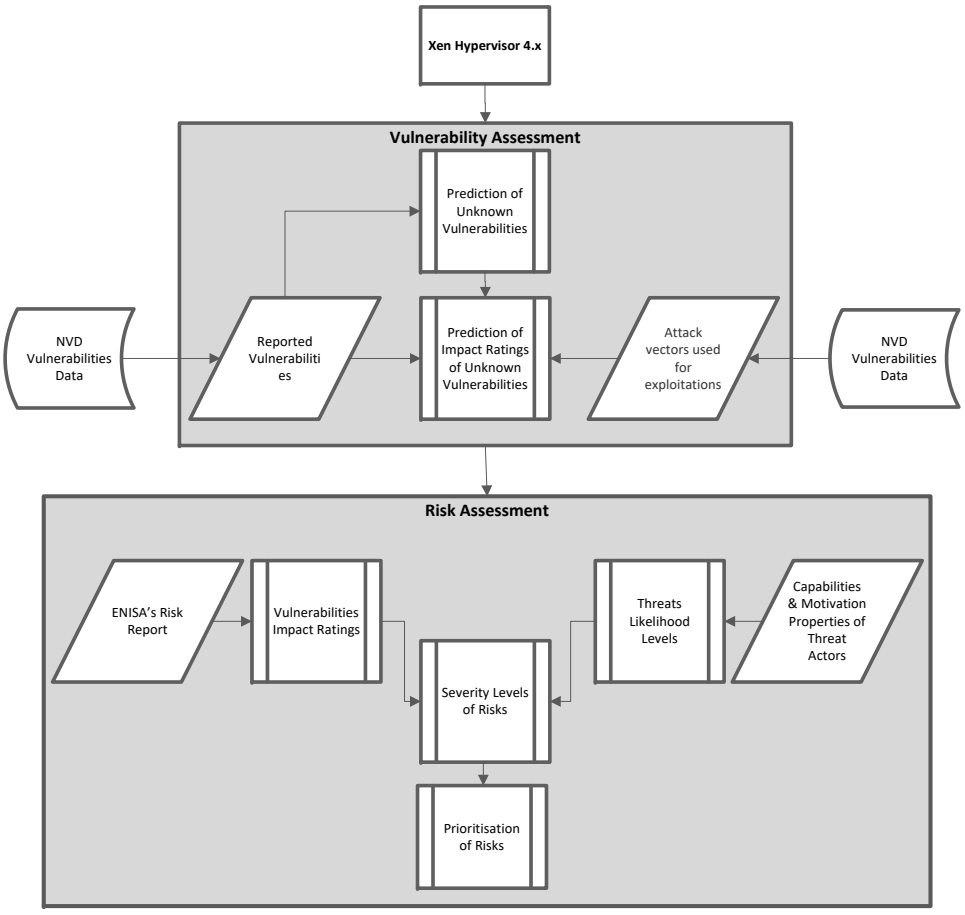

Figure 1. Vulnerability and Risk Assessment Process

Time (t)	Year	Quarter	Reported Vulnerabilities (D_t)	Time (t)	Year	Quarter	Reported Vulnerabilities (D_t)
1	2013	1	7	11		3	5
2		2	7	12		4	15
3		3	13	13	2016	1	5
4		4	16	14		2	16
5	2014	1	9	15		3	6
6		2	20	16		4	1
7		3	4	17	2017	1	17
8		4	12	18		2	5
9	2015	1	9	19		3	22
10		2	12	20		4	18

Table 1. Xen Reported Vulnerabilities

an attack graph. Later, the threat levels are propagated through the attack graph as physical, local, and network AV exploits. This provides the overall threat likelihood level to the target hypervisor. Qualitative threat likelihood levels are Severe, Substantial, Moderate, Low, and Negligible.

In the final step, risk severity levels to Xen are identified via a risk estimation matrix (Catteddu & Hogben, 2009). Using the matrix, the vulnerability impact ratings and threat likelihood levels are then mapped to determine severity levels. Eight risk types are listed, and their qualitative severity levels are determined as Low (1–3), Medium (4–6), and High (7–9).

Xen Vulnerability Assessment

Xen vulnerability assessment is performed, looking for unknown vulnerabilities in 2018 using TSWM. NIST (2017) reports 219 vulnerabilities from 2013–2017 where the most significant number of vulnerabilities are DoS related. This may reflect the type of hypervisor (bare metal) vs hosted hypervisors because the hosted version may benefit from the underlying OS protecting resource and memory management.

Prediction of Xen Unknown Vulnerabilities

To make the predictions, Xen 4.x reported vulnerabilities data in 2013–2017 retrieved from NVD (Table 1). Data are recorded quarterly and annualised. A growth trend is observed, and Level, Trend, and Seasonal factors estimated by deseasonalising actual data (D_t) using linear regression.

Deseasonalising D_t results in deseasonalised data that would have been there without seasonal fluctuations. P is the number of periods and as the actual data cycle repeats after every four periods (quarters), the value of $p = 4$.

$$\overline{D}_t = \frac{D_t - \left(\frac{p}{2}\right) + D_t + \left(\frac{p}{2}\right) + \sum_{i=t-1-\left(\frac{p}{2}\right)}^{i=t-1+\left(\frac{p}{2}\right)} 2D_i}{2p} \tag{1}$$

Equal weight is given to each season to average consecutive periods, p, of D_t. The average of D_t from $l+1$ to $l+p$ provides deseasonalised data for $l + ((p+1))/2$. p is even in this case, so deseasonalised data at a point between $l + (p/2)$ and $l + 1 + (p/2)$ is calculated. The deseasonalised data for $l + 1 + (p/2)$ is calculated by averaging the $D_t$values $l+1$ to $l+p$ and $l+2$ to $l+p+1$. Equation 1 is used to calculate, $\overline{D}_t$, for the period t, where p is even.

To perform the linear regression using the relationship between deseasonalised data and time(t), based on the change in D_t over time (Equation 2). $\overline{D}_t$ is the deseasonalised data over time t, L is the Level, and T is the rate of growth at t_0. L and T at t_0 are calculated using linear regression. $\overline{D}_t$ is a dependent variable and time, t, is independent. Using Table 1**Error! Reference source not found.**, initial Level, L_0, is 12.18 (intercept coefficient) and initial Trend, T_0, is -0.16 (X variable or slope). Therefore, $\overline{D}_t$ is calculated for any t.

$$\overline{D}_t = L + T \times t \tag{2}$$

D_t is the ratio to deseasonalised data, $\overline{D}_t$, is used to calculate the deseasonalised values for initialisation, $\bar{S}_t$ (Equation 3), where $\bar{S}_t$ for t.

$$\bar{S}_t = \frac{D_t}{\overline{D}_t} \tag{3}$$

After calculating $\bar{S}$ values for initialisation, S_t for a given time period is obtained by averaging $\bar{S}_t$ that correspond to similar time periods. For example, $p = 4$, so $\bar{S}$ are similar at time periods t, t_1, t_5, t_9, t_{13}, and t_{17}. Therefore, seasonal factors for these time periods are calculated as the average of five seasonal factors. S_t is obtained using Equation 4 for all the time periods in Table 1**Error! Reference source not found.**, $pt + i, 1 \le i \le p$, for given data cycles r.

$$\bar{S}_i = \frac{\sum_{j=0}^{r=1} S_j p + i}{r} \tag{4}$$

As actual data has 20 periods and periodicity of $p = 4$, then it has $r = 5$ seasonal cycles. Thus: $S_1 = 0.9$, $S_2 = 1.13$, $S_3 = 0.99$, and $S_4 = 1.21$.

Predictions

Before predicting for 2018, Level (L_t), Trend (T_t), and Seasonal factors $S_{5\ldots20}$ are calculated. Three smoothing parameters, α, β, and γ, are applied. α is used to calculate L_t (Equation 5), T_t is calculated using β (Equation 6), and γ is used to calculate S_t (Equation 7).

$$L_t = \alpha \frac{D_t}{S_t} + (1-\alpha)(L_{t-1} + T_{t-1}) \tag{5}$$

$$T_t = \beta(L_t - L_{t-1}) + (1-\beta)T_{t-1} \tag{6}$$

$$S_{t+p} = \gamma\left(\frac{D_t}{L_t}\right) + (1-\gamma)(S_t) \tag{7}$$

$L_0 = 12.18$ and $T_0 = -0.16$, and $S_{1\ldots4}$ are known. Therefore, TSWM is applied to predict the unknown vulnerabilities for the desired Time periods. The prediction for the first quarter of the year 2013 is obtained using Equation 2. 10.83 vulnerabilities are predicted for the first quarter of 2013. As a high value is predicted for the first quarter, L_t, T_t, and S_t are updated. Therefore, Level (L_1) is updated with an assumption that $\alpha = 0.10$, Trend (T_t) with $\beta = 0.20$, and S_5 is updated with an assumption that smoothing parameter $\gamma = 0.30$.

Calculating L_t, T_t, and S_t for all 20 periods, Equation 2 is used to predict unknown vulnerabilities for four quarters in 2018.

$$F21 = [L_{20} + (T_{20} \times 1)] \times S_{17} = 8.96$$
$$F22 = [L_{20} + (T_{20} \times 2)] \times S_{18} = 16.89$$
$$F23 = [L_{20} + (T_{20} \times 3)] \times S_{19} = 9.31$$
$$F24 = [L_{20} + (T_{20} \times 4)] \times S_{20} = 12.94$$

Xen Risk Assessment

Risk assessment is a process of risk management (Alturkistani & Emam, 2014; Kiran, 2014) that allows someone to determine the impact of vulnerabilities, the likelihood of threats, and severity levels of risks. Through risk evaluation, risks can be recorded and prioritised for decision making. Risk assessment can be quantitative (probability or proportion) or qualitative (normative or descriptive). It can also be inductive or

Vulnerability Impact	Threat Likelihood				
	Negligible	Low	Moderate	Substantial	Severe
Very Low	1	2	3	4	5
Low	2	3	4	5	6
High	4	5	6	7	8
Very High	5	6	7	8	9

Table 2. Risk Estimation Matrix

deductive (Cayirci, 2015). In this study, a qualitative inductive risk assessment process is applied to provide a risk estimation matrix (Table 2).

Vulnerability Impact Analysis

The impact of the loss of the factors confidentiality, integrity, and availability of security objectives for Xen vulnerabilities are rated as Very High, High, Medium, Low, and Very Low. We know that vulnerabilities often have similar exploitable threats (Albakri, Shanmugam, Samy, Idris, & Ahmed, 2014), for example, through cognitive bias, otherwise they pose no risk.

Vulnerability impact ratings are determined by using the CVSS Base Metric Group (Mell et al., 2006, 2007), where three scenarios and two TAs are presented to score Xen vulnerabilities.

Scenario 1 Presents a Privileged User (PU) TA that impacts availability and confidentiality by exploiting physical vulnerabilities, shutting down the Xen server, or stealing data. Vulnerability impact rating is Medium (6.8 Base score).

Scenario 2 A Normal User (NU) TA exploits a local hypervisor vulnerability by misusing privileged guest VM user space to impact confidentiality and integrity security objectives. A High impact rating (7.1 base score).

Scenario 3 A NU TA uses a network AV to exploit a vulnerability in the network stack of Xen to impact confidentiality, integrity, and availability security objectives. A High impact rating (8.5 base score).

Since it is not practicable to list a large number of possible scenarios, impact ratings are adopted from ENISA's risk framework (Table 3). Eight risk categories are identified, related to hypervisor vulnerabilities, with vulnerability and asset category subsets. There can be other risks and vulnerabilities but only common risks are presented in this paper. For example, Risk 1 is related to vulnerabilities V1 and V2, and assets A1, A2, A3, and A4. A High impact rating is determined because the exploitation of both V1 and V2 can result in the loss of service delivery and data that belong to customers. Also, the exploitation can affect the reputation of IaaS customer organisations.

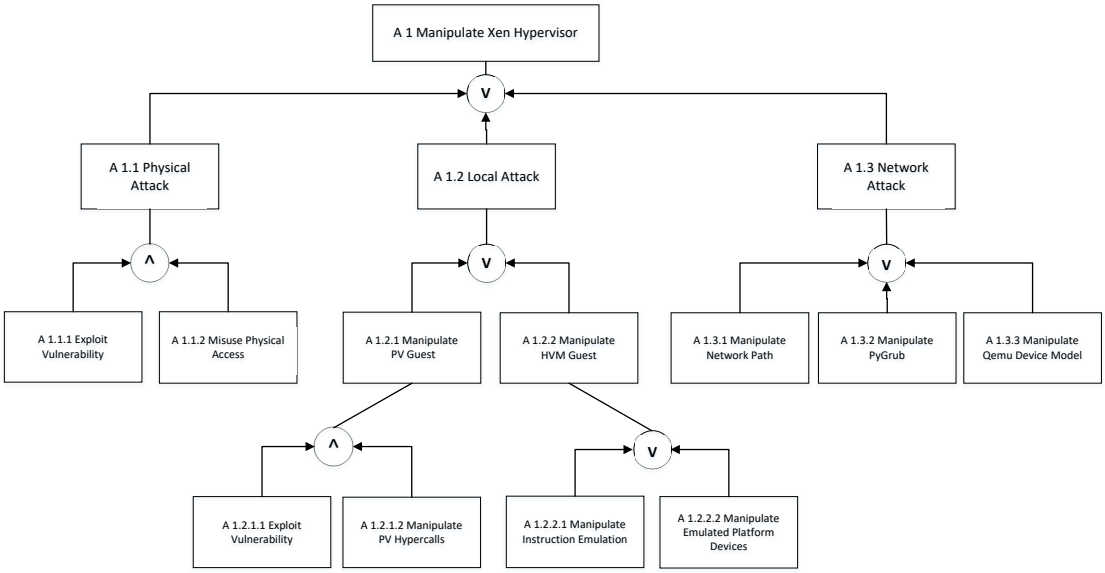

Figure 2. Xen Attack Graph

Risk No.	Risk Type	Vulnerabilities	Assets	Impact Rating
R1	Loss of Business Reputation Due To Co-Tenant Activities	V1, V2	A1, A2, A3, A4	High
R2	Isolation Failure	V1, V2, V3, V4	A1, A2, A3, A4, A5	Very High
R3	Malicious Insider	V5, V6	A1, A2, A3, A4, A5, A6, A7	Very High
R4	Intercepting Data In Transit	V3, V4, V5	A1, A2, A3, A5, A7, A8	High
R5	Data Leakage	V3, V4, V5	A1, A2, A3, A5, A6, A7, A9, A10	High
R6	Undertaking Malicious Probes Or Scans	V3, V4	A1, A4, A5	Medium
R7	Compromise Hypervisor	V1, V2	A2, A3, A4	Very High
R8	Privilege Escalation	V2, V5	A2, A3, A11, A12	High

Key:

[V1] Lack of resource isolation

[V2] Hypervisor code vulnerability

[V3] Possibility that internal (cloud) network probing will occur

[V4] Possibility that co-residence checks will be performed

[V5] Authentication, authorisation, and accounting (AAA) vulnerabilities

[V6] Inadequate physical security procedures

[A1] Company reputation

[A2] Personal sensitive data

[A3] Personal data

[A4] Service delivery

[A5] Customer trust

[A6] Employee loyalty and experience

[A7] Intellectual property

[A8] Backup or archive data

[A9] Credentials

[A10] Cloud service management interface

[A11] Access control

[A12] User directory - data

Table 3. Correlation of Risk, Vulnerabilities, Assets and Impact Ratings

Threat identification and Likelihood Analysis

With the potential attacker's goal determined, we next identify threats and determine their likelihood levels. These include impact confidentiality, integrity, and the availability of security objectives of virtual assets. The Xen attack graph (Figure 2) is a directed graph (V, E), where every node *n* in V is labelled with either a logical AND or OR, and the attack graph should have at least one source node. Source nodes in the graph that do not have incoming edges correspond to the source nodes of an attack tree.

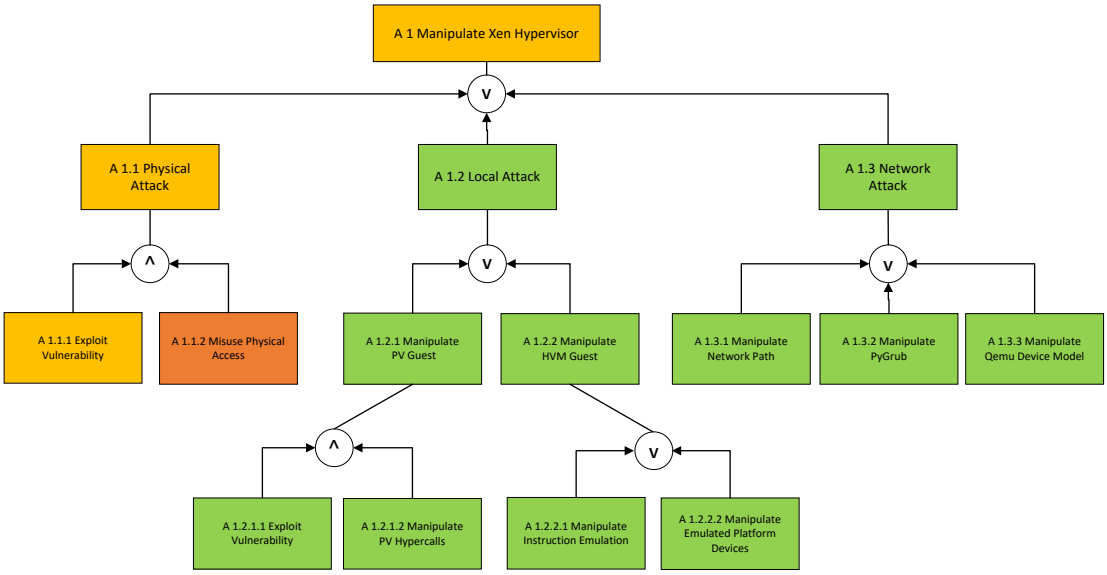

Figure 3. Threat Likelihood Levels from PU TA

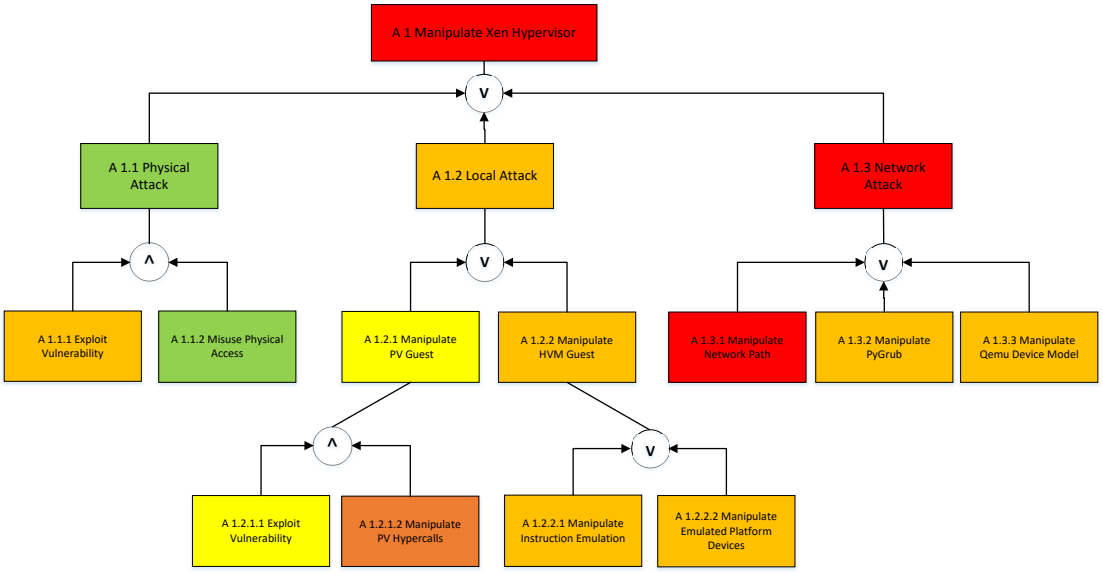

Figure 4. Threat Likelihood Levels from NU

The TA-specific source node assessment is performed by assigning a threat level to each source node. Later, threat levels through the graph (from the source node to root node) are propagated. Each node in the attack graph is either assigned or propagated with a threat level value which corresponds to the likelihood of an exploitation scenario. Threat likelihood levels are determined for exploitation scenarios: Physical, Local, and Network AVs. The overall threat likelihood level is determined at the root node of the graph that represents the likelihood of a threat to Xen hypervisor from a TA. The range is from Negligible (very unlikely) if the level is 1, to Severe (frequent) if the level is 5.

PU and NU are two TAs with capability and motivation properties (CESG, 2009) and have been assigned values from 1–5. These values are then combined in **Error! Reference source not found.** to calculate initial threat levels (Level 1 is Negligible, 2 is Low, 3 is Moderate, 4 is Substantial, and 5 is Severe). Initial threat levels are established based on analysis and literature when considering capabilities and motivations described in the HMG IS1 standard (CESG, 2009). Each TA are then assigned capability and motivation values (**Error! Reference source not found.**).

Motivation	Capability Level				
	Very little	**Little**	**Limited**	**Significant**	**Formidable**
Indifferent	Negligible	Negligible	Low	Low	Moderate
Curious	Negligible	Negligible	Low	Moderate	Substantial
Interested	Negligible	Low	Moderate	Substantial	Severe
Committed	Low	Low	Moderate	Severe	Severe
Focused	Low	Moderate	Substantial	Severe	Severe

Table 4. Properties of Threat Actors

Assigning Threat Levels to Tree Nodes

To assign threat levels to the tree nodes, threat likelihood levels that are based on capability and motivation properties (above) are assigned to each source node in **Error! Reference source not found.** so that both PU and NU TAs are analysed and these are shown in **Error! Reference source not found.** and Figure 4.

A 1.1.1 Exploit Physical Vulnerability: The threat level assignment starts with the source nodes connected to A 1.1 (Physical Attack). Threat levels are assigned to A 1.1.1 and A 1.1.2 source nodes. For the attack step, A 1.1.1 (Exploit Physical Vulnerability) minimal technical knowledge is required to exploit a physical vulnerability using physical access, for example, unplugging the power or network cable. But while both TAs are capable, the motivation value of the NU is lower because they typically have no direct access.

A 1.1.2 Misuse Physical Access: The CSP is responsible for physical security and normally the Xen host server is provided with adequate physical access protection. However, a systems administrator (PU TA) may have direct or physical access and may be influenced by a malicious party. As the only one present where the host server is installed and with sufficient privileges to access it, both capability and motivation values are high for the PU TA. Whereas a NU with malicious intent may be motivated to exploit a physical vulnerability, but in most cases NUs do not know the location of the CSP data centre.

Following on, initial threat levels for the remaining nodes are determined using **Error! Reference source not found.** and the source nodes in Figure 2**Error! Reference source not found.**. Due to a lack of available space, the assignment of threat levels to all the remaining source nodes are not presented in this paper.

Propagating Threat Levels

A 1.1.1 and A 1.1.2 source nodes (the child nodes of A 1.1) are assigned a threat level to achieve the propagation of threat levels to the root node. Threat levels are then propagated to the root node A 1 to establish threat likelihood levels. Threat levels for these nodes are further adjusted to achieve min/max conditions. The attack trees apply the min/max function for every logical AND relationship. This means that to achieve the attacker's primary goal, all sub-goals (the manipulation of Xen components) are required. Therefore, the sub-goal with the lowest threat level would determine the difficulty level of the attack scenario. Whereas for every logical OR relationship, it is the maximum threat level of the source node that is assigned to the respective child node. The OR relationship means that to achieve the attacker's goal, just one sub-goal is required. That is, the sub-goal with the highest threat level determines the difficulty level of the attack scenario. In the Xen attack graph, threat likelihood levels are red for Severe, orange for Substantial, dark yellow for Moderate, yellow for Low, and green for Negligible.

A 1.1 Physical Attack: The threat level from the source nodes A 1.1.1 and A 1.1.2 are propagated to A 1.1 (**Error! Reference source not found.**). This tree branch has no loop so the threat propagation is simple. Both A 1.1.1 and A 1.1.2 are connected to A 1.1 by an AND operator, so the threat level is minimal. Threat levels from A 1.1, A 1.2 and A 1.3 are again propagated to root node A 1 to determine the overall threat likelihood level to Xen.

A 1 Manipulate Xen Hypervisor: In this case the threat level is propagated to root node A 1. The child nodes, A 1.1–1.3, are linked by a logical OR and so the level of threat is at the maximum. Figure 3 illustrates the combined threat likelihood level of a PU TA from A 1. In this scenario, the overall threat likelihood level from a PU TA is moderate and the weakest point is a physical vulnerability.

By comparison, Figure 4 presents the combined threat likelihood level determined at the A 1 root node from a NU TA. In this case, if a NU exploits a network vulnerability by misusing his Guest OS (DomU) space to compromise Dom0 and Xen, then the overall threat likelihood level is Severe (it is red). We see that if a NU

exploits a vulnerability through Physical, Local, and Network AVs, A 1.1-1.3 are Negligible, Moderate and Severe respectively.

Risk Level Determination

Vulnerability impact ratings and threat likelihood levels are mapped to risk severity levels for Xen hypervisor and are determined using the risk estimation matrix (Table 3). Eight risks are considered from both PU and NU TAs (Table 5). Given this knowledge, we would expect that higher-level risks ought to be brought down to an acceptable level by improving the security controls and procedures.

Risk No.	Xen AV	Vul Impact Rating	Threat Likelihood	Level of Risk from PU TA	Risk No.	Xen AV	Vul Impact Rating	Threat Likelihood	Level of Risk from NU TA
R1	Local	High	Negligible	Medium	R1	Local	High	Moderate	Medium
R2	Local	Very High	Negligible	Medium	R2	Local	Very High	Moderate	High
R3	Physical	Very High	Moderate	High	R3	Physical	Very High	Negligible	Medium
R4	Network	High	Negligible	Medium	R4	Network	High	Severe	High
R5	Local	High	Negligible	Medium	R5	Local	High	Moderate	Medium
R6	Network	Medium	Negligible	Low	R6	Network	Medium	Severe	High
R7	Physical	Very High	Moderate	High	R7	Physical	Very High	Severe	High
R8	Local	High	Negligible	Medium	R8	Local	High	Moderate	Medium

Table 5. Risk Levels from PU and NU TAs

Comparison with Other Systems

To assess the vulnerability and risk process for generalizability, we have applied it to other important and commonly used open source infrastructure level systems, the Apache HTTP and Squid Proxy servers. From this assessment, the evaluation process did not present limitations and produced fairly accurate results. While we believe the overall risk assessment process is applicable and adaptable on open source software, due to inconsistent trend and seasonal factors, the Apache and Squid reported vulnerabilities data provides fluctuations. Thus, compared to Xen, the predicted results were not as accurate. We think that the attack trees and graphs used to determine threat likelihood levels are applied at a high level of abstraction and so they are limited to viable-threat-only vectors. Such vectors would obtain implicit information such as assets, vulnerabilities, and so forth.

During the Apache vulnerability assessment, for 2018 6.66 unknown vulnerabilities were predicted. Following the process, we determined that the predicted result is fair because there is some variance to the annual average reported vulnerabilities of 10. The impact ratings of vulnerabilities from three scenarios present physical threats as Medium, local threats as Medium, and network threats as High. The threat likelihood levels show that the most TA likely to exploit a network vulnerability is a Service Consumer (SC) TA. A SC TA poses a Severe likelihood level, whereas a PU TA is likely to exploit physical vulnerabilities, and a NU TA is likely to exploit local AVs. PU TA and NU TA pose Moderate and Substantial likelihood levels respectively.

By comparison, 9.13 unknown Squid vulnerabilities are predicted for 2018. The prediction is only fair because it is not close to the annual average of 5.06. The vulnerability impact ratings are physical threats as Medium, local threats as High, and network threat as Medium. In the case, a PU TA is likely to exploit a vulnerability through a physical AV and poses a Moderate likelihood level. Whereas, a NU TA and a new TA, the Indirectly Connected (IC) TA, pose Negligible and Substantial likelihood levels by exploiting vulnerabilities through local and network AVs respectively.

Discussion and Conclusion

The influence of PUs as TAs highlights that IaaS customers must ensure that CSPs follow sufficient hiring procedures. CSPs must perform proper background and security checks for new hires, and to prevent targeted approaches, staff profiles should not be public. However, responsibility for hypervisor security is goes beyond the CSP, for example CC service consumers need to be aware of creating AVs through poor programming. Another example, in a standard scenario, an IaaS NU has access to Xen and its assets via the guest VM, so a NU with malicious intent and misuses privileges to exploit the host OS and Xen poses a high threat severity level. Therefore, CSPs need to have adequate security controls and procedures in place to mitigate high severity risks.

We have found a need to provide IaaS customers with a realistic and accurate vulnerability and risk assessment platform. The vulnerability and risk assessment process can allow IaaS customers to prioritise risk by severity level and then make informed security decisions. Also, the process provides similar levels accuracy for other open-source applications. However, the data we have used is provided in good faith and a public domain, we do not think the process will yield as good results from proprietary systems because vendors may be reluctant to post vulnerabilities onto CVSS.

The accuracy of the time series predictions can be improved by including additional factors such as the type and frequency of reported vulnerabilities instead of using the number of reported vulnerabilities as the input dataset. The structured analysis approach developed here provides a broad threat modelling platform, so an optimised threat modelling technique can be used to ensure an in-depth threat analysis at a deeper technical level, with the ability to embed the knowledge from available catalogues.

References

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

Alturkistani, F. M., and Emam, A. Z. 2014. "A review of security risk assessment methods in cloud computing". In *New perspectives in information systems and technologies*, volume 1 pp. 443–453. Springer.

Bazargan, F., Yeun, C. Y., and Zemerly, M. J. 2012. "State-of-the-art of virtualization, its security threats and deployment models". *International Journal for Information Security Research*, (2:3/4), pp. 335–343.

Brohi, S. N., Bamiah, M. A., Brohi, M. N., and Kamran, R. 2012. "Identifying and analyzing security threats to virtualized cloud computing infrastructures". In *International Conference on Cloud computing technologies, applications and management*, 2012. pp. 151–155.

Catteddu, D., and Hogben, G. 2009. "Benefits, risks and recommendations for information security". *European Network and Information Security*.

Cayirci, E. 2015. "Models for cloud risk assessment: A tutorial". In *Accountability and security in the cloud* pp. 154–184. Springer.

CESG. 2009. "HMG IA standard no. 1 technical risk assessment". *CESG recommendation*.

Chhabra, S., and Dixit, V. 2015. "Cloud computing: State of the art and security issues". *ACM SIGSOFT Software Engineering Notes*, (40:2), pp. 1–11.

Dawoud, W., Takouna, I., and Meinel, C. 2010. "Infrastructure as a service security: Challenges and Solutions". The 7th international conference on Informatics and systems, 2010 pp. 1–8.

Elhage, N. 2011. "Virtunoid: Breaking out of KVM". *Black Hat* USA.

Fitó, J. O., and Guitart, J. 2014. "Business-driven management of infrastructure-level risks in cloud Providers". Future Generation computer systems, (32) pp. 41–53.

Hussain, M., and 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.

Hutle, M., Hansch, G., and Fitzgerald, W. 2015. "D2. 2 threat and risk assessment methodology". *Tunneling* and Underground Space Technology, (24:3), pp. 269–277.

Kiran, M. 2014. "A methodology for cloud security risks management". In *Cloud computing*. pp. 75–104. Springer.

Kortchinsky, K. 2009. "Cloudburst: A VMWare guest to host escape story". *Black Hat* USA.

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

Litchfield, A., and Shahzad, A. 2017. "A systematic review of vulnerabilities in hypervisors and their detection". In *Proceedings of the 23rd americas conference on information systems*.

Mell, P., Scarfone, K., and Romanosky, S. 2006. "Common vulnerability scoring system". *IEEE Security and Privacy*, (4:6).

Mell, P., Scarfone, K., and Romanosky, S. 2007. "A complete guide to the common vulnerability scoring system version 2.0". *First-forum of incident response and security teams* (1:1), p. 23.

NIST. 2017. "National vulnerability database (NVD)" [Internet web page]. Retrieved 31 December, 2017, from https://nvd.nist.gov/

Noel, S., Jajodia, S., Wang, L., and Singhal, A. 2010. "Measuring security risk of networks using attack graphs". *International Journal of Next-Generation Computing*, (1:1), pp. 135–147.

Reuben, J. S. 2007. "A survey on virtual machine security". Helsinki University of Technology, (2:36).

Roumani, Y., Nwankpa, J. K. & Roumani, Y. F. 2015. "Time series modeling of vulnerabilities". Computers & Security, 51, pp. 32–40.

Rutkowska, J., and Wojtczuk, R. 2008. "Preventing and detecting Xen hypervisor subversions". *Blackhat Briefings* USA.

Sabahi, F. 2011. "Virtualization-level security in cloud computing". In *IEEE 3rd international conference on Communication software and networks*. pp. 250–254.

Saripalli, P., and 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.

Shahzad, A., and Litchfield, A. 2015. "Virtualization technology: Cross-VM cache side channel attacks make it vulnerable". In *Proceedings of the Australasian conference on information systems*.

Shoaib, Y., and Das, O. 2014. "Pouring cloud virtualization security inside out". arXiv preprint arXiv:1411.3771.

Tanimoto, S., Hiramoto, M., Iwashita, M., Sato, H., and 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*. pp. 147–152.

Wang, H., Liu, F., and Liu, H. 2012. "A method of the cloud computing security management risk assessment". *Advances in Computer Science and Engineering*, pp. 609–618.

Wojtczuk, R. 2008. "Subverting the Xen hypervisor". *Black Hat* USA.

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