



UNIVERSITY OF LEEDS

This is a repository copy of *Multi-tenancy in cloud computing*.

White Rose Research Online URL for this paper:

<http://eprints.whiterose.ac.uk/80819/>

Version: Accepted Version

Proceedings Paper:

Aljehdali, H, Albatli, A, Garraghan, P et al. (3 more authors) (2014) Multi-tenancy in cloud computing. In: Proceedings of the 8th IEEE International Symposium on Service-Oriented System Engineering. 2014 IEEE 8th International Symposium on Service Oriented System Engineering (SOSE), 7-11 April 2014, Oxford, UK. IEEE , 344 - 351. ISBN 978-1-4799-2504-9

<https://doi.org/10.1109/SOSE.2014.50>

Reuse

Unless indicated otherwise, fulltext items are protected by copyright with all rights reserved. The copyright exception in section 29 of the Copyright, Designs and Patents Act 1988 allows the making of a single copy solely for the purpose of non-commercial research or private study within the limits of fair dealing. The publisher or other rights-holder may allow further reproduction and re-use of this version - refer to the White Rose Research Online record for this item. Where records identify the publisher as the copyright holder, users can verify any specific terms of use on the publisher's website.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Multi-Tenancy in Cloud Computing

Hussain AlJahdali, Abdulaziz Albatli, Peter Garraghan, Paul Townend, Lydia Lau, Jie Xu

School of Computing
University of Leeds
Leeds, United Kingdom

{schla, sc11a2a, scpmg, p.m.townend, l.m.s.lau, j.xu}@leeds.ac.uk

Abstract— As Cloud Computing becomes the trend of information technology computational model, the Cloud security is becoming a major issue in adopting the Cloud where security is considered one of the most critical concerns for the large customers of Cloud (i.e. governments and enterprises). Such valid concern is mainly driven by the Multi-Tenancy situation which refers to resource sharing in Cloud Computing and its associated risks where confidentiality and/or integrity could be violated. As a result, security concerns may harness the advancement of Cloud Computing in the market. So, in order to propose effective security solutions and strategies a good knowledge of the current Cloud implementations and practices; especially the public Clouds; must be understood by professionals. Such understanding is needed in order to recognize attack vectors and attack surfaces. In this paper we will propose an attack model based on a threat model designed to take advantage of Multi-Tenancy situation only. Before that, a clear understanding of Multi-Tenancy, its origin and its benefits will be demonstrated. Also, a novel way on how to approach Multi-Tenancy will be illustrated. Finally, we will try to sense any suspicious behavior that may indicate to a possible attack where we will try to recognize the proposed attack model empirically from Google trace logs. Google trace logs are a 29-day worth of data released by Google. The data set was utilized in reliability and power consumption studies, but not been utilized in any security study to the extent of our knowledge.

Keywords – Cloud Computing; Security; Multi-Tenancy; Attack Models; Cloud Data.

I. INTRODUCTION

Cloud Computing is recognized as currently one of the most popular technologies available; it can be seen as an instance of Computing as a Utility. In Computing as a Utility, customers utilize the concept of “pay-as-you-go” for applications, computing and storage resources [4] [5]. Along with the pay-as-you-go concept, the elasticity in upgrading or downgrading resources makes Cloud Computing a popular model for organizations [15]. Moreover, the cost effectiveness of Cloud Computing is encouraging its adoption; enterprises requiring a high level of elasticity and about to decide whether to build up their own IT infrastructure or to utilize Cloud infrastructure may find that using a Cloud infrastructure will give a better balance between cost and elasticity [5] [9] [15] [17].

Cloud Computing is defined as “a system, where the resources of a data centre is shared using virtualization technology, which also provide elastic, on demand and

instant services to its customers and charges customer usage as utility bill” [4]. Pay-as-you-Go and Elasticity along with On-Demand, Broad network access, Scalability and Virtualization are considered as the essential characteristics of Cloud Computing Model.

With the benefits of Cloud Computing come along challenges to the model; one of the most challenging of these aspects is security. Information Security refers to protecting information and information systems from unauthorized access, use, disclosure, disruption, modification, inspection, recording or destruction. Based on a study for the Cloud Security Alliance (CSA), there are seven top threats that organizations will face in adopting Cloud Computing [7]. These are Abuse and Nefarious Use of Cloud Computing, Insecure Application Programming Interfaces (API), Malicious Insiders, Shared Technology Vulnerabilities, Data Loss/Leakage, Account, Service and Traffic Hijacking and Unknown Risk Profile. In addition, another study by Gartner has also identified seven Cloud Computing security risks, which are Outsourcing Services, Regulatory Compliance, Data Location, Shared Environment, Business Continuity and Disaster Recovery, Hard Environment for Investigating Illegal Activity and Long Term Viability [8]. Moreover, a survey of Cloud providers by the International Data Corporation (IDC) in 2008 to study the obstacles or concerns for adopting Cloud Computing in enterprises showed that security as a concern came first with 88.5% of the votes, whilst availability; which is one of information security principles; came third with 84.8% of the votes [11] [17].

Such concerns are driven by Cloud nature of shared resources and Multi-Tenancy. The threat of data compromise increases in the Cloud, due to the increased number of parties leading to an increase in the number of points of access [2]. Also, delegating data control to the Cloud leads to an increase in the risk of data compromise where outsourced services bypass the personal, logical and physical security controls of a consumer. A number of concerns emerge regarding the issues of Multi-Tenancy and data remanance. Multi-Tenancy refers to resource sharing in Cloud Computing where any resource object is reusable in the Cloud infrastructure. Reusable objects must be carefully controlled and managed since they create a serious vulnerability and violate confidentiality through possible data leakage. Data leakage in this context may be caused by the fact that hardware in Cloud Computing is not separated; there is a good level of separation in Cloud Computing at the application and virtual layer but not enough in the hardware

layer [1]. Also, confidentiality could be breached due to the reusability of resource objects through data remanance, where a customer can request storage space from a Cloud provider and run a scan in order to search for sensitive data to other customers [1] [2] [6].

The most important challenge in studying security in Cloud Computing relies on the trade-off between security and cost, which is itself one of the important factors in shifting to Cloud Computing. Tim Watson, Head of the computer forensics and security group at De Montfort University notes:

“...although one provider may offer a wonderfully secure service and another may not, if the latter charges half the price, the majority of organizations will opt for it as they have no real way of telling the difference.” [4].

Also, George Wrenn, Security Solutions Director at Unisys recommends that customers must consider other factors more than price and top feature sets (i.e. feature sets will be different from one Cloud provider to another) before deciding to move critical systems and applications to Cloud [4]. From the previous quotations the trade-off between security and cost is obvious, where security is considered relatively costly.

The goal of our work is to understand Multi-Tenancy. Moreover, we will try to sense any suspicious behavior that may indicate to a possible attack based on a proposed attack model. The rest of this paper is organized as follows: section II highlights and identifies Multi-Tenancy as vulnerability. Section III focuses in Multi-Tenancy its definition, origins and its benefits. Section IV highlights the security challenges in dealing with Multi-Tenancy. Section V shows the scheme where system model, threat model, attack model and our approach in undertaking Multi-Tenancy is presented. Section VI will illustrate the reconstruction of the attack model from Google’s dataset. Finally, section VII and VIII will present the future work and conclude this paper.

II. LITTERATEUR REVIEW

Multi-Tenancy has been identified as a security issue in Cloud Computing by several researchers such as [1] who conducted a survey on security issues in service delivery models in Clouds and stated that Multi-Tenancy is a major Cloud Computing characteristic that may lead to confidentiality violation. [2] Also identifies Multi-Tenancy as a major threat to both confidentiality and privacy when talking about Cloud Computing security. In addition, [4] highlighted shared technology vulnerabilities – hence Multi-Tenancy – as one of the top threats to Cloud computing in a survey done on the existing literature. Moreover, [9] recognizes Multi-Tenancy as a new source of threat in Cloud Computing infrastructure.

From another point of view [10] links between Multi-Tenancy as a form of shared environment and the attraction of malicious activities in the Clouds. Intel IT Centre [13] generated a document of best practices on building secure Clouds; yet clearly highlights Multi-Tenancy and shared technology issues as security challenges for a Cloud environment. Where [14] in his work proposed a layered security approach for Cloud Computing, and states that

virtualization is one of the process hosting layer (i.e. servers) issues where competitors will have separate virtual machines in the same physical machine; hence Multi-Tenancy.

In [15] several areas were identified as danger in Clouds; under data governance the writer highlighted that Multi-Tenancy arrangements in Clouds are raising questions about data segregation. While NIST developed a report titled “Guidelines on Security and Privacy in Public Cloud Computing”; they identify Multi-Tenancy as of the security and privacy downsides in the Cloud [15]. In a totally different approach [17] interviewed five leading scientists from the cloud community; Raghu Ramakrishnan the Chief Scientist for Search and Cloud Platforms at Yahoo! was one of them, where his response to the question of “On a related note, for a graduate student starting a PhD, what would you say are the key fundamental challenges of cloud computing that should be addressed by new research in the field?” included Multi-Tenancy as a fundamental challenge of Cloud Computing. Again [18] raised questions in how Cloud Computing affecting security, privacy and trust; where he identifies Multi-Tenancy as one of the security issues.

Cloud Security Alliance (CSA) released a document titled “Security as a Service” [26] where they tries to define categories for services; they raised the question “How does one assure data isolation in a multi-tenant environment?”. Also, CSA in the same document stated that Multi-Tenancy is creating new targets for intrusion. In a study done by [19] to identify the challenges of security and privacy in Cloud Computing; Multi-Tenancy is recognized as one of the unique implications of security and privacy in Cloud computing. In the same direction [28] defines Multi-Tenancy as a major characteristic of Cloud Computing and a major dimension in the Cloud security problem that needs a vertical solution from the Software-as-a-Service (SaaS) down to Infrastructure-as-a-Service (IaaS). Where [20] highlight the fact that Multi-Tenancy may enable information leakage and increase attack surface which will affect the security of the Clouds. Also, [31], [32] and [33] considered Multi-Tenancy among the serious issues in Cloud security.

After highlighting Multi-Tenancy as a security concern in Cloud Computing, the need for a deep understanding of Multi-Tenancy is required in order to deal with it effectively.

III. MULTI-TENANCY

Multi-Tenancy is a natural result of trying to achieve economic gain in Cloud Computing by utilizing virtualization and allowing resource sharing [9] [15]. AS defined earlier, Multi-Tenancy refers to resource sharing in Cloud Computing, but such a definition is still general in the context of Cloud Computing, where Multi-Tenancy is seen differently from different service models.

In Software as a Service (SaaS), applications are provided as a service by the Cloud Service Provider (CSP) where the customer cannot monitor or control the underlying infrastructure; here, Multi-Tenancy means that two or more customers utilize the same service or application provided by the CSP regardless of the underlying resources [13].

In Infrastructure-as-a-Service (IaaS), where the customer is capable of provisioning computing, storing and




Figure 1: Multi-Tenancy Benefits' Tree.

networking resources and can control but cannot manage the underlying infrastructure, Multi-Tenancy occurs when two or more virtual machines (VMs) belonging to different customers share the same physical machine (PM) [15].

Multi-Tenancy has brought different arguments in Cloud Computing. While software developers see it as an opportunity, security experts see it as vulnerability [3] [12] [13] [15]. Even though security experts agree that Multi-Tenancy is a vulnerability that could lead to confidentiality being exposed, they vary in providing the solution for such vulnerability.

Whereas [12] suggests the elimination of the virtualization layer in order to prevent multi tenancy, [15] suggests that the provider should expose the risk of Multi-Tenancy to the customer and do nothing about it (i.e. give them the option of paying extra to avoid Multi-Tenancy). The first strategy seems very effective, but would eliminate vital benefits for Cloud providers such as VM mobility and financial gain due to resource sharing.

VM mobility is one of these benefits where providers can easily reallocate VMs to achieve better utilization and save power consumption. On the other hand, the second strategy will not enhance the Cloud security and customers especially enterprises are holding back investment in Cloud Computing because of security issues [1] [9] [19].

Moreover, current practice of UK enterprises is to deploy Private Clouds in order to cut costs and safeguard sensitive data [1]. We therefore identify that a solution securing Multi-Tenancy yet keeping its benefits is needed. So, a deep understanding of Multi-Tenancy is required in order to identify all the possible benefits brought to Cloud Computing because of Multi-Tenancy.

$$\text{Virtualization} + \text{Resource Sharing} = \text{Multi-Tenancy} \quad (1)$$

As equation (1) shows, in order for Multi-Tenancy to occur both virtualization and resource sharing must be allowed by the CSP. Fig. 1 shows all the identified possible benefits of Multi-Tenancy and by looking into the tree's leaves, it is easily recognized that the origin of the benefits could be linked either to virtualization, resource sharing or by combining both of them. For instance, separating the hardware failure from the software failure is achieved by virtualization. On the other hand, sharing the resource will increase the utilization which will lead to a reduction in cost by making the resource available for more than one customer. In other cases such as over provisioning and VM mobility, both virtualization and resource sharing will amplify their impact. VM mobility can contribute in

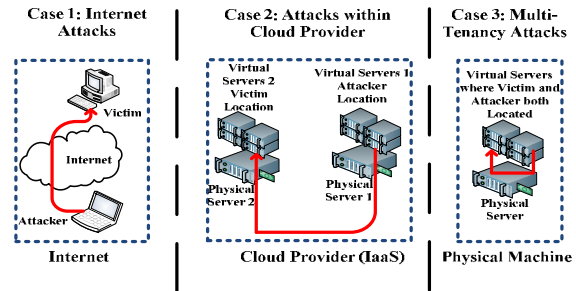


Figure 2: Difference between Multi-Tenancy and Traditional Cases.

maximizing the utilization of the infrastructure or reducing the power consumption by reallocating VMs into clusters and minimizing the number of servers used. Whereas over provisioning is considered one of the major features of Cloud Computing as it gives the opportunity for the CSP to seal more than the capacity of his infrastructure. These features are important for Cloud Computing and any proposed solution must add to them or at least try to keep them and not to eliminate any one of them.

IV. MULTI-TENANCY SECURITY CHALLENGES

What is unique about Multi-Tenancy in Cloud Computing is that both the attacker and the victim are sharing the same server (i.e. physical machine (PM)). Such a setup cannot be mitigated by traditional security techniques and measures, simply because it is not designed to penetrate inside servers and their monitoring techniques are limited to the network layer [19] [20].

To illustrate, Fig. 2 shows the different cases of attacker and victim locations and the networking between them. In case one, the attacker and the victim both are regular Internet users; in order to defend against such attacks, traditional network security techniques and devices are efficient.

In case two, both attacker and victim are customers in the same Cloud provider but each one of them is located on a separate server. This kind of setup is due to the utilization of the virtualization layer in the Cloud Computing Model; to secure such a setup, virtual network security devices and techniques must be implemented by Cloud providers [20].

Case three describes the problem that we intend to address in future work, where both the attacker and the victim are customers in the same Cloud and are sharing the same server. Such a situation is due to Multi-Tenancy; securing such a setup is not an easy task as network communication between the attacker's VM and the victim's VM is limited within the physical machine (PM). Therefore, traffic will not leave the physical machine, which is harder to be mitigated by virtual network security defenses as opposed to case two.

In order to secure such vulnerability, we must first answer the following question: how is Multi-Tenancy exploited? An answer can be found in [15], where an attack is generated over the Amazon EC2 Cloud to investigate data leakage. In order to carry out the attack, network probing is performed; following this, a brute force attack is generated to take advantage of the Multi-Tenancy effect by allocating the attacker's VM beside the victim's VM. The results show that

by spending just a few dollars, an attacker has a 40% chance to allocate his VM beside the victim's VM. After achieving Multi-Tenancy, a side channel attack – any attack takes advantage of the system characteristics – is generated to extract the data of the victims.

Obviously, any tenant can attack its neighbor because the type of attack that could be utilized, such as side channels, cannot be detected by the hypervisor or even the operating system.

So, there is no way to eliminate the Multi-Tenancy effect in order to keep its benefits yet the effect could be minimized and that what is this paper is trying to illustrate. Multi-Tenancy cannot be eliminated, yet a smart resource allocation technique will minimize the risk of Multi-Tenancy; in other words, a resource allocation technique will increase the level of difficulty of achieving Multi-Tenancy for customers, yet is easily managed by Cloud providers. What is interesting of Multi-Tenancy is that in order to achieve it for targeted victims, the attacker needs to invest an effort, time and cost. So, by making Multi-Tenancy difficult to be achieved by customers, we are restricting the number of potential attackers.

V. SCHEME

In this section a proposed system model, threat model and attack model will be demonstrated. Also, our approach in undertaking Multi-Tenancy security is presented. The main issue is to highlight Multi-Tenancy as an attack surface and try to control it by the proposed system model in the following.

A. System Model

Our proposed system model is shown in Fig. 3, where all the solid black links (i.e. 1c, 2c, 3c, 4c, 5c, 6c and 7c) represent control channels which are under Cloud provider control and responsibility. And all the dashed blue links (i.e. 1o, 2o and 3o) are operation channels where the letter (c) donates control and letter (o) donates operation. The separation of these channels is vital in order to enhance system security and implement security in depth. The red links reflects the attack path and the attacker interaction with the system. All the system components are defined below; following this, a description of the system flow is given.

- *Registration Unit*: the initial contact between the consumer and Cloud provider. Registration can be an online form or a contract signed between both parties. In this phase, all important information that will define the allocation mechanism should be gathered.
- *Verification and approval Unit*: this is a vital phase where the provider should verify and approve the information given by the customer. The importance of such a process is to protect the provider's image by avoiding any fraud possibility.
- *Control Database Unit*: is the location where the parameters and restrictions of the resource allocation technique are stored. There are two contacts to this Database; the first is made by the Cloud provider in order to store the customer information to be utilized by the allocation technique; the second is made by the

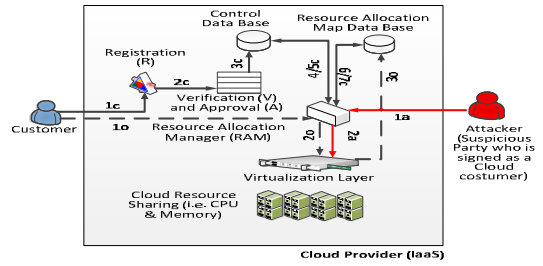


Figure 3: Proposed System Model.

resource allocation manager in order to extract the customers' resource allocation restrictions.

- *Resource Allocation Manager Unit (RAMU)*: is responsible for allocating resources following a customer request. The RAMU is the only system component to access the control database.
- *Resource Allocation Map Database Unit*: this Database is responsible of keeping updated records of resource allocation.

To understand this system better, we describe the scenario of a customer joining a Cloud provider. The process starts when a customer needs to utilize a public IaaS Cloud as its infrastructure. Firstly, the customer will register either online or by visiting the Cloud provider. Then, the provider will verify the information provided by the customer; official documents could be used for this purpose. If the customer passes the verification process, the provider will approve the process to the next stage. After the verification and approval processes, the customer's data will be stored in the control Database where all resource allocation restrictions will be specified by the CSP.

When the customer has been registered successfully and security restrictions are specified, the system is ready to be utilized by the customer. The customer's request to allocate resources shall be sent to the resource allocation manager unit (RAMU). Then, the RAMU will request the security restrictions from the control Database and the current resource allocation map from the resource allocation map Database in order to allocate the customer's resources in the proper location. And whenever a customer releases a resource the resource allocation map Database is updated immediately. In such a setup, an attacker will not have the chance to take advantage of the resource allocation mechanism, and the benefits of Multi-Tenancy are preserved.

An extra benefit of this system is that the control Database can have a different resource allocation method, whereby the system model will not be changed. This advantage will give the Cloud provider the opportunity to define security restrictions based on their business strategy. Also, it gives the provider the chance to implement their own resource allocation methods if needed. Moreover, it could be a security best practice to change the resource allocation method periodically in order to raise the system difficulty and make it hard to be predicted.

B. Threat Model

In this section we will describe our assumptions in regard to the environment. In the attack model shown in Fig. 4, an assumption of a secure hypervisor (the infrastructure




Figure 4: Attack Model.

management component) is made. Also, an attacker is assumed to know nothing about the Cloud provider infrastructure, which is more like a real situation. Moreover, the assumption that the Cloud provider is allowing Multi-Tenancy is made, which is a natural result of allowing resource sharing over virtualization; this is noticed in most Cloud implementations. The attack will take advantage of Multi-Tenancy regardless of any other Cloud component and feature. So, the attack is designed to exploit Multi-Tenancy and all the phases before being a Multi-Tenant does not take advantage of any Cloud well-known vulnerability such as the hypervisor vulnerability.

C. Attack Model

The Attack model we use is based on one of the scenarios utilized by [15] to explore information leakage in Amazon EC2 public Cloud. The nature of Information Security (InfoSec) for a specific vulnerability is that there could be a large number of attacks to exploit it [15]. However, one successful attack against a system will identify most of the possible vulnerabilities that can be utilized. Moreover, attacks vary in the sense of their behavior; for example, it is easy to detect any distributed denial of service (DDoS) attack and any attack consisting of port scanning due to the unexpected increase in traffic. Also, it is easy to identify viruses due to their unique signatures; whereas it is hard enough to detect iFrame attacks. An iFrame attack is an attack where an HTML code is embedded inside another HTML code as a frame in order to collect credit card information for instance. The following is a list of well-known attacks that could be utilized efficiently over the Cloud infrastructure:

- *Side Channel attack*: a side channel attack is any attack based on information gained from the physical implementation of a system. There are many side channel attacks known in the field; some of the well-known side channel attacks are timing attacks, power consumption attacks and differential fault analysis.
- *Brute Forcing*: brute forcing is an attack strategy or mechanism which could be applied over any kind of attack. It is one of the simplest strategies in order to build an attack but yet it is one of the most common used strategies. For instance if an attacker wants to find out a password of a system by utilizing a brute force strategy, the attacker will try every possible combination until the correct password is found. Therefore, brute forcing can be defined as running an attack operation

multiple times until a successful breach is achieved. Brute forcing is identified as one of the top ten attacks by the Data Breach Investigations Report (DBIR) where it forms 22% of data breach attacks [14].

- *Network Probing*: is a mechanism to find out the physical topology of a network that consists of IPs and servers connected in the network. Such information can be utilized to identify possible targets and to design an attack for a sub group in the network.

The attack model we propose is generated in three phases as shown in Fig. 4. In phase one, attacker register with the Cloud provider as a normal customer. In phase two, the attacker gathers information about the allocation technique and the Cloud infrastructure where network probing is utilized. The attacker can make sense of the allocation technique simply by requesting resources and then releasing them; this action will give the attacker knowledge of the allocation technique. Moreover, the attackers can take advantage of the information revealed by the Cloud provider about their infrastructure or any kind of systems or techniques they are using. After that, the attacker can utilize brute force techniques to generate VMs in order to achieve Multi-Tenancy. In phase three, after the attacker achieved Multi-Tenancy, a side channel attack is generated to extract the victim's data.

D. Approach

Once agreed that Multi-Tenancy is a vulnerability then we need to choose one of the well-known risk strategies in order to control and minimize its impact. There are four strategies to deal with risks; these four strategies are as follow:

- *Eliminating the risk.*
- *Mitigating the risk.*
- *Transferring the risk.*
- *Accepting the risk.*

Although eliminating the risk considered the most powerful strategy, it is not possible to apply it on Multi-Tenancy as most of the Cloud Computing benefits are linked to it, as illustrated previously.

So, in order to eliminate the risk of Multi-Tenancy we have to eliminate Multi-Tenancy and that can be achieved by either eliminating what makes Multi-Tenancy vulnerable or eliminating what forms Multi-Tenancy in the first place.

As mentioned previously in equation (1) that Multi-Tenancy is a natural result of allowing resource sharing over virtualization. So, in order to eliminate Multi-Tenancy we need to either eliminate the use of virtualization or disable resource sharing.

In both cases that is not acceptable as that will eliminate most of the Cloud Computing main drivers and marketing strength features as shown in Fig. 1.

The other direction is to try to eliminate what makes Multi-Tenancy vulnerable which is the possibility of taking advantage of the shared environment. Such possibility is valid because of the side channel attacks where side channel




Figure 5: Approach in Securing Multi-Tenancy.

attacks defined as taking advantage of the physical characteristics of the system.

So, Side channel attacks by their nature are unlimited in their number and evolve with time which makes it hard to eliminate all the possible side channel attacks. In addition, dealing with the known side channel attacks is not an easy task where the forms of such attacks are multiple.

For example, there are researchers dealing with the memory as the attack vector such as [40] where they propose to eliminate the shared memory attack. Although, they achieved the goal but their solution was a hypervisor dependent where it worked only over Xen.

Another form of attack is the timing side channel attacks and [41] in his proposal was able to eliminate three forms of timing side channel attacks. However, in order to eliminate these three forms of timing side channel attacks the CSP must sacrifice 2/3 of his infrastructure. In other words 2/3 of the infrastructure is overhead and the solution will not eliminate all the timing side channel attacks without even mentioning the other forms of side channel attacks.

From the previous two examples it is quite obvious that trying to eliminate all the side channel attacks will be at a very high cost if it is possible in the first place.

Furthermore, some of the known side channel attacks have a tight relationship where two forms of side channel attacks cannot be eliminated in the same time [18]. In other words if the vulnerability of side channel attack A is eliminated, the vulnerability of side channel attack B cannot be eliminated. So, if attack A is blocked that means attack B will be successful.

Since we cannot eliminate the side channel attacks and cannot tolerate the loss of major features due to the elimination of either virtualization or resource sharing then the eliminating the risk strategy is not acceptable and is not applicable when it comes to dealing with Multi-Tenancy risks. So, the second best strategy is to mitigate the risk and this is what we are after; where we try to balance between the benefits brought by Multi-Tenancy and the security putting in minds other factors such as performance and cost. Fig. 5 illustrates our approach where we highlighted Multi-Tenancy as vulnerability and a security concern in Cloud Computing and we illustrated that layer 1 cannot be eliminated either partially or completely. Also, layer 3 cannot be eliminated totally as described earlier.

So, the only angle left is how to go from layer 1 into layer 2. In other words how does Multi-Tenancy happen in the IaaS Clouds? And the answer is the resource allocation mechanism. And in order to mitigate the Multi-Tenancy risks the resource allocation mechanism must be controlled. Making the resource allocation mechanism a security aware mechanism will enhance the total security for the CSP and will minimize the surface attack. In addition, minimizing the probability of being a Multi-Tenant by controlling the resource allocation mechanism will have an impact on the underlying layer (i.e. layer 3) as layer 3 is dependent on layer 2.

VI. GOOGLE TRACE LOGS

To better understand the issues and challenges in developing and adopting the Cloud, an analysis on real Cloud data is a crucial step.

Google have recently released two sets of data (7-day and 29-day sets) [33]. These sets have been studied and analyzed by a number of researchers in the literature [44 - 48]. Those studies were focusing on resources utilization, scheduling, relations with Grid/HPC systems, scalability, cluster management and behavior of workloads but with little focus on user behavior, security and the patterns of the workloads.

To clarify, the users of these trace logs have been identified as Google engineers and services [33] [38]. Reference [39] has concluded that there is a dependence/relationship between resource utilization, number of tasks and user patterns. Another study by [33] which has examined Google trace logs in terms of workload characteristics stated that the most notable workload characteristic is heterogeneity. They state that such heterogeneity lead to complications in resource allocations and utilizations.

In addition, [35] conducted a comparative study between Google dataset and Grid/HPC systems, stating that Google workloads show that resource allocations are finer with respect to CPU and Memory than that of Grid/HPC systems. Reference [38] conducted a study on the workload characteristics of Google Dataset. They concluded that machines are continuously taken offline and online to combat system failures and to apply upgrades. Also, many of the submitted jobs are not latency sensitive as more jobs are killed before normal completion.

Therefore, having reviewed the related part of the literature and after examining the Google dataset, it is concluded that the Workload consists of many patterns, depending on the angle of attention. In this paper we highlighted Multi-Tenancy as vulnerability and provided in depth understanding related to different dimensions of Multi-Tenancy. So, we decided to empirically investigate the possibility of reconstructing the proposed attack model from the dataset released by Google. Such activity can be used as a monitoring tool where CSP can monitor some behavior that can be linked to well-known attack models.

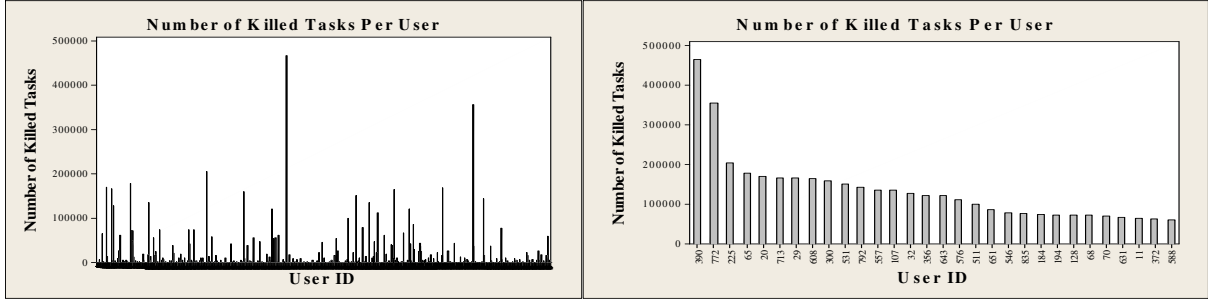


Figure 6: Number of Killed Tasks per User Where (A) is the lift figure and (B) is the right figure.

A. Reconstruct the Attack Model from the Dataset

In this section we will try to find out, is it possible to sense a suspicious behavior in the dataset, and to do so we will utilize the attack model proposed earlier and try to highlight if we can map it to the dataset either entirely or partially.

In order to find out a pattern of the attack model reflected on the dataset we will focus on phase two in Fig. 4. Since we do not have any information about phase one; which is the registration; in the dataset we will not look at it. In the case of phase two we will focus on brute forcing technique where it could be investigated based on Google's dataset.

So, in order to sense a brute forcing behavior we targeted the killed tasks in the dataset. In the dataset there were nine event types that any task can be tagged with. Between the nine event types there was a killed task event which represents a task was cancelled by the customer or a driver program. The killed task event is the only task event consists of human interaction where the customer could terminate the task.

So, in our analysis we decided to utilize the killed tasks to sense any brute forcing behavior. In our analysis we accepted the fact that cyber-attacks may be generated by humans or software and brute forcing is not an exception. The dataset consists of 25,000,000 task and 6,608,917 of the total number of tasks is tagged as killed tasks which represent 26.4% of the total number of tasks. The total number of customers in the dataset was 925 active customers. Only 725 of them committed a kill task event. Fig. 6 (A) shows all the customers who killed a task or more and Fig. 6 (B) shows the top thirty customers in terms of killing tasks.

We observed that most of the customers did not pass the threshold of killing tasks which is 200,000; only three customers passed that threshold. The customers IDs are 290, 772 and 225 where 290 notably killed over 450,000 tasks and 772 killed around 250,000 tasks. Also, we observed that the highest eight customers after the top three can be grouped together where they fall in the range between 150,000 and 200,000 killed tasks.

The following eight fall in the range between 100,000 and 150,000 killed tasks. The rest of customers committed a kill task event in a frequency below 100,000 times. In light of the above, we can highlight the customers 290, 772 and 225 as their behavior can give a strong indication of brut forcing technique. Such behavior can be linked to the

proposed attack model where customers' confidentiality can be violated.

VII. FUTURE WORK

Analyzing large datasets consumes time and needs a computational power in order to be analyzed, yet it is vital to increase our understanding of a system. In this paper we just run the first mile, yet we have been able to reconstruct a proposed attack model partially. Due to the computational limitation and time constrains we could not reconstruct the whole attack model. So, in the near future we will reconstruct the complete attack model with in depth analysis. Also, we will investigate empirically how restrictions in tasks affect Multi-Tenancy?

VIII. CONCLUSION

Multi-Tenancy is often seen as a benefit to Cloud providers; however, it comes with an associated security risk. When security comes first, a natural proposal is to eliminate this risk; [12] proposes the elimination of the virtualization layer in order to increase system security. However, the cost of such a change for existing systems (especially large Clouds) will be high. Also, the valuable feature of VM reallocation will not be possible in such a scenario, which will lead to performance degradation (i.e. low level of utilization of resources).

On the other hand, [13] shows Multi-Tenancy as an opportunity must be utilized without mentioning the security concerns related to it. Between those extremes, [15] identifies Multi-Tenancy as vulnerability yet suggests that Cloud providers expose it to customers without giving any solution to at least mitigate its risks. Such exposure to the problem without providing a real solution will make customers depart from Cloud providers.

This paper studies Multi-Tenancy in depth highlighting its origins, benefits and what is unique about it. Also, it proposes a system model and a resource allocation technique that will achieve the balance between both security and the benefits gained from Multi-Tenancy. In addition, we propose an approach in how to tackle Multi-Tenancy in a novel way in order to reach a balanced point between its benefits and Cloud security.

Moreover, this paper introduces security as a requirement when designing resource allocation techniques without affecting performance, power consumption, and cost. Finally, a proposed attack model is reconstructed partially from Google's dataset where three customers are identified as suspicious customers.

REFERENCES

- [1] S. Subashini, and V. Kavitha, "A Survey on security issues in service delivery models of cloud computing," *Journal of Network and Computer Applications* (2011).
- [2] Dimitrios Zissis, and Dimitrios Lekkas, "Addressing cloud computing security issues," *Future Generation Computer Systems* (2011).
- [3] H. Alaqrabi, Lu Liu, Jie Xu, Richard Hill, Nick Antonopoulos, and Yongzhao Zhan, "Investigation of IT security and compliance challenges in security-as-a-Service for cloud computing" (2012).
- [4] Md. Tanzim Khorshed, A.B.M. Shawkat Ali, and Saleh A. Wasimi, "A Survey on gaps, threat remediation challenges and some thoughts for proactive attack detection in cloud computing," *Future Generation Computer Systems* (2012).
- [5] Michael Armbrust, Armando Fox, Rean Griffith, Anthony D. Joseph, Randy H. Katz, Andrew Konwinski, Gunho Lee, David A. Patterson, Ariel Rabkin, Ion Stoica, and Matei Zaharia, "Above the clouds: a berkeley view of cloud computing" (Feb. 10, 2009).
- [6] Cloud Security Alliance, "Security guidance for critical areas of focus in cloud computing V2.1" (Dec. 2009).
- [7] Cloud Security Alliance, "Top threats to cloud computing V1.0" (March 2010).
- [8] Jon Brodtkin, "Gartner: seven cloud-computing security risks" (July 02, 2008).
- [9] Wayne A. Jansen, "Cloud hooks: security and privacy issues in cloud computing," *Proceedings of the 44th Hawaii International Conference on System Sciences* (2011).
- [10] David Teneyuca, "Internet cloud security: the illusion of inclusion," *SciVerse ScienceDirect* (2011).
- [11] Tharam Dillon, Chen Wu, and Elizabeth Chang, "Cloud computing: issues and challenges," *24th IEEE International Conference on Advanced Information Networking and Applications* (2010).
- [12] Eric Keller, Jakub Szefer, Jennifer Rexford, and Ruby B. Lee, "NoHype: virtualized cloud infrastructure without the virtualization," *The 27th Annual International Symposium on Computer Architecture* (June 19-22, 2010).
- [13] Afkham Azeez, Srinath Perera, Dimuthu Gamage, Ruwan Linton, Prabath Siriwardana, Dimuthu Leelaratne, Sanjiva Weerawarana, and Paul Fremantle, "Multi-Tenant SOA middleware for cloud computing," *IEEE 2nd International Conference on Cloud Computing* (2010).
- [14] Verizon RISK Team, "Data breach investigations report (DBIR)," (2012).
- [15] Prasad Saripalli, and Ben Walters, "QUIRC: a quantitative impactand risk assessment framework for cloud security," *IEEE 2nd International Conference on Cloud Computing* (2010).
- [16] T. Ristenpart, E. Tromer, H. Shacham, and S. Savage, "Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds", in *Computer and Communications Security (CCS)*, 2009.
- [17] F. Gens, "IT cloud services user survey, pt.2: top benefits & challenges," Oct. 2008. (<http://blogs.idc.com/ie/?p=210>).
- [18] Hagai Bar-El, "Introduction to side channel attacks,"
- [19] Ruoyu Wu, Gail-Joon Ahn, Hongxin Hu, and Mukesh Singhal, "Information flow control in cloud computing," (9-12 Oct. 2010).
- [20] Augusto Ciuffoletti, "Monitoring a virtual network infrastructure," (October 2010).
- [21] Intel IT Center, *Seven Steps for Building Security in the Cloud from the Ground Up* (September 2011).
- [22] Mehmet Yildiz, Jemal Abawajy, Tuncay Ercan, Andrew Bernoth, *A Layered Security Approach for Cloud Computing Infrastructure*, 10th International Symposium on Pervasive Systems, Algorithms, and Networks (2009).
- [23] S. Mansfielddevine, "Danger in the clouds," *Network Security*, vol. 2008, no. 12, pp. 9–11, Dec. 2008.
- [24] W. Jansen and T. Grance, "Guidelines on Security and Privacy in Public Cloud Computing."
- [25] G. Blair, F. Kon, W. Cirne, D. Milojicic, R. Ramakrishnan, D. Reed, and D. Silva, "Perspectives on cloud computing: interviews with five leading scientists from the cloud community," *Journal of Internet Services and Applications*, vol. 2, no. 1, pp. 2–9, Jun. 2011.
- [26] S. Pearson and A. Benameur, "Privacy, Security and Trust Issues Arising from Cloud Computing," *2010 IEEE Second International Conference on Cloud Computing Technology and Science*, vol. 8, no. 6, pp. 692–702, Nov. 2010.
- [27] Cloud Security Alliance, "Security as a Service, Defined Categories of Service," 2011.
- [28] G. I. Davida, D. L. Wells, and J. B. Kam, "Security and Privacy," *IEEE Concurrency*, vol. 8, no. 2, pp. 24–21, 2000.
- [29] M. Al Morsy, J. Grundy, and I. Müller, "An Analysis of The Cloud Computing Security Problem," 2010.
- [30] A. S. Ibrahim, J. Hamlyn, and J. Grundy, "Emerging Security Challenges of Cloud Virtual Infrastructure," 2010.
- [31] Z. Chen and J. Yoon, "IT Auditing to Assure a Secure Cloud Computing," *2010 6th World Congress on Services*, pp. 252–259, Jul. 2010.
- [32] S. Bleikertz, M. Schunter, C. W. Probst, and K. Eriksson, "Security Audits of Multi-tier Virtual Infrastructures in Public Infrastructure Clouds Categories and Subject Descriptors," pp. 92–102.
- [33] R. Chakraborty, S. Ramireddy, T. S. Raghu, and H. R. Rao, "Assurance Practices of Cloud Computing," pp. 29–27, 2010.
- [34] Charles Reiss, John Wilkes, and Joseph Hellerstein, "Google clusterusage traces : format + schema," 2012.
- [35] Sheng Di, Derrick Kondo, and Walfredo Cirne, "Characterization and Comparison of Cloud versus Grid Workloads," *IEEE International Conference on Cluster Computing*, 2012.
- [36] Charles Reiss, Alexey Tumanov, Gregory R. Ganger, Randy H. Katz, and Michael A. Kozuch, "Towards understanding heterogeneous clouds at scale: Google trace analysis," 2012.
- [37] Charles Reiss, Alexey Tumanov, Gregory R. Ganger, Randy H. Katz, and Michael A. Kozuch, "Heterogeneity and dynamics of clouds at scale," *Proceedings of the Third ACM Symposium on Cloud Computing*, 2012.
- [38] Liu Zitao, and Sangyeun Cho, "Characterizing Machines and Workloads on a Google Cluster," *41st International Conference on Parallel Processing Workshops*, 2012.
- [39] Ismael Solis Moreno, Peter Garraghan, Paul Townend, and Jie Xu, "An Approach for Characterizing Workloads in Google Cloud to Derive Realistic Resource Utilization Models," the 7th IEEE International Symposium of Service-Oriented System Engineering, 2012.
- [40] Francisco Rocha, Thomas Gross, Aad van Moorsel, "Defense-in-depth Against Malicious Insiders in the Cloud," *IEEE International Conference on Cloud Engineering*, 2012.
- [41] Peng Li, Debin Gao, and Michael K. Reiter, "Mitigating Access-Driven Timing Channels in Clouds using StopWatch," the 42nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks, 2012.