A Multi-Tenancy Cloud Trust Model using Quality of Service Monitoring: A Case of Infrastructure as a Service (IaaS)

Pascal M. Mutulu
School of Computing and Informatics
University of Nairobi, Kenya

Andrew M. Kahonge
School of Computing and Informatics
University of Nairobi, Kenya

ABSTRACT

Digitization and changes in technological trends have necessitated the need for enterprises to start or have plans to migrate their services to cloud computing environments. Therefore, benefiting from the many advantages that come with cloud computing. When a service level agreement is made between a cloud consumer and the cloud provider, the consumer usually is left with no choice but to trust the provider will deliver their bit. They have faith but lack a way of verifying or even doing QoS monitoring on their own. To solve this problem, we propose a multi-tenancy cloud trust model that uses QoS monitoring. It focuses on Infrastructure as a Service and, as seen in the results, assists cloud consumers in evaluating cloud service providers well before they purchase services. This prevents them from leasing already congested clouds or which do not meet their specifications. Cloud providers also stand to gain. A provider that honors their SLAs will tend to be more trusted than one who does not, leading to a better reputation.

Keywords

Cloud computing, Infrastructure as a service, Quality of Service, Monitoring, QoS monitor

1. INTRODUCTION

Most companies and organizations are going through digital transformations by automating their traditional business processes [1]. It states that companies not able to embrace the digitizing world may be victims of "digital Darwinism," and thus, enterprises that cannot adapt to technological trends may not survive. With all this digitization, the various services need to reside in servers; building a data center (DC) is costly, consumes much time, and requires enormous capital to maintain. This has necessitated these clients to look for third-party providers to offer them platforms to deploy their services.

The third-party providers, commonly known as "cloud providers," consume virtualization to create different clients' different instances. At a minimum, virtualization technology has host hardware, hypervisor, and virtual machines. The client's instances must not interfere with each other but share resources [2]. They should be segregated and appear as a physical server to the customer. This necessitates the need for multi-tenancy technology.

On acquiring such services as compute, storage, networking, the client and the Cloud Service Provider (CSP) agree on some service level agreements (SLAs) [3]. The cloud provider may commit that their cloud has all the cloud-computing characteristics such as redundancy, high availability, fault tolerance, optimum performance, among others. Since the cloud provider has most of the control, depending on the client's service, there is a need to confirm the provision of the agreements.

To enhance trust to the client that they are benefiting from all they purchased, there needs to be a third party means to confirm the same. This is the reason for coming up with this trust model to address that problem. It makes it possible for the cloud consumer to confirm the cloud platform status in real-time at any moment if they have access to the third-party QoS monitor.

The main objective was to come up with a multi-tenancy cloud trust model using QoS monitoring for IaaS. Specifically, reviewed the various trust models used in multi-tenancy clouds, came up with a multi-tenancy cloud trust model using QoS monitoring, developed a prototype of the proposed model, and finally evaluated the model.

2. RELATED WORK

2.1 Chains of trust in the cloud

This model focuses on customer verification of services through third-party professionals to foster trust through QoS monitoring and SLA verification [4]. They suggest a chain of trusts between the cloud provider, auditor, broker, cloud user, and the cloud service. It does not give control to the end customer to monitor any aspects on their own as most of the control in relations to trust is bestowed on the cloud auditor, such that as long as they have certified the cloud broker, cloud provider, and the cloud service, then the cloud user will also trust them.

The architecture of the cloud might change quite often. This might mean that the audit done by the cloud auditor might need some changes as soon as the architecture changes. This may not always be the case, as most audits are done yearly or after any significant changes.

2.2 Cloud Computing Service Security Strength Measuring Trust Model

This trust model measures cloud security and establishes a trust value [5]. It uses some considerations such as identity management, authorization, authentication, confidentiality, among others, to come up with the value. It only focuses on security in the cloud environment, and the parameters are evaluated through interaction with the cloud environment. Customers consume the trust value to evaluate the cloud vendor they ought to purchase. This model does not give the customers some monitoring level or a way of verifying the trust once they have acquired the cloud service.

2.3 Collaborative cloud services Authorization models in multi-tenancy environments

The research was about collaborative cloud services authorization models in multi-tenancy environments and suggested that trust between CSPs and cloud users is like the trust relations between organizations and their contracted outsourcing vendors [6]. They identify three independent organizations: the enterprise, the outsourcing company, and the auditing firm responsible for storage services, service coding, and reporting, respectively. They propose a mathematical model for authorization as a service. They do not seem to provide any tools employable to foster trust to the cloud consumers directly.

2.4 Trusted computing environment model (MTCEM)

As a countermeasure to cloud security risks, a discussion of a proposed Multi-tenancy trusted computing environment model (MTCEM) that implements the trusted computing groups (TCG) is done [7]. Trusted Computing Platform (TCP) is a set of principles, standards, and technologies that makes a data owner trust and holds accountable the underlying computing infrastructure where applications that create, store and change their data runs [8]. TCP comprises two assertions discussed below:

i. Transitive trust – This suggests that computing platforms might only adjust from a Core Root of Trust Measurement (CRTM). This includes hardware or even encrypted firmware certified by a certified body of specialists and thus deemed trustworthy. Implicitly trust is implied such that one level of initialization trusts the previous.

ii. Platform attestation – A computing policy displays to a third party that it is trusted. The other systems attest to the system's trustworthiness it interacts with and thus, in turn, considered reliable by other systems. The main challenge is how to express conventional reasonable and quantifiable metrics useful to show how trustworthy the system is.

A critical look at the trusted computing (TC) model discussed above presents some limitations and has some drawbacks as presented by the internet community. Professor Anderson (University of Cambridge) claims that it is more of Information Technology (IT) industry than for people. It might give providers much power making them come up with unfair policies.

GNU project founder and Free Software Foundation (FSF) president says that trusted computing may expose free operating software and free application to a risk that users may not have the capability to run them anymore [9]. Such criticisms raise some critical issues with trusted computing that may make it impossible to implement actual technology.

2.5 Subsequent Cross-Tenant Trust Model (CTTM)

They suggest a cross-tenant trust model (CTTM) in cloud computing. The model consists of unilateral trust relations that reflect access control needs by two different tenants: the trustor and the trustee. They suggest a multi-tenant authorization as a service (MTAaaS) to enable the implementation, as shown in figure 1 below [10].

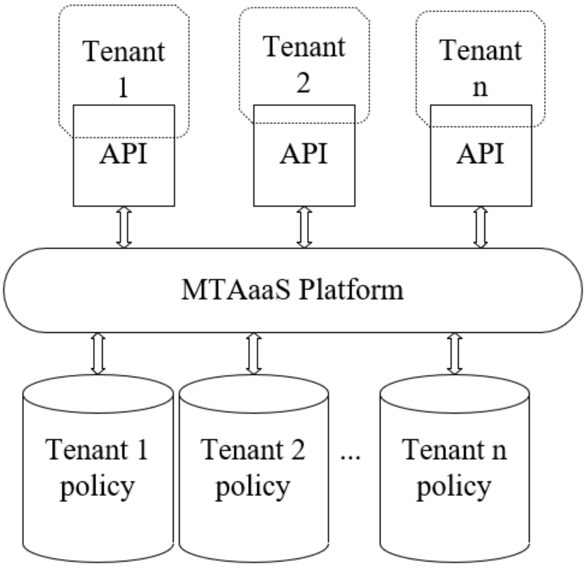

Figure 1: Cross-tenant trust model

Their crucial contribution is proposing a cloud-based MTAaaS where different tenants communicate to the MTAaaS platform using an application programming interface (API), which then gets the policies specific to the tenant and thus provides application-centric security.

Their work primarily focuses on the trust between the different tenants. Different tenants in a multi-tenant architecture need not know each other as essential; actualizing such a model could be difficult and does not contribute much trust between the CSP and the consumer. Cloud hardware providers already provide abstraction between tenants that is an essential requirement for such architectures.

2.6 QoS Monitoring

Provisioning of the appropriate resources to cloud workloads depends on the QoS requirements of such workloads [11]. In the IaaS, compute, and storage resources are offered at a fee. The resources may include and are not limited to CPU, memory, disk, storage, networks, and bandwidth, among many others. Cloud consumers select a cloud service that can offer services with an adequate QoS guarantee. In the cloud, QoS may entail the level of performance, availability, and reliability obtainable by an application, platform, or the infrastructure that hosts it. Sample QoS metrics or cloud SLA parameters include availability, throughput, response time, memory utilization, processing capacity, among others.

2.7 Cloud computing reference architecture

NIST identifies five major cloud actors. These key players make the complete architecture of the cloud ecosystem. Each is an independent element with its structure and thus works together through the user of well-predefined technologies. Under the cloud provider is where service models reside. The IaaS is usually the bottom layer below the PaaS and SaaS and seats directly above the abstraction layer. The actors are described below.

i. Cloud Consumer - A person or organization that maintains a business relationship and uses Cloud Providers' service.

ii. Cloud Provider - A person, organization, or entity responsible for making a service available to interested parties.

iii. Cloud Auditor - A party that can conduct an independent assessment of cloud services, information system operations, performance, and cloud implementation security.

iv. Cloud Broker - An entity that manages the use, performance, and delivery of cloud services and negotiates relationships between Cloud Providers and Cloud Consumers.

v. Cloud Carrier - An intermediary that provides connectivity and transport of cloud services from Cloud Providers to Cloud Consumers.

2.8 The Trust Model

Figure 2 below represents a prototype of the proposed trust model. The cloud computing reference architecture model has various actors. The plan was to introduce the third-party QoS monitor who can keep watch of the cloud service provider platform in real-time compared to the cloud auditor who occasionally audits the cloud.

Cloud consumers wishing to join the cloud, in addition to the audit reports done by the cloud auditor, can benefit from the QoS monitor by being able to monitor in real-time the cloud's overall performance. The third-party QoS monitor is built by a third party and connects to the cloud through a secure, high throughput direct connection link.

The overall prototype architecture of the proposed model, data flow from the cloud service provider to the cloud consumer, and how the various tools used are utilized to accomplish the solution are shown in the exact figure 2 below. At the cloud service provider, the OpenStack cloud platform is set up. It serves as infrastructure as a service cloud model platform. A virtual data center (vDC) is created for cloud consumers to create their virtual machines, images, and their virtual private clouds (VPC).

Broker is set up where third-party QoS monitors can connect to avoiding direct connection to the core cloud platform. The broker serves the sole purpose of mediating between the cloud platform and external networks. The broker has a module installed on the cloud platform from where it scraps metrics. The module is referred to as a node exporter. The tool used to work as the broker is Prometheus.

The third-party QoS monitor is once allowed to connect to the cloud service broker through a secure connection by use of application programming interfaces (APIs) and can get the metrics scraped from the cloud platform by the broker. The tool used to represent the third-party QoS monitor in this project is Grafana. It has the capability of visualizing the cloud platform status by displaying the various metrics. This is then shared with the cloud consumers. Various ways can be used to share the data, such as: provisioning them and giving them a web portal to log into or connecting through APIs or manually share the report with them inform of excel pdf or any other formats.

The cloud consumer uses the third-party QoS monitor's information to make trustworthy decisions related to service level agreements with the cloud provider. If allowed, they could also connect to the third-party QoS monitor with APIs.

3. METHODOLOGY

The exploratory research design was the overall research design strategy employed in this research. The Delphi method [12][13] was utilized to do data collection and analysis, population, and sampling, among other activities. The research design helped explore the concepts and techniques around trust in multi-tenancy clouds and thus able to develop a well-thought and researched conceptual framework from where a model is developed.

3.1 Population and Sampling

The population from which the sample is drawn consists of cloud experts or individuals who have been supporting cloud consumers technically. This is because such individuals know how the multi-tenancy cloud works, and due to the support they do to the cloud consumers, they know some pain points of these customers. A target of ten cloud experts either managing clouds directly or supporting cloud formed the population.

Purposive sampling was used in the selection of these individuals as there was a need to focus on their technicality and experience about cloud or using the cloud. One panel consisting of five individuals was to be deemed enough.

3.2 Data Collection

This being qualitative research and depending on the methodology, and the data collection technique consumed was interviews through open-ended questions and literature review. Since the respondents had to remain anonymous to avoid bias, they were contacted independently.

3.3 Data Analysis

The data analysis idea was to look for patterns within data and draw conclusions [14]. Data provided by the panelist was analyzed by comparing the common patterns among the various respondents. This helped in determining the QoS metrics to focus on as well as the best way to architect the model.

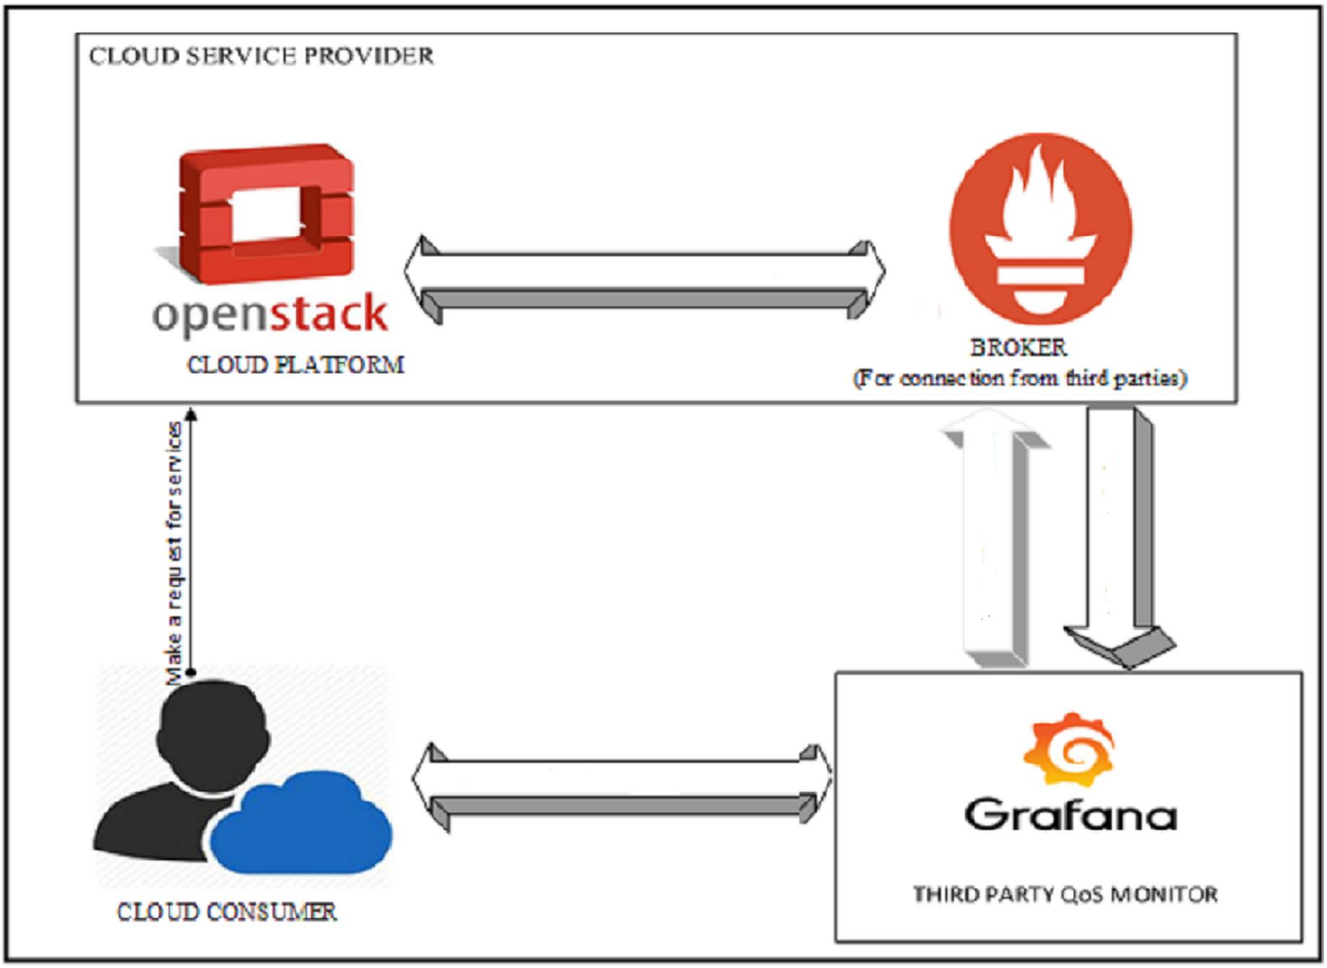

Figure 2: The overall prototype architecture of the proposed trust model

From the data collected through the open-minded interview sessions with the panelists, several conclusions were drawn. The first round of interviews was conducted before the prototype development began. It sought to establish cloud consumers' pain points in terms of QoS monitoring and acquisition of cloud services, institutions majorly consuming cloud services hosted locally, the best ways in which the trust model could be designed, and if any of them knew the existence of such a trust model in Kenya. The general conclusions drawn were as below:

i. Most cloud customers who consume cloud services within Kenya are primarily institutions that do not want their data to go outside the country. This is majorly contributed by policies within such institutions. This was according to four out of six respondents.

ii. According to two out of three panelists who were major cloud consumers, such customers in (i) choose to remain with a particular cloud consumer despite poor services as there were no enough options in Kenya, or they lacked insights about the other companies and feared they might experience the issues as with their current provider.

iii. Most cloud consumers lack visibility of their cloud service provider platform, making it difficult to make some decisions.

iv. None of the individuals interviewed knew of any existing platform in Kenya's CSP where they could verify the QoS metrics using a third-party QoS monitor.

v. Three of the cloud consumers said that such a trust model would help cloud consumers make some decisions related to cloud hosting services.

vi. They also expressed that CSP might not want a direct connection to their cloud platform by third parties as it might be a way of exposing them.

4. RESULTS

4.1 QoS Metrics Analysis

Bearing in mind that the resources allocated to a physical cloud platform are known as shown by the below sample figures 3,4, and 5, the cloud could be configured to assign logical capacity to cloud consumers that the cloud could physically accommodate possible. This means it is possible to assign fake resources like virtual CPU, ram, and hard disk to a consumer during provisioning and visible in their virtual data center. Unless they know how to confirm, then they would believe the suitable capacity has been provisioned.

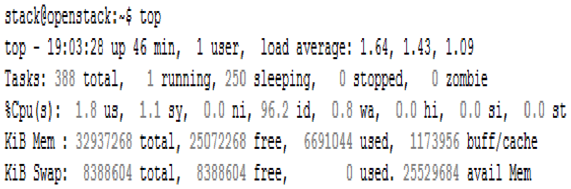

```
stack@openstack:~$ top
top - 19:03:28 up 46 min,  1 user,  load average: 1.64, 1.43, 1.09
Tasks: 388 total,   1 running, 250 sleeping,   0 stopped,   0 zombie
%Cpu(s):  1.8 us,  1.1 sy,  0.0 ni, 96.2 id,  0.8 wa,  0.0 hi,  0.0 si,  0.0 st
KiB Mem : 32937268 total, 25072268 free,  6691044 used,  1173956 buff/cache
KiB Swap:  8388604 total,  8388604 free,        0 used. 25529684 avail Mem
```

Figure 3: Physical, allocated RAM, and Swap.

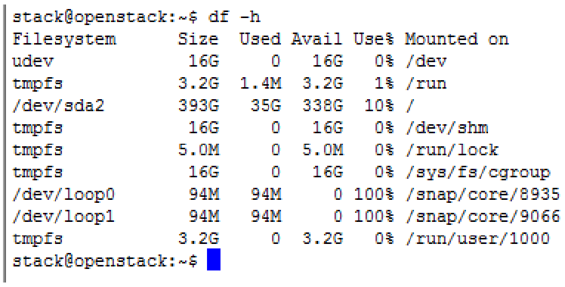

```
stack@openstack:~$ df -h
Filesystem      Size  Used Avail Use% Mounted on
udev             16G     0   16G   0% /dev
tmpfs           3.2G  1.4M  3.2G   1% /run
/dev/sda2       393G   35G  338G  10% /
tmpfs            16G     0   16G   0% /dev/shm
tmpfs           5.0M     0  5.0M   0% /run/lock
tmpfs            16G     0   16G   0% /sys/fs/cgroup
/dev/loop0       94M   94M     0 100% /snap/core/8935
/dev/loop1       94M   94M     0 100% /snap/core/9066
tmpfs           3.2G     0  3.2G   0% /run/user/1000
stack@openstack:~$
```

Figure 4: Physically allocated disk space and partitions

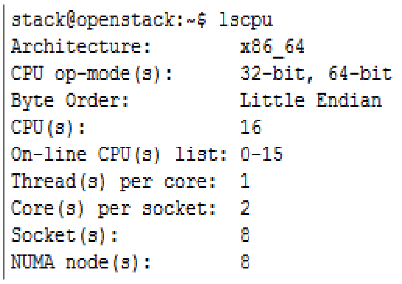

```
stack@openstack:~$ lscpu
Architecture:        x86_64
CPU op-mode(s):      32-bit, 64-bit
Byte Order:          Little Endian
CPU(s):              16
On-line CPU(s) list: 0-15
Thread(s) per core:  1
Core(s) per socket:  2
Socket(s):           8
NUMA node(s):        8
```

Figure 5: Physically allocated CPU

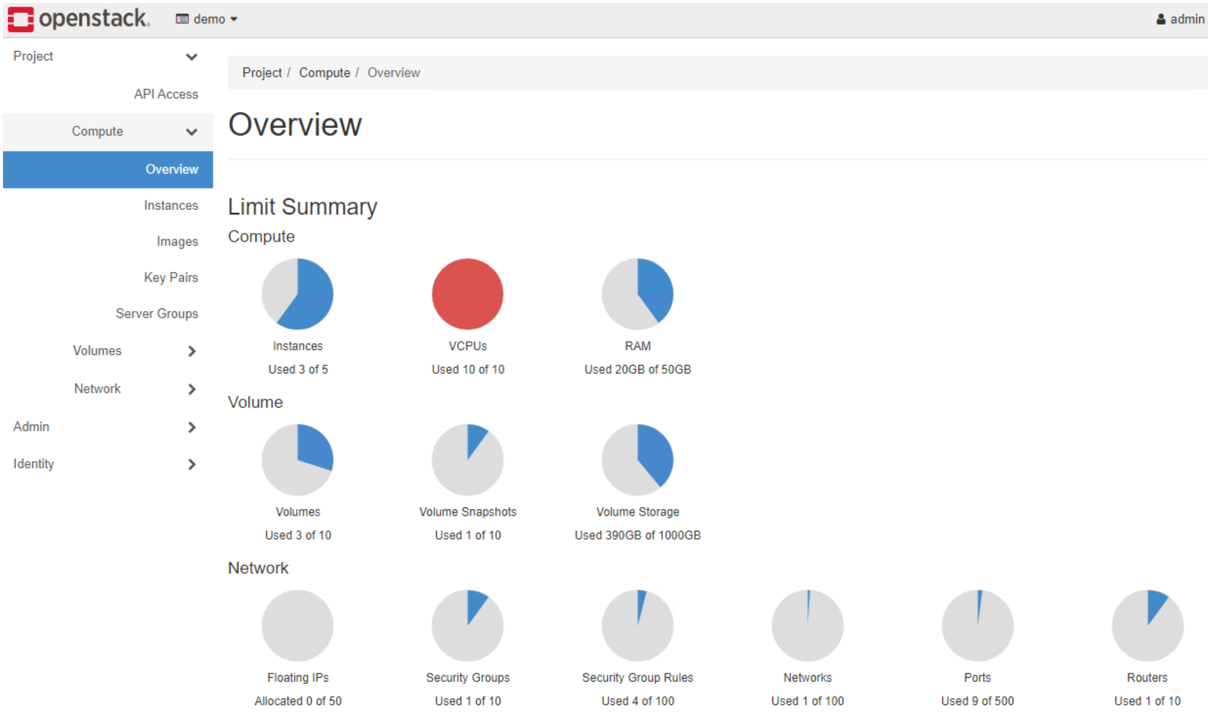

Project / Compute / Overview

Overview

Limit Summary

Compute

Instances Used 3 of 5

VCPUs Used 10 of 10

RAM Used 20GB of 50GB

Volume

Volumes Used 3 of 10

Volume Snapshots Used 1 of 10

Volume Storage Used 390GB of 1000GB

Network

Floating IPs Allocated 0 of 50

Security Groups Used 1 of 10

Security Group Rules Used 4 of 100

Networks Used 1 of 100

Ports Used 9 of 500

Routers Used 1 of 10

Figure 6: Sample resources allocated to a tenant

The above figures 3,4 and 5 show the resources allocated physically on the hardware running the cloud platform and represent the RAM, disk, and CPU, respectively.

Table 1: Metrics analysis

Physical RAM	31 GB
Physical disk on the partition hosting customers data	393GB
Tenant Demo allocated RAM	50 GB
Tenant Demo allocated HDD	1000 GB

Table 1 above shows sample metrics assigned to a cloud consumer (demo), as also shown in figure 6 below. Doing some analysis, it was realized that the user is assigned a RAM of 50 GB, whereas the physical cloud platform has a total of 31 GB RAM. Again, the user is allocated storage of 1000 GB, but the cloud platform has 393 GB. These are just sample metrics as it was possible to exaggerate many other metrics. With the QoS monitor, such illegalities could have been noticed firsthand.

4.2 Manipulation of metrics

Since open-source cloud platforms are highly configurable, it is possible to tune them to serve the purpose needed. Rogue CSPs cloud program it such that they steal small portions of resources from cloud consumers. The portions might almost be negligible, but they might affect applications efficiency.

For instance, in the cloud platform, one could edit consumer virtual data center metrics by just clicking the button and updating. A simple solution to this would be to set a threshold such that one cannot provide more than the cloud platform can physically accommodate, which may not always be the case.

5. DISCUSSION

As seen in the above results, CSPs can allocate logical resources to cloud consumers that even exceed what the physical infrastructure can accommodate. That was tested using various users in the prototype, and we were able to confirm that. The QoS monitor comes in to help cloud consumers confirm whether what they have been provisioned is available. This answers the first research question on the match between the provisioned and actual capacity provisioned.

Comparing the proposed model with the other models within related works it was found out that it is possible to give control to the cloud consumer to verify QoS metrics in real-time without having to go through the cloud auditor. One of the models closely relates to the proposed work of QoS monitoring by third-party professionals, mainly the cloud auditor, without giving the cloud consumers the ability to view metrics from any tool as audit reports might not be up to date. This makes the proposed trust model a more viable solution.

There is the greater possibility of commercializing the model. A cloud service provider wishing to show their commitment to meeting SLAs can pay a third-party QoS monitor and have consumers and potential clients access the free metrics. This is likely to boost sales on the CSPs side as it makes it easy for clients to verify cloud status before they purchase a service, thereby boosting their confidence and trust. Besides, the third-party QoS monitor would act as the cloud platform monitoring operations center and able to share unbiased feedback to the CSP where there are anomalies. The principle of "Trust but Verify" in the cloud would best fit this model and make it a source of trust using the QoS monitor.

6. CONCLUSION

Several trust models used in multi-tenancy clouds were identified, as discussed in related work. Most of them focused on trust among the different tenants or trust through third-party professionals. None focused on a third-party QoS monitor where cloud consumers could get metrics in real-time. This was a great opportunity for us to think through and come up with the developed model

Also, a multi-tenancy clouds trust model using QoS monitoring was developed. The model enables cloud consumers to go through the third-party QoS monitor to confirm in real-time the status of the cloud service provider platform status. This would help considerably once making the initial decision of which cloud platform to adopt.

Besides, a prototype to showcase the proposed model was developed. It captures the anomalies concerning the assignment of resources to cloud customers. The prototype was developed using readily available tools, which made it easier and cheaper to consume. Configuration and programming the tools to accomplish what was intended to achieve was successful.

6.1 Limitation

This system works well where cloud consumer is constrained to using CSP within a country. Since most of the full-fledged public cloud companies such as Microsoft Azure and Amazon web services do not have a presence in all the countries, then such consumers are left with few chances to choose from and hence a need to scrutinize what they have locally before they can leap in. Also, the model works best where the third QoS monitor and the CSPs trust each other first. CSP negotiation should be from a management level where the QoS monitor also acts as a monitoring entity for the said CSP, and hence they see it as a benefit and not a way of exposing them.

6.2 Further work

There is an opportunity to improve on the granularity of information given to the cloud consumer. This means that after the cloud consumer has already acquired the cloud service, they can continue using the third-party QoS monitor to show only their virtual data center that hosts their resource. This will enable them to see the overall status of the cloud and have more visibility to the resources allocated to them only. Cloud service providers can adopt the model. They could replicate it within their environment, making them more transparent to their discerning cloud consumers. As discussed, this could go a long way in helping them strengthen their reputation, which may lead to more sales and revenue growth.

7. REFERENCES

[1] Ismail, M.H., Khater, M., and Zaki, M., 2017. Digital Business Transformation and Strategy: What Do We Know So Far. *Cambridge Service Alliance, November*.

[2] AlJahdali, H., Albatli, A., Garraghan, P., Townend, P., Lau, L. and Xu, J., 2014, April. Multi-tenancy in cloud computing. In *2014 IEEE 8th International Symposium on Service Oriented System Engineering* (pp. 344-351). IEEE.

[3] Ansari, A., 2018. Service Level Agreement Governance For Cloud Computing. 10.13140/RG.2.2.11361.15206.

[4] Huang, J. and Nicol, D.M., 2013. Trust mechanisms for cloud computing. *Journal of Cloud Computing: Advances, Systems and Applications*, *2*(1), p.9.

[5] Shaikh, R. and Sasikumar, M., 2015. Trust model for measuring security strength of cloud computing service. *Procedia Computer Science*, *45*, pp.380-389.

[6] Tang, B., Sandhu, R. and Li, Q., 2015. Multi- tenancy authorization models for collaborative cloud services. *Concurrency and Computation: Practice and Experience*, *27*(11), pp.2851-2868.

[7] Brown, W.J., Anderson, V. and Tan, Q., 2012, September. Multitenancy-security risks and countermeasures. In *2012 15th International Conference on Network-Based Information Systems* (pp. 7-13). IEEE.

[8] Anderson, R., 2003. 'Trusted Computing' Frequently Asked Questions, https://www.cl.cam.ac.uk/~rja14/tcpa-faq.html accessed on 21st December 2019.

[9] Stallman, R, 2018. Can You Trust Your Computer?, https://www.gnu.org/philosophy/can-you-trust.en.html accessed on 21st December 2019.

[10] Tang, B. and Sandhu, R., 2013, August. Cross-tenant trust models in cloud computing. In *2013 IEEE 14th International Conference on Information Reuse & Integration (IRI)* (pp. 129-136). IEEE.

[11] Odun-Ayo, I., Ajayi, O. and Falade, A., 2018. Cloud Computing and Quality of Service: Issues and Developments.

[12] Skinner, R., Nelson, R.R., Chin, W.W. and Land, L., 2015. The Delphi Method Research Strategy in Studies of Information Systems. *Cais*, *37*, p.2.

[13] Rowe, G. and Wright, G., 2001. Expert opinions in forecasting: the role of the Delphi technique. In *Principles of forecasting* (pp. 125-144). Springer, Boston, MA.

[14] Oates, B.J., 2005. Researching information systems and computing. Sage.