Cloud Multi-Tenancy: Issues and Developments

Isaac Odun-Ayo
Covenant University, Ota, Nigeria
isaac.odun-ayo
@covenantuniversity.edu.ng

Sanjay Misra
Covenant University, Ota, Nigeria
sanjay.misra
@covenantuniversity.edu.ng

Olusola Abayomi-Alli
Covenant University, Ota, Nigeria
olusola.abayomi-alli
@covenantuniversity.edu.ng

Olasupo Ajayi
University of Lagos, Lagos, Nigeria
olasupoajayi@gmail.com

ABSTRACT

Cloud Computing (CC) is a computational paradigm that provides pay-per use services to customers from a pool of networked computing resources that are provided on demand. Customers therefore does not need to worry about infrastructure or storage. Cloud Service Providers (CSP) make custom built applications available to customers online. Also, organisations and enterprises can build and deploy applications based on platforms provided by the Cloud service provider. Scalable storage and computing resources is also made available to consumers on the Clouds at a cost. Cloud Computing takes virtualization a step further through the use of virtual machines, it allows several customers share the same physical machine. In addition, it is possible for numerous customers to share applications provided by a CSP; this sharing model is known as multi-tenancy. Though Multi-tenancy has its drawbacks but however, it is highly desirable based on its cost efficiency. This paper presents the comprehensive study of existing literatures on relevant issues and development relating to cloud multi-tenancy using reliable methods. This study examines recent trends in the area of cloud multi-tenancy and provides a guide for future research. The analyses of this comprehensive study was based on the following questions relating to recent study in multi-tenancy which are: what is the current trend and development in cloud multi-tenancy? Existing publications were analyzed in this area including journals, conferences, white papers and publications in reputable magazines. The expected result at the end of this review is the identification of trends in cloud multi-tenancy. This will be of benefit to prospective cloud users and even cloud providers.

UCC'17 Companion, December 5–8, 2017, Austin, TX, USA

ACM ISBN 978-1-4503-5195-9/17/12…$15.00
https://doi.org/10.1145/3147234.3148095

CCS CONCEPTS

• **Distributed Architecture**→Cloud computing • **Computing methodologies**→ Empirical studies**, Survey and overviews**

KEYWORDS

Cloud computing; multi-tenancy; virtualization; IAAS; PAAS; SAAS; CSP

1 INTRODUCTION

Cloud computing (CC) according to the National Institute of Science and Technology (NIST) is defined as a model that enables convenient and ubiquitous access to a shared pool of configurable computing resources that can be rapidly setup or torn down on users' demand with little or no service provider intervention [1]. Resources in CC could be at the hardware level: Infrastructure-As-A-Service (IAAS), at the software level: Software-As-A-Service (SAAS) or at a developer level: Platform-As-A-Service (PAAS) and deployed either as a private, public, community or hybrid model. In the IAAS model, hardware resources can be instantly allocated or released to customers through the use of Virtual Machines [2]. Automated resource allocation techniques such as Amazon's Elastic Cloud Compute are used by Cloud providers to achieve this [2]. Furthermore, easy management and better resource utilization are achieved by hosting multiple customers on the same Physical Machine (PM) using VMs. This is known as Multi-tenancy. CC and related technologies are advancing at such a drastic rate and will continue to impact the Information Technology world for many more years.

Cloud Service Providers (CSPs) are improving their services while Cloud users continue to understand and enjoy the benefits offered. CSPs use the SAAS to provide applications for the Cloud users over the Internet. Consequently, cloud user accesses and utilizes such applications using any device, anytime, anywhere without any worries about installations and system requirements. However, the only caveat is on the pay per use model of such applications. CSP use the PAAS to provide platform to enable Cloud users create and deploy their own applications. Such application are written in programming languages specified by the CSP who also provide relevant application programming and Web 2.0 interfaces, libraries, and scripts. CSP use the IAAS to offer compute resources and storage services to Cloud users. CSPs provide processors, memory

and storage infrastructure to Cloud users. With respect to the deployment models, private Clouds are fully owned, accessed and managed by an organization. They can either be hosted on premise or by a third party. Private Clouds are considered the most secure Cloud model [3] unlike Public Cloud that are owned by major CSPs who have large infrastructure and data centres spread across several geographical locations. In addition, small and big enterprises alike can utilize compute and storage resources offered by public Cloud. Public Clouds are considered less secured in comparison to Private Clouds. Community Clouds are owned by several organizations with shared and/or common interest. The infrastructure can be managed by the community or through a third party. Hybrid Cloud is a combination of either private, public or community Cloud. This model allows organizations maintain crucial services within their private data centre while migrating less sensitive data to the Cloud. Multi-tenancy is one of the core concepts of Cloud computing. This allows several users share some resources at the same time within the same location. Multi-tenancy refers to resource sharing in Cloud computing where any resource object is reusable in the Cloud infrastructure [4]; or a model that allows sharing of hardware infrastructure by various users from different organizations [5]. It is an approach to share an application instance between multiple tenants by providing every tenant a dedicated share of the instance, which is isolated from others with regards to performance and data privacy [6]. It allows making use of full economy of scale, by offering multiple Cloud customers customizable virtual dedicated hardware with shared instances of the same application and database [7]. Reusable objects must however be carefully controlled and managed since they can create a lead to vulnerability and violate confidentiality by providing avenue for possible data leakage. Also if not well managed, Multi-tenancy can result in stiff competition for resources among tenants [8]. Tenants are groups of users sharing the same view of an application. Tenants are groups of users sharing the same view of an application. Tenants are groups of users sharing the same view of an application. This view includes the data they access, the configuration, the user management and particular functionality. A multi-tenant application shares the same instance among different customers to reduce overhead cost. Handling different tenants within one application instance requires several modifications as every tenant need its own view. Multi-tenancy is a natural result of try to achieve economic gain in Cloud computing by utilizing virtualization and allowing resources sharing [7]. Multi-tenancy can be seen differently in different service types. In SAAS, applications are provided as a service by the CSP where the Cloud user cannot monitor or control the underlying infrastructure. In this case, multi-tenancy implies that multiple customers utilize the same service or application provided by the CSP regardless of the underlying resources [5]. In IAAS, customers are capable of dynamically provisioning computing, storage and networking resources and can control but cannot manage the underlying infrastructure. In this case multi-tenancy occurs when two or more VMs belonging to different Cloud users share the same physical machine.

The purpose of this paper is to examine multi-tenancy in Cloud computing. The rest of the paper is organized as follows: Section 2 examines related work, while Section 3 discusses various challenges in multi-tenancy. Section 4 highlights current trends in the industry. Section 5 concludes the paper and suggests future work.

2 RELATED WORK

[4] proposed a multi-tenancy authorization system with federated identity for Cloud-based environments using shibboleth. This approach utilized a tool know as shibboleth to facilitate the process of authentication, authorization and the implementation of the identity federation. The aim is to provide a more reliable means of association between a client and service provider. [6] proposed a multi-tenancy in Cloud computing with a focus on resource sharing on the Cloud. The paper proposed a model based on threats related to multi-tenancy and validated this model using real data from Google. [7] presented the architecture for enforcing multi-tenancy for Cloud computing environments. [8] presented an authorization system for multiple tenants in Cloud environments called Multi-Tenancy Authorization System (MTAS). MTAS was presented as an extension of the Role-Based Access Control (RBAC) with the introduction of trust among tenants. [9] proposed a multi-tenancy authorization models for collaborative Cloud services. The study focused on cross-tenant relationship, formalizing the MTAS by bolstering the inter-tenant trust and adding an administrative model. Experimental tests yielded favourable results. A cost-benefit analysis of multi-tenancy for SAAS applications was done by [10]. The paper looked at the merits and demerits of developing and deploying SAAS applications with respect to small and medium size enterprises. Ease of application and improved hardware utilization were amongst the highlighted benefits but scalability, security, coding complexity and maintenance cost were some of the. challenges. The paper concluded that ultimately a good architecture yields the best result.

[11] proposed a model for enhancing security in multi-tenant Cloud environments. Rather than focusing on resource or power efficient resource allocation to tenants, the work centred on resource allocation in a secure manner. It is believed that considering security from the resource allocation phase gives a more secured cloud. [12] analysed architectural concerns in multi-tenant SAAS applications. The study provides an understanding of the concept of multi-tenancy then discussed the existing multi-tenancy architecture and the challenges in their implementation. Some of the highlighted challenges include: performance isolation, persistency, QoS, differentiation and customizers. [13] presented, a hybrid of workload-aware resource reservation and NoSQL Multi-tenant storage scheme for Cloud environment. This model was called Argus and the hybridization proposed addressed the performance degradation associated with the use of NoSQL while Ngo et al. (2016) proposed a multi-tenant attribute-based access control for Cloud infrastructure services. The study focused on multi-Cloud providers providing services to multiple users. An access control model was proposed and implemented for multi-Cloud and

multiple clients. Espadas et al. (2013) proposed, a tenant-based resource allocation model for scaling SAAS applications over Cloud computing infrastructures. The unique nature of SAAS is being able to scale up or down applications as needed by the customers. However the limitation of the study as highlighted by authors is the charges deducted from customers for idle processor time and unused resources. In conclusion, the study proposes a model to optimize utilization of resources for SAAS multi-tenancy. Povedani-Molina et al. (2013) presented a highly adaptable and scalable monitoring architecture for multi-tenant Clouds called DARGOS. Knowing the availability and status of resources is important to the CSP, the author finally proposed an architecture that provides active resource monitoring information to the CSP. Unfortunately, using VMs, there is a limit to the number of tenants that can be hosted on a PMs due to high requirements for every VMs [10] [14]. However, with the use of multi-tenancy, several tenants can share the same software instance, thereby indirectly improving resource utilization which ultimately translates to lower overall cost of application.

3. MULTI-TENANCY CHARACTERISTICS AND ARCHITECTURE

3.1 Multi-Tenant Characteristics

Multi-tenancy has the following characteristics which are [10]:

3.1.1 Hardware Resources Sharing. In traditional single tenant software architecture, tenants have their own VMs, which they customize to their requirements. Unfortunately, using VMs, there is a limit to the number of tenants that can be hosted on a PMs due to high requirements for every VMs [10] [14]. However, with the use of multi-tenancy, several tenants can share the same software instance, thereby indirectly improving resource utilization which ultimately translates to lower overall cost of application

3.1.2 High Degree of Configurability. In a single tenant environment, every tenant has his own customized application instance; while in a multi-tenant set up all tenant share the same application instance, which appear to the tenants as a single dedicated one. Due to this, a key requirement of multi-tenant applications is the possibility to configure and customize applications to meet the widely varied tenants' needs. In multi-tenancy, configuration options must be integrated into the product design. In view of the high degree of configurability of multi-tenant software system, it may be necessary to run multiple versions of an application next to each other.

3.1.3 Shared Application and Database Instance A single tenant application may have many running instances and they may all be different from each other because of customization. In multi-tenancy, the differences no longer exist as the application is runtime configurable. This entails that in multi-tenancy, the overall number of instances will be much lower usually one, but the application may be replicated for scalability purposes. Consequently, deployment is much easier and cheaper, particularly in the area of deploying updates, as the number of instances which are affected by the deployment action is much lower. In addition, new data aggregation opportunities are opened because all tenant data is in the same place. Hence, user behaviours traces can be collected easily, which can help improve user experience. From the forgoing characteristics, multi-tenancy allows higher utilization of hardware resources. It enables easier and cheaper application maintenance. In addition, services are provided at a lower cost, with new data aggregation opportunities.

3.2 Multi-Tenancy Architectures in Managing Data.

Multi-tenancy is the defining characteristics of Cloud computing. The shared infrastructure changes the underlying economies of enterprise applications, allowing vendors maintain a single instance for thousands of customers. In multi-tenancy Cloud environment, multiple users using the same infrastructure can access and use an application. The application design must therefore distinguish between users to ensure that they do not share each other's data. There are three different methods for achieving multi-tenancy which are: using a database, using virtualization and through physical separation. In the case of SAAS, multi-tenancy is achieved via database and configuration with isolation provided at the application layer. So at the application layer, service providers must design and implement a specific class and then create an object of the class in a manner that services the need of multiple users in an effective way. Designing SAAS application in this way will solve many issue such as the need for data security, data separation and customized applications. Virtualization is another technology for achieving multi-tenancy especially for IAAS. Virtualization allows multiple copies of operations systems (VMs) run within a PM. These multiple VMs can then share the same physical hardware resources on the PM such as network card, disk storage. Though virtualization based multi-tenancy reduces costs and expenses, but compared to multi-tenancy using database technology, it is more costly.

Multi-tenancy can also be achieved through a dedicated technology that provides resources to tenants individually. This is known as multi-tenancy via physical separation. This option is by far the most expensive. Though CSPs offer this as an optional configuration to special customers, who want to use the Cloud service but do not want to share hardware resources with other customers.

The three approaches to managing multi-tenant data in the Cloud [16] are: storing tenant data on separate databases, which is the simplest approach to data isolation; housing multiple tenants on the same database, with each tenant having his own set of tables grouped into a schema created specifically for the tenant; and by using the same database, and same set of tables to host multiple tenants' data.

The general architecture for representing multi-tenancy for effective Cloud environment is at figure 1.

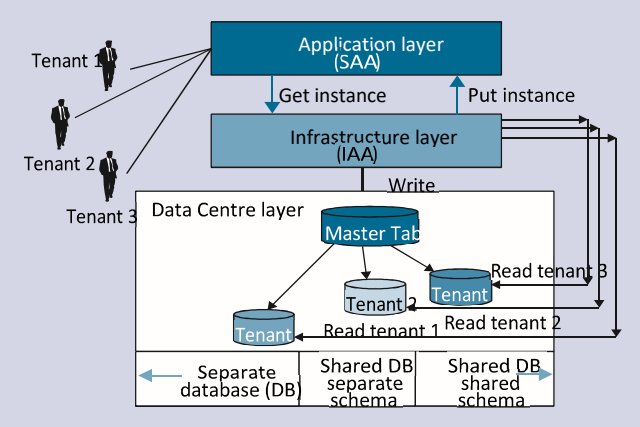

Figure 1: An Overview of a general Multi-Tenancy Cloud Architecture [5]

The architecture employs customer integration in three layers, which are the application, the infrastructure and the data-centre layer.

- The data-centre layer multi-tenancy is the most common and provides the highest level of security, if implemented correctly.
- Infrastructure layer dedicates one stack of software to a specific customer, with the possibility of deploying multiple stacks to each customer. The hardware requirement depends on the actual service used.
- Application layer, which involves both the software and infrastructure layer. This type of multi-tenancy can compromise security because application methods and database queries can access and store data from different user accounts. However, if implemented correctly, it can offer significant cost savings.

While multi-tenancy, on Cloud environments provides seemingly limitless scalability and an alternative to the expensive data centre infrastructure, it raises security and privacy issues because it hands the processing and storage task over to third parties. This requires building adequate security into every aspect of the SAAS application, as well as for every IaaS virtual service. This involves using filtering which provides an intermediary layer between a tenant and data source. Another is permission which uses access control lists. Finally is encryption which obscures every tenant's critical data.

3.3 Conceptual Multi-Tenant Architecture in Managing Applications

Multi-tenancy affects almost all layers of a typical application as shown in Figure 2.

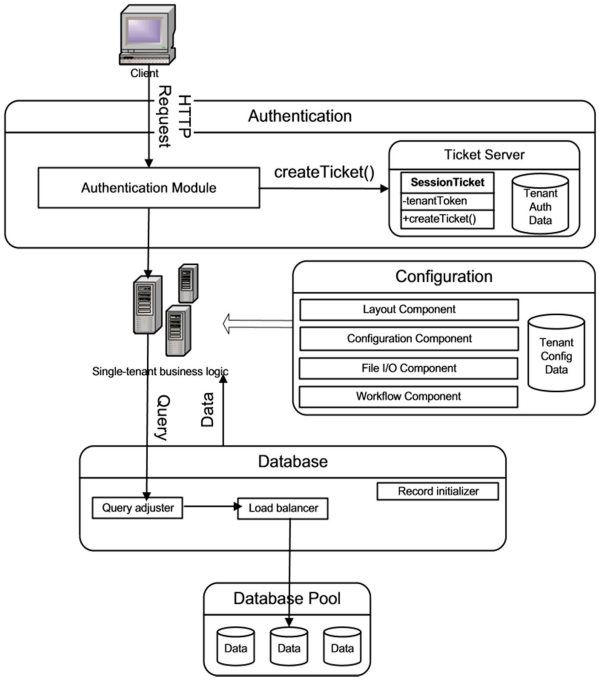

Figure 2. Architectural overview for multi-tenancy [10]

3.3.1 Authentication. In multi-tenancy, application and database instances are shared by numerous tenants, it is therefore imperative that separation of user data is implemented. User authentication and access control are used to achieve this, thus restricting tenants to only their own data or content.

3.3.2 Configuration. In order to give customers a feeling of a dedicated environment, multi-tenant applications require a lot of configuration. These configurations include the following: layout style, general configuration, file Input and output, and Workflow [10].

3.3.3 Database Data isolation is essential in multi-tenant environment as all tenants use the same application and database instance. Data isolation like authentication ensures that tenants only have access to their own data.

3.4 Security Challenges in Multi-tenancy.

Multi-tenancy in Cloud computing could lead to a situation whereby attackers and the victim share the same PM. This presents a new security challenge which traditional network focused techniques cannot mitigate. Figure 3 shows different cases of attacker and victim locations and the inter-connecting network.

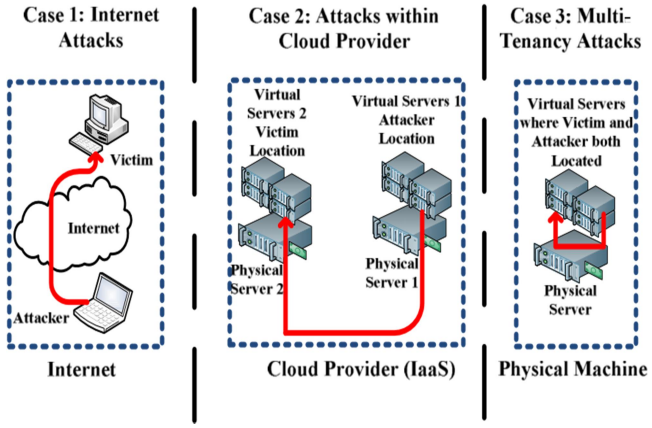

Figure 3: Differences between multi-tenancy and traditional attack cases [6].

Figure 3, case 2, shows a situation where both attacker and victim's VMs are hosted on different PMs but within the same CSP network. To mitigate this, virtual network security devices need to be put in place. Case 3 on the other hand, is one in which both the attacker and the victim share the same PM.

To illustrate how this can be achieved, an attacker can initiate a network probing for VMs in a PM. Once a target VM has been identified, a brute force attack can then be launched. This is one of the possible shortcomings of multi-tenancy and shows how by spending just a few dollars, an attacker can extract data from victims using a side-channel attack when hosted on the same PM. These attacks are often time not detectable by the PM's hypervisor or resident operating system.

4. INDUSTRY TREND IN PREVENTING MULTI-TENANT ATTACKS.

4.1 Multi-Tenancy in Private Clouds

Organizations are increasing their adoption of private Clouds. Building private Cloud implies offering a whole new model of service delivery from the data centre based on virtualization, orchestration, multi-tenancy and provisioning. IT planners are focusing on private IAAS or the delivery of dynamic compute, storage and networking to application development teams. For private Cloud IAAS to be beneficial, it must offer elastic, self-provisioned, on-demand, multi-tenant pools of shared compute, storage and networking [1]. A problem in today's data centre is that some aspects of networking are not compatible with private Cloud computing concepts. Specifically, the majority of network architectures assume a fixed relationship between device identification and physical address. Previously, it was uncommon for virtualization servers and storage to move around the data centre. Presently, virtual machines move for a number of reasons, such as load balancing, power management, maintenance and disaster recovery.

A key characteristic of the Cloud is multi-tenancy. In a private data centre, this means creating distinct virtual networks over a physical network so that human resource and finance for example can share the same server, storage and network resources. It is important to have controls in place to isolate one tenant from another. This is essential for business, security and compliance reasons. In addition, most IT organizations depend on virtual local area network VLAN as the core control to provide logic and multi-tenancy in the private Cloud. VLANs (802.1q) are effective for multi-tenancy by isolating one department's virtual server from another. There are issues with VLAN, because the maximum number is 4,096 technically. This may not be an issue for small data centres, but a large data centre with thousands of servers will have limitation issues. Also, hypervisors have limitations on how many host may be supported in certain configurations, which increases the number of VLANs significantly. VLAN management can be complex for network and virtualization administration in a Cloud environment. The process of tying a virtual VLAN to a physical VLAN is a manual process and managing more than 100 VLAN is a significant challenge.

4.2 Multi – Tenancy Security

Security in a multi–tenant environment must be addressed at both the SAAS and IAAS layers because of the services provided. The steps for ensuring security in a multi – tenant environment include VM segmentation, database segmentation and VM introspection

4.2.1 *VM Segmentation* Virtualization is the platform that underpins IaaS offerings. Central to virtualization platform is a specialized and optimized operating system called the hypervisor. The hypervisor in part serves to map traffic from the virtual machines to the underlying VM host hardware so the traffic can make its way through the data centre and out to the Internet and vice versa. The majority of security concern in the virtualized infrastructure relates to the co-residency of machines owned by different customers. Different clients on the same infrastructure with sensitive data and potentially different access policies are placed together. This can lead to unauthorized connection monitoring, unmonitored application login attempt, malware propagation and other forms of attacks.

Presently in the world of physical machines, network segmentation is used to ensure sensitive back-end services are well protected from the potentially vulnerable front-end available to the public. This has led to the development of the DMZ and tiered approaches to network centre design. Segmentation in a virtual environment is equally important because back-end database or application servers are valuable. VM isolation and segmentation is a primary requirement for VMs containing compliance related data such as personally identifiable information (PII)

4.2.2 Data Segmentation. Data segmentation is required in SaaS where tenants share a database infrastructure because tenants are souring the same application. Data of multiple tenants may likely be stored in the same database and may even share the same tables. A system for authentication and authorization of the access request is usually implemented so only certain review or fields are modifiable based on security policies. Encryption is also a primary security control to protect

data at rest so that if the database is compromised or data is stolen, it would be difficult to decipher the underlying data.

4.2.2 VM Introspection. VM Introspection (VMI) enables information gathering about virtual machines, virtual networks, security and virtual environment setting without the use of agents. The ability of malware to hide or disable from security agents is a security problem that has plagued the security industry for decades. VMI offers an interesting approach to leverage the hypervisor for an uncompromised inspection of VMs. VMI is basically a hypervisor based service that examines the internal state of a running virtually machine. Recent technologies have been commercialized that leverage VMI to provide high levels of segmentation and isolation for guest VMs or Cloud-services tenants. VMI provides rich details about the application and services that are installed on the virtual machine as well as its configuration.

5 CONCLUSION

Cloud computing is rapidly expanding and offering valuable services to Cloud users. There as various Cloud services such as SaaS, PaaS, and IaaS. Services such as application, compute resources and storage are provided by the CSP at price to Cloud users. Cloud deployment types such as private, public, community and hybrid Clouds are available in Cloud computing. A major characteristic of Cloud computing is multi-tenancy which enables the use of a single resources by multiple user from different places. Multi–tenancy has unique architecture based on the data or multi – tenant application. Multi – tenancy enables optimum use of resources on the Cloud, but also has security challenges.

ACKNOWLEDGMENTS

We acknowledge the support and sponsorship provided by Covenant University through the Centre for Research, Innovation and Discovery (CUCRID).

REFERENCES

[1] Mell, P. and Grance, T. (2011) "The NIST Definition of Cloud Computing", NIST Special Publication 800-145 White Paper. Retrieved from acm.org/citation.cfm?id=2206223

[2] I AWS LLC, "Amazon Elastic Compute Cloud (EC2)" Retrieved from http://aws. amazon.com/ec2

[3] Balasubramanian R. and Aramudhan M. (2012). Security Issues: Public vs Private vs Hybrid Cloud Computing. International Journal of Computer Application. 55(13) 35-41. .

[4] Xu, Y., Musgrave, Z., Nobel, B., and Bailey, M. (2014). Workload-Aware Provisioning in Public Clouds. IEEE Internet Computing, 18(4), 15-21, IEEE.

[5] Leandro, M., Nascimento, T., Dos Santos, D., Westphall, C., Westphall, C. (2012), 'Multi-Tenancy Authorization System with Federated Identity for Cloud-Based Environments Using Shibboleth', ICN2012, The 11th International Conference on Networks.

[6] AlJahdali, H., Albatli, A., Garraghan, P., Townend, P., Lau, L., Xu, J. (2014), "Multi-Tenancy in Cloud Computing", proceedings of the 8th International Symposium on Service-Oriented System Engineering, http://dx.doi.org/10.1109/SOSE.2014. 50.

[7] Fiaidhi, J., Bojanova, I., Zhang, J., Zhang, L., Kingdee, (2012), "Enforcing Multitenancy for Cloud Computing Environments", IT Pro 2012, IEEE Computer Society 1520-9202/12.

[8] Alcaraz-Calero, J., Edwards, N., Kirschnick, J., Wilcock, L., Wray, M. (2010), "Towards a Multi-tenancy Authorization System for Cloud Services

[9] Tang, B., Sandhu, R. and Li, Q. (2014), "Multi-tenancy authorization models for collaborative Cloud services", Concurrency and Computation: Practice and Experience Concurrency Computat.: Pract. Exper. 2015; 27:2851–2868

[101 Bezemer, C., Zaidman, A. (2010), "Multi-Tenant SaaS Applications: Maintenance Dream or Nightmare?" Accessed on 24 May 17

[11] AlJahdali, H., Townend, P., and Xu, J. (2013), "Enhancing Multi-Tenancy Security in the Cloud IaaS Model over Public Deployment", 7th International Symposium on Service-Oriented System Engineering, IEEE. .

[12] Krebs, R., Momm, C. and Kounev, S. (2011), "Architectural Concerns in Multi-Tenant SaaS Applications", Accessed on 24 May 17.

[13] Zeng, J., and Plale, B. (2016), "Argus: A Multi-tenancy NoSQL store with workload-aware resource reservation", Parallel Computing 58 (2016) 76 – 89.

[14] Ngo, C., Demchenko, Y., and Laat, C. (2016), "Multi-tenant attribute-based access control for Cloud infrastructure services", Journal of Information Security and Applications 27-28 (2016) 65–84.

[15] Espadas, J., Molina, A., Jiménez, G., Molina, M., Ramírez, R., Concha, D. (2013), "A tenant-based resource allocation model for scaling Software-as-a-Service applications over Cloud computing infrastructures", Future Generation Computer Systems 29 (2013) 273–28.

[16] Povedano-Molina, P., Lopez-Vega, J., Lopez-Soler, J., Corradi, A. Foschini, L. (2013), "DARGOS: A highly adaptable and scalable monitoring architecture for multi-tenant Clouds", Future Generation Computer Systems 29 (2013) 2041–205

[17] E-Guide (2015), 'Cloud Networking – The Next Frontier', TechTarget Networking Media Publication

[18] Robbie Higgins (2016) 'Securing a multi-tenant environment', Computer World Publication Mell, P. and Grance, T. (2011) "The NIST Definition of Cloud Computing", NIST Special Publication 800-145 White Paper. Retrieved from acm.org/citation.cfm?id=2206223.