Kaspersky - IT Security for the Next Generation - European Cup 2011

Title: Towards new access control models for Cloud computing systems
Category: 'In the Cloud' - Security
Author name: Gouglidis Antonios
E-Mail: agougl@uom.gr
Postal address: University of Macedonia, Economic and Social Sciences, 156 Egnatia Street, GR-540 06 Thessaloniki, Greece
City, Country: Thessaloniki, Greece
Name of the University: University of Macedonia, Department of Applied Informatics
Professor: Mavridis Ioannis, Assistant Professor, Ph.D., Dipl. Eng. in Computer Engineering and Informatics, mavridis@uom.gr

Abstract

Cloud computing is a composition of existing technologies viz. virtualization technology, disk storage, processors and so on, which gained considerable attention mostly from the enterprise. Cloud security is an active research area, due to the newly introduced SPI service model and the different deployment models that require the revision of several security concepts. Specifically, in this paper we give a brief presentation of Cloud computing and the terminology of access control concepts used in the Cloud. Additionally, we elaborate on the identification of access control's distinctive characteristics in the aforementioned systems. We use a conceptual categorization, which is a systems engineering methodology, in order to identify a series of characteristics for access control in the Cloud computing paradigm. Furthermore, we present a comparative review of two prominent access control models for the Cloud, namely the Role-based Access Control model (RBAC) and the Usage Control model ($UCON_{ABC}$). We anticipate this initiative to help for the definition of concrete access control requirements and the design and implementation of new access control models, in order to accelerate the adoption of Cloud technologies.

1. Introduction

The Cloud is a fairly new and emergent technology, which definition is a topic for discussion in several research papers (Foster et al., 2008). Nevertheless, Cloud computing is defined in (Mather et al., 2009) using five attributes viz. multitenancy, massive scalability, elasticity, pay as you go and self-provisioning of resources. These attributes successfully imprint the distinctive characteristics of the Cloud and differentiate it from similar technologies, as the Grid computing paradigm. As stated in (Gollmann, 2010), when new or old techniques are put to new in use, as in the Cloud, new security challenges arise. For that reason a series of security concepts require to be revised, so that to cope with the requirements of the above mentioned systems.

Privacy, trust and access control are some of the security concepts met in Cloud systems. In this paper, we will further examine the latter of the aforementioned. Access control is of vital importance in a Cloud environment since it is concerned with allowing a user to access a number of Cloud resources. An extensive research has been done in the area of access control in collaborative systems (Tolone et al., 2005, Zhang et al., 2008). Nonetheless, further examination is demanded, due to the partial or weak fulfilment of security requirements in the Cloud.

The aim of this document is to provide the reader with basic information about the models and the terminology used in Cloud computing systems and basic access control concepts. Furthermore, a series of thoughts are presented regarding the characteristics of Cloud computing systems. The methodology that is followed for the identification of the characteristics is based on the conceptual categorization presented in (Gouglidis and Mavridis, 2010). The value of this paper is to serve as an

initiative for further investigation of access control requirements in the area of Cloud computing, in order to assess the applicability of access control solutions in the Cloud infrastructure.

The structure of the remainder of this chapter is as follows. The next section provides a prerequisite terminology used in Cloud computing systems and in access control. In turn, we imprint the basic characteristics of Cloud computing systems. A comparative review is given of the Role-based Access Control model (RBAC) and the Usage Control model ($UCON_{ABC}$). Finally, we present our concluding remarks along with some future thoughts.

2. Background

This section provides fundamental information about the Cloud computing paradigm and access control. The provided information is mostly extracted from (Mather et al., 2009, Ravi S. Sandhu, 1994, Gollmann, 2010, Foster et al., 2008).

2.1. Cloud computing in a nutshell

As already mentioned, the definition of the Cloud is based on five attributes, which are able to capture its specificities. Multitenancy refers to the business model that is implemented by the Cloud, where a single shared resource can be used from multiple users. Massive scalability refers to the potential of the system to scale (i.e. increase) in resources. The on-demand and rapid increment or decrement of computing resources is translated as elasticity of the Cloud. Thus, when more storage space or bandwidth is required, this can be allocated, and vice versa. Pay as you go is the process of paying for the resources that are used. Last but not least, the users are provided with the ability to self-provision resources, namely storage space, processing power, network resources and so on.

The service model of Cloud computing is based on the SPI framework. SPI is an acronym that stands for Software-as-a-service (SaaS), Platform-as-a-service (PaaS) and Infrastructure-as-a-service (IaaS). Specifically the SaaS delivers to consumers the software that is used under a business model, namely the usage-based pricing. The PaaS offers the platform for the development of the applications, and lastly, the IaaS handles the provision of the required hardware, software and equipment, in order to deliver a resource usage-based pricing model.

Furthermore, the aforementioned service models are provided under three deployment models viz. public, private and hybrid Cloud. The public Cloud provision resources over the Internet and are accessible via a web application. A third-party operates as the host and performs all the required operations (i.e. management, security). The private Cloud provides the same functionality as the public deployment model, but in an internal and private network. This model requires the acquisition of the appropriate hardware and software. The hybrid model refers to the combination of the public and private deployment models. Usually, the latter model is used to keep sensitive data in the private network and deploy non-core applications in the public.

2.2. Access control in a nutshell

Access control's role is to control and limit the actions or operations in the Cloud systems that are performed by a user on a set of resources. In brief, it enforces the access control policy of the system, and at the same time it prevents the access policy from subversion. Access control in the literature is also referred to as access authorization or simply authorization.

A Cloud access control policy can be defined as a Cloud security requirement that specifies how a user may access a specific resource and when. Such a policy can be enforced in a Cloud system through an access control mechanism. The latter is responsible for granting or denying a user access upon a resource. Finally, an access control model can be defined as an abstract container of a collection of access control mechanism implementations, which is capable of preserving

support for the reasoning of the system policies through a conceptual framework. The access control model bridges the existing abstraction gap between the mechanism and the policy in a system.

Furthermore, in this section, a brief presentation of the reference monitor concept is given. This is mainly done because the application of the reference monitor concept is known to achieve high assurance access control mechanisms. Additionally, it provides guidelines for the design and implementation of secure computer systems (Ferraiolo et al., 2003).

The process of access control in any computer system guarantees that any access to the resources of the system conforms to its access control policy. The application of the abstract concept of the reference monitor is capable of providing the requirements that are posed from the access control process. The reference monitor operates as an access mediator between the subject's access requests and the system's objects. The accesses comply with the system's security policy. The reference monitor can be informed for the security policy of the computer system from an access control database. Moreover, all the security relevant transactions are kept into an audit file for security and traceability reasons.

The architecture of the reference monitor is the result of the application of three key implementation principles. These principles are the completeness, isolation and verifiability. Completeness requires from the reference monitor to invoke all the subject's references to an object and also to constitute it impossible to bypass it. The isolation principle assures that the reference monitor must be tamper-proof. This means that it must be impossible for an attacker to penetrate the reference monitor in a malicious way. Lastly, the verifiability principle appertains to the checking and validation of the system's security design through the use of software and system engineering techniques.

Nonetheless, the aforementioned reference monitor principles seem to be insufficient, especially in enterprise environments. This is mostly because the main objective of the reference monitor is the enforcement of each system's policy. Yet, it does not interfere with the articulation of a system's security policies. Thus, the principles of flexibility, manageability and scalability are introduced. The first principle assures that the access control policy of an enterprise can be enforced by the existing security system. The next refers to the ease of policy management and the latter requires from the security system to cope with the fluctuations in the number of the participating users and resources in a computer system.

The concept of reference monitor in open systems has been standardized with the X.812 access control framework (ITU-T, 1995). In brief, the main functions in X.812 are the Access Control Decision Function (ADF) and the Access Control Enforcement Function (AEF). The former component is responsible for the making of access control decisions. The decisions are made based on information applied by the access control policy rules, the context in which the access request is made, and the Access Control Decision Information (ADI). ADI is a portion in the Access Control Information (ACI) function, which includes any information used for access control purposes, including contextual information. Lastly, the AEF is responsible for the enforcement of the decision taken from the ADF.

Regarding access control policies, several of them have been introduced during the last decades, namely the Mandatory Access Control policies (MAC), the Discretionary Access Control policies (DAC) and the Role Based Access Control policies (RBAC). Each one of them serves specific security requirements in different working environments. As mentioned in the definition of the access control policy, a number of access control models are required and were developed in order for the policies to be represented by formal methods. Research on the MAC, DAC and RBAC has proven that an access control model, which can express the role based access control policies is also capable of enforcing both MAC and DAC policies (Ferraiolo et al., 2003). It is noteworthy that an attempt started along with the advancement of RBAC for the design of a series of Attribute Based Access Control models (ABAC). The ABAC model was mainly introduced to overcome a

number of RBAC's shortcomings and has also been proven capable of enforcing MAC, DAC and RBAC policies.

3. In the Cloud - Access control characteristics

The identification and definition of Cloud access control characteristics and requirements, namely the access control policy, greatly amplifies the design of a model and the implementation of a mechanism regarding access control. In order to appoint a series of characteristics regarding access control we use the conceptual categorization for Cloud systems proposed in (Gouglidis and Mavridis, 2010). Figure 1 depicts the four layers of the conceptual categorization. The entropy layer identifies requirements from the dispersion of the objects in a system and the assets layer from the type of shared objects within the boundaries of the entropy layer. The management layer defines requirements from policy management and the logic layer incorporates requirements that are not handled by the former layers. A set of core requirements for access control systems that are considered important for the Cloud environment, follows. The identification of the requirements incorporates also characteristics that are exposed by the three levels of the information security infrastructure in the Cloud viz. application level, host level, and network level, where applicable. These characteristics may vary depending on the use cases that need to be supported by a specific system.

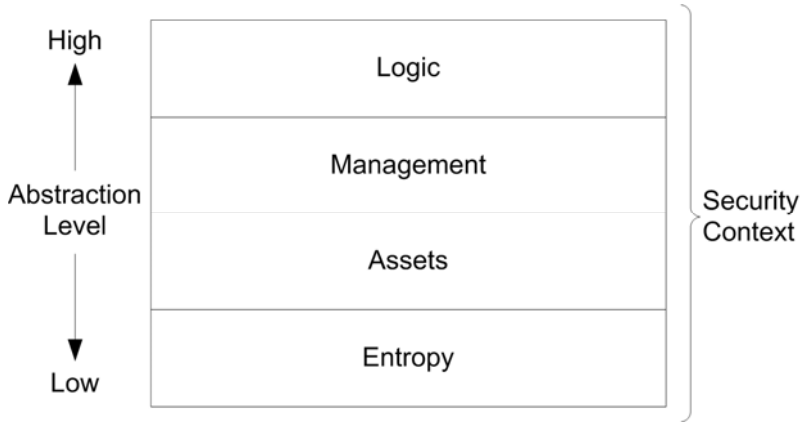

Figure 1. Conceptual categorization layer.

Entropy layer: Applications are provided to consumers as a set of services via the SaaS service model. Each application is in most cases accessible through a web interface. Usually the services are deployed under the same organization and thus under the same domain. However, the use of the public or hybrid deployment models requires the collaboration of services among the participating organizations. Therefore, the application's entropy level can be relative high, depending on the used deployment model. Additionally, the hosts that are used to provide the assets of the Cloud can also be characterized by their high dispersion when the public or hybrid model is deployed and low when the private deployment model is applied.

Assets layer: The assets in Cloud computing systems are of two type viz. software and hardware. The software is exposed as a set of services that can be realized by technologies such as the web services. Collaboration among services is applicable. Hardware resources can be CPU, storage space, network bandwidth and so on. Specifically, we recognize that the fine-grained sharing of any resource in a Cloud system includes a resource requestor and a provider. When a user requests access to an asset, access must be granted only if the requestor is a legitimate user and also authorized to access the specified asset. Furthermore, as described in the definition of the Cloud, multitenancy must be supported. Thus, when multiple consumers from different companies are using a Cloud service model, consumers and their data must be guaranteed that are protected from each other throughout the collaboration.

Management layer: The management of policies in a Cloud computing system is required to be centralized in most cases. However, if collaboration is required as in the public and hybrid Cloud deployment models, the management of policies requires to be distributed and applied among participating organizations. Moreover, each administrative user of an organization should administer the local policies of the organization. Additionally, administrators should run the policies in the collaboration that refer to resources of the administrator's organization. Furthermore, it must be guaranteed that no conflicts should exist among the policies of the individual organizations at the higher corporate level. Last but not least, the process of identifying policy violations should be automated, in all deployment models.

Logic layer: The main characteristic of the Cloud is the support of the business model that allows the provision of usage based pricing. Thus, quality of service policies (QoS) along with service level agreements (SLAs) must be supported, in order to provide to the consumers the agreed levels of quality. Resource providers should be able to define quality factors on their shareable resources. The quality factors concern the level of resource usage and can also be characterized as obligations that must be met from a provider when granting access to a resource requestor. For instance, quality factors could apply for setting disk quotas, memory or CPU utilization levels and so on and so forth. Furthermore, we identify the enforcement of the autonomy and security principle (Shafiq et al., 2005). The autonomy principle refers to the permission of an access under secure interoperation, if it is also permitted within the individual domain. The security principle pertains to the denial of an access under secure interoperation, if it is also denied within the individual domain. Furthermore, the principle of containment (Ravi Sandhu, 2008) that subsumes the principles of the separation of duties, least privilege and so forth, should be supported in each and among domains. The latter requirement greatly enhances the adoption of Cloud technologies in business organizations, where the existence of conflict of interest policies is presumed.

4. Discussion

In this section, the RBAC and the $UCON_{ABC}$ access control models are compared, as two of the most prominent access control models for the Cloud. The comparison is attempted with respect to the conceptual categorization for Cloud systems, with a view to specify a number of deficiencies in the aforementioned access control models. The criteria used throughout the comparison are based on the characteristics that were discussed and the evaluation is based on the level of fulfillment of the requirements by the access control models.

Table 1 illustrates the evaluation of the RBAC and $UCON_{ABC}$ models with respect to the entropy, assets, management and logic layers of the conceptual categorization.

Concerning the entropy layer, the characteristics that were defined, demand both the support of both a centralized and a distributed access control among different domains and the dynamic joining of new ones. The proposed standard RBAC model, as already seen, handles better centralized architectures and is rather weak in inter-domain collaborations. Such functionality is absent from the standard model. However, research in (Shafiq et al., 2005) has proven that RBAC can also be applied in multi-domain environments where distributed multiple organizations inter-operate. Yet, RBAC requires that all user domains must be known a priori, in order to access an object. On the contrary, the $UCON_{ABC}$ model, due to its support of attributes, can cope better with highly distributed environments. Furthermore, one of UCON's features is that it is possible to provide access to users in a collaborative environment without the need for them to be known by the resource a priori.

Access control models	**Conceptual categorization layers**			
	Entropy	**Assets**	**Management**	**Logic**
	Low /	Low /	Medium /	Medium

RBAC	Medium	Medium	High	
$UCON_{ABC}$	High	Medium	Low	Medium

Table 1. Comparisons between the different access control models.

In regard to the layer of assets, we mentioned that fine-grained access to resources should be supported. Additionally it should support obligations from the side of the resource provider. RBAC usually provides more course-grained access control to resources in contrast to $UCON_{ABC}$. Research has also been done in RBAC to extend it and to support finer-grained access control through the use of context (Tolone et al., 2005). Obligations are supported in $UCON_{ABC}$, but not in the notion demanded by the requirements. The notion of obligations is completely absent in RBAC.

RBAC supports improved administrative capabilities on the level of a domain in comparison to $UCON_{ABC}$. In more detail, RBAC can also provide management in a role-based fashion (Ferraiolo et al., 2003). However, a number of issues arise when it comes to inter-domain management of policies, and solutions are provided in existing literature (Shafiq et al., 2005). In contrast to RBAC, $UCON_{ABC}$ lacks administration.

Finally, the fulfillment of characteristics in the logic layer is fairly the same in both access control models. Features as the support of QoS or SLA rules is absent from both models. Nonetheless, RBAC supports the principles of separation of duties and least privilege better.

5. Conclusion

This paper introduced and explained the definition of Cloud computing environments, including associated concepts and characteristics. Access control models and authorization systems in the Cloud context are of vital importance due to their layered nature. Based on the results stemmed from our observations, we believe that the design and implementation of proper access control models for the Cloud computing paradigm is needed. Current access control models are not specifically designed to tackle the requirements of Cloud systems. By applying the conceptual categorization for the Cloud systems, we illustrated how to identify a list of basic access control's characteristics. In result, we expect the applied methodology to initiate further research for the definition of access control requirements in Cloud computing systems and moreover, to result in new access control models.

6. References

FERRAIOLO, D. F., KUHN, D. R. & CHANDRAMOULI, R. 2003. *Role-Based Access Control*, Artech House, Inc.

FOSTER, I., YONG, Z., RAICU, I. & LU, S. Year. Cloud Computing and Grid Computing 360-Degree Compared. *In:* Grid Computing Environments Workshop, 2008. GCE '08, 12-16 Nov. 2008 2008. 1-10.

GOLLMANN, D. 2010. Computer security. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2**,** 544-554.

GOUGLIDIS, A. & MAVRIDIS, I. 2010. On the Definition of Access Control Requirements for Grid and Cloud Computing Systems. *In:* DOULAMIS, A., MAMBRETTI, J., TOMKOS, I. & VARVARIGOU, T. (eds.) *Networks for Grid Applications.* Springer Berlin Heidelberg.

ITU-T 1995. X.812 Recommendation. *Data Networks and Open System Communications Security - Information Technology – Open Systems Interconnection – Security Frameworks for Open Systems: Access Control Framework.* ITU.

MATHER, T., KUMARASWAMY, S. & LATIF, S. 2009. *Cloud Security and Privacy: An Enterprise Perspective on Risks and Compliance*, Oreilly & Associates Inc.

RAVI S. SANDHU, P. S. 1994. Access Control: Principles and Practice. *IEEE Communications Magazine,* 32**,** 40-49.

RAVI SANDHU, V. B. Year. The ASCAA Principles for Next-Generation Role-Based Access Control. *In*, 2008. xxvii-xxxii.

SHAFIQ, B., JOSHI, J. B. D., BERTINO, E. & GHAFOOR, A. 2005. Secure Interoperation in a Multidomain Environment Employing RBAC Policies. *IEEE Trans. on Knowl. and Data Eng.,* 17**,** 1557-1577.

TOLONE, W., AHN, G.-J., PAI, T. & HONG, S.-P. 2005. Access control in collaborative systems. *ACM Comput. Surv.,* 37**,** 29--41.

ZHANG, X., NAKAE, M., COVINGTON, M. J. & SANDHU, R. 2008. Toward a Usage-Based Security Framework for Collaborative Computing Systems. *ACM Trans. Inf. Syst. Secur.,* 11**,** 1--36.