SERVICE MODEL DRIVEN VARIATIONS IN SECURITY MEASURES FOR CLOUD ENVIRONMENTS

Ramaswamy Chandramouli
National Institute of Standards & Technology, Gaithersburg, MD, USA

ABSTRACT

With the increasing adoption of cloud computing service models – Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS), proper implementation of adequate and appropriate security protection measures has become a primary area of concern. In an enterprise computing architectural stack, the components in all layers are owned or controlled by a single entity – the enterprise. However, in cloud service environments, the control of the various layers is split between the cloud provider and the cloud subscriber based on the cloud service model. In this paper, we analyze the impact of this difference in control on the set of actual security protection measures that can be implemented in the various layers for different cloud service models. We also point out the value of access to lower layers for implementing protection measures for components at a higher layer.

KEYWORDS

Cloud Service Models, Security Protection, Cloud Provider, Cloud Subscriber, Virtual Machine, Hypervisor

1. INTRODUCTION

There are now three generally accepted cloud service models [1] – Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS) and two major players – the cloud provider and the cloud subscriber. The set of layers over which each of these players have control varies with the cloud service model or environment. The various layers in a cloud service architectural stack based on slight variations from the model given by cloud security alliance [2] is given in Figure 1 below. While there is general consensus on the threats and security goals for each layer [3], what is not clear is the impact of the difference in control over the various layers in different cloud service models (or environments) on the set of security protection measures that can be implemented. Some researchers [4] feel that this difference in control has negligible impact on actual security measures. However we beg to differ and in this paper, we go on to describe in detail the variations in security protection measures for each layer in the three cloud service models – SaaS, PaaS and IaaS.

Going up from the bottom, in all three cloud service models (SaaS, PaaS and IaaS), the facility, networking infrastructure, server (host) hardware and the resource abstraction layer (consisting of a hypervisor since virtualization is the most common abstraction mechanism used) are entirely under the control of the cloud service provider. Out of these 4 layers, since the facility, networking infrastructure and server hardware layers are common to all IT infrastructures, and hence the security protection measures for these three layers are beyond the scope of this paper. Going up one more layer in the stack, we find that although the resource abstraction layer (virtualization layer) is not common to all IT infrastructures, it is invariably deployed in all three cloud service models and is entirely under the control of the cloud provider. Hence we exclude this layer as well from our discussion. The organization of the rest of the paper is as follows. In section 2, we discuss the variation in security protection measures at the VM layer followed by variations in protection measures at the middleware layer in section 3. A detailed discussion of the various aspects of security protection for the application layer is the topic for section 4. Section 5 provides the conclusions.

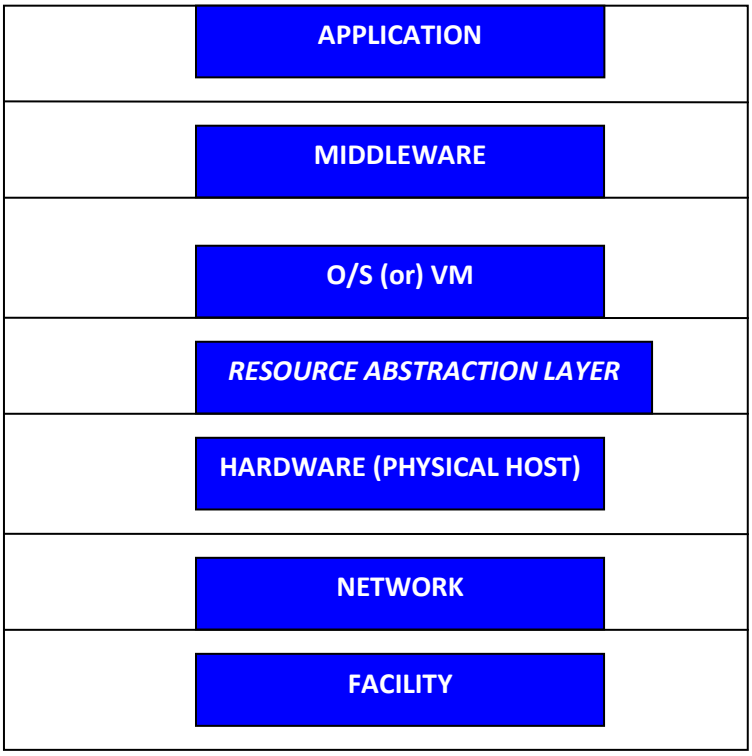

Figure 1. Cloud Service Layers

2. VARIATIONS IN VIRTUAL MACHINES (VMs) PROTECTION MEASURES

In an IaaS cloud service model, the VM is the unit of rental for the cloud subscriber and hence the subscriber has full control over it. In PaaS or SaaS, the VM itself is an optional artifact (although in this paper we assume it to be an essential one) depending upon whether the cloud provider uses a virtualized infrastructure or not. The security focus areas for a VM consists of: (a) VM-SF1: The Images from which running VM instances are launched should be securely configured, (b) VM-SF2: Malicious and unwanted traffic flowing in and out of the VM must be filtered out based on source/destination IP Address and Port as well as the network protocol type and (c) VM-SF3: The VMs themselves must be protected against viruses, malware and other threats [5]. For the third focus area, there are two options [5]: a centralized solution or a decentralized solution. The variations in protection measures for these focus areas based on the cloud service environment are discussed below.

2.1 VM Protection Measures in SaaS, PaaS and IaaS Cloud Services

Taking up the case of VM-SF1, we find that in SaaS and PaaS cloud environments, VM images are created by the cloud provider and are usually configurations made up of hardened pared-down version of O/S with minimal features (often called golden images) along with ensuring that other software such as web servers that are part of that image are also free of vulnerabilities. In the case of IaaS, it is the cloud subscribers who can either create their own images or choose to use any of the template images provided by the cloud provider. If they choose to use any of the external VM images (cloud provider or any other source) they have to ensure: (a) the VM Image is certified to be free of malware and contains the secure version of the O/S with latest patches and (b) the VM Images offered are digitally signed so as to provide assurance regarding their origin and authenticity. As regards VM-SF2, there are no variations in security measures due to cloud service models and the usual implementation in all three cloud services is the use of software-based firewalls. As regards VM-SF3, two types of solutions are possible - a virus/malware protection solution based in each individual VM as well as a centralized solution based in a separate dedicated VM. In SaaS and PaaS models,

a centralized solution using a security appliance hosted on a hardened dedicated VM is used for performing virus detection and remediation on all VMs hosted on a hypervisor or virtualized server. In an IaaS cloud scenario, the cloud subscriber has only control over the VMs he/she has subscribed and hence protection measures for virus/malware has to be implemented in a decentralized way at each of the VMs (with or without an agent-based technology) he/she has rented.

3. VARIATIONS IN MIDDLEWARE PROTECTION MEASURES

In a SaaS environment, the middleware is developed (or procured), deployed and used entirely by the cloud provider. In the case of an IaaS environment, the scenario is the same except that the entity involved here is a cloud subscriber instead of a cloud provider. The protection measures for middleware in these two environments consist, therefore, of following essentially the best practices needed for any software development lifecycle. However, the middleware in the case of a PaaS environment, though developed and deployed by the cloud provider, is offered as a subscriber-usable application. Hence the following additional security protection measures are needed: (a) All the Middleware elements should be certified by an independent third party to be malware free, (b) The Middleware should be designed such that it will only accept communication through an encrypted channel [6], (c) The architecture of the Middleware layer should be such that the chances of security exploits due to misconfiguration of the middleware components are minimal and (d) Middleware that are provided exclusively for security management (as opposed to ones that are provided for supporting application integrity/performance - e.g. middleware that provides functions such as identity verification, authorization/access control, input data validation, event logging etc) must have maximum trust by carrying independent accredited third party validations/certifications.

4. VARIATIONS IN APPLICATION LAYER PROTECTION MEASURES

Secure development and deployment practices are required universally for all applications – whether provided as part of a cloud service or hosted for a client or for an internal enterprise use. Based on this observation, certain security focus areas within the overall application lifecycle are: They are: (a) Secure application development environment (b) Robust Authentication and Authorization of application users, (c) Distribution of applications based on sensitivity levels and (d) Protection of application data. The variations in protection measures associated with each of the security focus areas among the three cloud service environments are discussed below:

4.1 Secure Application Development Environment

In SaaS and IaaS the tools needed for application development are owned and used by the same entity. Hence the usual protection measures for secure deployment of these tools apply. In the case of PaaS, the tools are owned by cloud subscriber but are used by cloud subscriber. Hence special protection measures are needed. They are: (a) Ensure that the development tools and library codes (compile time and run time) provided by PaaS cloud providers have been certified to be free of malware and (b) Ensure that the PaaS provider provides digitally signed versions of these tools and library code and that all cloud subscriber developers start using these tools and libraries after verification of the associated digital signatures.

4.2 Authentication and Authorization of Users

Robust authentication and granular authorization are key first line defenses in application protection. The key requirement here is that the cloud subscriber would like to enforce the same authentication policies and authorization policies that they enforce in their internal data centers even if their application is hosted in the cloud. In an IaaS environment, the applications are owned, controlled and hosted on the VMs controlled by the cloud subscriber who then has the choice to implement any suitable well established authentication and authorization mechanism of choice. In the case of Paas, the applications (which in this case are development tools and libraries) are owned, controlled and hosted on resources controlled by cloud provider and hence the same situation as we saw in the case of IaaS cloud subscriber applies. When it comes to SaaS environment,

since the volume of the user base is likely to be large, SaaS providers often would like to provide support for a standardized claims/assertions (e.g., SAML) based authentication and an associated identity federation protocol (e.g., SAML or WS-Federation) instead of managing its own identity verification services.

4.3 Distribution of Applications

In a SaaS environment, where applications and the resources to host them are fully under the control of cloud provider, the applications can be distributed in such a way that applications at the same security level are hosted on a particular virtualized host [7]. However in an IaaS cloud offering that offers a multi-tenant environment, applications of different sensitivity levels will be running on a given virtualized host depending upon work load efficiencies algorithms used by the cloud provider. Hence cloud subscribers, in order to protect their applications, have to deploy VM-level protections described in section 2 of this paper. In the case of PaaS environment, the software pack (i.e., development toolkits & libraries) does not have pronounced sensitivity levels and hence this type of protection measure need not be considered.

4.4 Protection of Application Data

In SaaS and IaaS environments, data belonging to multiple cloud subscribers are going to be sharing the same storage resources and hence each subscriber needs to protect its own data (at rest) through some form of encryption. Further, in the case of an IaaS environment, a subscriber may be hosting an application with multiple tiers with network traffic between them sharing the cloud provider's internal network along with other subscribers [8]. In this instance, the cloud subscriber has to protect its data in transit through some form of encryption especially if the data is of a sensitive nature, the eavesdropping of which may have security implications.

5. CONCLUSIONS

In this paper, we described the variations in security protection at each of the top 3 layers of cloud services – Application, Middleware and VM layer, due to differences in the entity that controls each layer. We also saw that for any layer, a more efficient security protection measure can be provided if the same entity controlled the layer below it. Since the network, hardware and resource abstraction layers in all cloud service models are controlled by cloud provider, it has available to it more efficient protection measures. This knowledge can be leveraged by a cloud subscriber to ensure that a given cloud provider has these efficient protection measures deployed in its service and also explore the possibility of a cloud subscriber providing selective access to those lower level APIs (e.g., introspection API of the hypervisor) to enable the cloud subscriber to customize/enhance some security protection measures to meets its special business and regulatory compliance needs.

REFERENCES

[1] Mell, P., and Grance, T.,, 2011, *A NIST Definition of Cloud Computing, NIST SP 800-145*, http://csrc.nist.gov/publications/PubsDrafts.html#SP-800-145

[2] Cloud Security Alliance, 2009, *Security Guidance for Critical Areas of Focus in Cloud Computing, v2.1*, www.cloudsecurityalliance.org/csaguide.pdf

[3] Cloud Security Alliance, 2010, *Top Threats to Cloud Computing V 1.0*, http://www.cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf

[4] Grobauer B., et al, 2011, Understanding Cloud Computing Vulnerabilities, *IEEE Security and Privacy*, Vol 9, No 2, pp 50-57.

[5] Spring J., 2011, Monitoring Cloud Computing by Layer, Part 1, *IEEE Security and Privacy*, Vol 9, No 2, pp 66-68.

[6] Ranabahu A., and Maxmilien E.M., 2009, A Best Practice Model for Cloud Middleware Systems, *Proceedings of the Best Practices in Cloud Computing*: Designing for the Cloud Workshop, ACM Press, pp 41-51

[7] Spring J., 2011, Monitoring Cloud Computing by Layer, Part 2, *IEEE Security and Privacy*, Vol 9, No 3, pp 52-55

[8] Kaufman, L.M., Data Security in the world of cloud computing, *IEEE Security and Privacy*, Vol. 7, No. 4, 2009, pp. 61-64