Contents lists available at SciVerse ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

A novel intrusion severity analysis approach for Clouds

Junaid Arshad *, Paul Townend, Jie Xu

School of Computing, University of Leeds, Leeds, LS2 9JT, UK

ARTICLE INFO

Article history:
Received 8 November 2010
Received in revised form
15 March 2011
Accepted 5 August 2011
Available online 16 August 2011

Keywords:
Cloud computing
Cloud security
Intrusion severity analysis
Machine learning

ABSTRACT

Cloud computing presents exciting opportunities to foster research for scientific communities; virtual machine technology has a profound role in this. Among other benefits, virtual machine technology enables Clouds to offer large scale and flexible computing infrastructures that are available on demand to address the diverse requirements of scientific research. However, Clouds introduce novel security challenges which need to be addressed to facilitate widespread adoption. This paper is focused on one such challenge—intrusion severity analysis. In particular, we highlight the significance of intrusion severity analysis for the overall security of Clouds. Additionally, we present a novel method to address this challenge in accordance with the specific requirements of Clouds for intrusion severity analysis. We also present rigorous evaluation to assess the effectiveness and feasibility of the proposed method to address this challenge for Clouds.

1. Introduction

The advent of internet technologies such as, Service Oriented Architectures (SOA), has significantly influenced the methods used in e-Science along with the emergence of new computing paradigms to facilitate e-Science research. Cloud computing is one such emerging paradigm which makes use of contemporary virtual machine technology. This collaboration between internet and virtual machine technologies enables Cloud computing to emerge as a paradigm with promising prospects to facilitate the development of large scale, flexible computing infrastructures that are available on demand to meet the computational requirements of e-Science applications. Related to this, Cloud computing has witnessed widespread acceptance mainly due to compelling characteristics such as live migration, isolation, customization and portability, thereby increasing the value attached with such infrastructures. The virtual machine technology has had profound role in it. Amazon [1], Google [2] and GoGrid [3] represent some of the commercial Cloud computing initiatives whereas Nimbus [4] and OpenNebula [5] represent academic efforts to establish a Cloud.

Cloud computing has been defined in different ways by different sources; however, for the purpose of the research described in this paper, Clouds have been defined as follows: *a high performance computing infrastructure based on system virtual machines to provide on-demand resource provision according to the service level agreements established between a consumer and a resource provider.*

A Cloud computing system representing the above definition has been presented in Fig. 1. A system virtual machine, as described in this definition, serves as the fundamental unit for the realization of a Cloud infrastructure and emulates a complete and independent operating environment. Within the scope of this paper, the Cloud platforms focused on satisfying computation requirements of compute intensive workloads have been defined as *Compute Clouds* whereas those facilitating large scale data storage as *Storage or Data Clouds*. For the rest of this paper, the terms *Cloud computing* and *Clouds* have been used interchangeably to refer to our definition of compute Clouds.

As with any other emerging paradigm, different models of Cloud computing have been proposed to harvest its benefits. These are *Infrastructure as a Service (IaaS)*, *Software as a Service (SaaS)* and *Platform as a Service (PaaS)* [6]. With regard to these models, the Cloud computing system defined earlier and illustrated in Fig. 1 resembles IaaS and therefore, inherits the characteristics of this model of Clouds. From the definition of Cloud computing presented earlier, the term Cloud computing has been used to refer to IaaS model of Cloud computing for the rest of this paper.

However, security underpins the extensive adoption of Cloud computing. Related to this, virtualized computing systems such as Clouds, introduce novel challenges for the development of enabling mechanisms, in general and security in particular which require dedicated efforts for their solution [7]. The emphasis of this research is to investigate security issues due to the ability of virtualization to host multiple different computing environments on a single physical resource. Therefore, an intrusion detection and severity analysis system, residing in the most privileged domain, is required to monitor multiple virtual machines with possibly diverse security requirements. This requires identification of the

* Corresponding author. Tel.: +44 113 343 1707; fax: +44 113 343 5468.
E-mail address: sc06ja@leeds.ac.uk (J. Arshad).

0167-739X/$ – see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.future.2011.08.009

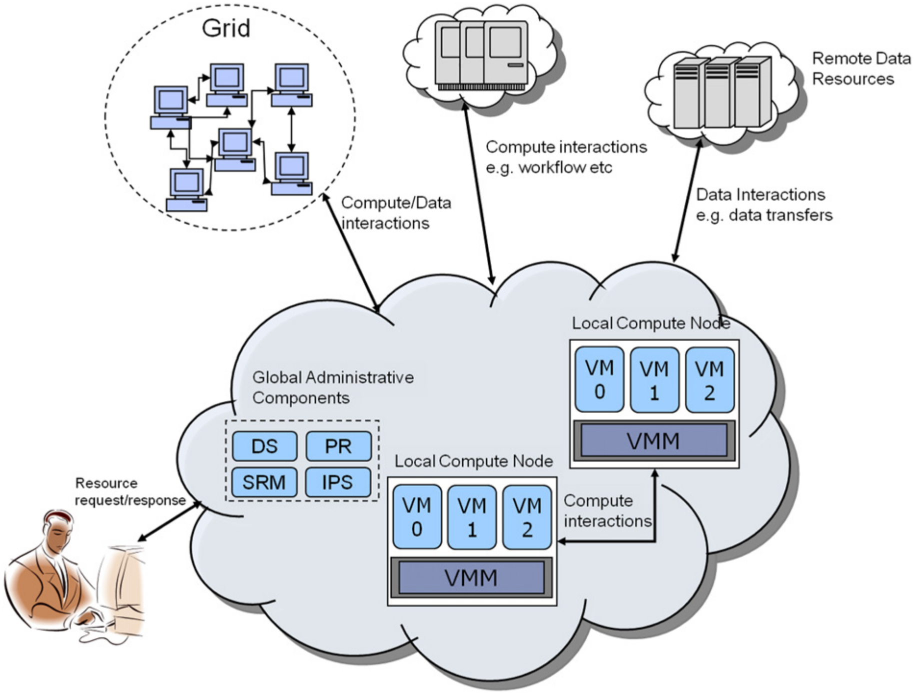
Grid
Remote Data Resources
Compute interactions e.g. workflow etc
Data Interactions e.g. data transfers
Compute/Data interactions
Local Compute Node
VM 0
VM 1
VM 2
VMM
Global Administrative Components
DS
PR
SRM
IPS
Local Compute Node
Compute interactions
Resource request/response

Fig. 1. A Cloud computing system.

fact that a particular malicious attempt can have a different degree of impact on different applications. Within the context of this research, this has been defined as the *Level of Severity (LoS)* of a security breach for a particular application. There can be different implications of this concept such as the invoking of an appropriate recovery mechanism based on the LoS of a particular attack. To the best of our knowledge, this is *the first effort to identify the significance and application of this concept* for Cloud computing.

The efforts described in this paper are a part of the system described in [8] and are focused on establishing a severity analysis component for the overall architecture. As part of the methodology to address the severity analysis problem, it is believed that the severity problem can be treated as a special case of the traditional classification problem as it essentially involves segregating intrusion trails into different categories. Furthermore, machine learning techniques have been used to perform the intrusion severity analysis. The parameters used for this analysis include security requirements for guest virtual machines along with SLA state and the frequency of attack on a particular security requirement.

This paper is organized as follows. Section 2 describes the problem and the state-of-the-art related to it. This is followed by a description of system and fault models for the proposed system in Sections 3 and 4, respectively. Section 5 describes the proposed methodology to address the intrusion severity problem. This is followed by two critical aspects of the proposed method i.e. classification techniques used as part of the proposed solution and the security quantification to achieve customized operation in Sections 6 and 7, respectively. Section 8 presents a detailed explanation of the various aspects of evaluation followed by a mention of the conclusions and potential future work in Section 9.

2. Problem definition and related work

The research presented in this paper is considered to be related to severity analysis of intrusions in general and for Cloud computing in particular. Furthermore, similarities can be held with traditional network based systems as well. Therefore, the existing literature in these domains has been explored to draw a comparative analysis of the proposed approach with contemporary approaches.

With respect to Cloud computing, to the best of our knowledge, we are the first to identify the intrusion severity analysis problem for such systems. However, there have been efforts in traditional systems to address this problem. An obvious example of such systems is network intrusion detection system (NIDS) where an intrusion detection system is usually deployed at a border node to look after a whole network of computers. As part of the network, different sub-domains or clusters can have varying security requirements. Our experience with such systems has been that the problem of potential varying impact of an intrusion is addressed by defining customized security policies for such groups of nodes. However, there are certain defining differences between such systems and the virtual machine based systems such as Clouds. First, the policies tend to be static, largely due to the static nature of the monitored systems. This is because the groups of nodes tend to have stable security requirements which have been established overtime based on experience with such systems. However, with Clouds, the monitored virtual machines are added and removed dynamically. Furthermore, the security requirements of individual virtual machines are also envisaged to be diverse, aggravating the problem. Second, the security policies in traditional systems are designed and managed by a system security administrator, with some input from the users, who is responsible for the security of the whole system. This human intervention can become the weak link to realize customization and on-demand operation offered by Clouds. Additionally, it also affects the intrusion response time thereby affecting the overall security of the system.

With respect to the use of intrusion severity to select optimal response, intrusion response systems [9] use different metrics to achieve this objective. Schnackenberg et al. [10] and Porras and Neumann [11] represent two such approaches which use a severity metric. However, this metric is proposed to be computed by a human administrator through an offline analysis at the policy definition stage. In this case, severity is usually calculated based