Mitigating Security Threats to Large-Scale Cross Border Virtualization Infrastructures*

Philippe Massonet[1], Syed Naqvi[1], Francesco Tusa[2], Massimo Villari[2],
and Joseph Latanicki[3]

Centre d'Excellence en Technologies de l'Information et de la Communication
{Syed.Naqvi,Philippe.Massonet}@cetic.be
Università degli Studi di Messina, Facoltà di Ingegneria
{mvillari,ftusa}@unime.it
Thales
Joseph.Latanicki@thalesgroup.com

Abstract. Cloud Computing is being a computation resources platform where it is possible to make up an environment flexible and scalable able to host any kind of services. In Cloud Computing, virtualization technologies provide all the needful capabilities to deploy services and run applications in an easy way. Furthermore, large-scale cross border virtualization infrastructures present promising landscape to cope with the ever increasing requirements of modern scientific and business applications.

The large-scale cross border virtualization infrastructures can be seen as a federation of heterogeneous clouds. We present pragmatic analysis of the potential threats posed to the emerging large-scale cross border virtualization infrastructures. We have taken into consideration both *internal* and *external* threats to these infrastructures. We also drive the discussion considering a real model of cloud. In particular an *infrastructure cloud* is briefly presented; a useful scenario where to assess security threats and apply security solutions, that is the European Project, RESERVOIR.

Keywords: Cloud Computing, Security Architecture, Threats modelling, Virtualization infrastructure.

1 Introduction

Currently available cloud architectures do not strongly address security necessities [1,2]. Security has to be considered as an integral part of the development process rather than being later addressed as an add-on feature. The conception of a comprehensive security model requires a realistic threat model. Without such a threat model, security designers risk wasting time and effort implementing safeguards that do not address any realistic threat. Or, just as dangerously,

* The research leading to the results presented in this paper has received funding from the European Union's seventh framework programme (FP7 2007-2013) Project RESERVOIR under grant agreeement number 215605.

D.R. Avresky et al. (Eds.): Cloudcomp 2009, LNICST 34, pp. 73–82, 2010.

they run the risk of concentrating their security measures on one threat while leaving the underlying architecture dangerously exposed to others.

In this paper, we drive the discussion considering a real model of cloud. In particular an *infrastructure cloud* is briefly presented, where it is possible to assess the security aspects through a meaningful scenario, that is the Resources and Services Virtualization without Barriers (RESERVOIR) [3]. The RESERVOIR platform presents concepts as virtualization infrastructure, VEEs, dynamic deployment, elastic and autonomic systems where all actions must to be performed in a secure way. Furthermore the dynamic management of computational resources among sites represents the main challenge to cope by the RESERVOIR cloud computing middleware.

Afterwords a brief description of RESERVOIR, we present a detailed analysis of the threats to large-scale cross border virtualization infrastructures. These threats are broadly classified into two major categories namely *internal threats* and *external threats* so as to complement the DolevYao threat model [4]. We also present some mitigating techniques to cope with these threats and position them with the existing solutions.

The paper is organised as follows: Section 2 surveys related works; Section 3 briefly covers RESERVOIR basic concepts, explaining its architecture, entities and stockholders involved. Section 4 presents all the threats that a cloud infrastructure may suffers by attackers. Sections 5 explains how to face the threats previously highlighted, providing some solutions, case by case. Section 6 finally concludes the dissertation.

2 Related Works

The term *Cloud Computing*, has recently become popular together with *Web 2.0*. Since such paradigm is mostly new, there are dozens of different definitions for Cloud Computing and there seems to be no consensus on what a Cloud is: the paper [5] aims to compare and contrast Cloud Computing with Grid Computing from various angles, explaining the essential characteristics of both. According to the authors, Cloud Computing is not completely a new concept; it has intricate connection to the existing Grid Computing paradigm and other relevant technologies. This paper offers a good starting point to identify the different kind of issues involved in cloud computing: the ones related to security represented a valid basis for our research.

Paper [6] refers to the threats analysis of those scenarios involving general computer systems: attackers and defenders both strive to gain complete control over them. To maximise their control, both attackers and defenders have migrated to low-level, operating system code. This paper assumes the perspective of the attacker, who is trying to run malicious software and avoid detection. By means of the proposed approach, the authors hope to help defenders to understand and defend against the threat posed by a new class of rootkit, called VMBR (Virtual Machine based rootkit), which install a virtual machine monitor underneath an existing operating system. As our main paper topic, the one of this work refers to the study of internal threats involved in the execution of virtual machines. Differently from our case, the study is not strictly related to Cloud Computing environments.

3 RESERVOIR - An Example of Large Scale Cross Border Virtualization Infrastructure

Nowadays, all the commercial cloud infrastructures do not provide any detail of whole components compounding their systems. As we already highlighted, in order to overcome to these limitations and survey however these type of cloud infrastructures, we performed our assessment on the RESERVOIR cloud scenario. In this section we briefly describe the RESERVOIR architecture (many more details are presented in [3]), hence we will opportunely address the security issues of a federation of infrastructure providers in the cloud computing context.

RESERVOIR will introduce an abstraction layer that will allow to develop a set of high level management components that are not tied to any specific environment. This abstraction involves a federation of heterogeneous physical infrastructures. As shown by Figure 1 (reference architecture), in RESERVOIR, more sites (site A and site B) can share physical infrastructure resources on which service applications can be executed. All the entities depicted by the picture are explained just below.

Every site is partitioned by a virtualization layer into virtual execution environments (VEEs). These environments are fully isolated runtime modules that abstract away the physical characteristics of the resource and enable sharing. The virtualized computational resources, alongside with the virtualization layer and

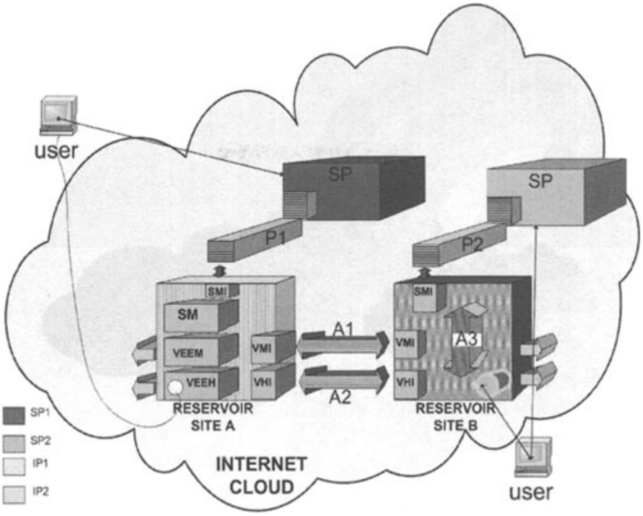

Fig. 1. RESERVOIR reference architecture: a federation of heterogeneous physical infrastructures

all the management enablement components, are referred to as the VEE Host. A service application is a set of software components which work to achieve a common goal. Each component of such service application is executed in a dedicated VEE. These VEEs are placed on the same or different VEE Hosts within the site, or even on different sites, according to automated placement policies that govern the site. Neither Service Provider (SP) nor final User are aware of the real mapping between service application and hardware resources. In RESERVOIR's model, there is a separation between SP (e.g. ebay, or Salesforce) and Infrastructure Providers (IP - Amazon, Google, Flexiscale, etc.). SP are the entities that understand the needs of particular business and offer service applications to address those needs. SPs do not have the computational resources needed by these service applications, instead, they lease resources from a cloud, which provides them with a seemingly infinite pool of computational resources.

RESERVOIR clouds installed on each site present three different layers (see Figure 1 RESERVOIR Site A) described as follows:

- Service Manager (SM): it is responsible for the instantiation of the service application by requesting the creation and configuration of VEEs for each service component, in agreement with SP performed with a shared manifest.
- Virtual Execution Environment Manager (VEEM): it is responsible for the placement of VEEs into VEE hosts.
- Virtual Execution Environment Host (VEEH): it represents a virtualized resource hosting a certain type of VEEs. VEEM issues generic commands to manage the lifecycle of VEEs, and VEEHs are responsible for translating these commands into commands specific to the virtualization platform abstracted by each VEEH.

4 Security Threats to RESERVOIR Infrastructure

In this section we assess the security issues raising in RESERVOIR architecture, highlighting those involved in a federation of infrastructure providers in the cloud computing context. We underline that the added value of our dissertation is not given by a simple threats classification, given that the work provides the gathering of more security concerns, with a complete (360 degrees) perspective of Cloud Computing environments.

In order to take decisions about the RESERVOIR security architecture, information security, policy creation and enforcement, an analysis of the various kinds of threats facing the RESERVOIR architecture, its applications, data and information systems is required. Moreover, in order to identify all the possible threats to federations of heterogeneous physical infrastructures, we provide a simple classification: 1) *within* a RESERVOIR site for all the interactions among VEEM, VEEH, and SM; 2) *across* the RESERVOIR sites for the SLA based VMI interactions between the VEEMs of different RESERVOIR sites; 3) *outside* the RESERVOIR sites for the interaction between SM and SP (SMI). Actually, the threats reported in item 1 and 2 are quite similar. The communication can be affected by the same type of threats. The vulnerability appears during the communication between entities and it is also present in all the network interfaces. The communications can be categorised as follows: *horizontal communication* (parallelepipeds P1 and P2, arrows A1 and A2); *vertical communication* (vertical arrow A3).

The endpoints in the horizontal communication are both SMs with SPs and RESERVOIR sites (i.e Site A and B), while in the vertical communication the entities involved are SMs, VEEMs and VEEH in each site (i.e Site A or B). Horizontal communication exposes endpoints toward External Threats. The communications occur throughout Internet since there is an high level of risk. Vertical communication is the subject of Internal Threats. The SMI, VMI and VHI interfaces are located in External Threats.

4.1 External Threats

The Internet represents the same origin of threats for the communication across the RESERVOIR sites (VMI-VHI interfaces) and outside the RESERVOIR sites for the SMI interface (e.g. injection, identity theft and spoofing).

All the interfaces could be also exposed different attacks (e.g. denial of service, flooding and buffer overflow). These kind of threats are aimed toward provoking a *system crash*, leading to the inability to perform ordinary functions. All the interfaces (SMI, VMI and VHI), are affected by the same issues, but we have to underline the solutions in some cases are different. Considering the VMI and VHI interfaces, the RESERVOIR system administrator has the full capability to manage security policies and to apply them on both the sides (endpoints of site A and site B). Hence in RESERVOIR it is possible to select an its own security framework. While in the case of communication between SM and SP (SMI), the RESERVOIR cloud has to use a common security framework shared with many different partners. Since, it is necessary to solve the same issues under two different perspective views.

4.2 Internal Threats

RESERVOIR site has a logical representation with three different layers, but these layers can be compounded by one or more hardware components. Figure 2 gives an overview of these entities and relative mapping with a simplified view of the hardware. First of all, it is possible to split the site in two different virtual zones: *control and execution zone*. In the *control zone* there are: Service Manager (SM), VEEM (in bridge configuration between control and execution zone), network components (router, switch, cable, etc.), SMI/VMI interfaces and VHI internal interface.

In the *execution zone* instead there are: VEEH, VEEM (in bridge configuration between control and execution zone), VHI internal interface: VHI, network components (router, switch, cable, etc.), network storage: NAS, databases, etc and VHI/User Internet access interfaces.

The *control zone* can be considered a trusted area. Some threats can appear through the interfaces SMI and VEEM, since they fall into the same cases of

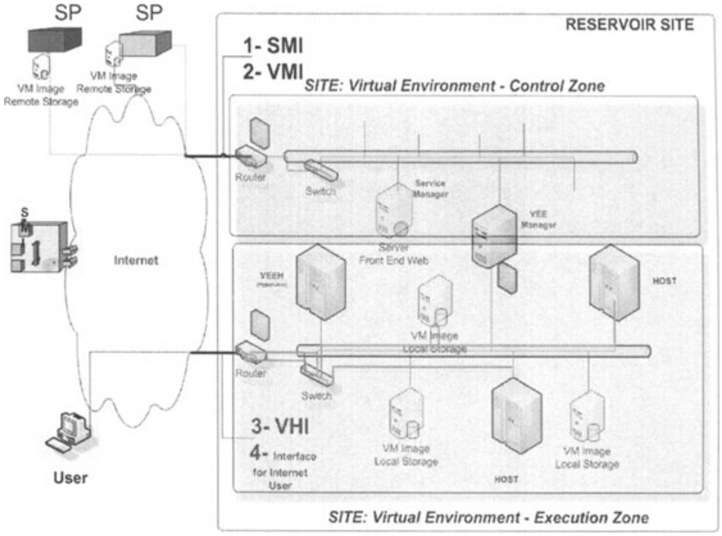

Fig. 2. RESERVOIR site: internal representation

external threats. The firewall located next to the router increases the trustworthiness. In this zone the weak ring of the chain is represented by the VEEM. It is the bridge between two areas, and it allows to exchange data among the zones. Figure 2 shows a firewall close to the VEEM, added to prevent any attacks from the execution area. The zone with high level of risk is represented by the *execution zone*. It can be considered as Demilitarised Zone (DMZ). This area shares has all the hardware components. The hypervisor (VEEH) uses the network, storage, CPU and ram (host) to load and execute all the VEEs. To better explain the role of each component it can be useful to evaluate chronologically all the phases necessary to execute a Virtual Execution Environment: VEEH, once all the requirements from VEEM are received, it downloads the VM Image from the SP, stores the Image into the NAS, it performs the setup configuration and executes the VM. The internal threats related with these phases can be classified as: 1) authentication/communication of SPs and other RESERVOIR site; 2) misbehaviour of service resource allocation due to malicious manifests; 3) data export control legislation: on an international cloud or between two clouds; 4) fake command for placement of VEEs and compromising data integrity of Distributed File System (NFS, SAMBA, CIFS); 5) Storage Data compromising (fake VEE image); 6) data privacy compromising; 7) hypervisor and OS security breaking; 8) data partitioning between VEE.

To avoid any fraudulent access, the VEEH has to verify *authentication/ communication* of SPs and other RESERVOIR sites. Thus is the same behaviour analysed for all the communications in external threats. Relatively to later group of threats (3,4,5 - 6,7,8) RESERVOIR site has to guarantee different types of isolation, that is: runtime isolation, network isolation and storage isolation.

Runtime isolation resolves all the security problems with the underlying OS. The hypervisor has to provide all the solutions.

Network isolation is addressed via the dynamic configuration of network policies; virtual circuits that involve Routers and Switches (Virtual LAN) (See figure 2, there are more virtual circuits with different colours).

To avoid fake VEE image loading and do not compromise data privacy, *storage isolation* has to be performed and secure protocols has to be used. Protocols like NFS, SAMBA, CIFS are not secure. Virtual Execution Environment, downloaded from any generic SP, can expose the infrastructure toward back door threats, spoofing threats and malicious code execution (virus, worm and Trojan horse). The RESERVOIR site administrator needs to know at any time the state of threats, with a strong monitoring of the *execution zone.*

5 Mitigating Techniques for Security Threats

This section presents some security techniques that could be used to mitigate some of the security threats described in the previous section. It is by no means a complete and detailed description of the RESERVOIR security architecture that is required to cover all of the threats described in the previous section. This section does not argue on the isolation needed at hypervisor level (VEEH) (*runtime isolation*). These type of threats could meaningful compromise the whole architecture and they have to be treated in a careful way. Paragraph 5.6 highlights a possible solution able to reduce, and even remove all the risks related to *runtime isolation.*

5.1 Centralised or Decentralised PKI: Cross Certification?

One of the key security issue in a virtualized architecture is the identification/authentication of all the different elements which build up a Cloud. To be able to identify and authenticate such elements, one solution is to use a Private Key Infrastructure (PKI) based on certificates controlled by a Certification Authority (CA). But two solutions are available, a centralised or a distributed architecture. Another issue is raised by the fact that every architecture provider will have its own PKI. To solve this issues, one could use a cross certification process which will permit the use of every agreed CA certificates in the cloud, but this process is quit painful to run due to legal aspects. Another solution would be to create a root CA and then the PKI becomes fully centralised. This solution brings new issues such as, who is going to manage and run this root CA.

The choice of centralised or distributed PKI also depends on the centralised or decentralised cooperation between RESERVOIR sites. In the case of centralised cooperation a virtual organisation could be formed by relying on a unique certification authority. The virtual organisation could then provide authentication and access control for all RESERVOIR sites: cooperation would only be authorised between RESERVOIR sites that are members of the same virtual organisation.

However, in the case of decentralised cooperation between sites that form a loosely coupled federation, a distributed PKI architecture is more adequate. In this approach each site is responsible for establishing and managing trust relationships with other RESERVOIR sites. A potential security architecture for RESERVOIR could supports multiple certification authorities. This architecture introduces certification authorities (CA) and a new component for each site, an LDAP slave server. CA entities can be external, e.g. Verisign or Digital Signature Trust Company, some sites can have their own RESERVOIR certification authorities.

The LDAP server represents the entity where it is possible to publish certificates of service providers (SP1, SP2, SP3 etc. etc.), service managers (SM site A, SM site B, SM site C etc.), VEEM (VEEM site A, VEEM site B, VEEM site C etc.), as well as relationships between sites and VEEH (VEEH of site A, VEEH of site B, etc.) and relationships between VEE and service providers (VEE1 belong to SP1, VEE2 belong to SP2, VEE3 belong to SP3 etc.). In a Master/Slave configuration each site has a consistent copy of all information.

5.2 Ciphering: Communications, Data, Customer Data in the Management

One of the major threats in a virtualized architecture is about the communications and data confidentiality. Many technical solutions are available, such as Secure Socket Layer (SSL), IPSEC... One has to be careful to use the right algorithm and the right key length to be sure of the robustness to the solution. Speaking of keys, some issues raise. Who is delivering keys, how are they distributed? A good way is to use the TPM component which is mainly built for this purpose. It could be used also to generate keys to ciphered data, but what about the key recovery process issue. How to recover the key used to cipher data when this key has been lost.

5.3 Virtual or Physical Firewalls

Obviously, there will be firewalls in a virtualized architecture, but we can use physical or virtual one. Physical firewalls are well known and described. Some of them are certified and we know a lot about their security. Some virtual firewall are now available, and it seem more elegant to use them in a virtualized architecture. On both type of firewall, an issue is raised about their management. Some new threats should be taken into account. A simple human error could brake the full isolation (this threat exists also in a standard architecture). In that case traceability of the administration activity should be available to be able to build organisation processes to avoid such errors. This traceability which could available to the Cloud service provider as to the user, could be a good way to inspire confidence in a Cloud Computing architecture.

5.4 Virtual Switches: VLAN in the Architecture

Virtual LAN Network (VLAN) technology is well use, accepted in the IT world and can be used in a virtualized architecture. As for firewall some virtual switches begin to be available in these architecture. These VLAN can be used to isolate networks, but again as for firewalls the administration issue has to be solve and traceability is a possible way to help to solve it.

5.5 Securing Migration of VEEs

The security of migration of VEE between different RESERVOIR sites that have different security policies must be addressed. One approach to securing migration is to use security profiles. The service provider that submits a service manifest to a primary RESERVOIR site also needs to provide a required security profile. Submission to the RESERVOIR site would only be authorised if the required security profile matches the infrastructure security profile of the primary RESERVOIR site. Migration of VEEs to a destination site would only be authorised if the required security profile matches the destination security profile.

A security profile is defined in terms of security features found at each site such as the use of HTTPS, a firewall, an encrypted file system, a VPN tunnel or a VLAN. Security profiles is ordered from less secure to more secure. This ordering between security profiles provides the basis for comparing and matching security profiles.

5.6 Mitigating Techniques through the OpenTC Solution

Considering the architecture presented previously, many threats may be derived by the compromising of *runtime isolation.* The risks are carried out by the fact that a malicious software (malware) can be execute at VEEH level. These *malwares* could be installed either inside the VEEs or in between of hypervisors and hardware. Latest type of threats are well recognised in [6]. The authors underline the possibility to install a malware able to change the boot sequence. In our cloud platform, we don't have to make an in-dept introspection of hypervisors' functionalities. But, the architecture needs to monitor the hypervisor's behaviour and verify its authenticity and integrity.

Therefore, our cloud implementation we are developing, has to guarantee isolation at VEEH level and it has to be able to avoid the probability that a malicious software gains the control of a site. In order to mitigate these threats, we identify a set of capabilities based on Trusted Computing (TC), and in particular through its open source implementation: OpenTC.

Trusted Computing is an effort to bring some of the properties of closed, proprietary systems to open, commodity systems. This is done using a combination of hardware and software components. Furthermore, these components allow to check and enforce the integrity of a system, and authenticate itself to remote systems. The hardware block that provides trustiness to whole system is called Trusted Platform Module (TPM), that is tamper-resistant and has an embedded private key. This component is able to assure the identification of all the hardware or software components of the architecture, but it has to be available on all the equipments which is not always the case. Although TC is controversial as the hardware is not only secured for its owner, but also secured against its owner as well, we think, its feature may really increase the trustiness in Cloud Computing.

6 Conclusions and Perspectives

We have presented a pragmatic analysis of a range of potential threats to the emerging large-scale cross border virtualization infrastructures. The focal point of this work was cloud computing architectures. In the detailed presentation of these threats and their impact on the overall functioning of clouds is elaborated. We have also explored various security solutions to effectively address the security requirements of virtualization infrastructures. It is important to remember that security is a process, the threat picture is always changing, and threat analysis needs to be continuously updated. In other words, virtualization infrastructure should be subject to constant review and upgrade, so that any security loophole can be plugged as soon as it is discovered.

We are working on a comprehensive security model for a reference architecture of Cloud deployment. We plan to use this threats analysis in defining various core functionalities of the eventual security solutions.

References

1. Amazon Web Services: Overview of Security Processes, http://s3.amazonaws.com/aws_blog/AWS_Security_Whitepaper_2008_09.pdf
2. Comprehensive review of security and vulnerability protections for Google Apps, http://www.google.com/a/help/intl/en/admins/pdf/ds_gsa_apps_whitepaper_0207.pdf
3. Juan Caceres, R.M., Rochwerger, B.: Reservoir: An architecture for services, the first issue of the reservoir architecture document (June 2008), http://www.reservoir-fp7.eu/twiki/pub/Reservoir/Year1Deliverables/080531-ReservoirArchitectureSpec-1.0.PDF
4. Dolev, D., Yao, A.C.: On the Security of Public Key Protocols. In: Proceedings of the IEEE 22nd Annual Symposium on Foundations of Computer Science, pp. 350–357 (1982)
5. Foster, I., Zhao, Y., Raicu, I., Lu, S.: Cloud Computing and Grid Computing 360-Degree Compared. In: Grid Computing Environments Workshop, GCE 2008, November 2008, pp. 1–10 (2008)
6. King, S.T., Chen, P.M., Wang, Y., Verbowski, C., Wang, H.J., Lorch, J.R.: Subvirt: Implementing malware with virtual machines. In: SP 2006: Proceedings of the 2006 IEEE Symposium on Security and Privacy, Washington, DC, USA, pp. 314–327. IEEE Computer Society, Los Alamitos (2006)