A Multi-Layered Architecture for a Secure Virtualization Environment

Amir Soltani Nezhad
Universidade de Lisboa
Faculdade de Ciências
LaSIGE
amir@lasige.di.fc.ul.pt

António Casimiro
Universidade de Lisboa
Faculdade de Ciências
LaSIGE
casim@di.fc.ul.pt

Paulo Verissimo
Universidade de Lisboa
Faculdade de Ciências
LaSIGE
pjv@di.fc.ul.pt

***Abstract*—Virtualization is a way to efficiently use and distribute system resources. There are two ways to design a virtualization environment: using hardware or software. Although most commodity CPUs support virtualization, they are able to only support one single layer of virtual machines. To cope with this limitation, one promising approach is to use recursive virtualization, which can also be exploited to increase the security level of a virtualization environment. In this paper we propose an architecture applying the idea of *defense-in-depth* to enhance the security of virtualization systems.**

***Keywords*- Recursive; Virtualization; Hypervisor; Security; Defense-in-Depth**

I. Introduction

Cloud computing providers use virtualization as an efficient way of deploying resources. Virtualization is thus widespread and millions of users are served by this technology everyday. In consequence, the need for protecting users against attackers is increasingly obvious.

Defense-in-depth is an advanced security technique. Using this technique, we are able to place diverse and successive defending layers and mechanisms in the way of adversaries that wish to attack a system. The task of the adversary becomes more difficult, since it must then defeat layers one by one, and with potentially different attack methods. To apply the idea of defense-in-depth in virtualization environments, we need to be able to stack several virtual machines as layers. Unfortunately, current hardware does not support this kind of layering. Therefore, we need to emulate it by software.

Recursive virtualization is an approach that lets a virtual machine recursively create another layer of virtual machines and as such, might serve the purpose of supporting the defense-in-depth technique. However, recursive virtualization classically results in a huge overhead, which has been an obstacle to its acceptance. To deal with this problem, [3] introduced an efficient way of implementing recursive virtualization by managing recursive VMs in the root hypervisor rather than repeating the support for nested virtual machines in every layer. The approach makes it feasible to implement this type of nested virtualization.

In this paper, we present an architecture combining the ideas of defense-in-depth and recursive virtualization to make virtualization environments much more resilient to attacks.

II. Threat Model

In this section, we describe the threat model considered. We assume that a user virtual machine can be malicious, and may attempt to compromise other VMs and the whole system. Example attacks include injecting malicious code or triggering a bug in the hypervisor by a malicious VM. An attacker may also try to reach and compromise operating systems in intermediate layers of the recursive virtualization stack. Attacks may eventually affect integrity and availability of our multi-layer architecture. In this work we do not consider hardware security. In fact, we assume that malicious code cannot tamper with hardware, platform BIOS and other important hardware parts.

Therefore, our ultimate goal is to increase the protection of the main hypervisor and other user VMs, against attacks originating from malicious user VMs. We do so by interposing layers of *defensive VMs* between the user VMs layer and the hypervisor.

III. Related Work

The theoretical foundations of recursive virtualization were first addressed by Popek and Goldberg [4]. However, given that high overhead has always been the side effect of recursive virtualization, to the best of our knowledge, our proposal is novel in using recursive virtualization to enhance security of virtualized systems. There is already a considerably large number of papers dealing with security in virtualization and hypervisors, but from different perspectives. For example, [6] discusses a small hypervisor (micro hypervisor) named NOVA, that cannot be exploited easily by an adversary. This is because the small piece of code can be easily verified, making bug removal easier than with large hypervisors. The authors propose to move to the user level many functions that are not necessary at lower levels. In consequence, the trusted computer base (TCB) becomes very small and hard to be exploited by attackers. Another notable example is SecVisor [5]. SecVisor is a tiny hypervisor that prevents attackers from injecting arbitrary code into some OS kernel. The idea behind this work is that a CPU should just run kernel-mode code that has been verified by SecVisor. Also, nobody can modify the kernel-mode code, but SecVisor.

IV. A Secure Architecture for a Recursive Virtualization Environment

At this stage, we only focus on proposing an architecture for virtualization environments, as depicted in Figure 1. We would like to choose a tiny hypervisor for the lowest (and main) layer of our architecture, and then customize it such that it contains the capabilities introduced in [3] for a low overhead recursive virtualization system. We believe that this change will not be huge and, therefore, will not make the hypervisor code large and unverifiable. SecVisor for example, would be a good candidate for deploying the root-layer of our architecture, given its characteristics described earlier.

On top of this root layer we are able to have several virtual machines. In regular non-recursive systems, one would normally create user VMs with the desired guest operating systems just above the root hypervisor. What we propose in this work is to add several *defensive VMs* between the root hypervisor and those user VMs. Those defensive VMs will feature diverse mechanisms implementing defense-in-depth against attacks on the hypervisor, making it harder for attackers to take control of the system. The number and functionality of these defending layers is adjustable, depending on the level of protection desired for our virtualization environment. Cloud providers can leverage on this versatility, creating different configurations with incremental security.

Protection of the intermediate layers is also of concern. By using SecVisor or any similar hypervisor as our root layer, we can also detect any changes in the kernels of those intermediate layers, resulting from direct attacks on the defensive VMs, and thus increasing the baseline security of our recursive virtualization system. Likewise, diversity will be used not only in the defense mechanisms themselves, but also in the structure of the intermediate layers, where we plan to use various simple operating systems along with different small and verified hypervisors such as SecVisor and NOVA.

Whilst using small hypervisors in our architecture shrinks the attack plane by reduction of the TCB footprint, it also means that we cannot expect them to provide all the functionalities of commodity hypervisors normally required by user VMs. We solve this problem by introducing a user-level virtual machine monitor (VMM) as the top level, as shown in Figure 1. Resilience against attacks on the VMM becomes a concern, which may be addressed, for example, by isolation mechanisms that confine compromises of an individual VM to the corresponding VMM component [1].

V. Conclusion

In this paper, we propose the idea of defense-in-depth to elevate the security degree of a virtualized environment. For this purpose, we resort to recursive virtualization, taking advantage of a recently proposed scheme that obviates the well-known exponential overhead problem [3]. Defense in depth is implemented by a stack of *defensive virtual machines*, which implements diverse mechanisms

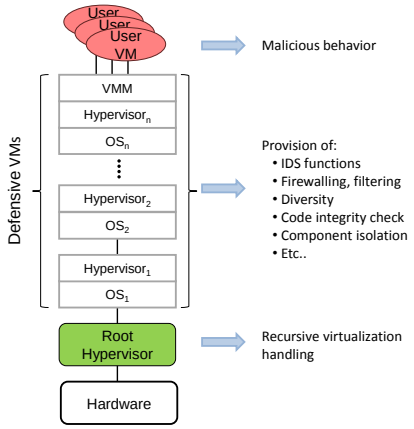

Figure 1. An Architecture for Recursive Virtualization

which make it difficult for an attacker to compromise the root hypervisor and other user VMs. The defense virtual machine stack is also protected by several measures. As future work, we are going to implement and materialize this architecture, evaluate the security of the virtualization system, and compare it to other secure virtualization environments in the literature.

Candidate mechanisms for the defensive virtual machines are for example firewalls or network or host-based intrusion detection systems [2] as well as isolation mechanisms such as wrappers.

Acknowledgment

This work was partially supported by the EC, through project TCLOUDS (FP7-257243), and by FCT, through project TRONE (CMU-PT/RNQ/0015/2009) and the Multiannual and CMU-Portugal programmes. We warmly thank Bernhard Kauer for several fruitful discussions.

References

[1] P. Colp, M. Nanavati, J. Zhu, W. Aiello, G. Coker, T. Deegan, P. Loscocco, and A. Warfield. Breaking up is hard to do: security and functionality in a commodity hypervisor. In *SOSP*, pages 189–202, 2011.

[2] T. Garfinkel and M. Rosenblum. A virtual machine introspection based architecture for intrusion detection. In *NDSS*, 2003.

[3] B. Kauer, P. Verissimo, and A. Bessani. Recursive virtual machines for advanced security mechanisms. In *Proceedings of the 2011 IEEE/IFIP 41st International Conference on Dependable Systems and Networks Workshops*, DSNW '11, pages 117–122, 2011.

[4] G. J. Popek and R. P. Goldberg. Formal requirements for virtualizable third generation architectures. In *SOSP*, page 121, 1973.

[5] A. Seshadri, M. Luk, N. Qu, and A. Perrig. Secvisor: a tiny hypervisor to provide lifetime kernel code integrity for commodity oses. In *SOSP*, pages 335–350, 2007.

[6] U. Steinberg and B. Kauer. Nova: a microhypervisor-based secure virtualization architecture. In *EuroSys*, pages 209–222, 2010.