COMPUTERS & SECURITY 62 (2016) 177–192

Available online at www.sciencedirect.com

ScienceDirect

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

Computers & Security

Exploring infrastructure support for app-based services on cloud platforms [1]

Hai Nguyen [a], *Vinod Ganapathy* [a,*], *Abhinav Srivastava* [b], *Shivaramakrishnan Vaidyanathan* [a]

[a] *CS Department, Rutgers University, 110 Frelinghuysen Road, Piscataway, NJ 08854, USA*
[b] *AT&T Labs-Research, 1 AT&T Way, Bedminster, NJ 07921, USA*

ARTICLE INFO

Article history:
Received 26 February 2016
Received in revised form 29 June 2016
Accepted 27 July 2016
Available online 2 August 2016

Keywords:
Cloud computing
Cloud apps
App stores
Virtualization
Security
Introspection

ABSTRACT

Major infrastructure-as-a-cloud (IaaS) providers have recently been building marketplaces of "cloud apps," which are VMs pre-installed with a variety of software stacks. Clients of cloud computing leverage such markets by downloading and instantiating the apps that best suit their computing needs, thereby saving the effort needed to configure and build VMs from scratch.

We posit that the notion of cloud apps as defined by these marketplaces is nascent and does not allow apps to leverage the benefits of virtual machine (VM) introspection technology developed over the past decade. We envision a marketplace of apps that can interact with client VMs in a rich set of ways to provide a number of services that are currently supported only by cloud providers. This allows clients to deploy services such as VM introspection-based security tools and network middleboxes on their work VMs without requiring the cloud provider to deploy these services on their behalf.

This paper presents models to support such a marketplace of expressive cloud apps. We present a study of the design space of these models to understand their performance and deployment tradeoffs. We also consider the design of a permissions-based framework to contain untrusted third-party cloud apps. Finally, we demonstrate the utility of our models by building and evaluating a number of security tools built as cloud apps.

1. Introduction

Modern cloud computing platforms such as Amazon EC2 and Windows Azure have become popular over the last several years. These platforms host client computations on a shared computing infrastructure while ensuring that mutually-distrusting clients cannot affect the confidentiality or integrity of each other's computations. The key enabling technology is virtualization. By running client computations within virtual machines (VMs) and controlling access to the hypervisor layer, cloud providers ensure that clients are isolated from each other.

In parallel with these developments, the security and computer systems research communities more broadly have been investigating the use of virtualization to offer novel services to VMs. For example, prior research has developed a variety of virtual machine introspection (VMI) (Chen and Noble, 2001)

[1] This paper is a revised and expanded version of a vision paper that appeared in the 2012 ACM Cloud Computing Security Workshop [Srivastava and Ganapathy, 2012].

* Corresponding author.
E-mail addresses: hdn11@cs.rutgers.edu (H. Nguyen), vinodg@cs.rutgers.edu (V. Ganapathy), sv330@cs.rutgers.edu (A. Srivastava), abhinav@research.att.com (S. Vaidyanathan).

http://dx.doi.org/10.1016/j.cose.2016.07.009
0167-4048/

based security tools such as sophisticated malware detectors using VMI (Garfinkel and Rosenblum, 2003; Payne et al., 2007; Srivastava and Giffin, 2008). For example, using VMI one can build as rootkit detectors for operating systems executing within the VM; rootkit detectors cannot be built using traditional in-VM methods. Similarly, network-function virtualization allows the creation and deployment of networked services such as network intrusion detection systems (NIDS) and firewalls to be deployed on VMs.

In this paper, we explore the possibility of offering such enhanced services as *apps* on a cloud platform. That is, a client should be able to download an "app" that implements such services (such as a NIDS or a rootkit detector), and use the app on his VMs executing on the cloud. While the motivating use-cases in this paper are all *security apps*, *e.g.*, VMI-based services that enhance the security of client VMs, the concepts that we develop apply more broadly to any cloud-based service.

Contemporary cloud platforms offer a nascent notion of cloud apps and cloud app marketplaces (NetEx; Managing VMWare vApp – VMWare vSphere 4 ESX and vCenter Server; The Cloud Market: Complete Catalog of Amazon EC2 Images; Cisco – SourceFire). These "apps" are primarily VMs installed with pre-configured software stacks for a variety of standard workflows, such as Web servers and database servers. They primarily benefit clients who lack the expertise or manpower to perform detailed configurations. For example, Amazon allows publishers to create and publicly offer VM images, known as *Amazon Machine Images* (AMIs) (AWS MarketPlace – Find and Buy Server Software and Services that Run on the AWS Cloud), that execute on EC2. AMIs that offer a variety of standard software stacks (*e.g.*, LAMP or SQL database software) are now available, and customers can directly simply instantiate them to their specific domains.

Unfortunately, the notion of apps supported by contemporary cloud providers does not support our vision of services as apps. We envision a cloud app market where apps (implemented as VMs) offer standard utilities such as firewalls, NIDS, storage encryption, and VMI-based security tools. Cloud apps can also implement a number of other non-security-related utilities, such as memory and disk deduplication, and network middleboxes such as packet shapers and QoS tools. Clients can leverage these utilities by simply downloading the appropriate cloud apps, and linking them suitably with their work VMs.

The key challenge in realizing this vision on current cloud computing environments is that such interaction between VMs is disallowed because of security considerations. On a cloud platform such as Amazon EC2 or Microsoft Azure, client VMs – even those executing on top of the same hypervisor on a single physical machine – are completely isolated from each other. Thus, one client VM cannot intercept the I/O or memory state of another VM, even if that VM belongs to the same client. The only exception to this isolation is the *privileged management VM*, a per-node VM controlled by the cloud provider that supervises the execution of client VMs. The management VM oversees all I/O from client VMs and can inspect and modify the memory and CPU state of all VMs executing on that platform. While it is possible to implement VMI-based services or I/O interception within the management VM, it requires the cloud provider to deploy the services on behalf of the client. As a result, the current model fails to realize the full power of an app-based approach – *e.g.*, wherein clients can download and deploy VMI-based services on other VMs that they own.

The primary contribution of this paper is in exploring how we can change the security model in today's cloud computing infrastructures to support our vision of cloud apps. As we discuss in this paper, a variety of design options are possible, and our goal is to perform a comprehensive exploration of the design space rather than committing to any one model. We therefore make the following contributions:

- In Section 3, we present a number of motivating examples of cloud apps to illustrate our vision. These apps range from standalone apps to ones that involve complex system- and network-level interactions with other VMs. We present these examples to identify the key challenges in enabling such an app model atop a cloud platform.
- In Section 4, we present a detailed exploration of the design space for supporting cloud apps. We consider various models of virtualization that are supported by contemporary cloud platforms, and offer designs for each of these models to support cloud apps. We also develop the notion of cloud app permissions, which clients can use to reason about and control the behavior of apps developed by third parties (just as in mobile app markets).
- We have implemented each of our designs atop the KVM hypervisor (Section 5). In Section 6, we demonstrate the utility of our model by building and evaluating a number of security-related cloud apps and evaluate their performance for various points in the design space.

2. Threat model

We assume a standard cloud computing model, where a cloud provider, *i.e.*, an entity such as Amazon or Microsoft, provides computing infrastructure. Clients rent resources from the cloud provider, and run their virtual machines on the cloud provider's hardware.

To support cloud apps, we assume that the apps are hosted on a marketplace, which is supported either by the cloud provider or a third-party. Clients download these apps and apply them to their VMs to obtain access to various services. The main security problems that we must address are:

(1) How do we ensure that a cloud app downloaded by the client can access that client's VMs alone?
(2) How do we confine the cloud apps so that they can access the CPU, memory and I/O state of a client's VMs but only to the extent that they need to achieve their advertised goals.

To achieve these goals, we assume that the *cloud provider is trusted*, and does not intentionally violate the security and privacy of clients. Thus the computing platform, consisting of the physical hardware and the hypervisor, is trusted. To a certain extent, clients can verify their trust in the cloud provider using trusted hardware, *e.g.*, using attestation based on the TPM). We also assume that the cloud provider's infrastructure is equipped with IOMMU units to enable I/O virtualization. The design points