

Towards Trusted Cloud Computing

Nuno Santos

Krishna P. Gummadi

Rodrigo Rodrigues

MPI-SWS

Abstract

Cloud computing infrastructures enable companies to cut costs by outsourcing computations on-demand. However, clients of cloud computing services currently have no means of verifying the confidentiality and integrity of their data and computation.

To address this problem we propose the design of a *trusted cloud computing platform* (TCCP). TCCP enables Infrastructure as a Service (IaaS) providers such as Amazon EC2 to provide a closed box execution environment that guarantees confidential execution of guest virtual machines. Moreover, it allows users to attest to the IaaS provider and determine whether or not the service is secure before they launch their virtual machines.

1 Introduction

Companies can greatly reduce IT costs by offloading data and computation to cloud computing services. Still, many companies are reluctant to do so, mostly due to outstanding security concerns. A recent study [2] surveyed more than 500 chief executives and IT managers in 17 countries, and found that despite the potential benefits, executives “trust existing internal systems over cloud-based systems due to fear about security threats and loss of control of data and systems”. One of the most serious concerns is the possibility of confidentiality violations. Either maliciously or accidentally, cloud provider’s employees can tamper with or leak a company’s data. Such actions can severely damage the reputation or finances of a company.

In order to prevent confidentiality violations, cloud services’ customers might resort to encryption. While encryption is effective in securing data before it is stored at the provider, it cannot be applied in services where data is to be computed, since the unencrypted data must reside in the memory of the host running the computation. In Infrastructure as a Service (IaaS) cloud services

such as Amazon’s EC2, the provider hosts virtual machines (VMs) on behalf of its customers, who can do arbitrary computations. In these systems, anyone with privileged access to the host can read or manipulate a customer’s data. Consequently, customers cannot protect their VMs on their own.

Cloud service providers are making a substantial effort to secure their systems, in order to minimize the threat of insider attacks, and reinforce the confidence of customers. For example, they protect and restrict access to the hardware facilities, adopt stringent accountability and auditing procedures, and minimize the number of staff who have access to critical components of the infrastructure [8]. Nevertheless, insiders that administer the software systems at the provider backend ultimately still possess the technical means to access customers’ VMs. Thus, there is a clear need for a technical solution that guarantees the confidentiality and integrity of computation, in a way that is verifiable by the customers of the service.

Traditional trusted computing platforms like Terra [4] take a compelling approach to this problem. For example, Terra is able to prevent the owner of a physical host from inspecting and interfering with a computation. Terra also provides a remote attestation capability that enables a remote party to determine upfront whether the host can securely run the computation. This mechanism reliably detects whether or not the host is running a platform implementation that the remote party trusts. These platforms can effectively secure a VM running in a single host. However, many providers run data centers comprising several hundreds of machines, and a customer’s VM can be dynamically scheduled to run on any one of them. This complexity and the opaqueness of the provider backend creates vulnerabilities that traditional trusted platforms cannot address.

This paper proposes a *trusted cloud computing platform* (TCCP) for ensuring the confidentiality and integrity of computations that are outsourced to IaaS ser-

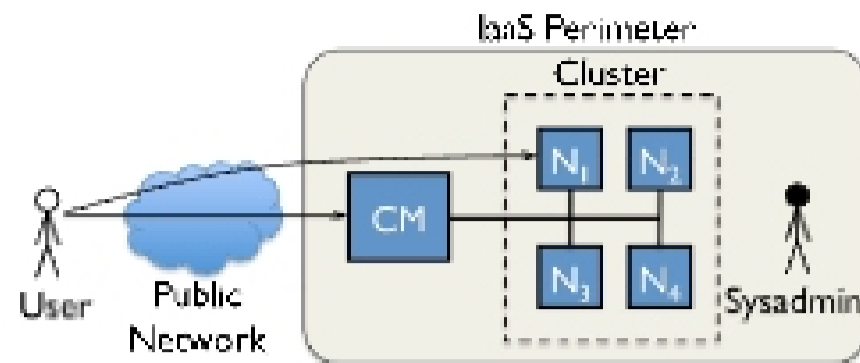


Figure 1: Simplified architecture of Eucalyptus.

vices. The TCCP provides the abstraction of a closed box execution environment for a customer’s VM, guaranteeing that no cloud provider’s privileged administrator can inspect or tamper with its content. Moreover, before requesting the service to launch a VM, the TCCP allows a customer to reliably and remotely determine whether the service backend is running a trusted TCCP implementation. This capability extends the notion of attestation to the entire service, and thus allows a customer to verify if its computation will run securely.

In this paper we show how to leverage the advances of trusted computing technologies to design the TCCP. Section 2 introduces these technologies and describes the architecture of an IaaS service. Section 3 presents our design of TCCP. Although we do not yet have a working prototype of TCCP, the design is sufficiently detailed that we are confident that a solution to the problem under discussion is possible.

2 Background

2.1 Infrastructure as a Service

Today, myriads of cloud providers offer services at various layers of the software stack. At lower layers, Infrastructure as a Service (IaaS) providers such as Amazon, Flexiscale, and GoGrid allow their customers to have access to entire virtual machines (VMs) hosted by the provider. A customer, and user of the system, is responsible for providing the entire software stack running inside a VM. At higher layers, Software as a Service (SaaS) systems such as Google Apps offer complete online applications than can be directly executed by their users.

The difficulty in guaranteeing the confidentiality of computations increases for services sitting on higher layers of the software stack, because services themselves provide and run the software that directly manipulates customer’s data (e.g., Google Docs). In this paper we focus on the lower layer IaaS cloud providers where securing a customer’s VM is more manageable.

While very little detail is known about the internal organization of commercial IaaS services, we describe (and base our proposal on) Eucalyptus [6], an open source IaaS platform that offers an interface similar to EC2. Fig-

ure 1 presents a very simplified architecture of Eucalyptus. This system manages one or more clusters whose nodes run a virtual machine monitor (typically Xen) to host customers’ VMs. Eucalyptus comprehends a set of components to manage the clusters. For simplicity, our description aggregates all these components in a single *cloud manager* (CM) that handles a single cluster; we refer the reader to [6] for more details.

From the perspective of users, Eucalyptus provides a web service interface to launch, manage, and terminate VMs. A VM is launched from a virtual machine image (VMI) loaded from the CM. Once a VM is launched, users can log in to it using normal tools such as ssh. Aside from the interface to every user, the CM exports services that can be used to perform administrative tasks such as adding and removing VMIs or users. Xen supports live migration, allowing a VM to shift its physical host while still running, in a way that is transparent to the user. Migration can be useful for resource consolidation or load balancing within the cluster.

2.2 Attack model

A sysadmin of the cloud provider that has privileged control over the backend can perpetrate many attacks in order to access the memory of a customer’s VM. With root privileges at each machine, the sysadmin can install or execute all sorts of software to perform an attack. For example, if Xen is used at the backend, Xenaccess [7] allows a sysadmin to run a user level process in Dom0 that directly accesses the content of a VM’s memory at run time. Furthermore, with physical access to the machine, a sysadmin can perform more sophisticated attacks like cold boot attacks and even tamper with the hardware.

In current IaaS providers, we can reasonably consider that no single person accumulates all these privileges. Moreover, providers already deploy stringent security devices, restricted access control policies, and surveillance mechanisms to protect the physical integrity of the hardware. Thus, we assume that, by enforcing a security perimeter, the provider itself can prevent attacks that require physical access to the machines.

Nevertheless, sysadmins need privileged permissions at the cluster’s machines in order to manage the software they run. Since we do not precisely know the praxis of current IaaS providers, we assume in our attack model that sysadmins can login remotely to any machine with root privileges, at any point in time. The only way a sysadmin would be able to gain physical access to a node running a customer’s VM is by diverting this VM to a machine under her control, located outside the IaaS’s security perimeter. Therefore, the TCCP must be able to 1) confine the VM execution inside the perimeter, and 2) guarantee that at any point a sysadmin with root privi-

leges remotely logged to a machine hosting a VM cannot access its memory.

2.3 Trusted Computing

The Trusted Computing Group (TCG) [10] proposed a set of hardware and software technologies to enable the construction of trusted platforms. In particular, the TCG proposed a standard for the design of the *trusted platform module* (TPM) chip that is now bundled with commodity hardware. The TPM contains an endorsement private key (*EK*) that uniquely identifies the TPM (thus, the physical host), and some cryptographic functions that cannot be modified. The respective manufacturers sign the corresponding public key to guarantee the correctness of the chip and validity of the key.

Trusted platforms [1, 4, 5, 9] leverage the features of TPM chips to enable *remote attestation*. This mechanism works as follows. At boot time, the host computes a measurement list *ML* consisting of a sequence of hashes of the software involved in the boot sequence, namely the BIOS, the bootloader, and the software implementing the platform. The *ML* is securely stored inside the host’s TPM. To attest to the platform, a remote party challenges the platform running at the host with a nonce n_U . The platform asks the local TPM to create a message containing both the *ML* and the n_U , encrypted with the TPM’s private *EK*. The host sends the message back to the remote party who can decrypt it using the *EK*’s corresponding public key, thereby authenticating the host. By checking that the nonces match and the *ML* corresponds to a configuration it deems trusted, a remote party can reliably identify the platform on an untrusted host.

A trusted platform like Terra [4] implements a thin VMM that enforces a *closed box* execution environment, meaning that a guest VM running on top cannot be inspected or modified by a user with full privileges over the host. The VMM guarantees its own integrity until the machine reboots. Thus, a remote party can attest to the platform running at the host to verify that a trusted VMM implementation is running, and thus make sure that her computation running in a guest VM is secure.

Given that a traditional trusted platform can secure the computation on a single host, a natural approach to secure an IaaS service would be to deploy the platform at each node of the service’s backend (see Figure 1). However, this approach is insufficient: a sysadmin can divert a customer’s VM to a node not running the platform, either when the VM is launched (by manipulating the CM), or during the VM execution (using migration). Consequently, the attestation mechanism of the platform does not guarantee that the measurement list obtained by the remote party corresponds to the actual configuration of the host where the VM has been running (or will be run-

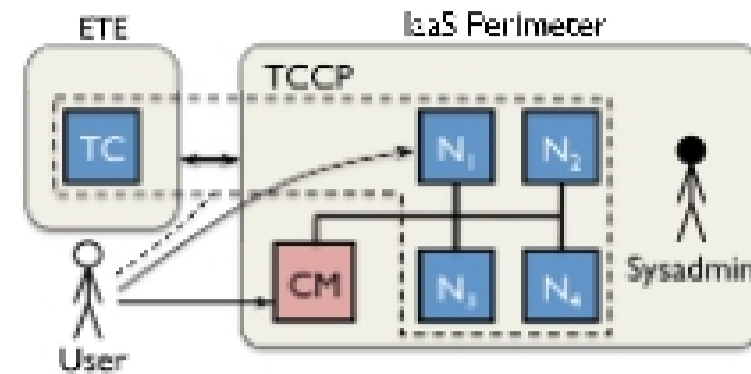


Figure 2: The components of the trusted cloud computing platform include a set of *trusted nodes* (N) and the *trusted coordinator* (TC). The untrusted *cloud manager* (CM) makes a set of services available to users. The TC is maintained by an *external trusted entity* (ETE).

ning in the future). Therefore, the TCCP needs to provide a remote attestation that guarantees the immutability of the platform’s security properties in the backend.

3 Trusted Cloud Computing Platform

We present the *trusted cloud computing platform* (TCCP) that provides a *closed box* execution environment by extending the concept of trusted platform to an entire IaaS backend. The TCCP guarantees the confidentiality and the integrity of a user’s VM, and allows a user to determine up front whether or not the IaaS enforces these properties. Next section gives an overview of TCCP, and Section 3.2 presents a detailed design.

3.1 Overview

TCCP enhances today’s IaaS backends to enable *closed box* semantics without substantially changing the architecture (Figure 2). The trusted computing base of the TCCP includes two components: a *trusted virtual machine monitor* (TVMM), and a *trusted coordinator* (TC).

Each node of the backend runs a TVMM that hosts customers’ VMs, and prevents privileged users from inspecting or modifying them. The TVMM protects its own integrity over time, and complies with the TCCP protocols. Nodes embed a certified TPM chip and must go through a secure boot process to install the TVMM. Due to space limitations we will not go into detail about the design of the TVMM, and we refer the reader to [5] for an architecture that can be leveraged to build a TVMM that enforces local *closed box* protection against a malicious sysadmin.

The TC manages the set of nodes that can run a customer’s VM securely. We call these nodes *trusted nodes*. To be trusted, a node must be located within the security perimeter, and run the TVMM. To meet these conditions, the TC maintains a record of the nodes located in the security perimeter, and attests to the node’s platform to verify that the node is running a trusted TVMM