



Creating Trust in the Cloud

DEPLOY WORKLOADS IN THE CLOUD WITH
CONFIDENCE BY USING UBUNTU[®] SERVER
AND INTEL[®] XEON[®] PROCESSORS



Overcoming Cloud Security Obstacles

Cloud computing offers compelling opportunities for more cost efficient, flexible computing. However, security is a top-of-mind concern for IT professionals, especially in highly regulated industries such as government and financial services. After all, when computing resources reside in fluid pools, how do you keep a compromised system from infecting other systems in that environment? How can you be sure that sensitive workloads run in protected and trusted environments to satisfy compliance requirements? How can you confidently move workloads between systems without fear of spreading malware or exposing your images to stealthy threats? And how can you be sure that user devices, such as PCs, smart phones, tablets, and other emerging gadgets, access the data and workloads in the cloud in a secure manner?

As organizations plan to take advantage of the benefits of cloud computing—whether on public, private, or hybrid clouds—they must take equal care to secure their workloads and sensitive data. Traditional approaches, such as anti-malware software, might not catch threats before they spread to multiple systems across your cloud environment. Software-based blacklisting solutions are also challenged by newer, more sophisticated stealth attacks, such as escalation-of-privilege attacks, which can take control of a system and might remain undetected. In addition, these blacklisting solutions are always one step behind the current threats because they are dependent on reactively identifying and disseminating current threat information to users.

More advanced heuristic-based anti-malware solutions provide better protection but cannot fully secure the underlying hypervisors and BIOS on which your software and virtual machine images reside. Even when effective, anti-malware software is not designed to verify the integrity of the underlying platforms, so it does not assist with compliance needs in highly regulated industries.

A Deeper Approach to Secure Cloud Computing



A more comprehensive approach to security in the cloud is to enable trust from the silicon up, which strengthens security and compliance. A stack that combines hardware optimized for security with open-source cloud software dramatically simplifies the process of creating trust and verifying platform integrity in the cloud. This paper describes such a stack, composed of:

- Ubuntu® Server 12.10 from Canonical
- OpenStack* (the Folsom release, which includes integrated OpenAttestation* [OAT] software development kit [SDK])
- Servers powered by select Intel® Xeon® processors on which Intel® Trusted Execution Technology (Intel® TXT) is enabled¹

With this stack, Intel and Canonical bring differentiated and advanced security capabilities to open-source cloud deployment. The solution relies on tamper-resistant, hardware-based protections with mechanisms for verifying platform integrity. This hardware-based approach helps you enjoy the benefits of cloud computing with a higher level of confidence in the security of your systems and workloads.

This higher confidence grows from trust: you can select hosts for inclusion in trusted compute pools, verify that the components of the hosts' launch stack have not been compromised, and then ensure that workloads you choose run only on hosts in trusted compute pools. Service providers can take advantage of these capabilities to offer secure public cloud services, while IT organizations that use the solution to build private clouds can more easily automate the assignment of sensitive data and workloads to trusted compute pools.

What Are Trusted Compute Pools and How Do They Help?



Trusted compute pools are collections of computing platforms on which the launch process has been measured and that have been verified to be trustworthy. Intel TXT and OAT verify trust status. A trustworthy platform, whether it is in a public cloud or your own data center, is trustworthy because the integrity of those pre-operating system launch components has been verified. By definition, trusted compute pools can only comprise systems with launch stacks that have not been compromised.

No security measure can guarantee complete protection from determined attackers. However, you can significantly increase your level of confidence if you can be assured that critical elements of the launch environment—such as firmware, BIOS, and hypervisor—have not been tampered with. These elements execute before the operating system and its anti-malware drivers load, so they are a blind spot in your view of the enterprise security landscape. In a cloud environment, where computing and storage assets are managed as flexible pools, the state of these critical elements can be even more opaque.

This trustworthiness is established through a measured launch environment (MLE) and an attestation procedure that compares signatures from a host's boot process with signatures that are known to be valid. If the boot process contains unknown elements, or known elements behave uncharacteristically, the host is marked as untrusted and is not included in a trusted compute pool. This approach strengthens protection against attacks that target lower levels of the launch environment, including the BIOS and hypervisor.

Trusted compute pools enable additional benefits:

- They improve security for virtualized data centers and cloud environments because the environment controller (such as OpenStack) isolates platforms with unknown elements or elements that have unexpected traits.
- They support compliance by design by providing verifiable audit trails of a sensitive workload's execution environment.
- They enhance operational agility when combined with intelligent management software that can automate workload assignment, such as Juju®, the cloud orchestration tool from Canonical, and OpenStack.

Intel and Canonical Approach to Trusted Compute Pools



Canonical is a leader in secure cloud solutions, providing a supported route to a cloud infrastructure you can trust. In addition, the OpenStack release includes the Open Attestation (OAT) software development kit (SDK), which provides the critical third-party verification of platform integrity. Canonical and Intel provide an integrated trusted compute pools capability that uses OpenStack and OAT with Ubuntu Server (12.10). This combination of software and Intel hardware technologies enables IT organizations to build more secure clouds and to deploy workloads to those clouds with greater confidence.

This journey to trusted compute pools consists of three phases, which this paper describes in greater depth.

- **Prepare for trust:** IT professionals enable Intel TXT and the Trusted Platform Module (TPM) on the hosts and populate the whitelist with the known good state of the hosts.
- **Establish trust:** When the host powers on, Intel TXT measures the firmware, BIOS, bootloader, hypervisor, and other elements of the launch stack. Values are compared against known-good values stored in the host's TPM. If the values do not match, launch control policies enabled on the host can prevent the host from booting.
- **Verify:** The OpenStack scheduler queries OAT to verify the trust status of every host in the pool.

These phases lay the foundation for trust in the cloud. A user (application owner or server administrator) can now request a secure virtual machine instance from within the OpenStack dashboard and verify that it will only run on a host that is part of a trusted compute pool.

ESTABLISH TRUST WITH TRUSTED BOOT

The concept of trust relies on a verifiable chain of integrity that is forged from the silicon up through the firmware, BIOS, bootloader, operating system, hypervisor, and other elements of the computing stack. That chain is anchored in a tamper-resistant root in the hardware—the TPM—in which launch measurements, or hashes, are stored in platform configuration registers (PCRs). The environment controller (operating system kernel or hypervisor) compares those measurements against known good values and detects variations which can indicate a compromise of one or more launch components.

The environment that is launched as a result of this verification process is called the measured launch environment (MLE), and it is the foundation on which trusted compute pools are built. In the Ubuntu and Intel trusted compute pools solution, the MLE is provided by Trusted Boot (tboot) and Intel TXT.² Tboot is an open source, pre-kernel/pre-hypervisor module that works with Intel TXT to perform measurements and then enable launch control policies based upon those measurements. Tboot has been included in Ubuntu Server since version 11.10 and can be configured with tools provided through command line interfaces (CLI).

Tboot verifies launch components using provisioned Intel TXT hashes and enables launch control policies that can prevent booting of systems whose components cannot be verified. Intel® Virtualization Technology for Directed I/O (Intel® VT-d), embedded on the chipset, provides further protection.³ It allows the MLE to protect itself and other software—such as guest virtual machines—from unauthorized device access to memory. Intel VT-d can be used to block access to specific physical memory pages, and the protection is enforced for all direct memory access (DMA) to the protected pages.⁴

VERIFY TRUST WITH OPEN ATTESTATION (OAT)

When a measured and verified system launch is complete and the platform is ready for workloads, it is not yet considered a trusted host for the purposes of trusted compute pools. The platform must first be certified by a third party in a process called attestation. In the solution this paper describes, this process is performed by Open Attestation (OAT). OAT is an open source SDK that verifies host integrity using the remote attestation protocol defined by the Trusted Computing Group (TCG).⁵ OAT can be added and configured as part of Ubuntu Server.

ATTESTATION FLOW

When creating a trusted compute pool environment, administrators install OAT agents on target hosts and configure the certificates and keys so that the agent can authenticate with the OAT server using secure communications. Administrators then provision the OAT whitelist tables with the PCR values that are known to be valid for the hosts that will be included in the trusted compute pool. These values may be obtained from the platforms' manufacturer or from an initial platform launch that is known to be secure, such as a first launch after shipment from the factory.

When configuration and the OAT whitelist table provisioning are complete, OAT is ready to begin verifying the trust status of hosts. Attestation begins when an OAT agent on a measured host—called an “attesting host”—requests appraisal from the OAT server. The OAT server's appraiser calls for the attesting host's PCR values and sends a nonce that the agent will use to cryptographically sign its responses. The attesting host retrieves its PCR values from the TPM, and sends the response to the OAT appraiser. The appraiser invokes the OAT privacy certificate authority to verify the attesting host's signature and confirms the PCR values against the whitelist table. If the signature is valid and the PCR values match the whitelist table values, then the host is trusted and the appraiser can give a “trusted” response to any attestation queries about that host.

BUILD ON TRUST WITH OPENSTACK

Now that you have established and verified trusted hosts, Intel and Ubuntu with OpenStack can help you take advantage of them, intelligently and automatically. Hosts verified as trusted are added to trusted compute pools and OpenStack with

Ubuntu can intelligently manage them. Data center administrators can then require that confidential data or sensitive workloads run on these trusted hosts that are better controlled and have had their configurations more thoroughly evaluated through Intel TXT, OAT, and OpenStack.

This granular manageability is made possible by the OpenStack scheduler and trusted filter. Intel contributed enhancements to the OpenStack scheduler that enable the scheduler to detect the characteristics of cloud platforms. With that information, the scheduler can better determine which platforms are acceptable candidates for a workload that requires trust.

For example, when requesting a virtual machine from within the OpenStack management dashboard, you or an application owner can specify that workloads should execute only on trusted hosts. The scheduler then invokes a “trusted filter,” which causes it to exclude from consideration any host that the OAT server cannot verify as trusted.

A Clear Path to Trust in the Cloud



If concerns about trustworthy hosts have limited or prevented your use of the cloud, Ubuntu Server and Intel technologies provide solutions to alleviate those concerns. When your cloud platforms run on servers powered by select Intel Xeon processors, the hardware foundation is in place on which you can build trust in the cloud. Ubuntu Server simplifies the process of building trusted compute pools by bringing together necessary components, such as tboot, OpenStack, and OAT.

The joint solution lets you deploy sensitive data and workloads to the cloud with confidence, knowing that they are running on hosts that launched without suspicious anomalies and on which the integrity of the platform firmware, BIOS, hypervisor, and operating system code has been verified. OpenStack then simplifies the process of taking advantage of secure platforms by automatically excluding untrusted hosts from cloud trusted compute pools. Such visibility and control also help your compliance with operational and audit requirements.

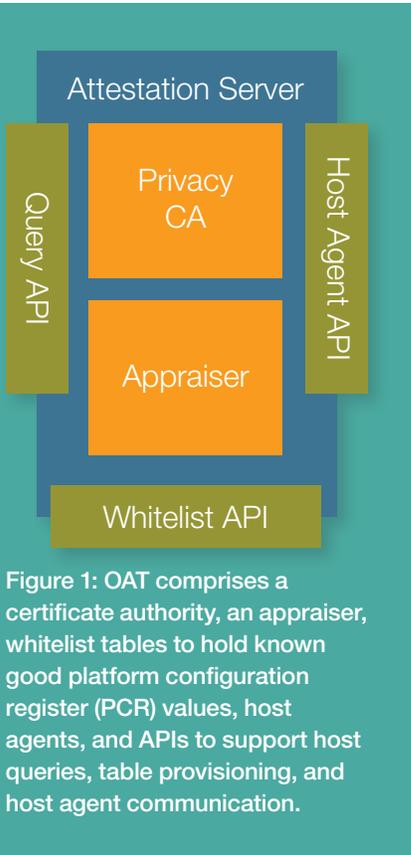


Figure 1: OAT comprises a certificate authority, an appraiser, whitelist tables to hold known good platform configuration register (PCR) values, host agents, and APIs to support host queries, table provisioning, and host agent communication.



Intel® Trusted Execution Technology (Intel® TXT) can help establish trust and compliance for cloud workloads.

www.intel.com/txt



Ubuntu® and OpenStack* software simplify your route to a secure, open cloud.

www.ubuntu.com



¹ No computer system can provide absolute security under all conditions. Intel® Trusted Execution Technology (Intel® TXT) requires a computer with Intel® Virtualization Technology, an Intel TXT-enabled processor, chipset, BIOS, Authenticated Code Modules and an Intel TXT-compatible measured launched environment (MLE). Intel TXT also requires the system to contain a TPM v1.s. For more information, visit www.intel.com/go/inteltxt. For a current list of server manufacturers and models that support Intel TXT, see <http://www.intel.in/content/www/in/en/architecture-and-technology/trusted-execution-technology/trusted-execution-technology-server-platforms-matrix.html>.

² Intel® TXT, the host TPM, and Intel® Virtualization Technology and Intel® VT-d must be enabled in the BIOS for tboot to work properly.

³ Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, and virtual machine monitor (VMM). Functionality, performance, or other benefits will vary depending on hardware and software configurations. Software applications may not be compatible with all operating systems. Consult your PC manufacturer. For more information, visit <http://www.intel.com/content/www/us/en/virtualization/virtualization-technology/hardware-assist-virtualization-technology.html>.

⁴ For more information on Intel® VT-d and DMA protection mechanisms, see Chapter 1.10 of the Intel TXT Software Development Guide, <http://www.intel.com/content/www/us/en/software-developers/intel-txt-software-development-guide.html>.

⁵ Developed by Intel, the SDK Intel was built on the NSA's National Information Assurance Research Laboratory (NIARL) Host Integrity at Startup to measure and report status for host platforms which contain a Trusted Platform Module (TPM). The implementation takes advantage of Infrastructure Work Group Integrity Report Schema Specification available at http://www.trustedcomputinggroup.org/resources/infrastructure_work_group_integrity_report_schema_specification_version_10/. The OpenAttestation SDK supports web API for third-party software to integrate and access web-based attestation to support cloud usage models.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

Copyright © 2013. Intel Corporation. All rights reserved.

Ubuntu, Juju, and Canonical are registered trademarks of Canonical Ltd.

*Other names and brands may be claimed as the property of others.

Printed in USA 0413/MS/PRW/PDF Please Recycle 328889-001US