SOLUTION BRIEF
Intel® Trusted Execution Technology
Cloud Computing

Trusted Geolocation in the Cloud

By implementing an automated hardware root of trust, organizations can monitor and enforce geolocation restrictions and policies, ensuring that their workloads in the cloud are deployed on trusted hardware in known locations.

> "NIST is playing a critical role in providing federal agencies with objective insight and guidance on securely moving applications to the cloud."
>
> *– Jason Kimrey*
> *Director of Federal Sales, Intel*

Trusted Security in the Cloud

Shared cloud computing offers agility and flexibility to organizations, transparently using whatever resources are available to process workloads. However some organizations are reluctant to use this technology because of security concerns.

The National Institute of Standards and Technology (NIST) in the United States, working with Intel, VMware, and RSA Archer, conducted a proof of concept (PoC) that demonstrates trusted geolocation in the cloud using Intel® Trusted Execution Technology (Intel® TXT). Trusted geolocation allows organizations to determine and control the location of the hardware that processes and stores their sensitive information and applications in the cloud. The resulting NIST publication provides a blueprint the general security community can use to validate and implement the described PoC.

The benefits of cloud computing are achieved by migrating workloads to available compute resources. In a secure cloud system, when workloads from different organizations land on a single server, they are segregated so that they do not interfere with each other, gain access to each other's sensitive data, or otherwise compromise the security or privacy of the workloads.

However, some organizations have concerns involving the location of the cloud servers. Depending on the system, workloads may migrate from cloud servers located in one country to those in another. Each country may have its own set of laws for data security, privacy, and other aspects of IT. Because the requirements of these laws may conflict with an organization's policies, it may be necessary to ensure that workloads use only cloud servers physically located in a specific country. This process involves determining the server's location—known as *geolocation*.

Addressing the Need for Trust

A platform that is not trustworthy places the workload at risk of compromise and cannot provide assurance that the claimed geolocation of the cloud server is accurate. A key component of geolocation is the hardware root of trust, which is an inherently trusted combination of tamperproof hardware and firmware that maintains the integrity of the platform and the geolocation information—making it a trusted platform. Geolocated platforms in which there is a hardware root of trust are aggregated into trusted compute pools, segregating them from untrusted resources—resulting in trusted geolocation.

ROOT OF TRUST

Geolocation enables identification of a cloud server's approximate location by adding that information to the server's root of trust. "The hardware root of trust is seeded by the organization with the host's unique identifier and platform metadata stored in tamperproof hardware. This information is accessed using secure protocols to assert the integrity of the platform and confirm the location of the host."

(from "Trusted Geolocation in the Cloud: Proof of Concept Implementation," NIST, sec. 2.1).

Trusted compute pools allow IT to gain the benefits of the dynamic cloud environment while still enforcing higher levels of protections for more critical and sensitive workloads. These compute pools are created by designating trusted compute resources that meet the specific and varying security requirements of users and placing those resources in a segregated area of the cloud. Access to this area is controlled so that only appropriate applications are deployed there, and audits on that portion of the cloud are enabled so users can verify compliance.

Creating Trusted Compute Pools with Trusted Geolocation

The process of establishing a trusted compute pool with trusted geolocation in a cloud has three main stages, as shown in Figure 1. During stage 1, each compute platform must be attested as trustworthy, and a safe hypervisor launch must be ensured. During stage 2, the cloud system must ensure that workload migration occurs only between trusted resources. Trusted geolocation is ensured at stage 3, with continuous monitoring and enforcement of geolocation restrictions.

Stage 1: Platform Attestation and Safe Hypervisor Launch

Attestation is the process of providing a digital signature of a set of platform configuration registers (PCRs) to the platform. This stage provides a basic assurance of platform trustworthiness and enables faster detection of security issues. Stage 1 has three steps.

1. **Configure the server.** Set up the cloud server platform as being trusted, including configuring the hardware, BIOS, and hypervisor.
2. **Verify the hypervisor.** Before each hypervisor launch, verify the trustworthiness of the cloud server platform set up in step 1.
3. **Continually monitor the hypervisor.** During execution, frequently repeat the measurements in step 2 to continually assure trustworthiness. These measurements then become an ongoing part of a continuous monitoring process.

Stage 2: Trust-based Secure Migration

Ensure that workloads are deployed and then migrated only among trusted server platforms within the cloud. Stage 2 has two steps.

1. **Deploy to the trusted platform.** Apply the verification tests established in step 3 of stage 1, and deploy a workload to only those platforms deemed trustworthy.
2. **Migrate to trusted platforms.** After deploying a workload, ensure that it migrates to only hosts with comparable trust levels. This is determined by applying the verification tests from step 3 of stage 1, on both the workload's current server and the server to which the workload is migrating. Both servers must pass their audits in order for the migration to occur.

Figure 1. The three stages of establishing a trusted compute pool with trusted geolocation.

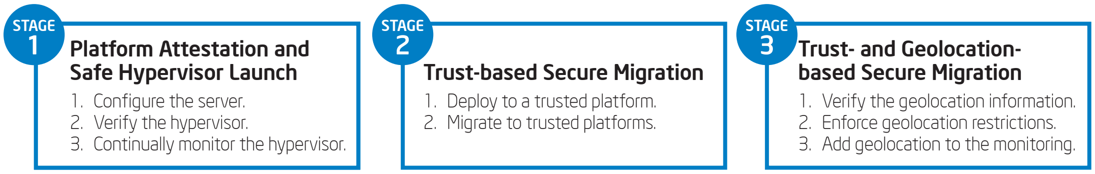

STAGE 1

Platform Attestation and Safe Hypervisor Launch

1. Configure the server.
2. Verify the hypervisor.
3. Continually monitor the hypervisor.

STAGE 2

Trust-based Secure Migration

1. Deploy to a trusted platform.
2. Migrate to trusted platforms.

STAGE 3

Trust- and Geolocation-based Secure Migration

1. Verify the geolocation information.
2. Enforce geolocation restrictions.
3. Add geolocation to the monitoring.

Stage 3: Trust- and Geolocation-based Secure Migration

Ensure that workloads migrate only to trusted server platforms while also taking geolocation restrictions into consideration. Stage 3 has three steps.

1. **Verify geolocation information.** Ensure that any platform to be included in the trusted geolocation pool had its geolocation set as part of its initial configuration in step 1 of stage 1. This process is a cryptographic hash within the hardware cryptographic module in BIOS. Ensure that the geolocation information can be readily verified and audited.
2. **Enforce geolocation restrictions.** Before deploying or migrating a workload, add a geolocation check to the pre-deployment and pre-migration verification in steps 1 and 2 of stage 2.
3. **Add geolocation to monitoring.** Add geolocation checks to the continuous monitoring put in place in step 3 of stage 1 to ensure trustworthiness of the platforms. This process should audit the geolocation of the cloud server platform against geolocation policy restrictions.

The NIST Proof of Concept Implementation

The PoC implementation designed by NIST used commercial off-the-shelf products provided by Intel, VMware, and RSA Archer*. The implementation assumed the use of homogeneous cloud servers, which have the same platform architecture and hypervisor type and reside in the same cloud with a single management console.

Trust Starts with the Hardware and the Hypervisor

A necessary component of hardware root of trust and any trusted compute pool is Intel® TXT, which is supported by the Intel® Xeon® processor 5600 series as well as the more recent Intel® Xeon® processor E3, E5, E7 family. Intel TXT is a set of enhanced features in the microprocessor, chipset, I/O subsystems, and other platform components designed to provide trusted boot and to protect sensitive information from software-based attacks. When coupled with an enabled OS, a hypervisor, and enabled applications, these capabilities provide confidentiality and integrity of data in increasingly hostile environments.

It is important that the hypervisor used in the trusted pool has implemented Intel TXT support. The NIST PoC implementation uses VMware ESXi*, as shown in Figure 2. Hosts are initially configured with VMware vCenter* using the vSphere* Client. If the ESXi hypervisor detects Intel TXT-enabled hardware, the hypervisor undergoes a measured launch. During the launch, the BIOS and virtual machine monitor components are measured (cryptographically) and extended into the server Trusted Platform Module's (TPM) PCRs. The measurement values stored in the TPM are accessible through vCenter using the VMware Web Services SDK and cached in the vCenter database. Intel has provided the Intel TXT plug-in, a reference plug-in to vCenter, which provides basic attestation and access to the measurements for a given ESXi Server from within vCenter.

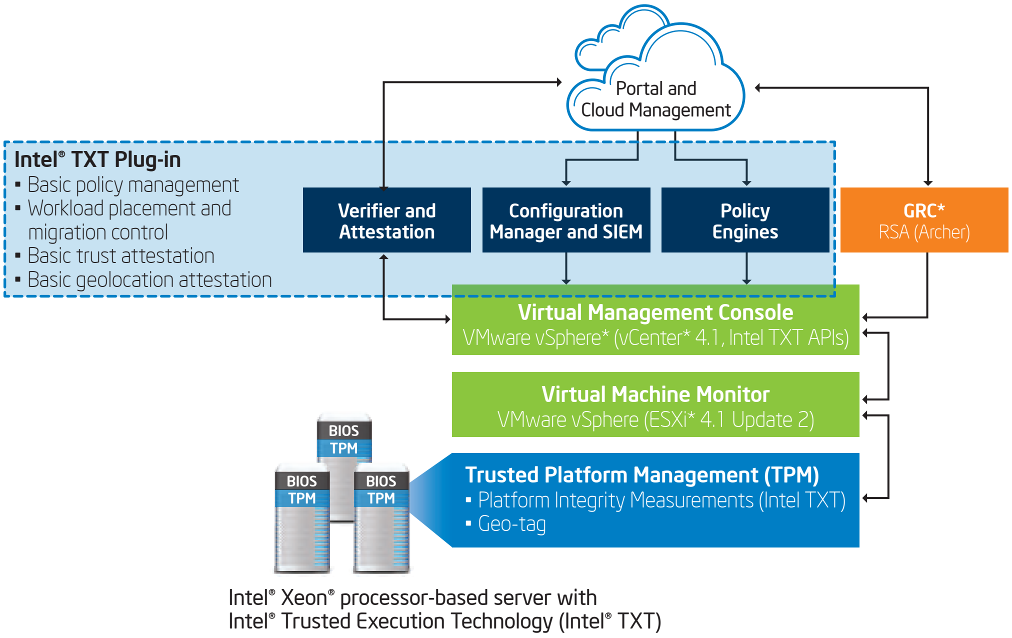

Portal and Cloud Management
Intel® TXT Plug-in
▪ Basic policy management
▪ Workload placement and migration control
▪ Basic trust attestation
▪ Basic geolocation attestation
Verifier and Attestation
Configuration Manager and SIEM
Policy Engines
GRC*
RSA (Archer)
Virtual Management Console
VMware vSphere* (vCenter* 4.1, Intel TXT APIs)
Virtual Machine Monitor
VMware vSphere (ESXi* 4.1 Update 2)
Trusted Platform Management (TPM)
▪ Platform Integrity Measurements (Intel TXT)
▪ Geo-tag
BIOS
TPM
BIOS
TPM
BIOS
TPM
Intel® Xeon® processor-based server with Intel® Trusted Execution Technology (Intel® TXT)

Figure 2. Trusted Cloud Solution Reference Design for the National Institute of Standards and Technology.

Monitoring and Measuring in the Trusted Compute Pool

NIST's PoC implementation used the RSA Archer GRC* program to verify and monitor the trustworthiness of each platform in the trusted compute pool. This program provided a dashboard-based interface for monitoring the mix of trusted and untrusted platforms in a cloud along with their IP addresses, system validation measurements, and measurement statistics. Measuring each server's characteristics frequently, such as at intervals of every five minutes, helped to achieve a continuous monitoring solution for the servers.

Conclusion

By adding geolocation restrictions to the trusted compute pool in a cloud, trusted geolocation is an important solution to the security concerns experienced by those in federal, medical, and large corporate IT security. The NIST PoC implementation used commercially available off-the-shelf products from Intel, VMware, and RSA Archer that are available through open procurement from any hardware OEM. VMware has released ESX 5.1, which is the building block for future stages of this solution. New solutions are available from organizations such as HyTrust and Virtustream. Building on the VMware ESX 5.1 release, revision 2 of this PoC implementation will include the changes shown in Figure 3.

The NIST publication "Trusted Geolocation in the Cloud: Proof of Concept Implementation" provides a thorough overview of the implementation as well as specific steps and detailed notes that can serve as a blueprint or template for organizations in the general security community that want to validate and implement the described PoC.

Figure 3. Phase 2.0 2012 Trusted Cloud Solution Reference Design for the National Institute of Standards and Technology.

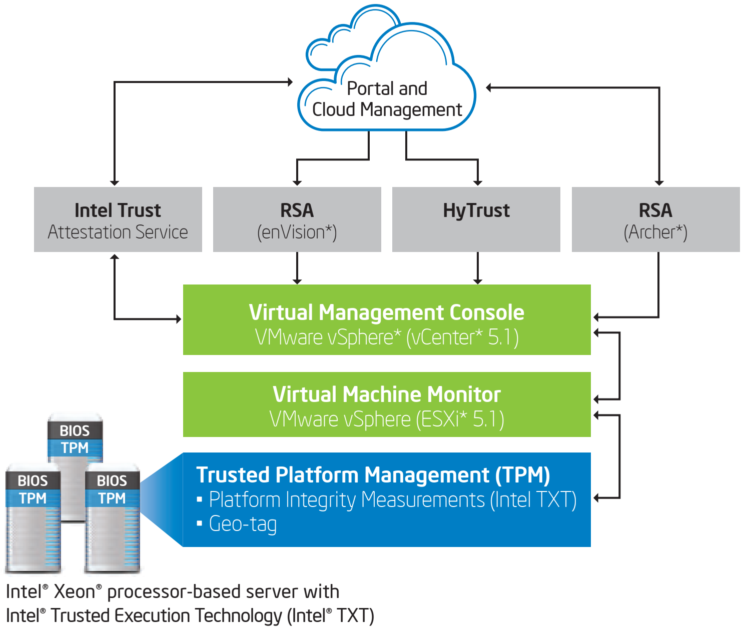
Portal and Cloud Management
Intel Trust
Attestation Service
RSA
(enVision*)
HyTrust
RSA
(Archer*)
Virtual Management Console
VMware vSphere* (vCenter* 5.1)
Virtual Machine Monitor
VMware vSphere (ESXi* 5.1)
Trusted Platform Management (TPM)
▪ Platform Integrity Measurements (Intel TXT)
▪ Geo-tag
BIOS
TPM
BIOS
TPM
BIOS
TPM
Intel® Xeon® processor-based server with
Intel® Trusted Execution Technology (Intel® TXT)

To read about the NIST PoC, download the NIST publication "Trusted Geolocation in the Cloud: Proof of Concept Implementation" (final version expected to be available in March) at: **http://csrc.nist.gov/publications/drafts/ir7904/draft_nistir_7904.pdf**

For more information about Intel® Trusted Execution Technology and the root of trust, visit **www.intel.com/go/inteltxt**

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/technology/security

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families: Go to: Learn About Intel® Processor Numbers

Trusted Platform Module (TPM): The original equipment manufacturer must provide TPM functionality, which requires a TPM-supported BIOS. TPM functionality must be initialized and may not be available in all countries.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

*Other names and brands may be claimed as the property of others. Printed in USA 0213/SLIO/KC/PDF ♻ Please Recycle 328168-001US