

TRUSTED GEOLOCATION IN THE CLOUD

The National Cybersecurity Center of Excellence's building blocks are example cybersecurity implementations that address technology-adoption gaps across sector boundaries.

PROBLEM

While cloud computing offers businesses and other organizations cost savings and flexibility, these shared resources can introduce security and privacy challenges. Enterprises that use cloud services want to be assured that:

- the cloud compute platform hosting their workload has not been modified or tempered
- sensitive workloads on a multi-tenancy cloud platform are isolated within a logically defined environment from the workloads of competing companies
- workload migration occurs only between trusted clusters and within trusted data centers
- cloud servers are located in their preferred regions or home countries so that the cloud provider is subject to the same data security and privacy laws

Unfortunately, traditional geolocation, the process for asserting the integrity of the cloud compute hardware and enforcing the physical location of an object, is based on operational security control: The method is not secure and does not lend itself to automation and scaling of the cloud computing platform.

SOLUTION

To meet these business needs and help accelerate the adoption of cost-saving cloud technologies, the National Cybersecurity Center of Excellence (NCCoE) is collaborating with Intel and RSA on the Trusted Geolocation in the Cloud building block. This automated "hardware root of trust" determines the integrity of the compute hardware and restricts the workloads to cloud servers within a location. The hardware root of trust is a tamper-proof combination of hardware and firmware—deployed by a cloud service provider, business or organization using cloud services—with a unique identifier for the cloud server host and metadata about the server platform. Using secure protocols, a business can access this information to find out if the platform is still as it was when first deployed, determine the location of the cloud server, and enforce geolocation-based restrictions.

This building block relies on trusted compute pools: segregated portions of the cloud with security requirements that match the value and sensitivity of the workloads they contain. We have demonstrated the use of RSA Archer commercial enterprise and risk management solutions and dashboard reporting to assess these pools, or compute nodes, by:

- measuring them at launch time
- monitoring them continuously
- determining their physical location and levels of compliance

This proof of concept will help businesses feel confident in the security of virtual workloads in a hosted infrastructure and cloud technology. The Trusted Geolocation in the Cloud implementation has been published as draft NIST Interagency Report 7904.

The National Cybersecurity Center of Excellence at the National Institute of Standards and Technology addresses businesses' most pressing cybersecurity problems with practical, standards-based solutions using commercially available technologies. The NCCoE collaborates with industry, academic and government experts to build modular, open, end-to-end reference designs that are broadly applicable and repeatable.

LEARN MORE ABOUT NCCOE
Visit <http://nccoe.nist.gov>

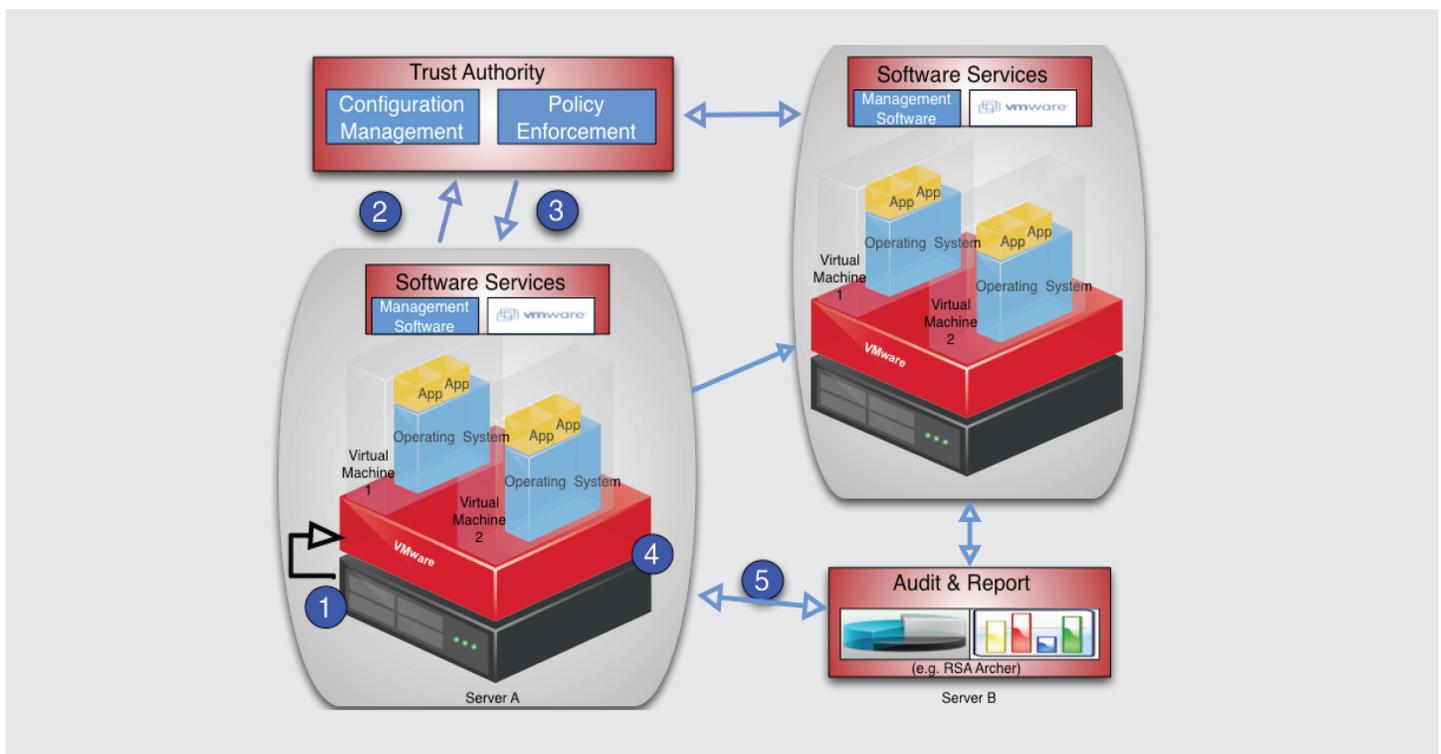
CONTACT US
nccoe@nist.gov
301-975-0200

STAGE 0	STAGE 1	STAGE 2
Platform attestation and safer hypervisor launch	Trust-based homogeneous secure migration	Trust-based and geolocation-based homogeneous secure migration
<ol style="list-style-type: none"> 1. configure a cloud server platform as being trusted 2. before each hypervisor launch, verify (measure) the trustworthiness of the cloud server platform 3. during hypervisor execution, periodically audit the trustworthiness of the cloud server platform 	<ol style="list-style-type: none"> 1. deploy workloads only to cloud servers with trusted platforms 2. migrate workloads on trusted platforms to homogeneous cloud servers on trusted platforms; prohibit migration of workloads between trusted and untrusted servers 	<ol style="list-style-type: none"> 1. have trusted geolocation information for each trusted platform instance 2. provide configuration management and policy enforcement mechanisms for trusted platforms that include enforcement of geolocation restrictions 3. during hypervisor execution, periodically audit the geolocation of the cloud server platform against geolocation policy restrictions

HIGH-LEVEL ARCHITECTURE

Generic steps performed in the operation of the solution:

1. Server A performs an Intel Trusted Execution Technology (TXT) measured launch, with Intel TXT populating the platform configuration register (PCR) values.
2. Server A sends a Trusted Platform Module (TPM) quote to the Trust Authority. The TPM quote includes signed hashes of the BIOS, TBOOT, ESX and geotag values.
3. The Trust Authority verifies the signature and hash values and sends an authorization token to Server A.
4. Server A's management layer executes a policy-based action (in this case, a virtual machine transfer to Server B).
5. Server A and Server B get audited periodically based on their PCR values.



Visit <http://csrc.nist.gov/nccoe/Building-Blocks/Trusted-Geolocation-in-the-Cloud.html>

Any mention of commercial products is for information only; it does not imply recommendation or endorsement by NCCoE or NIST.