ITL BULLETIN FOR FEBRUARY 2016

IMPLEMENTING TRUSTED GEOLOCATION SERVICES IN THE CLOUD

Michael Bartock, Karen Scarfone,[1] and Larry Feldman,[2] Editors
Computer Security Division
Information Technology Laboratory
National Institute of Standards and Technology
U.S. Department of Commerce

Introduction

Organizations that use Infrastructure as a Service (IaaS) cloud computing technologies sometimes need to take the geographic locations of cloud servers into account. For example, each country has its own laws regarding data security and privacy; some of these laws may require organizations to ensure that data on cloud servers remain within national borders. Enforcing geographic boundaries for IaaS cloud workloads requires new approaches. To address these challenges, the Information Technology Laboratory has released NIST Internal Report (NISTIR) 7904, *Trusted Geolocation in the Cloud: Proof of Concept Implementation*. This publication proposes the use of a hardware root of trust method for maintaining the integrity of geolocation information for cloud servers.

Challenges with Geolocation Services in the Cloud

Shared cloud computing technologies are designed to be highly agile and flexible, transparently using whatever resources are available to process workloads for their customers. This makes the use of cloud computing a major cost savings for many customers when compared with traditional data center technologies. However, as discussed above, there are often serious security and privacy concerns with allowing unrestricted workload migration. As a result, an organization may want to address these concerns by restricting which cloud servers it uses based on the geographic location of each server. This might be as broad as excluding servers in certain countries, or it might be as narrow as only allowing servers in one legal jurisdiction. It might also involve restricting workloads so that servers, and the data they contain, are physically located in the same country of origin as their data.

Being able to restrict the physical location of cloud workloads obviously depends on determining the approximate location of each cloud server. This process is known as geolocation. Geolocation can be accomplished in many ways, with varying degrees of accuracy, but traditional geolocation services are not secured. This leaves them susceptible to tampering and other malicious actions that could produce false results and allow workloads to be migrated to unacceptable locations. A related problem is that traditional geolocation services are enforced through management and operational controls that cannot be automated and scaled. In cloud environments, which are highly dynamic, the lack of automation is a huge hindrance and does not allow for seamless workflow migration. Therefore, traditional geolocation services cannot be trusted to meet cloud security needs; new trusted geolocation services are needed.

[1] Karen Scarfone is a Guest Researcher from Scarfone Cybersecurity.
[2] Larry Feldman is a Guest Researcher from G2, Inc.

Overview of the Proposed Solution

The primary purpose of NISTIR 7904 is to propose a potential solution for establishing and using trusted geolocation services. This solution has been tested through a proof of concept implementation designed to determine whether or not the potential solution: a) addresses the weaknesses inherent in traditional geolocation services; and b) meets the requirements to enforce workload locations in cloud environments. The publication provides sufficient details about the proof of concept implementation so that organizations can reproduce it if needed for their own purposes. It is important to note that the proof of concept implementation presented in NISTIR 7904 is only one possible way to solve the challenges of geolocation services in the cloud.

The proposed solution for enforcing and monitoring geolocation services is to establish a hardware root of trust method. A *hardware root of trust* is an inherently trusted combination of hardware and firmware that maintains the integrity of the geolocation information and the platform. The hardware root of trust is seeded by the organization, with the physical server's unique identifier and platform metadata stored in tamper-resistant hardware. This information is accessed by management and security tools using secure protocols to confirm the physical location of the server.

This solution also verifies the integrity of the cloud server's platform before allowing a workload to be migrated to it. Obviously, if the server's integrity has been compromised, workloads should not be migrated to it because of the high probability of compromise of the workloads themselves. In terms of geolocation, the server's integrity is particularly important because without it, there is no assurance that the server is actually in the physical location that it claims to be.

Stages for Solution Implementation

Implementing the entire solution for enforcing and monitoring trusted geolocation services in the cloud can be accomplished in stages. Each stage builds on the previous stage(s) to eventually reach the ultimate goal, which is to be able to enforce geolocation requirements for all workload migration within a cloud computing environment. NISTIR 7904 defines three stages:

- **Stage 0: Platform Attestation and Safer Hypervisor Launch.** This stage ensures that the cloud workloads are run only on trusted server platforms. Key steps within this stage include the following:
    - *Configure each cloud server platform as being trusted.* This includes the server's hardware and hypervisor; it also includes the host operating system if the hypervisor is running on top of one instead of directly on the hardware. If the hardware, host operating system, or hypervisor has already been compromised, obviously that needs to be fully addressed before proceeding with this stage.

- *Measure the trustworthiness of the cloud server platform before each hypervisor launch.* This involves verifying the configurations set in the previous step to ensure that the assumed level of trust in still in place.

- *Periodically audit the trustworthiness of the cloud server platform during hypervisor execution.* This is basically the same as the previous step, but it is performed on a regular basis during execution.

- **Stage 1: Trust-Based Homogeneous Secure Migration.** This stage allows cloud workloads to be migrated among homogeneous trusted server platforms within a cloud. For the purposes of this publication, homogeneous cloud servers are those that have the same hardware architecture and the same hypervisor type, and that reside in the same cloud with a single management console. The goal for this stage is to ensure that a workload is migrated only to servers that have the same level of security assurance as the source server. Key steps within this stage include the following:

  - *Ensure that a cloud server has a trusted platform before deploying a workload to it.* This would rely on the steps from stage 0, measuring platform trustworthiness before allowing a workload to be placed on the cloud server.

  - *Migrate workloads on trusted platforms to homogeneous cloud servers on trusted platforms; prohibit migration of workloads between trusted and untrusted servers.* If a workload has been deployed to a trusted platform, the level of assurance can be sustained only if it is migrated to hosts with comparable trust levels. So this is built upon stage 0 performed on both the workload's current server and the destination server. Only if both servers pass their audits can the migration be permitted to occur.

- **Stage 2: Trust-Based and Geolocation-Based Homogeneous Secure Migration.** This stage allows cloud workloads to be migrated among homogeneous trusted server platforms within a cloud, taking into consideration geolocation restrictions. Key steps within this stage include the following:

  - *Have trusted geolocation information for each trusted platform instance.* This information would be stored within the cloud server's cryptographic module (as a cryptographic hash within the hardware cryptographic module) so that it could be verified and audited readily.

  - *Provide configuration management and policy enforcement mechanisms for trusted platforms that include enforcement of geolocation restrictions.* This builds upon and enhances stage 1 by adding a geolocation check to the server to which the workload is migrated.

- *During hypervisor execution, periodically audit the geolocation of the cloud server platform against geolocation policy restrictions.* This is built upon stage 0, but it is specifically auditing the geolocation information against the geolocation policies to ensure that the policies are not being violated.

Conclusion

NISTIR 7904 explains the security and privacy challenges inherent in using IaaS cloud computing technologies that may cross geographic boundaries. For various reasons, including compliance with security and privacy laws and regulations, an organization may need to restrict in which geographic locations its data may be placed. In IaaS cloud environments, migrating a workload within a cloud requires the implementation and use of an automated solution for verifying the integrity of a cloud server and its geographic location, determining if this location is acceptable for the workload based on the workload's geolocation policy, and permitting migration of the workload only after these checks are passed.

The trusted geolocation service proposed in NISTIR 7904 is based on a hardware root of trust. This involves using trusted hardware and firmware for each cloud server, configuring each server platform to be trusted, and measuring platform trustworthiness before migration, as well as periodically during hypervisor execution, to ensure that the platform stays trustworthy. Applying geolocation policies on top of this allows automated enforcement and verification of compliance with geolocation requirements for all workloads on homogeneous trusted server platforms within a cloud. The proof of concept implementation presented throughout the rest of NISTIR 7904 shows how the proposed trusted geolocation service can be implemented on real-world cloud architectures to demonstrate that the service is a feasible solution for achieving trusted geolocation in the cloud.

In addition, this publication provides the following:

- Details on hardware-based security functions used to establish and enforce trusted geolocation services;
- Supplementary information on setting up your own proof of concept implementation of the trusted geolocation services solution; and
- The security controls from NIST Special Publication (SP) 800-53 Revision 4 and the subcategories from the NIST Cybersecurity Framework that most affect trusted geolocation services in cloud computing environments.

Additional Resources

NIST SP 800-53 Revision 4, Security and Privacy Controls for Federal Information Systems and Organizations

NIST SP 800-125, *Guide to Security for Full Virtualization Technologies*

NIST SP 800-147B, *BIOS Protection Guidelines for Servers*

Draft NIST SP 800-155, *BIOS Integrity Measurement Guidelines*

Framework for Improving Critical Infrastructure Cybersecurity, Version 1.0

ITL Bulletin Publisher: Elizabeth B. Lennon
Information Technology Laboratory
National Institute of Standards and Technology
elizabeth.lennon@nist.gov