Towards Trusted Location Specific Information for Cloud Servers

Leo Hippeläinen and Ian Oliver

Nokia Bell Labs, Finland
E-mail: {leo.hippelainen; ian.oliver}@nokia-bell-labs.com

Received 8 September 2017;
Accepted 10 October 2017

Abstract

Every physical datacenter is located somewhere on the globe. It is subject to the local legislation, including data protection related laws. A cloud service can be delivered from a set of datacenters in several locations. Responsibilities of the service provider include ensuring that legal and agreed constraints are respected also by its subcontractors, for example, those providing cloud computing resources. Several countries have data protection legislation that restrict sharing copies of sensitive data to locations that do not have compliant legislation. This paper presents ideas to dependably detect location specific information, like the legislation properties, of the current physical host server executing a service.

Keywords: Datacenter design, Trusted cloud geolocation, Data sovereignty, Privacy, Confidentiality, Data integrity, Data protection.

1 Introduction

During summer 2015 a government administration in a European country outsourced its databases to an international cloud service provider to save costs [1, 2]. The provider migrated the databases and their administration to another country. About one year later it was discovered that there were insufficient precautions in place to protect sensitive national data from being

Journal of ICT, Vol. 5_1, 1–38.
doi: 10.13052/jicts2245-800X.511

disclosed to unauthorized parties. To prevent similar cases, not only general awareness of data protection issues must be improved but also means are needed to support surveillance of geographical locations of services running on computing clouds.

Knowing geographical location or jurisdiction of a servicing cloud datacenter is important to cloud customers that own sensitive data [3] and to privacy enforcement authorities [4]. The data sovereignty concept pinpoints the applicable data protection legislation, which impacts on data security. On one hand, data protection laws restrict legal geographical locations of data instances to those that have compliant legislation [5]. On the other hand, laws protecting societies against threats, like terrorist acts, may allow authorities to legally investigate suspicious foreign data stored in or bypassing territory of their country.

Geographical issues are caused by the basic characteristic of cloud computing: smooth migration of computing and storage workloads within and among the datacenters. With an ordinary cloud based service the actual location of the hosting server is not an issue as long as service level agreement (SLA) terms become fulfilled and customer gets what she/he has subscribed for. However, application with location bound data shall not be freely migrated across datacenters located in different jurisdictions.

Mainstream geolocation methods are focusing on verifying geographical location of mobile clients, not servers. Consequently, traditional positioning methods for roaming targets are not well suited for attesting location of servers, which are mostly stationary. In case of cloud computing, hardware is not migrated but workloads.

Network Function Virtualization (NFV) enables telecommunication (telco) networks increasingly being implemented using regular datacenter hardware. Security and dependability requirements for telco are often more demanding than those for traditional datacenter applications. Special attention must be paid on incorporating mechanisms into NFV to support high availability computing and data protection at levels which are sufficient for the criticality of telco networks [6].

A classic example of the requirements of geolocation related to service provision and data storage can be found in the case of Lawful Intercept (LI) where access to data is granted based on a Judge's (or similar authority) orders [7]. In this case, access and subsequent storage and processing of data is wholly within a nation's physical borders. For telecommunications operators of all kinds, this places strict requirements on the physical placement of computing resources, including virtual ones.

In an ideal case, a cloud customer could fully trust that the cloud provider is honest and open about where his servers are located and how they are allocated. However, in case of any doubt, trustful, real-time verification of the claimed geographical location of the service providing host computers by the customer is intriguing. This is the problem we try to resolve.

In general, data records are stored to several locations during their lifetime. They have not only the primary storage location but also replica locations. During processing, there are copies in work memory and CPU registers. Moreover, data is transmitted between the storage devices and processors, and between datacenters, in case data server and compute server are in different locations. Adversaries can make illegal copies of data in transit.

In this paper, we extend upon previously published work [5], which focused on investigating the legal obligations to knowing the location of data in cloud computing applications and on applicability of existing technology to resolve the problem. In the earlier work, we also identified the key stake holders and design constraints. In this paper, we provide requirements and use cases for a location specific information delivery system and elaborate technical solution to deliver dependable site-specific information for cloud computing servers and applications. Previously published work [8] is a shorter version of an earlier, unpublished edition of this paper.

The rest of this paper is organized as follows: Section 2 introduces terms and concepts of the domain. Section 3 documents needs of the stakeholders and subsequent requirements for dependable location specific information systems. Section 4 derives prerequisites that are necessary for resolving the challenge. Section 5 presents a set of potential solutions, and Section 6 presents evaluation results. Finally, Section 7 has conclusions of the work and recommends future efforts.

2 Terms and Technologies

2.1 Concepts

To be able to discuss about this domain we must share the basic concepts. This section presents annotated list of the essential concepts.

2.1.1 Trust

There are many definitions for trust. We can define trust intention as "the extent to which one party is willing to depend on the other party in a given situation with a feeling of relative security, even though negative consequences are

possible" [9]. In computer systems context, the "intention" and "feelings" should be understood as those of the system developers. Correspondingly, essential "given situations" and criteria for "relative" should be documented in system specifications.

Usually trust between systems is based on some evidence about the identity of the other party and its commitment to act honestly. Various evidence categories, for example cryptographic key certificates, must be identified and specified at system development time.

In this paper word "dependable" is used as a synonym to "trustable", meaning something you want to trust on. Depending on the context, trust can be defined less generally, like in trusted computing.

2.1.2 Trusted computing

Trusted computing is defined as the use of a computer when there is confidence that the computer will behave as expected [10]. Consistency is based on knowing exactly, what are the hardware configuration and executing software binaries. This is achieved with chain of trust, which is explained as the next concept.

Trusted computing may be implicitly associated with correct behaviour of the system in all situations. However, the guarantee covers only that the behaviour is known and consistent, not that it is always correct. Trusted computer technology does not guarantee that executing programs are free from vulnerabilities but at least the exact variants of the binaries and, hopefully, their problems are known. Good development quality control practices should be used to create as correct as possible software for highly trusted systems.

2.1.3 Chain of trust

In the trusted computing context, the chain of trust refers to mechanisms which establish trust on the next system layer through validation performed by the already running hardware and software [11]. The root of the chain (root-of-trust) is permanently stored to computer motherboard so that it cannot be changed. Similar approach can be applied to security certificates.

With computer systems, the root-of-trust can a piece of initial boot code, which is run in a special processor mode to check that firmware is not tampered with. Firmware then checks the master boot record from disk before starting it, etc. Checks are usually based on measurements of the binaries and configuration data, which in practice means computing a hash value of the binary entity and comparing the result with known good values. Application layer binaries can be checked, for example, using digital signatures, which are also based on hash functions.

In trusted computing, the root-of-trust can be implemented using hardware conforming to TPM standard [12]. One such implementation is Intel TXT [13].

2.1.4 Attestation

Attestation is the act of showing or providing evidence that something is true. Attestation in the context of trusted computing is used to verify that the computed hash values are good, i.e., the measured entity is what it claims to be. Attestation can be organized as a separate virtual function which is executed in a server of the cloud computing system. Obviously, attestation server itself must comply with trusted computing requirements.

2.1.5 Cloud computing

Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management efforts or service provider interaction [14].

Cloud computing is typically more cost efficient than in-house datacenters because computing resources can be shared between several customers and the customers can pay only for those resources they actually need. Although administration of a cloud datacenter is more complex the advantages, especially costs and flexibility, enable it to win market share over conventional datacenters. Support for geographical location specific information is one of the new kind of problems that cloud computing creates when it is applies to processing and storage of sensitive data.

2.1.6 Jurisdiction

Jurisdiction refers to the territory within which a court or government agency may properly exercise its power [15]. Each country or a state of a federation can have its own laws, regulations and authorities. The geographical location of the datacenter site defines the jurisdiction in which the datacenter is and, consequently, which laws are applied to its operation. In case of digital information, the interesting laws are those regulating data protection and privacy.

2.1.7 Data sovereignty

Data sovereignty is the concept that information, which has been converted and stored in binary digital form, is subject to the laws of the country in which it is located. Many of the current concerns that surround data sovereignty

relate to data that is stored in a foreign country from being subpoenaed by the host country's authorities or some malicious actors, because prevalent laws do not set prohibiting enough punishments or because monitoring is not strong enough [16].

2.1.8 Data residency

Data residency refers to the physical or geographic location of an organization's data or information. Like data sovereignty, data residency also refers to the legal or regulatory requirements imposed on data based on the country or region in which it resides [17].

2.1.9 Trusted location specific data

Trusted geolocation refers to knowledge about geographical location of a server or datacenter, for which there is sufficient evidence to establish trust that the location information is correct. The evidence must cover both the source of the information and the information delivery chain of hardware and software components towards the inquirer. Global location coordinates can be mapped to the name of the corresponding country or even properties of the legislation in the country. These pieces of information could be made directly available per datacenter site as location specific information.

2.2 Basic Entities

In this section, the very basic definitions related to a datacenter are listed and explained from the viewpoint of this paper.

2.2.1 Datacenter

A datacenter is a facility used to house computer systems and associated components, such as telecommunications and storage systems, backup power supplies, redundant data communications connections, environmental controls (e.g., air conditioning, fire suppression) and various security devices [18]. In this paper, a datacenter simply denotes a set of interconnected servers and storage systems with communication channels to the Internet.

2.2.2 Datacenter site

The physical place of the datacenter. A site has a geographical location, which can be identified with global coordinates (longitude, latitude) or by the name of the location, which must be unique within the name space.

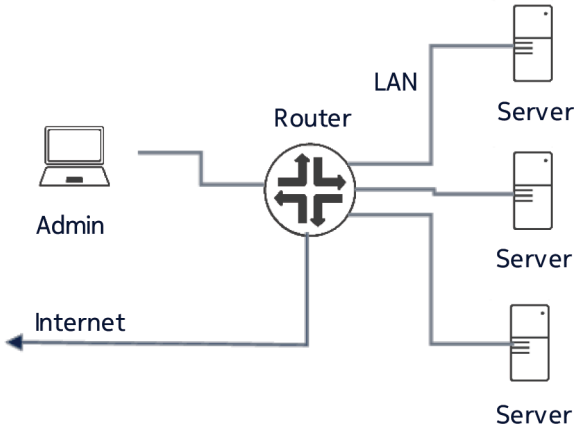

Figure 1 Elements of a simplified datacenter.

2.2.3 Local Area Network

A local area network (LAN) is a fast data transmission network to interconnect computers in a small local area, like a datacenter. Each computer has a physical adapter, a network interface card (NIC), which connects the computer via a data cable to a LAN router or switch. A computer can have one or several NICs that are connected to the same or different LANs.

A router is a cross-connection of LAN cable ports. It forwards incoming data packets to outgoing ports based on the destination address in the data packet header. Usually a LAN has several interconnected routers and other network equipment. Based on this definition Software Defined Networking (SDN) can be considered as a router technology.

2.2.4 Server

A server is a computer with processor cores, random access memory, LAN connectivity and control circuitry. Often a server also has mass memory (e.g., disk drive), a display port and sockets for options, like Trusted Platform Module (TPM) chip. More about server categories in the next subsection. Servers can be optimized for computing or providing storage space.

2.3 Stakeholders

There are various parties who are interested in knowing the geographical location of cloud servers. The most essential ones and their interests are listed in this section.

2.3.1 Cloud Service Customer

A cloud service customer (CSC) employs compute and storage resources made available by the contracted cloud service providers for executing cloud based services implemented by the customer's application software. Cloud service customer may have legal obligations to take care of that the service is provided by servers within a certain legislation.

2.3.2 Cloud Service Provider

A cloud service provider (CSP) makes available a set of compute and data storage servers to host application software. The hardware can be in several datacenter sites at different geographical locations on the globe. Cloud provider wants to maximize profit and sales of his/her cloud resources by offering competitive value for the money to cloud customers. Cloud provider also wants to share computing workloads optimally to the servers across his/her datacenter network.

2.3.3 Cloud Service End User

A cloud service end user (CSEU) is a human or an automated actor who consumes functions of a service provided by a cloud customer. While utilizing the service the end user may enter proprietary information to the system.

2.3.4 Third-Party Auditor

A third-party auditor (TPA) is an independent auditor trusted by both cloud provider and cloud customer who have business relationship. External Auditor can perform formal audits to datacenter sites and produce audit reports about conformance of cloud provider's services and processes against applicable laws, regulations and SLAs. External Auditor appreciates easiness of auditing geographical locations of physical servers of Cloud Provider.

3 Requirements for Location Specific Data Support

3.1 Data Protection Legislations

The wide-spread adoption of cloud computing services, as well as new approaches to data storage including object storage, have broken down traditional geopolitical barriers more than ever before. In response, many countries have regulated new compliance requirements by amending their current laws or enacting new legislation that requires customer data to be kept within the country the customer resides [16, 19].

Verifying that data exists only at allowed locations can be difficult. It requires the cloud customer to trust that their cloud provider is completely honest and open about where their servers are hosted and adhere strictly to service level agreements (SLAs) [16].

Per OECD [20] over sixty countries had adopted by year 2011 data protection or privacy laws that regulate transborder data flows. By year 2014 the number of countries has increased over 100 [21].

European Union (EU) General Data Protection Regulation (GDPR) [22] defines rules for the protection of the fundamental rights and freedoms of natural persons regarding the processing of personal data. Article 1 says that "The protection of natural persons in relation to the processing of personal data is a fundamental right." It sets limitation to where data can be stored and when it should not be kept anymore.

CDPR Articles 33 and 34 mandate data owner to report personal data breaches to supervisory authorities and to the subject without undue delay.

GDPR and OECD recommendations [23] encourage transborder data flows, if all countries on the way have compatible legislation. On the other hand, governmental surveillance laws in some countries are in conflict with GDPR and thus it is not acceptable to store and process the personal data of EU citizens in those countries.

In Australia and New Zealand, The Privacy Amendment Act [24] allow personal information to be transferred, but the sender must ensure the recipient will comply with the Australian Privacy Principles (APPs). Even then the Australian sender remains liable for the recipient's behavior in this context.

In the Peoples Republic of China, the Cybersecurity Law was enacted by the Standing Committee of the National People's Congress of China on November 7, 2016 [25]. According to the law, operators of key information infrastructure (e.g. communication) must retain, within the territory of China, critical and personal information which they collect and produce during their operations in China. They may still be able to transmit this information overseas, but only after undergoing and passing a security review. These operators are also required to undergo a network safety assessment at least once a year.

In Russia, personal data localization requirements implemented by the Amendments to the Personal Data Law as of September 2015 mandate all personal data of Russian citizens to be stored in databases that reside in the territory of the Russian Federation [26]. Personal data can still be duplicated to servers outside Russian borders, if other Russian laws regarding personal data are followed. The law does not restrict remote access to databases located

in Russia. However, the legislation is not quite exact and leaves room for interpretation.

In the United States of America (USA), there is no single, comprehensive federal (national) law regulating the collection and use of personal data [27]. However, there are many government policies and regulations that deal specifically with data privacy and residency issues. The Health Insurance Portability and Accountability Act (HIPAA) is a data privacy and security law designed to protect medical information. Payment Card Industry Data Security Standard (PCI DSS) is a set of policies to secure credit and debit cardholder information. [17] Compromised private data must be informed to victims.

From the legislations, we can summarise these needs:

- All data in transit or at rest must be routed and stored in geographical areas which have legislation that is compatible with that of the origin of the data.
- In certain countries, at least one copy of the protected local data must reside within borders of that country.
- Possible data breaches must be detected quickly.

Because in all legislations the responsibility of conforming data protection laws is by the data owner, we must convert the needs to system requirements that help the sensitive data owner to monitor application and data locations. The requirements will be listed in the context of the cloud service customer in Section 3.3.1.

3.2 Commercial Constraints

3.2.1 Per server cost impact

A datacenter site can have thousands or even millions physical servers, and a cloud system may extend to several datacenter sites. Price of a server is assumed to be in range of 500 to 10,000 euros. Even a small cost increase multiplies to considerable total amount, especially if existing servers need to be equipped with new parts. Thus requirement:

- REQ-1.1: Support for dependable location specific information should not add anything to the manufacturing costs of a trusted server hardware.

3.2.2 Provisioning costs

Each physical server is commissioned for use. Setting configuration parameter should be automated to avoid costly and error prone manual steps. Support for dependable location specific information quite probably adds some settings

or secrets to the onboard configuration data of the server but they should be automated.

- REQ-1.2: Dependable location specific information should not add need for new manual steps in the new server deployment process.

3.3 Stakeholder Needs

This section outlines requirements from the perspectives of the various stakeholders. Let's start with the stakeholder which is probably the most interested in trustable location specific information: the cloud service customer.

3.3.1 Cloud service customer needs

Cloud service customer has some geographical location sensitive application. The application executables may be under embargo ruling but most often it is the processed data that should not be stored or transporter via noncompliant jurisdictions. Even if the service provider can be trusted there still can be justification for extra assurance effort. Unintentional human errors may occur, which can cause SLA terms becoming violated or open opportunities for data breaches.

The legal obligations explained in Section 3.1 can be mapped to following requirements:

- REQ-2.1: A cloud application must be able to undeniably detect in real time location specific information applicable to the employed physical server.
- REQ-2.2: It must be possible to check afterwards from log records which physical servers were allocated for running the application of a cloud customer.
- REQ-2.3: A cloud application must be able to undeniably detect in real time the employed physical transmission connections and equipment.
- REQ-2.4: It must be possible to check afterwards from log records which physical connections were used for transferring a certain piece of sensitive data.

If a cloud software application can read a piece of location specific information then the data owner can see it by running a suitable software in the server. Integrating this information inquiry functionality to every location sensitive application gives certainty that retrieved information indeed applies to the service the application provides.

Typically, several cloud customers share resources of cloud datacenters. The workloads are scheduled to servers hosts by the cloud computing manager.

To discover wrong allocations, it would be useful if the customers could audit datacenter sites and drop a certificate to the sites qualified for their use. These certificates become pieces of location specific information. Then at run-time, a location sensitive application should be able to check if the datacenter or server is certified and report possible violations.

- REQ-2.5: It should be possible to mark a datacenter site to be certified for storing and/or processing protected data of a certain application.

3.3.2 Cloud service provider needs

An honest cloud service provider has need to satisfy geographical restrictions agreed with its cloud customers. Behaving differently would harm future business. If the cloud provider has more than one datacenter site, he wants to optimize physical server resources across server pools. To use the location information as an input to resource scheduling, the location of every server must be known at runtime. REQ-2.1 above can be utilized for creating a server inventory, which then is used for workload scheduling.

The service provider doing everything right is still not enough. The cloud service customers and auditors must also be convinced that the agreements are indeed fulfilled. The cloud provider must have mechanisms and procedures in place to prove that geographical constraints are and were conformed to. The evidence in practice can be implemented as trusted logging, which implies that cloud provider must support implementations of REQ-2.2 and REQ-2.4 above.

To support implementing REQ-2.3 the cloud service provider should maintain real-time inventory of active transmission connections. This is probably challenging and in practice there may be more cost-efficient means to protect data in transit, like encryption.

3.3.3 Cloud service end user needs

If the cloud customer is a SaaS provider, then the critical data is owned by end users of the application. Now the end users become interested in where the data is physically stored. Thus, the cloud service provider must make possible for the end users to detect server locations. However, this case does not add technical challenge compared to REQ-2.1 . . . REQ-2.4 above.

3.3.4 Third-party auditor needs

If there is lack of trust between cloud provider and cloud customer, an external auditor can be contracted to produce necessary testimonial about geographical

locations of servers. This delegates the challenge of acquiring dependable location information to an external auditor.

The auditor faces an obvious problem: How to audit, with reasonable effort, the locations of thousands of physical servers in several datacenters? The exact cabin positions are not needed. It is enough to enlist in which datacenter each server resides. Nevertheless, if location specific information is stored to each server, correctness of that must be audited.

- REQ-3.1: An independent auditor must be able to verify datacenter site and other location specific data associated with each trusted physical server.

Third-Party Auditor can document the site information of servers at the time of the audit. However, servers can be removed or added or moved between datacenters after the audit. The server inventory database must be maintained preserving the same level of dependability as what the auditor achieved in his initial audit.

- REQ-3.2: The location specific information attestation database must be maintained in real time in a dependable manner.

Again, also the intersite transmission connections and cables should be audited and listed, including their geographical routes.

- REQ-3.3: Transmission connections and routes between datacenters must be listed and audited from jurisdictions point of view. The list must be maintained if transmission contracts are changed.

3.4 Requirements from Location Information Cheating Mitigation

A dishonest cloud provider may try to report false location specific information to hide his slippage from the SLA terms [28]. The provided solutions shall prevent or, at least, reveal all tricks a dishonest cloud provider can come up with to mislead its cloud customers and independent auditors. A complete list of cheating patterns would be useful but research of this topic is still in its infancy. Here are some examples:

3.4.1 Location specific information forged

For example, if location evidence is based on a GPS receiver, signal feed to it can be spoofed so that the datacenter appears to be elsewhere than it actually is [29].

Another technique could be that the software functions delivering location specific information are not performing according to specifications and return misleading data.

Third method could be to provision wrong location information. The location specific information can be stored to non-volatile memory of the server hardware, like into a TPM chip register. If the data is set wrongly, the server reports wrong location. The same applies if the configuration data is overwritten by a piece of malware. The latter case can be mitigated by using trusted computing techniques, but still it might be possible if there are vulnerabilities.

If the software stack that deliver responses to location specific information inquiries are faulty or compromised, the inquiring application receives wrong data. Again, trusted computing mechanisms provide an efficient mitigation.

We must insist that a server shall not be able to pretend that it is at a different site than it physically is. All attempts to seduce protected data to illegal jurisdiction using this cheat pattern must be detected and prevented.

- REQ-4.1: Attempts to make a server on a remote site to look like it is located on the local site must be detected.

3.4.2 Migrated server not reconfigured

If location specific information is stored to the server hardware, for example in to a TPM PCR, then that information must be reconfigured if the server is moved to another datacenter. Failing to do this makes the server to look like it is still in its original location. This can be caused by a mistake or a deliberate action.

Usually physical servers are not moved between datacenter sites once having been commissioned. Nevertheless, it can still happen during the lifetime of a server.

- REQ-4.2: If a server is moved to another site, the move is detected and the server is associated with its new datacenter site.
- REQ-4.3: When a server is retired, its location information must be removed.

3.4.3 Migrate back to legal server at the time of location check

The location inquiry is implemented as an API function call, which the client software can call when it wants to retrieve a location specific attribute. It may be possible that the cloud service provider can detect these invocations and migrate the application to a legal server during the time of the inquiry. REQ-2.2 can support revealing also this kind of cheating pattern.

3.4.4 Mobile datacenter

Physical servers are not expected to change datacenter site very often after being commissioned to a site. However, there are small datacenters assembled into a sea container. It is possible to move the whole datacenter.

Prospects to provide dependable location specific information in the case of mobile datacenter are less good than with stationary datacenter site. Perhaps the easiest solution is to declare a mobile datacenter always as untrustable and avoid migrating there any location critical workloads.

4 Prerequisites to the Implementation Platform

In addition to specifying requirements for the dependable location specific systems, assumptions should be made regarding the underlying platform. Otherwise the solution would not be feasible.

4.1 Trusted Software

A dependable data storage is not enough to guarantee that the data is propagated intact through the software stack to the requestor. Server software may have vulnerabilities that enable malicious program code being smuggled in to the server. Therefore, location specific information cannot be considered dependable if malicious code could have changed it. To increase trust in software Trusted Computing Group has created TPM standard [12] which is also published as standard ISO/IEC 11889:2015. Most modern servers have support for TPM chip, either version 1.2 or 2.0. TPM 2.0 API can also be implemented using ARM TrustZone [30] and Intel PTT [31], which have a hardware based root-of-trust and can guarantee that only approved binaries and data are used.

- PRE-1: Only trusted computing servers shall be used for executing services that need dependable location specific information.

4.2 Middleware Enhancements

Data records stored in the repositories of a data server are written to and retrieved from using middleware like a database engine. Consequently, there is always some piece of software which must run in the same server as where the data resides. If that piece of software could also respond to inquiries of location specific data then there are means to find out where the active databases are located.

- PRE-2: Middleware that saves and retrieves data to and from physical storage devices shall be made to support inquiry operation for resolving location specific information from its physical execution environment.

4.3 Server Identifiers

To create allocation log records each server must be identifiable by software and by a human observer. The identifier can be, for example, a serial number or associated PKI key, which remains the same during lifetime of the server. The identifier must be available at run-time to software. If the identifier is changed, the server is considered to be a different one.

- PRE-3: Each physical server has computer and human readable unique identifier.

4.4 On-Site Auditing

An auditor, either third-party or a representative of the cloud service customer, must be allowed to visit a datacenter site and personally verify, for example, that a certain server hardware exists at that site.

- PRE-4: Independent auditor can visit the datacenter premises to observe its geographical location and seal some computer readable evidence to the site. This can be dedicated configuration of physical servers or security module at the site.

4.5 Logging Mechanism

Unfortunately, misbehaviour cannot always be detected in real time with reasonable costs. In many cases it is enough if the case can be investigated by combining evidence gathered at runtime. For example, a location inquiry may raise suspicion that illegal cloud server is being randomly used. By investigating log records of that time should give confidence whether or not there is enough reason for punishments.

- PRE-5: Trustable logging mechanism is provided by the underlying cloud computing platform.

5 Methods of Location Detection and Specific Data

In general, location detection problem can be split to two parts: 1) detecting that a group of servers is at the same site and 2) having trustable mechanism to knowing location of the site.

An earlier survey of geolocating techniques [5] lists many alternatives but most do not qualify, because they add to server hardware manufacturing costs (REQ-1.1).

The solution to the locating problems should be such that it cannot be influenced by any stakeholder. To convince critical customers geographical trust should be based on real world phenomena, like speed of light or physical proximity.

For example, we could use Bluetooth wireless communication to enable servers in a datacenter to connect to the same personal area network. If they can then we know that they must be at the same site because Bluetooth signal has maximum range of 400 meters [32]. Unfortunately, Bluetooth has too low limit for the maximum number of nodes to be useful for detecting proximity of servers [33]. Neither do existing servers have Bluetooth radios built in.

5.1 Persistent Onboard Configuration Data

Location information can be stored as configuration data to the persistent memory of a physical server, like a register in onboard TPM chip. The data value must be configured at commissioning phase which may require new manual steps to the server's configuration script. This may violate REQ-1.2.

A suspicious person can think that if data can be written once it may be rewritten later again with a different value. Even if the data is proven to be unmodifiable, the whole server can be moved and, consequently, data becomes outdated and cause REQ-3.2, REQ-4.2 and REQ-4.3 to fail. But as long as the server stays in its original location, location information is valid and dependable.

Our conclusion is that server's onboard configuration data alone cannot offer fully trustable evidence of the location. This solution is also expensive to audit in case the datacenter has hundreds or more servers (REQ-3.1).

Willingness to trust onboard configuration data can be increased if several independent data instances can be compared in real time. For example, there can be means to reliably detect that a group of servers exists at the same site and they all have the same location specific configuration information. Now unauthorized change of the configuration data in one server can be detected. Nevertheless, this is not enough for highly critical location sensitive applications.

A physics based root for establishing geographical trust is needed to reach trust level beyond that achievable with configuration data alone.

This is analogous to having a hardware based root of trust for facilitating trusted computing [12].

5.2 Round-Trip Time Measurements

Physical distance can be approximated by measuring round-trip time (RTT) from sending a request to receiving the corresponding response. Special distance-bounding protocols have been developed for this purpose [34]. By measuring distance from two or more known locations ("landmarks") to the datacenter of interest we can find its position using trilateration. However, global positioning inaccuracy can be in range of 1000 km due to transmission and computing delays and congestion.

RTT measurements can be used also for checking if two servers are in the same datacenter site, because LAN is usually faster than connections between datacenters. Even though the connection cables are not direct lines between servers and the accuracy depend on the OSI layer of the measurement, by defining a reasonable threshold, it should be possible to do the grouping. Still, there can be fast fiber connections able to hide the fact that a set of servers resides in another site. We must trust the independent auditor to review the reasonable RTT threshold value and to discover possible direct fibers, which bypass normal routed network connections.

RTT threshold value can be used to collect hard evidence that a server is physically near some known landmark, i.e., they are connected to the

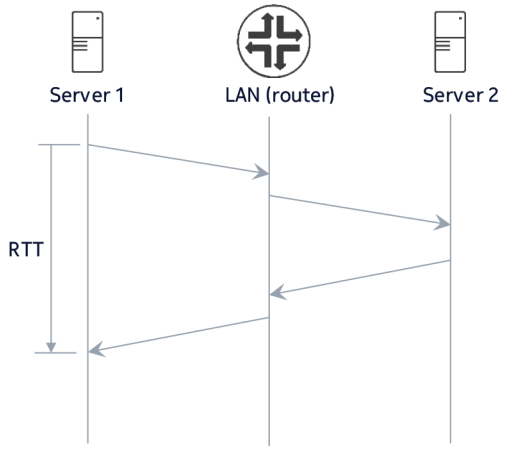

Figure 2 Round trip time between Server 1 and Server 2.

same LAN. This contributes to implementing REQ-2.1 and REQ-3.2 partially, but not in real-time. Server moving or forging location (REQ-4.1 and REQ-4.2) can be supported. Neither additional hardware (REQ-1.1), nor any new manual steps (REQ-1.2) are necessary. A database is needed to fully support of REQ-3.2. A removed server is detected (REQ-4.2 and REQ-4.3) when RTT measurements cannot anymore reach the server.

Server proximity data or map can be calculated from the RTT measurements. This implies maintenance challenge so that, for example, decommissioned servers are cleaned away from the data. Obviously, RTT measurements can and should run in the background continuously to keep proximity data updated.

5.3 Network Topology Discovery

Datacenters typically use wired LAN to connect the servers to each other and outside world. LAN topology can be detected using Link Layer Discovery Protocol (LLDP) [35] and Broadcast Domain Discovery Protocol (BDDP) or with proprietary equipment manufacturer specific protocols [36].

Each network layer can have different topology. It should be possible to see from the link layer topology, which servers are probably at the same site. For more dependable results, RTT measurements needed to find out connection link lengths.

Topology discovery can add value by discovering active transmission links and equipment, like routers (REQ-2.3). Detection is not real-time but delayed to the next discovery round.

Topology discovery should also be integrated with network management functions, like SDN (software defined network) controller, to get notifications concerning topology changes (REQ-4.2). SDN controller should also produce dependable event logs to document changes to flow tables (REQ-2.4). But alone topology discovery cannot support these.

Network implementations based on SDN employ controllers and switches themselves can be virtual machine instances. Even if dedicated equipment is used, they are software and firmware based. If trust on geographical location information is based on mapping network topology from the forwarding and routing tables of the network gear, the dependability of the discovered information should be evaluated. For the time being commercial network devices are not utilizing trusted computing approaches. Consequently, mission critical computing cannot assume that plain text data transmission is dependable nor that information derived from topology discovery can be fully trusted.

5.4 Attestation Service

Attestation service is often used in trusted computing for validating server integrity and to detect unauthorized changes [13]. Good reference (hash) values known by the attestation service originate from software build process and they are delivered to the attestation service using trustable means.

Attestation service could be extended to offer also geographical location specific information applicable to the datacenter servers. A possible result from a site audit is a database of physical servers with their identifiers and locations. If this database is made available to a trusted attestation service which serves all sites of the cloud provider, applications can inquire location of a server from the attestation service. The inquirer needs to know the identifier of the server of interest, which we assume being available at runtime (assumption PRE-3 in Section 4.1.3).

Trustworthiness of attestation service data depends on the source(s) of information. Attestation data could be prepared and delivered by the cloud management. If there are doubts that cloud service provider is not completely honest, we may not want to trust attestation data from this source either.

Data from independent auditor can be assumed to be dependable. However, changes made after the audit must be updated to the attestation service database. It is not practical that the auditor continuously audits the changes and as there is no other dependable source for updates, we can conclude that attestation service is not a possible solution for dependable location specific information.

As already mentioned the server database must be actively maintained. The challenge is that all changes should also be verified by an independent auditor or otherwise the dependability of data will degrade. This is a task that involves human effort and, because of that, maintaining data dependability becomes cumbersome and expensive.

5.5 Seismic Acceleration Measurements

Detecting seismic vibrations would provide an undisputable indication of the detectors being at the same site or different sites. It is assumed that the vibration patterns caused by seismic activity is unique as a function of time per a site. Seismic waves can originate from different sources, natural or artificial. It is well known how different frequencies propagate through the ground. The challenge is to filter seismic accelerations from the vibrations caused by other sources, like fans and read/write heads seeking over disk surface.

Even though datacenter servers do not have acceleration sensors it can be possible that some disk drives do have them. Sensors are used for protecting

the drive against rapid accelerations or to monitor mechanical wear out of the drive. It could be possible to utilize the acceleration sensors to record also seismic vibrations at the datacenter premises. All servers of a site should experience similar wave formats whereas servers at another site are exposed to different seismic wave patterns.

It may even be possible to find out global position of a seismometer by comparing its recordings to the list of signals originating from known centers of seismological events, like earthquakes. It is probably also possible to verify presence of servers at a site by verifying that auditor controlled artificial seismic activity is detected by the accelerometers in the servers.

5.6 GeoProof

GeoProof [28, 37] combines proof-of-storage (POS) protocols with the distance-bounding protocol to verify that certain data exists at a certain datacenter. GeoProof architecture entities include a third-party auditor (TPA), a tamper proof verifier device (V) and a cloud data server (P) Figure 3. The intention is to prove it still has the data file(s) saved to it. The TPA communicates with the V, which is located in the datacenter and connected to its LAN. Geolocation of V is assumed to be dependably known. RTT measurements are used for checking that the distance from V to P is short enough for them being at the same site.

GeoProof is designed to provide a geographic assurance for the data owners, that their data remains in the same physical location specified in the SLA. Thus, it offers solution to a different problem set up than what is outlined in this paper.

GeoProof supports location assurance only if the runtime datacenter is already known. In a case the cloud provider has several datacenter sites within

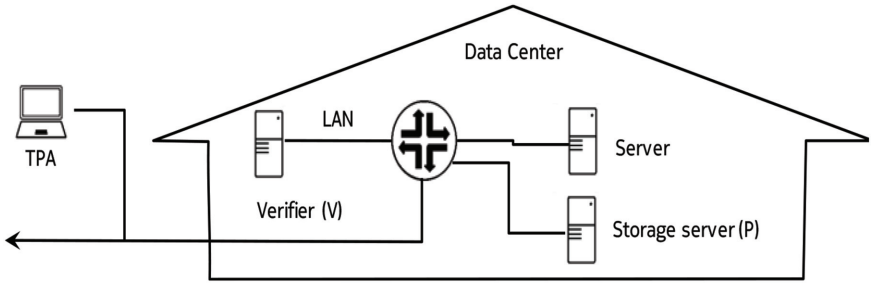

Figure 3 GeoProof architecture.

qualified jurisdiction, it can be possible to migrate customer's files among these datacenters without violating SLA. GeoProof does not offer help to locating the data replicas. Nor does it necessarily reveal an active replica located in a lower cost datacenter if the legal on can prove that it has a copy, too.

GeoProof assumes that there is a verifier device but leaves open how to assure its geographical location with high confidence. Primary purpose of the device is to root the geographical location and to verify proof of storage (POS) or proof of retrieval (POR) from the storage server. To fulfill (REQ-3.1) an auditor can use the verifier device to check if a server is on the same site or not.

5.7 Site Anchor

An independent auditor could, as an alternative to checking existence of servers at a datacenter site, observe a special site anchor device being deployed to a datacenter site. A site anchor is a small trusted computer (in TCG terms) with enough protected memory to store location specific data, like geographical coordinates, name of the site, or identifier of the jurisdiction or legislation at that site. The information is stored to the device under supervision of a trusted auditor.

Figure 4 depicts site anchor based system for storing datacenter specific information. The figure is drawn similar to GeoProof's Figure 3 but behaviour

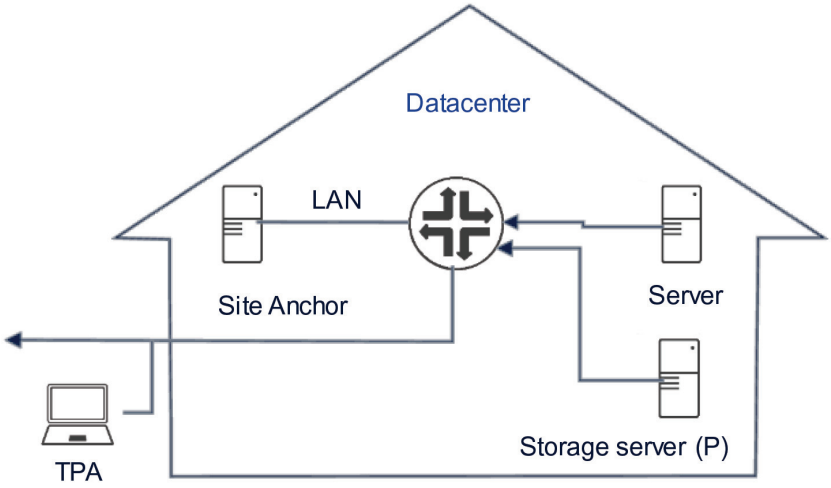

Figure 4 Site anchor architecture.

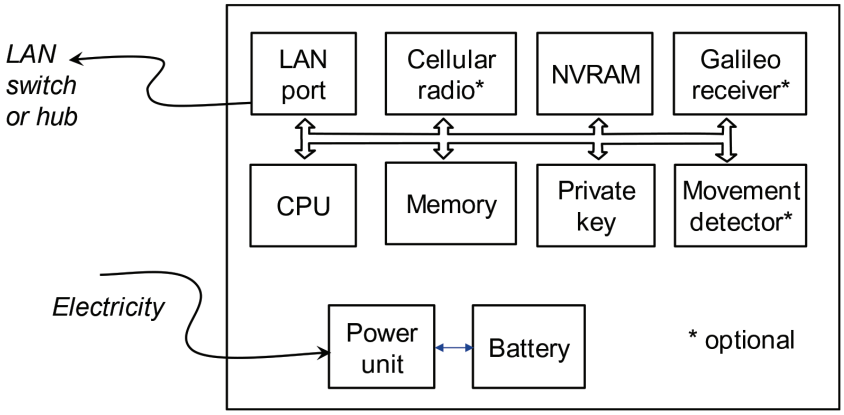

Figure 5 Example functional blocks of a site anchor.

is different. Site anchor device primarily serves the servers, not third-party auditor. Applications can inquire location specific information from the site anchor via an API library running in each of the trusted servers of the cloud.

When new servers are added, they should have configuration data which enables them to communicate with a local site anchor. The needed parameters need not be site specific, they can be cloud provider or customer specific instead. The communication with the anchor should be ciphered with dedicated keys. The main benefit is that servers of a datacenter site need not know the site-specific information, but they retrieve it from the anchor. LAN routing tables should be configured to direct inquiries to the local anchor, which can be assigned the same local address in all datacenters.

A site anchor should be commissioned and sealed to the physical structures of site in such a manner that it cannot be moved or reconfigured after the auditor leaves the premises. A site anchor device can be low cost hardware with support to trusted computing, for example, a credit card sized ARMv8 computer with LAN connection and power over Ethernet.

To avoid directing location specific data inquiries to a site anchor at another site, RTT measurement should be made by the inquiring server. Too slow responses should be sent again a reasonable number of times and if still failing, an exception should be raised to alarm management.

A single location anchor per site can be an unreliable solution. However, typical high-availability design patterns can be employed to increase reliability.

6 Results and Discussion

6.1 Evaluation of the Solution Alternatives

Table 1 summarises fulfilment of requirements per technical solutions that are introduced in Section 5.

Requirement related to location specific information of the transmission equipment and connections (REQ-2.3, REQ-2.4 and REQ-3.3) are mostly out of reach for every proposed solution candidate. Site anchor proposal fulfils all the rest of the requirements.

6.2 Findings Related to Location Specific Data

In literature, geolocation coordinates and attesting them have received major attention. However, the longitude and latitude numbers are seldom useful as such. In most cases the coordinates must be mapped to some other space, like legislation properties or site name, before conclusions can be made regarding the geographical location.

If there are means to know dependably the location of an entity, there must be evidence based on laws of physics to dependably associate that location knowledge to other entities, e.g., server computers in the proximity. Otherwise, there leaves room for doubt which equals to lack of trust.

Trusted computing principles must be applied to software that is processing proximity evaluation and passing onwards location specific information to the inquirer. Data signing can be used to allow validation of the propagated information and thus allow some untrusted middlemen in the chain of passing the data. However, RTT kind of proximity measurements must be made by trusted software.

Based on our studies it seems possible that geographical location of datacenters and servers in them can be dependably assured. For locating data items, there must be means to associate the physical location of the server

Table 1 Summary of requirement fulfilment

REQ#:	1.1	1.2	2.1	2.2	2.3	2.4	2.5	3.1	3.2	3.3	4.1	4.2	4.3
Config.data	+	+	+	+	–	–	(+)	+	(+)	–	–	–	+
RTT	+	+	–	+	–	–	–	–	–	–	+	+	+
Topology	+	+	–	+	(+)	–	–	–	–	–	–	–	+
Attest svc	+	–	+	+	(+)	–	(+)	+	(+)	–	(+)	(+)	–
Seismic	(+)	+	–	+	–	–	–	–	–	–	+	+	–
GeoProof	+	+	–	+	–	–	–	–	–	–	+	+	+
Anchor	+	+	+	+	–	–	+	+	+	–	+	+	+

computer with the physical storage device. One possible method is to make to data manager software that accesses the storage device to respond to inquiries regarding local physical location specific information. If the management software can be trusted, it implies that the location specifics of data can be trusted, too.

7 Conclusion

Legal restrictions to geographical location of data processing are driven by governments willing to protect their citizens and enterprises against data breaches, which may become easier if data is stored or transmitted at or via locations that do not have as strict data protection legislation as the domestic one.

Need for a technical solution to attest dependably geographical location of cloud servers becomes more and more important as governments and enterprises move their database to cloud computing environments. This tendency is driven by smaller costs of cloud computing compared to in-house datacenters.

A set of requirements is derived from legal, commercial and stakeholder needs. Data breaches due to data sovereignty can happen when data is at rest in a storage server, being computed and while in transit over transmission connections. Protecting data in transit is a problem domain and cloud processing of data is another. These two should be solved separately. Technical solutions covered in this paper are within cloud computing domain and try to solve the problem of supporting geographical location specific information in cloud datacenters.

None of the proposed solutions alone is adequate but they can be integrated to a sufficient solution. Based on the reasoning in this paper, then best solution candidate for further studies is a combination of site anchor and RTT measurements. This can be extended with an attestation service if its maintenance can be taken care of within cost and time budget. Also, LAN topology discovery can be used to cross-check that server locations and RTT measurements are parallel to what is known about LAN connection in the datacenter.

The location anchor could be equipped with seismic sensors, trustworthy Galileo satellite receiver or cellular network based positioning method to facilitate fully automatic location reference. Even though radio signals are not assumed to reach antennas inside a datacenter, satellite and cellular

radio antennas can be wired to the location anchor device with reasonable costs. To create trust to the geolocating system, it may be necessary to have independent auditor to pay a visit to the datacenter site to assure, that the location anchor is properly deployed and secured.

Acknowledgements

The authors would like to thank Professor Tuomas Aura of Aalto University. This work was made under the Finnish Dimecc Cyber Trust Program.

Appendix – Fulfilment of Requirements

Following subsections contain detailed analysis about how the methods introduced above fulfil the listed requirements. The detailed analysis is presented as tables, one per requirement. A plus sign marks compliance and minus non-compliance. Parenthesis denote partial compliance.

Fulfillment of Commercial Constraints

REQ-1.1 lowers economic threshold to deploy dependable location specific information support into use.

REQ-1.1	Support for dependable location specific information should not add anything to the manufacturing costs of a trusted server hardware.
Configuration data	+ Already available mechanism, especially TPM registers can be used.
Round-Trip Time	+ Sufficient accuracy expected to be achievable with software based measurements.
Network Topology	+ Topology discovery does not need special hardware.
Attestation	+ Possibly dedicated attestation server is a good idea but no additional hardware is required to every server.
Seismic	(+) Existing server hardware is sufficient only if the accelerometers already present in the servers (or their hard disk units) can be used. Requires further research.
GeoProof	+ A special verifier device is required in every datacenter but no additional hardware is required to every server.
Site Anchor	+ A special site anchor device is required in every datacenter but no additional hardware is required to every server.

REQ-1.2 focuses on avoiding new effort while commissioning new server hardware. Assuming server configuring is already automated, additional parameters require changes only to the configuring script but the actual commissioning effort remains the same. However, if the configured data has some site-specific elements the correct site settings may need to be manually selected or set. Connecting the new server to the datacenter LAN can make the server to register itself to the cloud manager, or in some systems a human operator must activate the membership.

REQ-1.2	Dependable location specific information should not add need for new manual steps in the new server deployment process.
Configuration data	– If site specific information is in the persistent memory of each server the deployment process must differ while setting destination site specific values.
Round-Trip Time	+ Probably a new server can automatically register itself to the RTT measurement polling list.
Network Topology	+ Topology discovery can be implemented without server specific configuration.
Attestation	– Server commissioning needs server specific information to be added to the attestation database by an authorized person.
Seismic	+ No server specific settings needed.
GeoProof	+ Server is associated with location via RTT based proximity to the verifier device. External bookkeeping may be necessary to maintain site specific information.
Site Anchor	+ Server is associated with location via RTT based proximity to the site anchor device. Site specific information must be set once to the anchor device by an authorized person.

Fulfillment of Cloud Service Customer Needs

REQ-2.1 is the primary requirement for dependable location specific information support. Real-time in this context can, at least, refer to an application running on a cloud server being able to invoke API operation, which returns wanted information, like name of jurisdiction or geographical location. The API library can retrieve the information in various ways, depending on the underlying solution alternative. If only server identification is available at the processor, it must be used to address the wanted information from a lookup table maintained by the cloud system in a trustworthy manner.

REQ-2.1	A cloud application must be able to undeniably detect in real time location specific information applicable to the employed physical server.
Configuration data	+ Wanted information can be retrieved from persistent memory of the cloud server.
Round-Trip Time	– A server can be associated with a cluster of servers but the association is not necessarily real time enough. However, RTT as a method does not provide any storage space for location specific information.
Network Topology	– A server can be associated with a cluster of servers but the association is not necessarily real time enough. However, topology discovery as a method does not provide any storage space for location specific information.
Attestation service	+ Attestation database can be made to support location specific information. The search key can be a unique identifier of the server. Database maintenance procedures required to keep the information real time.
Seismic	– A server can be associated with a cluster of servers but the association is not real time. Seismic location discovery as a method does not provide any storage space for location specific information.
GeoProof	– A server can be associated with a cluster of servers but the association is not necessarily real time enough. However, GeoProof as a method does not provide any storage space for location specific information.
Site Anchor	+ Information can be inquired from the local site anchor.

REQ-2.2 facilitates cloud manager allocation decisions becoming inspectable afterwards. It can be assumed (PRE-5 in Section 4.1.5) that cloud management operations produce log records for this purpose. Secure logging mechanisms should be used.

REQ-2.2	It must be possible to check afterwards from log records which physical servers were allocated for running the application of a cloud customer.
Configuration data	+ Possible changes to location specific data should be recorded with time stamps.
Round-Trip Time	+ RTT measurement events should produce log records.
Network Topology	+ Topology discovery events should produce log records.
Attestation service	+ Attestation database events should produce log records.
Seismic	+ Seismic acceleration measurements should be continuously recorded. Postprocessing this data together with recordings from other sites and observatories is the normal procedure to find out the global location.
GeoProof	+ Action performed for the third-party auditor should produce log records. Preferably also location specific data inquiries should produce log records.
Site Anchor	+ Action performed via commissioning console should produce log records. All log records must be sent from the site anchor to a trusted storage server.

REQ-2.3 is like REQ-2.1 but for the transmission connections instead of cloud servers. Scheduling transmission resources is not as much under control of the cloud service provider as are the computing and storage resources. However, log records should be created from all decisions that may be of interest to cloud service customers. The provider can also require from the transmission subcontractor to make their log records available.

REQ-2.3	A cloud application must be able to undeniably detect in real time the employed physical transmission connections and equipment.
Configuration data	– This method is not aware of transmission connections.
Round-Trip Time	– RTT measurements aim to discover transmission connection line lengths. However, RTT as a method is not aware of transmission connections actually used by the cloud applications.
Network Topology	(+) Network topology discovery may record also transmission equipment that implement the connections. However, topology discovery as a method is not aware of transmission connections actually used by the cloud applications.
Attestation service	(+) Attestation database could store location specific information also for the transmission equipment and lines. Transmission management service could consult the database while making routing decisions. There must be some method, for example topology discovery, that puts transmission location specific data to the attestation service.
Seismic	– Seismic method cannot locate transmission connections nor is aware of connections actually used by the cloud applications.
GeoProof	– GeoProof cannot locate transmission connections nor is aware of connections actually used by the cloud applications.
Site Anchor	– Site anchor cannot locate transmission connections nor is aware of connections actually used by the cloud applications.

REQ-2.4 is like REQ-2.2 but for logging the transmission events instead of cloud servers.

REQ-2.4	It must be possible to check afterwards from log records which physical connections were used for transferring a certain piece of sensitive data.
Configuration data	– This method has no visibility to transmission events.
Round-Trip Time	– RTT method should create log records from its measurement actions which can be compared later with log records from transmission events. However, RTT method could be a subscriber of transmission connection change events.

Network Topology	– Topology discovery method should create log records from its actions which can be compared later with log records from transmission events. Topology discovery could be a subscriber of transmission connection change events.
Attestation service	– Attestation service is not aware of connections events.
Seismic	– Seismic method is not aware of connections events.
GeoProof	– GeoProof method is not aware of connections events.
Site Anchor	– Site anchor method is not aware of connections events.

REQ-2.5 extends use of location specific information support to marking a datacenter as validated per cloud service customer basis. In this approach, the criteria for validation need to be known only by the customer and its authorized auditor.

REQ-2.5	It should be possible to mark a datacenter site to be certified for storing and/or processing protected data of a certain application.
Configuration data	(+) Persistent configuration data could be used for this purpose but then the validation must be done per server basis. This does not easily scale up to multitude of servers and customers.
Round-Trip Time	– RTT method does not store site specific data.
Network Topology	– Topology discovery method does not store site specific data.
Attestation service	(+) Attestation database could be used for this purpose but then the validation must be done per server basis. This does not easily scale up to multitude of servers and customers.
Seismic	– Seismic method does not store site specific data.
GeoProof	– GeoProof method does not store cloud application accessible site-specific data.
Site Anchor	+ Customer and site-specific data can be stored to a site anchor. An anchor has limited memory capacity but there can be several anchors per site.

Fulfillment of Auditor Needs

Main point of REQ-3.1 is to worry about providing technical means for an auditor to inspect large amount of physical cloud servers in a short period of time.

REQ-3.1	An independent auditor must be able to verify datacenter site and other location specific data associated with each trusted physical server.
Configuration data	+ An auditor visiting a site must check persistent location specific configuration data item of each trusted cloud server one by one. This should be possible with clever auditing software running in the laptop of the auditor while it is connected to the LAN of the datacenter.
Round-Trip Time	– RTT method does not store site specific data.
Network Topology	– Topology discovery method does not store site specific data.
Attestation service	+ When the location specific information is stored to an attestation service database, the auditor must verify that physical servers of the site and their data in the database are correct and that there are no servers listed to the site that actually resides elsewhere. The challenge is to provide human observable indication that couple a listed server with a physical server at the site in a trustable manner. One possibility is use run RTT tests from the auditor's computer while it is connected to the LAN of the site.
Seismic	– Seismic method does not store site specific data.
GeoProof	– GeoProof method does not store cloud application accessible site-specific data.
Site Anchor	+ Using RTT measurements over site LAN the auditor can validate that the anchor is at the site. Software in the trusted servers can check that the site anchor is at the same LAN as the server. By running an application is each server the auditor can check that the server returns expected location specific information as a response to an inquiry.

REQ-3.2 demands that the location specific information is maintained in real time. What "real time" actually means should be specified more accurately. For most cloud based services an update delay of one minute should not be a problem. After all, location specific information of the servers should change less often than once a year.

REQ-3.2	The location specific information attestation database must be maintained in real time in a dependable manner.
Configuration data	(+) Changing persistent configuration data is not a technical challenge except when there are security features around it. The data can be quickly altered in a server but how to be sure that the change is authorized. One possibility is that only a third-party auditor can perform the update but this is a clumsy and expensive solution. If a server is moved to another site, RTT or some other reliable mean is needed to verify the new proximity. In addition, authorization keys are needed to prevent unauthorized access to critical configuration data.

Round-Trip Time	– RTT method does not store site specific data.
Network Topology	– Topology discovery method does not store site specific data.
Attestation service	(+) An authorized actor should be able to update the data in the attestation service database. Also, RTT or some other reliable mean is needed to verify the new proximity.
Seismic	– Seismic method does not store site specific data.
GeoProof	– GeoProof method does not store cloud application accessible site-specific data.
Site Anchor	+ A moved cloud server connects to the local site anchor of the new site and gets from there the location specific data that is applicable at the new location. Same procedure applies if a new server is added to the datacenter site. No manual intervention is required.

REQ-3.3 is corresponding requirement to cover transmission connections as REQ-3.1 and REQ-3.2 are for cloud servers.

REQ-3.3	Transmission connections and routes between datacenters must be listed and audited from jurisdictions point of view. The list must be maintained if transmission contracts are changed.
Configuration data	– This method has no visibility to transmission events.
Round-Trip Time	– RTT method could compare latest measurement results with earlier results and raise flag if there are such changes that could be caused by, for example, new route. However, RTT measurement do not offer detailed information of a change so that transmission location database could be updated.
Network Topology	– Topology discovery method could compare latest discovery findings with earlier results and raise flag if there are such changes that could be caused by, for example, new route. However, the observations do not offer detailed information of a change so that transmission location database could be updated.
Attestation service	– Attestation service is not aware of connections events.
Seismic	– Seismic method is not aware of connections events.
GeoProof	– GeoProof method is not aware of connections events.
Site Anchor	– Site anchor method is not aware of connections events.

Fulfillment of Cheating Mitigation

Location related cheating methods are dependent on the implementation of the support for location specific data. Thus, this topic should be revisited once there is more information available regarding the implementation.

REQ-4.1 prevents dishonest cloud service provider to spoof location specific information.

REQ-4.1	Attempts to make a server on a remote site to look like it is located on the local site must be detected.
Configuration data	– It is realistic to assume that the owner of a datacenter can get authorized access to the configuration data of the servers and change it. Even trusted computing mechanism may not be enough to prevent forging data if the hash covering configuration data is also changed.
Round-Trip Time	+ RTT measurements can reveal if a server is connected to the same physical LAN or not.
Network Topology	– If a server is connected to a LAN switch of another datacenter topology discovery of the LAN does not reveal this.
Attestation service	(+) If the attestation service is under control of a third-party auditor we should be able to trust that the location specific data is correct. It can be possible that the auditor is misled by a cleverly dishonest datacenter owner. Forging attestation service data is less difficult if the attestation service is under control of a dishonest party.
Seismic	+ Natural seismic waves cannot be manipulated but it may take some time before latest accelerometer results are analysed.
GeoProof	+ GeoProof relies on verifier device which is proven to reside inside a datacenter. The device can perform RTT measurements which can detect if a server is further away than in the datacenter LAN.
Site Anchor	+ A server gets the location specific data from the local site anchor. Because site anchor cannot be reconfigured the only possibility to forge server's location specific information is to make it to connect to an anchor of another site. RTT measurements should prevent this, especially if they are performed by the anchors and the servers.

REQ-4.2 brings to focus possibility that a server is moved without reconfiguration to another site. This can reveal a weak point if the onboard data contain location specific information.

REQ-4.2	If a server is moved to another site, the move is detected and the server is associated with its new datacenter site.
Configuration data	– When location specific data is stored locally to the persistent memory of a server and the server is moved, the server still reports the according to the old location in the new site. Additional checks are needed for continuously monitoring for inconsistencies, for example, that all servers of a site report the same location. This will consume CPU cycles mostly in vain because server moves are rare and only small fraction of them forget configuration update.

Round-Trip Time	+ RTT measurements can reveal if a server is connected to the same physical LAN or not.
Network Topology	– If a server is connected to a LAN switch of another datacenter topology discovery of the LAN does not reveal this.
Attestation service	(+) If the unique access key value to retrieve the information, like a serial number, is bound with the server motherboard, attestation service may continue responding with outdated data, unless the data is updated as part of the migration process. In general, if the attestation service is under control of a third-party auditor we should be able to trust that the location specific data is correct. It can be possible that the auditor is misled by a cleverly dishonest datacenter owner. Forging attestation service data is less difficult if the attestation service is under control of a dishonest party.
Seismic	+ Natural seismic waves cannot be manipulated but it may take some time before latest accelerometer results are analysed.
GeoProof	+ GeoProof relies on verifier device which is proven to reside inside a datacenter. The device can perform RTT measurements which can detect if a server is further away than in the datacenter LAN.
Site Anchor	+ A server gets the location specific data from the local site anchor. Because site anchor cannot be reconfigured the only possibility to forge server's location specific information is to make it to connect to an anchor of another site. RTT measurements should prevent this, especially if they are performed by the anchors and the servers.

REQ-4.3 covers right to be forgotten if a server is taken off from service. This is not a critical requirement because a missing server probably raises alarm flags when workload is scheduled for it to process and the server does not respond.

REQ-4.3	When a server is retired, its location information must be removed.
Configuration data	+ Configuration data goes out of reach when the server is disconnected.
Round-Trip Time	+ RTT measurements cannot be made with a missing server which should raise alarm flag and trigger removal of the server from server inventory.
Network Topology	+ Topology discovery cannot reach a missing server and it gets dropped from topology maps.
Attestation service	– Attestation service, unless promptly updated, continues reporting the last location of the retired server.
Seismic	+ No new seismic measurements arrive from the retired server which should raise alarm flag and trigger removal of the server from server inventory.
GeoProof	+ A missing server cannot anymore be verified to own the required database causing failure in data residency test.
Site Anchor	+ No server specific information stored.

References

[1] Anonyms (2017). *Government under Fire After Transport Agency Data Breach.* Available at: http://sverigesradio.se/sida/artikel.aspx? programid=2054&artikel=6740394 [accessed July 28, 2017].

[2] Anonyms (2017). *Swedish Authority Handed Over 'Keys to the Kingdom' in IT Security Slip-up.* Available at: https://www.thelocal.se/20170717/swedish-authority-handed-over-keys-to-the-kingdom-in-it-security-slip-up [accessed July 26, 2017].

[3] Palad, N., and Michalas, A. (2014). *One of Our Hosts in another Country: Challenges of Data Geolocation in Cloud Storage.* Sweden: Swedish Institute of Computer Science, 1–6.

[4] Anonyms (2013). *The OECD Privacy Framework.* Available at: https://www.oecd.org/sti/ieconomy/oecd_privacy_framework.pdf

[5] Hippeläinen, L., Oliver, I., and Shankar, L. (2016). *Survey of Cloud Server Geolocating Techniques.* Available at: http://fruct.org/publications/fruct19/files/Hip.pdf

[6] Ravidas, S., Lal, S., Oliver, I., and Hippeläinen, L. (2017). *Incorporating Trust in NFV: Addressing the Challenges.* Paris: IEEE Xplore Digital Library.

[7] Anonyms (2017). *Lawful Interception.* Available at: http://www.etsi.org/technologies-clusters/technologies/lawful-interception [accessed March 14, 2017].

[8] Hippelainen, L., Oliver, I., and Lal, S. (2017). "Towards dependably detecting geolocation of cloud servers," in *Proceedings of the 11th International Conference on Network and System Security, NSS 2017,* Helsinki, Finland, August 21–23, 2017 (Helsinki: Springer International Publishing), 643.

[9] McKnight, D. H., and Chervany, N. L. (1996). *The Meanings of Trust.* Available at: http://www.misrc.umn.edu/workingpapers/fullpapers/1996/9604_040100.pdf [accessed March 2, 2017].

[10] Kittleson, N. (2012). *Trusted Computing Overview.* Available at: https://scap.nist.gov/events/2012/itsac/presentations/day2/4oct_11am_kittleson.pdf [accessed March 15, 2017].

[11] Wilkins, R., and Nixon, T. (2016). *The Chain of Trust.* Available at: http://www.uefi.org/sites/default/files/resources/UEFI%20Forum%20White%20Paper%20-%20Chain%20of%20Trust%20Introduction_Final.pdf [accessed March 14, 2017].

[12] Trusted Computing Group (2016). *Trusted Platform Module Library Specification, Family "2.0", Level 00, Revision 01.38.* Available at: https://trustedcomputinggroup.org/tpm-library-specification/ [accessed March 2, 2017].

[13] Futral, W., and Greene, J. (2013). *Intel Trusted Execution Technology for Server Platforms.* New York City, NY: Apress.

[14] Mell, P., and Grance, T. (2011). *The NIST Definition of Cloud Computing.* Gaithersburg, MD: National Institute of Standards and Technology.

[15] Anonyms (2016). *Jurisdiction.* Available at: https://www.law.cornell.edu/wex/jurisdiction [accessed November 22, 2016].

[16] Anonyms (2013). *Definition: Data Sovereignty.* Available at: http://whatis.techtarget.com/definition/data-sovereignty [accessed September 18, 2016].

[17] Rouse, M. (2015). *Definition: Data Residency.* Available at: http://searchcloudcomputing.techtarget.com/definition/data-residency

[18] Anonyms (2016). *Data Center.* Available at: https://en.wikipedia.org/wiki/Data_center [accessed September 26, 2016].

[19] Determann, L., Bekeschenko, E., and Perevalov, V. (2015). *Residency Requirements for Data in Clouds—What Now?* Available at: http://www.globalequityequation.com/files/Uploads/Documents/Equity%20 Equation/Residency%20Requirements%20for%20Data%20in%20Clouds%20–%20What%20Now.pdf [accessed September 2016].

[20] Kuner, C. (2011). *Regulation of Transborder Data Flows under Data Protection and Privacy Law: Past, Present and Future*. Paris: OECD Publishing.

[21] DLA Piper (2016). *DLA Piper's Data Protection Laws of the World Handbook.* Available at: https://www.dlapiperdataprotection.com/ [accessed September 2016].

[22] Anonyms (2016). *Regulation on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC*. Washington, DC: General Data Protection Regulation.

[23] Anonyms (2013). *The OECD Privacy Framework.* Available at: https://www.oecd.org/sti/ieconomy/oecd_privacy_framework.pdf

[24] Anonyms (2012). *Privacy Amendment (Enhancing Privacy Protection) Act.* Federal Register of Legislation.

[25] Anonyms (2016). *Final Cybersecurity Law Enacted in China.* Available at: https://www.huntonprivacyblog.com/2016/11/08/final-cybersecurity-law-enacted-china/ [accessed November 14, 2016].

[26] Gallia, A. L., McLoughlin, L. P., Khaskelis, A. S., and Voltchenko, M. A. (2015). *Russian Federation: Russia's Personal Data Localization Law Goes Into Effect.* Available at: http://www.mondaq.com/russianfederation/x/435890/Data+Protection+Privacy/Russias+Personal+Data+Localization+Law+Goes+Into+Effect

[27] Jolly, I. (2016). *Data protection in United States: Overview.* Available at: http://uk.practicallaw.com/6-502-0467 [accessed September 14, 2016].

[28] Albeshri, A. A., Boyd, C., and Gonzalez Nieto, J. (2012). *Geoproof: Proofs of Geographic Location for Cloud Computing Environment.* Macau: IEEE, 506–514.

[29] Hambling, D. (2017). *Ships Fooled in GPS Spoofing Attack Suggest Russian Cyberweapon.* Available at: https://www.newscientist.com/article/2143499-ships-fooled-in-gps-spoofing-attack-suggest-russian-cyberweapon/ [accessed August 14, 2017].

[30] ARM Limited (2017). *ARM Security Technology Building a Secure System using TrustZone Technology.* Available at: http://infocenter.arm.com/help/ index.jsp?topic=/com.arm.doc.prd29-genc-009492c/CACGCHFE.html [accessed May 2017].

[31] Intel (2014). *Strengthening Security with Intel Platform Trust Technology.* Available at: https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/enterprise-security-platform-trust-technology-white-paper.pdf [accessed August 30, 2017].

[32] Coupe, C., Hegenderfer, S., and Schmidt, S. (2016). *Debunking the Myth: Bluetooth Range.* Available at: https://www.bluetooth.com/∼/media/files/marketing/range%20white %20paper.ashx?la=en

[33] Anonyms (2011) *How Many Active Bluetooth Devices Can I Reliably Detect in a Single Space?* Available at: http://electronics.stackexchange.com/questions/21991/how-many-active-bluetooth-devices-can-i-reliably-detect-in-a-single-space [accessed September 14, 2016].

[34] Brands, S., and Chaum, D. (1993). "Distance-bounding protocols," in *Proceedings of the Workshop on the Theory and Application of Cryptographic Techniques,* Berlin: Springer.

[35] LAN/MAN Standards Committee (2009). *Station and Media Access Control Connectivity Discovery,* New York, NY: IEEE Computer Society.

[36] Ochoa Aday, L., Cervelló Pastor, C., and Fernández Fernández, A. (2015). *Current Trends of Topology Discovery in Open Flow-based Software Defined Networks.* Available at: http://upcommons.upc.edu/handle/2117/ 77672 [accessed November 18, 2016].

[37] Albeshri, A., Boyd, C., and Nieto, J. G. (2013). Enhanced GeoProof: improved geographic assurance for data in the cloud. *Int. J. Inform. Sec.* 13, 191–198.

Biographies

Leo Hippeläinen, MScEE, is a graduate from Helsinki University of Technology 1978. Over 40 years of software systems experience spanning from 8-bit Z80 systems and Nokia's DX200 through to modern day NFV, SDN and trusted cloud systems. He has been involved with many EU and TEKES projects including Celtic Plus SEED4C concentrating on high-integrity computing for telecommunication system. He is currently a senior security researcher at Nokia Bell Labs with particular interesting in geographically trusted cloud computing and IoT. He is also pursuing licentiate degree at Aalto University.

Ian Oliver is a security specialist at Bell Labs working on Trusted and High-integrity Network Function Virtualisation for 5G Networking, blockchain and the semantics of privacy. He holds a research fellow position at the University of Brighton working on semantics and diagrammatic reasoning. He is the author of the book Privacy Engineering: A data flow and ontological approach.