Physical Attack Protection with Human-Secure Virtualization in Data Centers

Jakub Szefer, Pramod Jamkhedkar, Yu-Yuan Chen and Ruby B. Lee
Department of Electrical Engineering
Princeton University
{*szefer, pjamkhed, yctwo, rblee*}*@princeton.edu*

***Abstract*—Cloud computing-based data centers, which hold a large amount of customer data, are vulnerable to physical attacks and insider threats. Current protection and defense mechanisms for security of data held in data centers are either completely physical (sensors, barriers, etc.) or completely cyber (firewalls, encryption, etc.). In this paper we propose a novel cyber-physical security defense for cloud computing-based data centers against physical attacks. In our system, *physical* sensors detect an impending physical/human attack which triggers *cyber* defenses to protect or mitigate the attack. The key to the cyber defenses is that in cloud computing data centers the data is loosely coupled with the underlying physical hardware, and can be moved/migrated to other physical hardware in the presence of an attack. In this paper we propose a model for coupling such cyber defenses with physical attack-detection sensors. We further describe a preliminary architecture for building such a system with today's cloud computing infrastructure.**

***Keywords*-physical attacks, data center, moving target defense**

I. Introduction

We present a human-aware, self-adapting approach to autonomously and proactively defend code and data executing on servers inside the data center. Like water utilities, the electricity grid, or SCADA systems, data centers are cyber-physical systems (CPS) and share many of the characteristics of these systems. They have a physical component comprising of computing infrastructure, servers, etc., and a cyber component consisting of the code and data executing or stored in the data center.

While in the past, data centers were isolated and each enterprise or company would have its own, now they are being increasingly used to provide computation and storage as a utility –"cloud" computing being a prime example of such use of data centers. Another trend is that the operator or owner of a data center is a separate party from the customers who are using the data center as a utility to perform their computation. As cloud computing becomes prevalent, data centers will become prime targets for attackers to steal information or damage computing infrastructure to compromise availability. Surprisingly, physical attacks compromising the security of information technology (IT) infrastructures have been identified as one of the most overlooked aspects of data center security [1], there has been no, or limited, cyber-physical defenses proposed. Moreover, the defense mechanisms in data centers have typically followed isolated approaches that are either entirely in the physical space, or entirely in the cyber space.

To protect data centers against physical attacks, numerous physical prevention and detection mechanisms are implemented today [2]. These include security locks, crash barriers, two-factor authentication, surveillance cameras, security guards, etc. These protection and defense mechanisms remain entirely in the physical space, and despite their existence, physical breaches and insider threat attacks have continued to occur [3], [4]. Cyber-level protection mechanisms including persistent encrypted and hashed storage, e.g., [5] provide sound confidentiality protection against physical attacks, but can incur tremendous overhead in a data center handling vast amounts of data.

In this work we propose a cyber-physical defense against physical attacks in data centers. The key to the defense is that in today's data centers, customers' data and code reside and execute within virtual machines (VMs). Virtual machines are containers for code and data, and because of the underlying virtualization software, the code and data are not tied down to the physical system that they are executing or stored on [6]. VM migration is a common technique used to move a VM from one physical machine to another [7]. Virtualization allows VMs to be easily deleted, encrypted or moved – which forms the trio of our defense strategies. Our defense strategies are in the spirit of *moving target defense*, where code and data are moved to avoid the attacker. However, the power of virtualization and moving target defense has always been used only against software-level (cyber) attacks confining its applicability only in the cyber space [8]. We extend these ideas to protect VMs, and the code and data contained in them, from physical attacks.

In this paper, we propose a cyber-physical security framework for cloud computing data centers that combines the security mechanisms in cyber and physical spaces, and exploits the power of virtualization to provide dynamic security against physical attacks. We protect against *human* attackers who can use physical access (illegitimate or legitimate as in the case of insider attacks) to extract code or data from the compute infrastructure inside the data centers. We *secure*

978-1-4673-2266-9/12/$31.00 ©2012 IEEE

code or data through the use of delete, encrypt or move cyber defenses. We leverage physical intrusion detection systems to warn of an impending physical attack and trigger these defenses. The key component of today's data centers which we use is *virtualization*. Hence we term our approach *human-secure virtualization*.

The contribution to the field and the novelty of this work lies in:

- the presentation of how to integrate cyber and physical defenses to protect code and data executing or stored in data centers from physical attacks, and
- the specification of a defense framework against physical attacks based on t_{detect} to t_{attack} time window during cuber which defenses can be activated.

This paper presents details of our approach for data center defenses against physical attacks. We have ongoing work on implementing the presented scheme in OpenStack [9], a data center management software.

The remainder of the paper is organized as follows. Section II describes the human-aware data center defense framework. Section III presents a preliminary system architecture for implementing the framework using today's cloud computing software. Section IV lists related work and we conclude in Section V.

II. Human-Aware Data Center Defense Framework

Our work is based on a contemporary computing paradigm where computation is performed inside virtual machines (VMs). A VM is a software implementation of a computing environment, within which an operating system and programs can be installed and run (Figure 1). VMs are enabled by a hypervisor that virtualizes the underlying physical hardware. On a given physical machine a hypervisor can simultaneously run multiple VMs, which share the underlying hardware via virtualization. A VM can be moved between physical machines by means of VM migration.

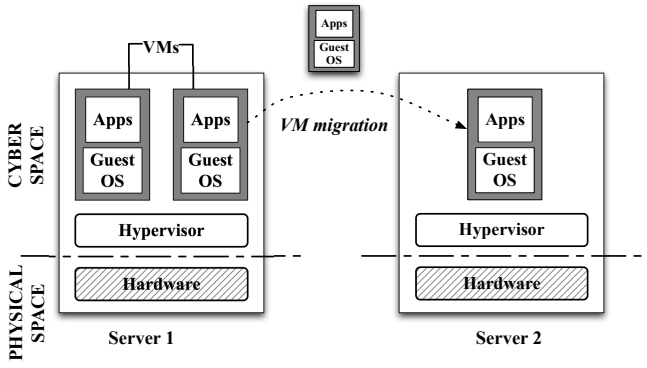

Figure 1. Virtualization software, the hypervisor, is used to partition the physical resources of servers for virtual machines (VMs), allowing the usage of a data center as a utility. Virtualization allows separation of the physical components of the servers from the cyber components which are the VMs and which can be migrated among the physical servers.

A. Threat Model and the Attacks

We aim to protect against physical attacks on compute, storage and networking equipment in a data center setting. This work is orthogonal to, and complements, work on defenses of these entities from cyber attacks. The main object to protect are the VMs running on the physical servers in a data center, which contain code and data belonging to cloud customers. By extracting the contents of the VMs, an attacker can gain valuable information, including proprietary or sensitive data belonging to cloud customers.

The human threat we are worried about is a physical attacker inside a data center who can gain physical access to the servers where the VMs are running. An attack can be carried out by different types of individuals including outsiders, maintenance individuals, insider threat, etc. We assume that once an attacker has physical access to a machine, at that point in time the attack can be considered successful.

Attack detection is performed by physical intrusion detection systems present in data centers. Our proposed system migrates or scrubs all code and data from devices when an intruder is detected, and we do not attempt to distinguish whether a physical intruder is a friend or foe. In particular, this allows us to defend against the insider threat problem, where a legitimate data center employee may use his or her position to physically access the hardware and perform a physical attack.

We do not protect against software attacks on the virtual machines or the management software; this is an orthogonal problem. We assume correct implementation and execution of defenses employed within a data center. There is already a lot of computer security work on software defenses [10]. We also assume correct implementation and operation of physical security sensors.

B. Physical Sensors and Security

An interesting feature of the data centers, unlike some of the older cyber-physical systems such as water utilities, is the existence of a plethora of physical security sensors and intrusion detection mechanisms. Data centers routinely include motion sensors, cameras, electronic locks on doors, etc. [11]. These have been installed to allow only authorized personnel to enter, and for post-mortem investigation if a breach has occurred (e.g., use of video footage to see who entered a restricted area). These sensors provide an infrastructure that can signal an alert about an impending physical attack before it is carried out. Such a warning system provides various degrees of response times for defense mechanisms to be triggered. We propose to leverage these readily available detection systems to create the human-aware, self-adapting protection of virtual machines running in the data center.

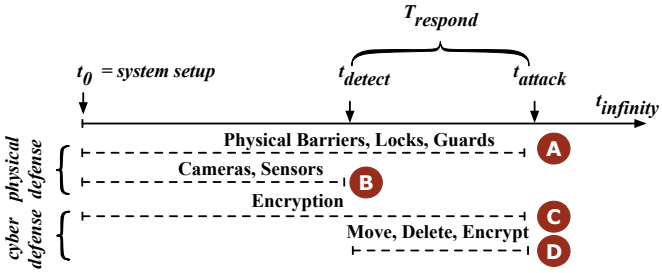

Figure 2. Defense strategy timeline.

C. Attack Detection and Defense Timeline

Our defense framework against physical attacks is based on the time difference between attack detection and the actual attack, as shown in Figure 2. We consider two timestamps in our framework: time of detection (t_{detect}) denoting the time at which an attack is detected, and time of attack (t_{attack}) denoting the time at which the attacker has physical contact with the computing equipment.

The physical sensors mentioned previously are used to trigger a warning at t_{detect}. While the exact t_{attack} is not known, we can designate a lower bound (worst case). After detecting intrusion, the location of the sensor and the location of the equipment can be used to determine how much time would be needed for an attacker to gain physical access to the piece of equipment. As an example, a recent industry white paper notes that "if someone has physical access to your servers, it takes about 60 seconds to take full control of the servers and start shipping your data out the door." [12]. In this scenario, t_{attack} would be $t_{detect} + 60\ sec$.

Figure 2 shows different security mechanisms and their relationship with these two time stamps. Mechanisms shown in case *A* are preventive mechanisms that operate wholly in the physical space, and their effectiveness ends after an attack occurs. The goal of these measures is to delay t_{attack}, possibly forever (i.e., to shift t_{attack} as far right as possible). Mechanisms shown in case *B* are detection mechanisms that operate in physical space, and their defensive purpose ends after an attack has been detected; although they can be utilized for evidence gathering as the attack proceeds.

The goal of these mechanisms is to detect an attack as early as possible. Persistent encryption mechanisms shown in case *C* assume that an attacker can attack without any warning ($t_{detect} = t_{attack}$), and data is kept encrypted all the time. Such mechanisms, even though effective, can incur significant overhead in a large data center.

In our defense framework, we capitalize on the fact that effective detection mechanisms (case *B*) coupled with effective protection mechanisms (case *A*), both operating in the physical space, will induce a significant delay between t_{detect} and t_{attack}. The physical detection mechanisms (case *B*) would produce a warning at t_{detect} that would trigger appropriate cyber defense mechanisms, based on the expected time to respond ($T_{respond}$).

D. Defense Framework

Based on the inputs from physical detection systems, such as sensors, and the estimated time to respond, it is possible to trigger various types of response mechanisms in the cyber-space. We have identified three primary types of defenses or response actions, which have been used or could be used against physical attacks to protect virtual machines:

1) *Delete:* In this case, code and data are deleted from the computing equipment so that when the attacker has access to the computing equipment, the data is no longer there.
2) *Encrypt:* In this case, code and data are encrypted within the computing equipment, so that when the attacker has access to the computing equipment, the attacker cannot make use of the data since it is in encrypted form.
3) *Move:* In this case, code and data are moved away from the computing equipment so that when the attacker has access to the computing equipment, the data is no longer there.

It is possible to use a combination of these strategies based on the available response time, type of attack, the security goals to be achieved, type of compute or storage infrastructure and the types of requests from the cloud customers. Each response strategy has its own limitations in terms of the time taken to carry out the response, practical feasibility, cost, the security protections it offers (in terms of confidentiality, integrity, availability, privacy, etc.) and the after effects of the response. Table I summarizes the capabilities and computation cost (in terms of time) required for the three primary defense mechanisms. The defense mechanisms discussed below are a small set of the range of possibilities that can be developed from these three basic strategies. As more types of attack detection inputs become available, these defense strategies will be refined.

Table I
DEFENSE MECHANISMS AND THEIR ANALYSIS.

Defense	Confidentiality	Availability	Integrity	Cost
Delete	Yes	No	No	Low
Encrypt	Yes	Limited	No	Medium
Move	Yes	Yes	Yes	High

In the case of *Delete*, the time to carry out the response is comparatively lower, confidentiality is safeguarded, however availability is lost after the data has been deleted. It is possible to use *Delete* with replication (at a remote site), which would ensure availability, but at the cost of significant space and computation overhead due to the replication process.

In the case of *Encrypt*, confidentiality is maintained, and so is availability (unless the attacker destroys or steals

the storage equipment). However, the cost of encryption is higher in terms of computation and time, and there may not be sufficient time to respond with encryption, once an attack has been detected. Although, encryption strategies can be effective in the case of scenarios where there is sufficient time *a priori*, such as maintenance schedules. Encryption can also be used *on-demand*, based on specific requests from cloud customers for their individual VMs.

In the case of *Move*, confidentiality, integrity and availability are all preserved after the attack (even if the attacker destroys or steals the equipment). Moving large amounts of data out of a data center may not be possible for shorter response times. However, the *Move* defense can be optimized with clever combinations of physical barriers, detection mechanisms, compartmentalization of the data center, *on-demand* services and priority-based data categorization. For example, certain sections of a data center can be made "high-security" areas, where high-priority (determined by cloud customer demands) data can be moved once an attack has been detected. Cloud providers can charge the customers for such costly and sophisticated defense mechanisms.

III. System Architecture

In this section we describe the initial design details of our system architecture, the main components of which are:

- the physical sensors and the physical security monitors which gather information from the sensors,
- the management infrastructure software which has been modified to accept input from physical security monitors to trigger the defensive measures, and
- the actual compute and networking nodes which carry out the defensive measures.

A. Components of the Architecture

The management infrastructure receives sensor data via the physical security monitor, as shown in Figure 3. The physical security monitor collects all the data from the attached sensors, formats it and converts it to the standard API (application programming interface) calls to the management infrastructure. The physical security monitor can implement drivers for the myriad of sensors, abstracting the details from the data center management software.

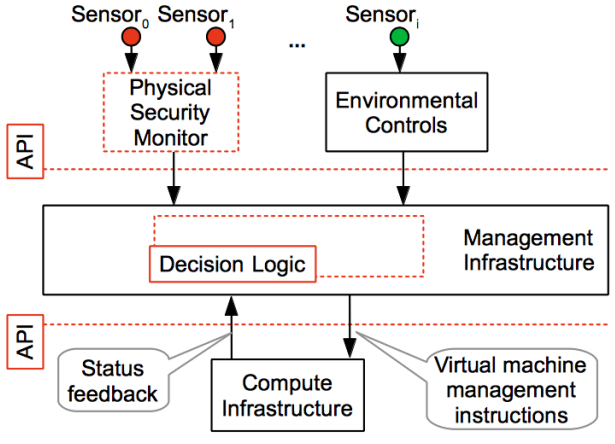

Figure 3. The components of our proposed system. The three new key components are shown in dashed lines: Physical Security Monitors, API and Decision Logic.

Also, the management infrastructure software could obtain data from environmental sensors already used in the data centers. Alternatively, for example, the climate control infrastructure could be augmented with special modules to send the relevant environmental data (e.g., humidity or "door-open" sensor data) to the data center management software via our API.

Once received, the data would be processed by our new decision logic components inside the management infrastructure. Unlike today's management software, the physical security-aware manager needs to be able to track the physical security of the servers which involves a number of new duties:

1) track physical sensor locations and current data from sensors,
2) track compute and networking equipment locations,
3) map which sensor's data affect which compute equipment (based on the location of the equipment and the sensor), and
4) map how sensor data changes affect security (or if certain changes do not affect security).

At the minimum, we suggest that physical intrusion sensors be installed in each server room. This includes "door-open" sensors at the entrances, cameras with face recognition in the room, or pressure sensors in the floor tiles. More sophisticated sensors can be installed, such as humidity sensors to sense the presence of a physical intruder. Many of these are already present in the data centers [13], [14] and other sensors can easily be added. Ideally, however, these would be incorporated at data center design time. Moreover, many of the sensors can be shared with the climate control systems (e.g., "door-open" sensor or humidity sensors) to limit the cost of our design.

The location of the sensors is a function of the data center floor plan and server room layout. Ideally the server type and location can be automatically extracted from the CAD (computer aided design) files, which describe the data center and server room facilities. To map the location of the servers, equipment and the sensors, the CAD files can be automatically read by the management software. A system architect will be needed to specify the scope of each sensor. For example, a camera can only see the servers in the same room where it is located and only in the direction it faces. Special features, such as the auto-pan feature for a camera, also need to be specified. Given this metadata about sensors and the two types of CAD files, the management software can automatically create and maintain an accurate view of the data center.

B. Computing Defense Strategies

Once the management software is aware of the physical status of the data center, it can act upon any unexpected (or expected as in the case of scheduled maintenance) changes in the environment. The environmental changes, such as a person opening a door to a server room, signal a potential impending attack. To deal with such attacks, we have proposed the three strategies of delete, encrypt or move. The decision logic, shown in Figure 4, in the management infrastructure is responsible for calculating the actual defense strategies and triggering them upon a physical intrusion. This is the "magic sauce" of any deployment which will realize the defenses of virtual machines against physical attacks in data centers. The inputs to the decision logic include: physical inputs about the physical infrastructure (data center layout, location and type of sensors, etc.), inputs from the security monitors reporting the current status from the sensors, as well as inputs about the load of the networks, availability of the servers or the security requirements for VMs, etc. The inputs are obtained via the API and may require specific interpreter modules so each input can be properly interpreted.

Scheduled events are one interesting input to the decision logic. Any regular maintenance will be scheduled and the information about the schedule can help the defense. For example, when a scheduled server maintenance is coming up, the management infrastructure has essentially a large t_{detect} to t_{attack} time window, during which the defenses can be activated early on. Then, by the time the scheduled server maintenance occurs and a technician is working on a server, all sensitive information has been deleted, encrypted or moved per our proposed cyber defenses.

The physical infrastructure database is generated automatically from the CAD files described earlier. The database ensures that the management software has current and up-to-date view of where the compute equipment is physically located (e.g., as servers are relocated, the database is updated). The cyber infrastructure status is the information obtained from the software running on the different pieces of compute-related equipment. For example, a server can report that the anti-virus software found a virus on the system, and hence it can no longer be trusted to carry out the defenses.

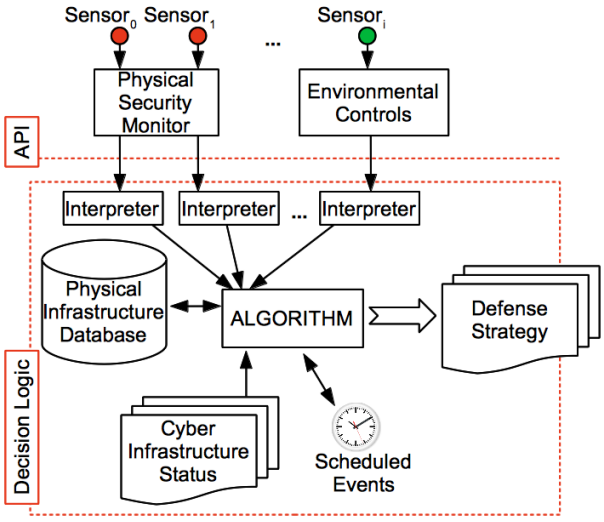

Figure 4. Decision logic.

C. Enacting Defense Strategies

Given these inputs, the algorithm inside the decision logic outputs the defense strategy that is to be taken at each time instant. As described in the previous section, the defense strategy will be a mix of defensive measures depending on the available time for the defense to be activated. The defense strategies are translated into actual virtual machine management instructions (bottom of Figure 3) so that the compute equipment can carry out the needed measures.

As more and more inputs are received, the defense strategy will be refined so the system is self-adapting to the current situation. For example, if a person is detected in a server room, the compute infrastructure closest to the entrance will be defended first as that's what the attacker can reach most quickly. As sensors pinpoint which server cabinet the attacker is trying to open, the focus will move to these exact servers. Such feedback can ensure that the code and data are continuously protected.

Also, the decision logic outputs can specify reconfiguration of the sensors. The reconfiguration can result in better detection (e.g., cameras can be set to send higher resolution video capture images to aid better face recognition). Finally, when the physical attack is over or an "all clear" condition is met, the algorithm outputs instructions for resuming normal operation on the equipment that was under attack.

IV. Related Work

Physical attacks have been one of the most overlooked aspects of information technology (IT) security [1]. Fortunately, in data centers numerous physical measures are applied during the design, construction, and operation of data centers to ensure the security of data center infrastructure from physical attacks [2]. Works have explored barriers, alarms, entry control, contraband detection, CCTV (closed-circuit television) surveillance, etc. [15] for protecting data centers. These measures are typically aimed at ensuring authorized access to computing infrastructure, detecting security breaches, and a recovery plan in the event of a security breach. They also mainly focus on environmental, not human, factors such as fire security or failures of supporting utilities (e.g., power utility) [16].

Ensuring authorized access to data centers typically includes measures like security locks, biometric authentication, compartmentalization and isolation of secure areas,

placing crash barriers, sign-in books, two-factor authentication, etc. These mechanisms are often complemented with mechanisms for detecting physical intrusions such as sensors, cameras, surveillance guards, etc. These approaches have largely remained in the physical realm, meaning that the attacks and their solutions are both handled in the physical space. Despite these measures, physical breaches have been a common occurrence, including insider threats.

Cyber-level approaches against physical or hardware attacks have been in the form of secure server architectures that have focused on persistent encryption of data outside the processor boundary [17], [18], [5] or special cryptographic co-processors [19]. While these approaches can provide sound security against physical attacks, implementing them on a large scale within a data center will affect performance and lead to significant amounts of redundancy, given that physical attacks are not very frequent.

Virtualization in cloud computing provides unique advantages in terms of protecting customers' code and data from potential security threats from hackers. The capability to "move" code and data around in the presence of software attacks has been well exploited in the form of *moving target defense* [8]. Until now, however, moving target defense has only been used to protect data from software-level attacks. This means that the power of virtualization and moving target defense has largely remained in cyber-space.

Distinct from past work, our framework builds a cyber-physical defense system which leverages physical detection mechanisms and cyber defenses to provide cost-efficient data center security from physical attacks.

V. Conclusion

We have presented a novel method for protecting against physical attacks in data centers. Rather than hardening the hardware or erecting physical barriers for attackers, we used three methods of delete, encrypt or move to protect code and data inside virtual machines. Our approach is based on the observation that there is a time delay between attack detection and actual attack, which can be used in the data center setting to activate the defenses. We show how *physical* sensors can be used to warn of impending attack and trigger a *cyber* defense. Such human-aware defenses protect against physical attacks, including the insider threat which can be especially important in data centers where the data center operator is typically not the owner of the code and data running on the server.

References

[1] K. J. Higgins, "The 10 Most Overlooked Aspects of Security," November 2006. [Online]. Available: http://www.darkreading.com/security/application-security/208808177/the-10-most-overlooked-aspects-of-security.html

[2] S. Heare, "Data Center Physical Security Checklist," SANS Institute, Tech. Rep., December 2001. [Online]. Available: http://www.sans.org/reading_room/whitepapers/awareness/data-center-physical-security-checklist_416

[3] R. Blau, "NYPD civilian worker busted in mass cop-ID theft," New York Post, March 4, 2009. [Online]. Available: http://www.nypost.com/p/news/regional/item_8BX3Qmj7PnfxPOdMPZTr8K

[4] J. R. Muir and M. Partner, "Stolen Drives and Servers, Dont Think it Cant Happen in Your Data Center," Trusted Strategies LLC, September 2007.

[5] D. Champagne and R. B. Lee, "Scalable architectural support for trusted software," in *Proceedings of the International Symposium on High Performance Computer Architecture*, ser. HPCA, Jan. 2010.

[6] R. Buyya, "Market-Oriented Cloud Computing: Vision, Hype, and Reality of Delivering Computing as the 5th Utility," in *Proceedings of the International Symposium on Cluster Computing and the Grid*, ser. CCGRID, May 2009.

[7] R. J. Masti, "On the security of virtual machine migration and related topics," Master's thesis, Department of Computer Science, ETH Zurich, 2010.

[8] S. Jajodia, A. K. Ghosh, V. Swarup, C. Wang, and X. S. Wang, Eds., *Moving Target Defense: Creating Asymmetric Uncertainty for Cyber Threats*. Springer, 2011.

[9] *OpenStack Compute: An Overview*, OpenStack. [Online]. Available: http://openstack.org/downloads/openstack-compute-datasheet.pdf

[10] W. Dawoud, I. Takouna, and C. Meinel, "Infrastructure as a Service Security: Challenges and Solutions," in *Proceedings of the International Conference on Informatics and Systems*, ser. INFOS, March 2010.

[11] S. D. Scalet, "19 Ways to Build Physical Security into a Data Center," CSO Security and Risk, Data Protection How-tos.

[12] C. Almond, "A Practical Guide to Cloud Computing Security," A white paper from Accenture and Microsoft, 2009. [Online]. Available: http://www.lmtnet.com/blog/wp-content/uploads/2010/03/guide-to-cloud-computing.pdf

[13] "Motion sensing as a proxy for Data Center "freshness"," Redwood Systems. [Online]. Available: http://www.redwoodsystems.com/node/1178

[14] "Sensaphone Remote Monitoring Solutions." [Online]. Available: http://www.sensaphone.com/data_center.php

[15] L. J. Fennelly, *Effective Physical Security*. Butterworth-Heinemann, 2003.

[16] M. E. Whitman and H. J. Mattord, *Principles of Information Security*. Cengage Learning, 2011.

[17] R. B. Lee, P. Kwan, J. P. McGregor, J. Dwoskin, and Z. Wang, "Architecture for protecting critical secrets in microprocessors," in *Proceedings of the International Symposium on Computer Architecture*, ser. ISCA, June 2005.

[18] J. S. Dwoskin and R. B. Lee, "Hardware-rooted trust for secure key management and transient trust," in *Proceedings of the Conference on Computer and Communications Security*, ser. CCS, 2007.

[19] J. Dyer, M. Lindemann, R. Perez, R. Sailer, L. van Doorn, and S. Smith, "Building the IBM 4758 secure coprocessor," *Computer*, vol. 34, no. 10, pp. 57 –66, October 2001.