jamk.fi

Enhanced physical access detection using environmental sensor information from existing devices

Jani Malkamäki

Master's thesis
November 2018
School of Technology, Communication and Transport
Master's Degree Programme in Information Technology
Cyber Security

jamk.fi

Description

Author(s) Malkamäki, Jani	Type of publication Master’s thesis	Date November 2018
		Language of publication: English
	Number of pages 53	Permission for web publication: x

Title of publication
Enhanced physical access detection using environmental sensor information from existing devices

Degree programme
Master’s degree programme in Information Technology

Supervisor(s)
Kotikoski, Sampo

Assigned by
Ryynänen, Mikko

Abstract

Data centers and the different types of equipment inside racks in these facilities are a critical part of any modern organization’s ICT infrastructure. Data centers are used to store the most important electronic information an organization owns. These facilities are deemed as a high security risk and must be protected accordingly. Many layers of physical and information security controls are deployed to protect the different critical assets inside the data center. Even when this layered defense is implemented, some residual risks may be left behind. Physical security on rack level is in many occasions handled using traditional locking mechanisms which cannot generate information when or where these rack environments have been accessed. This information would be important in the mitigation of this residual risk as it could help answer the question where and when access has been made to these critical rack environments.

The objective was to find a suitable method for detection when these rack environments were physically accessed. The objective was met by creating a method that consists of collection and analysis of device temperature information inside the rack environments. Examining how this temperature information would behave in the situation when a rack door was opened. From the information analysis, a temperature baseline and thresholds were detected which can be used to detect access made to rack environments.

The result was a security control solution that was able to collect, analyze and give trusted indications when an access had been made to the monitored rack environments. Other conclusion was that this method could be with ease extended to monitor the organization’s other similar rack environments.

Keywords/tags (subjects)
Data center, rack environment, physical security, information Security, data analysis

Miscellaneous (Confidential information)

jamk.fi

Kuvailulehti

Tekijä(t) Malkamäki, Jani	Julkaisun laji Opinnäytetyö, ylempi AMK	Päivämäärä Marraskuu 2018
		Julkaisun kieli: English
	Sivumäärä 53	Verkkojulkaisulupa myönnetty: x

Työn nimi
Enhanced physical access detection using environmental sensor information from existing devices

Tutkinto-ohjelma
Master’s degree programme in Information Technology

Työn ohjaaja(t)
Kotikoski, Sampo

Toimeksiantaja(t)
Ryynänen, Mikko

Tiivistelmä

Konesalit sekä niissä toimivat erityyppiset laitteet ovat kriittinen osa jokaisen nykyaikaisen organisaation ICT-rakennetta. Näissä datan varastoinnin kannalta tärkeimmissä kohteissa säilytetään usein yrityksen arvokkaimpia tietoja, joten turvallisuudesta ei voida tinkiä. Useita päällekkäisiä tietoturvakontrolleja asetetaan suojaamaan näiden konesalien kriittisiä laitteita ja informaatiota. Vaikkakin konesalit ovat pääsääntöisesti hyvin suojattuja, saattavat jotkin niiden osista olla suojattuna muita heikommin. Useasti laitekaappien ovet ovat suojattu perinteisemmillä lukitusjärjestelmillä, koska laitekaapeille pääsyä suojaa tarkemmin vartioitu kehä kuten konesalin ovet. Perinteisten lukitusjärjestelmien käyttäminen laitekaappien lukitukseen voi kuitenkin synnyttää riskin, jos ulompi suojauskehä on läpäisty. Perinteiset lukkojärjestelmät eivät pysty keräämään tietoa siitä, milloin tai missä laitekaappien ovet ovat avattu ja pääsy laitteistoon on mahdollistunut.

Tehtävänä oli siis löytää tapa, jonka avulla tämä laitekaappien ovien avaaminen voitaisiin tunnistaa. Tätä varten kehitettiin tunnistusmenetelmä, joka hyödynsi tietoteknisten laitteiden keräämiä lämpötilatietoja. Näiden lämpötilatietojen perusteella luotiin malli, joka pystyi tunnistamaan lämpötilatiedoissa tapahtuvia poikkeavuuksia.

Tunnistetut poikkeavuudet yhdistettiin tilanteisiin, jolloin laitekaapin oli avattuna. Lopputuloksena mallin perusteella rakennettiin tietoturvakontrolli, joka pystyi keräämään ja menetelmää hyödyntäen tuottamaan hyvin luotettavan sekä reaaliaikaisen indikaation, milloin laitekaappien ovia oli avattuna. Tämä kontrolli oli myös erittäin helposti laajennettavissa suojaamaan muita samantyyppisiä laitekaappeja ja niiden sisältämiä kriittisiä laitteistoja.

Avainsanat (asiasanat)
Konesali, Data analyysi, Fyysinen turvallisuus, Tietotuvallisuus, Laitekaappi

Muut tiedot

ACKNOWLEDGMENTS

I would like to give a special thanks to my wife Tarja. Without her support this thesis would never been written.

Contents

Acronyms

CCTV	Closed-circuit television
CRAC	Computer Room Air Conditioning
DAR	Data at Rest
DMA	Direct Memory Access
ETL	Extract Transform Load
HPE	Hewlett Packard Enterprise
HIDS	Host Based Intrusion Detection System
ILO	Insight Lights Out
IDS	Intrusion Detection System
LCP	Liquid Cooling Pack
UPS	Uninterruptible Power Supply
OA	Onboard Administrator
PC	Personal Computer
STD	Standard Deviation

Figures

Tables

1 Introduction

This thesis focuses on data center physical security, addressing the threats that may affect these facilities and assets stored in them. Introducing an innovative way of doing physical access detection at rack level.

1.1 Data centers

Data centers or computer complex as they were called in the past are facilities where most of the organizations' important electronic data and equipment are stored. The first computing and storage equipment demanded plenty of physical space as they were usually very large installations. To house the massive amount of different equipment needed to run these first computers, facilities such as data centers and computer complexes were created. The first pioneers in the field of computer science had to work inside these computer complexes. The access to the computing power of these large mainframes, different types of terminals were used. These terminals were located near the actual mainframe machines, because there was no modern network infrastructure available yet.

With the introduction of personal computers (PCs) by IBM in the year 1981 the situation how and where work was done started to change. The work that previously had to be done inside these computer rooms could now be conducted on these personal computers elsewhere. The creation of more modern and robust network infrastructure changed the situation even more. These new networks created the possibility to use these servers or mainframes over the network even from the other side of the world. The facilities were still used to host mainframes, servers and other types of equipment; however, there was no longer a need to access them inside these computer complexes. This led to the situation where most if not all authorized people left these facilities, the doors and equipment cabinets were locked, and lights were turned off. A modern data center was born.

Currently, these facilities are still accessed, and people do some work there but less often compared to the past. At the present, these facilities are accessed mainly by maintenance personnel responsible for different maintenance tasks related to these data centers. Terminals, equipment and racks inside data centers did not vanish, but are used by administrators and other maintenance personnel working on different types of maintenance tasks which focus on keeping machines and infrastructure running without interruptions. Most of these tasks can be carried out remotely over the network without the need to access the data center at all. Example of a computer complex from the past. The Real Time Computer Complex developed by IBM for the Gemini program IBM (2018).

Figure 1. Real Time Computer Complex, Gemini program

1.2 The negative shift in physical security

From security perspective a negative shift affecting physical security was introduced when people left these computer complexes. The situation changed so that people were not watching over people to the same degree as in the past. Less people with access to these facilities leads to the situation where they will have less eyes on them, which led to the situation where mitigation of this risk had to be addressed. The created gaps are mitigated by creating a layer of different types security controls around these facilities. Controls such as CCTV cameras, security guards and electronic locking mechanisms are designed to protect the information assets from unauthorized physical access Oriyano (2014, 93-101). The arrangement of a set of physical security controls is so-called perimeter defense Messina (2018). While many new physical security controls are introduced to enhance physical security in and outside these facilities, racks and cabinets inside these facilities are still mainly protected by traditional locking mechanisms. These traditional locking mechanisms will provide some level of security, however standard rack and cabinet door function mainly in the principle of keeping dust out not people.

Using traditional locking to physically secure the equipment cabinets may lead to the situation where it is impossible to answer the question when the cabinets holding information and technical assets were accessed. From perspective of an audit trail, real-time security monitoring and post analysis many questions are left unanswered, which will make the investigation and detection of a possible security breach more difficult. Key management plays a significant role when using traditional locking in protecting organization assets. When key management is done correctly, it may in to some degree mitigate the risks concerning these rack environments.

2 The objective of this thesis

As described in the previous chapters there may be some residual risks left even inside a highly protected facility such as data center. For example, using traditional physical security controls, e.g. traditional locking mechanisms in the rack environments should be mitigated somehow.

2.1 The objective

The main objective for this paper was to study and create a working solution for physical access detection focusing on these data center IT rack cabinets, trying specifically to answer the questions where and when these cabinets where accessed. Who and why could not be answered using the data created by the security solution.

The security solution should be able to detect and have the possibility to give real time alarms to other security monitoring systems or used as such. The speed of the detection was not as relevant as the detection accuracy, the amount of false positive detections should be as low as possible and every access made to the rack should be detected.

2.2 Challenges meeting the objective

The main challenge was to find some type of information that is already present in the system which could be used to detect access to the rack environments. One way of meeting the goal of the objective was to focus on the information collected by equipment environmental sensors that are present on all modern devices. Devices such as servers normally use this information to protect themselves from overheating. A hypothesis was created whit the prediction that the same sensor information used to protect these devices from physical damage could be used in an alternative way, using this information in the way it could answer the questions where and when these racks housing these devices have been accessed.

Another challenge on creating the desired solution required gaining more knowledge in different fields of studies including thermal dynamics, data analytics, physical security and more precisely, anomaly detection.

One of the most common rack types in the organization’s data centers was used to create and evaluate the proposed solution. All the testing and information gathering conducted in this thesis took place in one of these equipment racks. The challenge in this part was that the specific rack was still used for production purposes and no negative impact to production should be created from the testing and data collection made by the solution.

The challenge facing the design of solution was that it had to be constructed so that the end solution could be deployed in production after the completion of the evaluation test phase. To meet up with this goal, all the necessary steps to create a working solution included the installation of the metric-based analysis environment where components such as the Elastic stack, Grafana analysis software had to be installed. One of the most challenging parts of the work was the creation of the data collection queries using Logstash and Linux shell scripts.

2.3 Research questions

The main research questions are as follows:

1. What are the limitations of the currently implemented traditional physical security controls possibly affect the effectiveness of real time physical access detection?

2. Could these limitations pose a risk for the information assets?

3. Can the residual risks left by traditional locking mechanisms be mitigated using the temperature information collected from the device temperature sensors?

Literature could be used to answer the first two research questions; however, the third research question could only be answered by testing it. A case study was created to address the third question as it was seen to be the most suitable research method.

3 Security policies controls and limiations of these controls

This chapter takes a short look on different types of policies and security controls used to protect assets such as IT rack cabinets. Weaknesses these traditional physical controls e.g. CCTV cameras and locking mechanisms are also addressed as well as the exploitation of these weaknesses.

3.1 Security policies and risk management

Organizational risk management processes are set in place to determinate the level of security in which assets are to protect. Kim & Solomon define risk management as the process of identifying, assessing, prioritizing and addressing risks (Kim & Solomon 2014, 128). Risk management is used to evaluate the level of risk that may affect an asset. Any private or governmental actor who handles information that must be protected has to implement security policies, physical and informational security controls. When these polices, or controls are broken, it may lead to great harm for the organization or the information owner. Any data breach or downtime on services of any organization or government will have a negative impact on their reputation and may lead to financial losses.

Physical security of a data center and server room can be deployed in a wide variety of ways. Perimeter defense is normally used in the scenarios when, as stated by Messina (2018), the facility is deemed to be a high security risk. An external perimeter may be considered an additional layer of protection. Perimeter defense consist of security controls e.g. walls, doors, locks, security guards and other physical methods of protecting the assets residing in the inner perimeters of a data center (Oriyano 2014, 93-101).

Physical controls protecting assets are designed based on the physical security policies that for most parts only focus on the physical security. The physical security policies dictate, for example what kind of locking should be used, how high the wall should be around the perimeter and where guards should be placed. Physical security controls such as security guards and CCTV monitoring are normally responsible for intrusion detection in the physical layer and are separated from technical and informational security controls such as firewalls, malware defense and intrusion detection systems.

3.2 Limitations of physical and information security controls

Traditional physical security controls such as locks and their effectiveness are hard to measure because they will not give information about their effectiveness on preventing unauthorized access. For example, a traditional lock on a door will not collect metrics on how many times or when it has been opened or if someone had access to the physical assets it is protecting.

Key management plays a significant role when traditional locking is used. However, it will not specifically answer the question where this key was used to gain access. In a case of traditional key management is used to administer the access to the rack environments, the security of the information assets completely relies on this key management process its policies. If a key has been stolen or misplaced, the security of the physical and information assets will depend on how this key management process works and how long it takes for this process to detect this missing key. In many occasions, this will mean a delay in detection and will create a long-lasting negative effect on the security of the protected assets.

Even if using a more advanced locking mechanism there can be limitation as described by Williamson in the following way. Traditionally, access to individual racks has always been protected by key-based systems with manual access management. In some instances, data center owners have turned to a more advanced coded key system, however even this approach provides little in the way of security and no record of who has accessed the data center, making the collation of accurate audit reports practically impossible (Williamson 2018). If the electronic locks or electronic locks with biometric identification have possibility to relay information to other security systems, they can be used to answer all the questions of who, when and where. These types of more advanced locking mechanisms are usually installed only on the outer perimeters of the data center, such as data center door or server room doors but not on every individual rack cabinet.

Other physical controls such as CCTV cameras can be useful deterrent and a tool for higher level of physical security (Messina 2018). A physical security breach is detected by a burglar alarm CCTV camera system or other means relies on a security person who monitors and dispatch appropriate personnel to investigate the possible breach. The delay from detection to mitigation can give the attacker a long time to do misconduct on the targeted information and physical assets. The possibility of this control to provide real-time protection for the assets depends on how the CCTV monitoring is conducted and how the well the personnel are trained to do their job. If there are no security personnel monitoring these cameras around the clock this will make camera recordings usable only in post study of the possible incident (Messina 2018). This gap can be reduced by introducing better physical security controls, e.g. having more security personnel near the information assets and monitoring the environments 24/7. The arrangement where security personnel are stationed around the clock can be costly and not meaningful in all situations. Considerations on the placement of CCTV cameras focus on protection of perimeter entrances and critical access points (Oriyano 2014, 98). However when CCTV cameras are only installed or face the outer perimeters, monitoring the inner perimeter of the facility will be impossible. It will be hard to identify the specific assets that were compromised in the inner perimeters of, for example the data center.

These few examples show some of the weaknesses that traditional physical security controls may have. Not all is lost because these controls supplement each other, and they can be improved by merging even more information from different security controls to get a more comprehensive security situation picture where incident detection speeds and quality can be higher. When this security information is spread through multiple security layers with no information exchange conducted between these layers protecting everything everywhere with advanced security controls will come with a hefty price tag.

Different layers of the information security are descripted usually in the following manner as described by the InfoSec stating that successful organizations have layers of security.

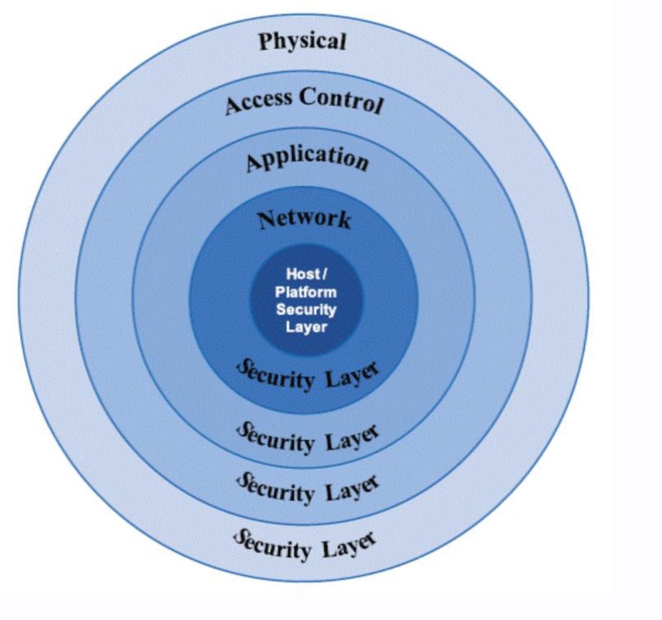

Figure 2 Different security layers (The InfoSec Handbook 2018)

As mentioned, it is the author’s opinion that one of the main and possible the worst limitation is the fact that physical security layer and the controls protecting this layer rarely provide security information to the information security layers. As this physical access detection information is not available to the informational security layer these security controls focus on the electronic information. There should be more focus on how the physical security layer and controls could be designed so that they would have the possibility to communicate with the information security layer. When these different layers start talking to each other in real-time, a higher level of security is possible.

there is a low amount of information flow currently available from the physical security layers to the other layers. The information that is gathered in the designed security solution is fundamentally residing in the physical layer and will be used to enhance physical access detection. The information that is collected and used to detect changes at the physical level, will be used to protect the information assets on the lower layers of the information security framework. This was one of the key findings in this thesis.

3.3 Technical information security controls on host level

As this work focuses mainly on the physical security level, information security controls that are discussed that are also used to protect assets against the host level attacks. Existing informational security controls on host level are made of wide variety of tools like host-based intrusion detection system (HIDS). These tools are designed to detect intrusion attempts at the host level and when correctly implemented can greatly reduce the different risks posed on the information assets.

While risks can be reduced by implementing these advanced information security controls, they cannot be used to protect equipment such as backup devices. Backup devices e.g. tape libraries may hold the same critical information that the servers hold. These tape devices can be one point where the critical information is extracted undetected. Stolen or missing tapes from these devices may be left unnoticed depending again on how the devices are monitored. Encryption is one way to mitigate the effect that an organization would be affected through this kind of data theft. Data encryption and network traffic encryption plays a critical part in any security aware organization. The encryption of information should also be a normal a practice in every organization.

Again, nothing is perfect as many encryption products protect data on the end-point machines use encryption data at rest (DAR) type of security, which means that data is open after encryption keys are provided by a user who possesses these keys. There are methods that include direct memory attacks (DMA) where using the suitable ports such as FireWire or Thunderbolt can lead to direct access to the machine memory where encryption keys are stored. These keys then can be extracted from the memory and then be used to unencrypt the encrypted content. Sometimes even an unlocked terminal may be found and then the attacker does not even have to steal the keys but just go and steal the open data.

These types of attacks demand the physical access to the machine and should never be possible in a secure site e.g. a data center. While not probable, these scenarios are possible and give one more reason for having some monitoring of the physical access implemented to rack environments.

3.4 Missing audit trail

If there are limitations on the collection of physical access information, this may affect the creation of a good audit trail. Audit trail basically all the information that is available for the security personnel when doing post forensic work on a security incident. The main goal is to create a step-by-step story of an incident: what has happened and how.

When one first discovers a computer crime, one must document exactly what events occurred (Easttom 2018). If there is the reason to believe that misconduct has taken place inside a data center or computer room, it may be in some cases difficult and time consuming to precisely identify what equipment or information may have been affected by the incident. The size of the data center, location of data center will affect the time it takes to find the equipment that may have been part of the incident, as time is usually a critical factor in mitigation of the possible after effects that different kinds of data breaches will have. What was stolen and how will it affect the company; decide what measures must be taken when a possible data breach has happened. Speed of the investigation is critical when carrying out a security analysis, as time goes by some of the information may be lost. This is the reason why physical access to environments should be collected.

3.5 Misconduct of personnel and insider threads

In some cases, the threat does not come from an outside source. Misconduct or accidents caused by personnel may lead to a situation where something has gone wrong, e.g. unplugging a wrong power cord or accidentally turning off a wrong server. These accidents still happen from time to time. Depending on the time it takes to correct the mistake, a situation may arise when some downtime to services will be introduced. When detecting what specific rack cabinets are accessed in a specific point in time, the proposed solution could also be used to mitigate the risk where an unintentional accident could happen. The solution should be uses to check that the maintenance personnel are accessing the right cabinet and not relying that they know that it is the right one.

One of the worst treats that an organization face is a Insider threat. It manifests in the form of a person working for the company or a subcontractor who has access to secure information and can even exploit the knowledge of the organization security controls to steal information undetected. The actor in some cases may even have access to the datacenters and this way have access to the organization's most important information. When gaining this access, the actor could have the possibility to extract secure information without any physical or information security controls even noticing this taking place, e.g. stealing backup tapes, hard disks or extracting information from the servers.

Insiders can cause a great deal of misconduct and no organization would like to acknowledge this threat exists, at least in their own organizations. As stated by (Easttom 2015, 147), case studies demonstrate that industrial espionage is not some exotic fantasy dreamed up by a paranoid security expert. It is an unfortunate, but quite real, aspect of modern business. While normally attackers try to access to data thought the network infrastructure cases of insiders preaches are still a considerable risk to information assets. As the figure 3 illustrates outsiders were found to be responsible for 45 percent of the attacks recorded in 2014, 55 percent of attacks were carried out by those who had insider access to organizations' systems. (IBM 2015)

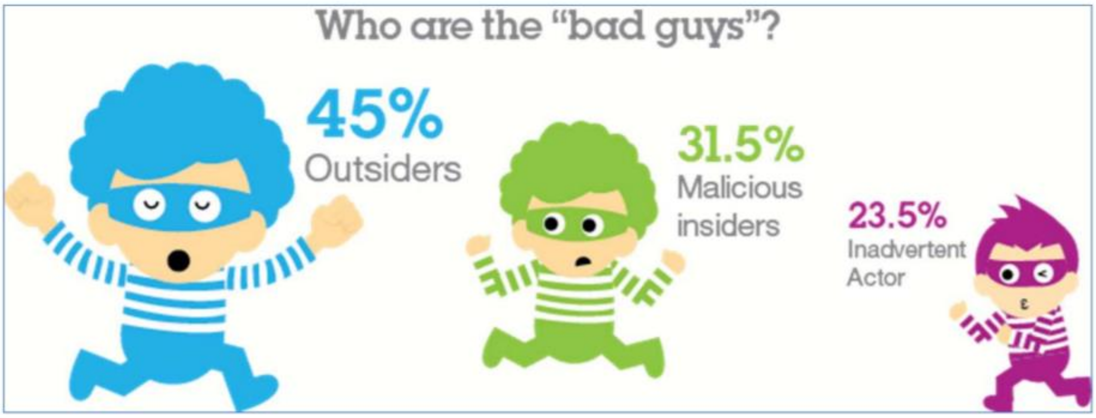

Figure 3 Who are the bad guys? (IBM 2015)

European Union Agency for Network and Information Security (ENISA) has also listed the top threats, and insider the threats are listed in the top 10 in years 2016 to 2017. The insider threat is also at the top 10 and the trend has been stable. Figure 4 created by ENISA illustrates to different top 10 trends and the change in ranking.

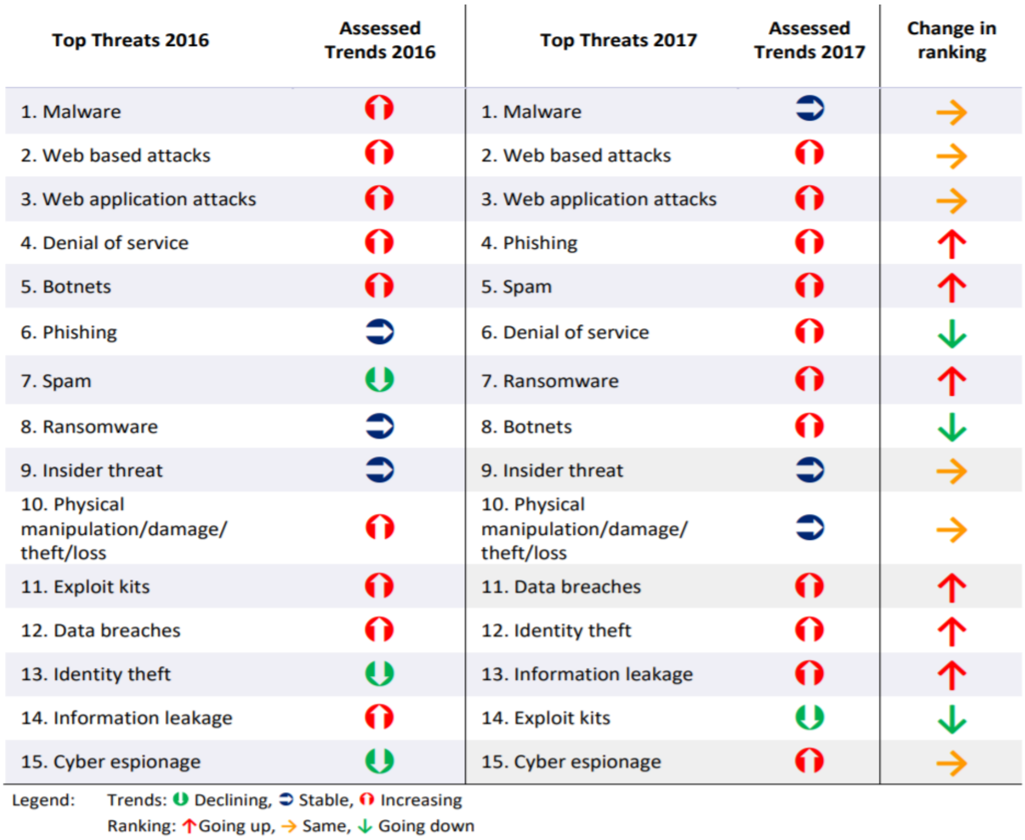

Top Threats 2016	Assessed Trends 2016	Top Threats 2017	Assessed Trends 2017	Change in ranking
1. Malware	⇧	1. Malware	⇨	→
2. Web based attacks	⇧	2. Web based attacks	⇧	→
3. Web application attacks	⇧	3. Web application attacks	⇧	→
4. Denial of service	⇧	4. Phishing	⇧	↑
5. Botnets	⇧	5. Spam	⇧	↑
6. Phishing	⇨	6. Denial of service	⇧	↓
7. Spam	⇩	7. Ransomware	⇧	↑
8. Ransomware	⇨	8. Botnets	⇧	↓
9. Insider threat	⇨	9. Insider threat	⇨	→
10. Physical manipulation/damage/ theft/loss	⇧	10. Physical manipulation/damage/ theft/loss	⇨	→
11. Exploit kits	⇧	11. Data breaches	⇧	↑
12. Data breaches	⇧	12. Identity theft	⇧	↑
13. Identity theft	⇩	13. Information leakage	⇧	↑
14. Information leakage	⇧	14. Exploit kits	⇩	↓
15. Cyber espionage	⇩	15. Cyber espionage	⇧	→

Legend: Trends: ⇩ Declining, ⇨ Stable, ⇧ Increasing
Ranking: ↑Going up, → Same, ↓ Going down

Figure 4 Figure 1 ENISA Top threats 2016-2017

3.6 Section summary

Information and physical security layers dependent on each other and have their reason for existence. While technical information security controls have evolved and keep on evolving to be better in protecting information assets the physical security controls have stayed more traditional as they are seen to be efficient in the protection of the physical layer. In the author’s opinion this difference in how the layers are protected creates a gap between the layers that cannot be fixed without more real-time information exchange between layers. Understanding the methods and controls used in physical and information security layers will help control creators combining these two domains. There are many reasons behind why there is not enough communication between these two layers.

One of the main reasons is the separation of the knowledge and duties of the different professionals responsible of protecting these different security layers. Understandably, professionals working with the physical security controls are not allowed to access the information systems that they are protecting, and the other way around the owners of the information assets may not have access to the monitoring and control systems used by the security personnel. If this information exchange does not take place one must be aware of the gap this creates. Sharing real time situation information between these different actors and security controls could be possible; however, it is challenging to create.

When looking at rack level security it is still common that many of the physical security controls are traditional type and the access controls are completely reliant on the fact that there is trust between the people accessing these racks and the information owner. In other words, a written policy document protects these information assets from unauthorized physical access.

As Williamson states, as data center threats continue to increase, technologies such as electronic-locking swing handles will become even more in demand as companies seek greater protection for their business-critical data. Therefore, security systems must become more intelligent and integrated in order to be affordable to business owners for their peace of mind. The industry has an increasing need for access solutions that combine protection and monitoring functions to combat the often-overlooked threat of physical data breaches (Williamson 2017). To get to this goal of affordable security solutions, solutions like proposed by this paper should be studied more. Every security control does not need to be expensive and hard to manage and a layered defense constructed from multiple security controls can create an even harder perimeter to penetrate. It is crucial to have the best combination of solutions protecting the organization’s assets both on informational and physical security domains.

4 Racks and cooling

Short introduction on data centers and rack environments and the way cooling is managed in these racks is discussed in this chapter. This section opens the concept and us of environmental sensors. This information is relevant for the reader as it explains what information was used to create the physical access detection solution.

4.1 Different types of racks and cooling

Racks, IT racks or IT cabinets are essentially a cabinet designed for storing IT equipment such as servers, network equipment and disk storage units. IT racks are commonly divided into multiple different types including open frame racks, closed rack cabinets and transportable racks. Fully enclosed cabinets or racks are the most common rack type available (The Server Rack FAQ 2007). They are cooled in different ways; the cooling has a critical function as all the equipment inside these cabinets generate heat. The heated air must then be transferred out from the racks and cooled down using a cooler to a lower temperature that can be reused for cooling purposes. This creates a flow of air to flow inside these racks. The circulation or airflow moves the cold air through the rack air conditioning system (CRAC) to the front of the equipment and then is sucked back by the equipment fans. Figure 5 illustrates how this rack air flow in a raised floor cooling solution.

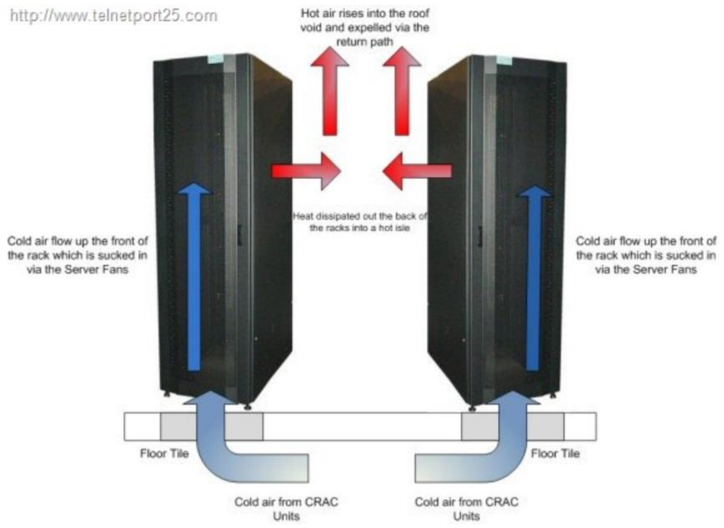
http://www.telnetport25.com
Hot air rises into the roof void and expelled via the return path
Heat dissipated out the back of the racks into a hot isle
Cold air flow up the front of the rack which is sucked in via the Server Fans
Cold air flow up the front of the rack which is sucked in via the Server Fans
Floor Tile
Floor Tile
Cold air from CRAC Units
Cold air from CRAC Units

Figure 5 HP raised floor and rack airflow (telnetport25)

In-Rack is the most precise cooling available, as the rack and the air conditioner operate in a close relationship with one another. Cold air has no choice but to pass through the servers; hot air has no choice but to pass through the heat exchanger (42U Data Center Solutions 2015).

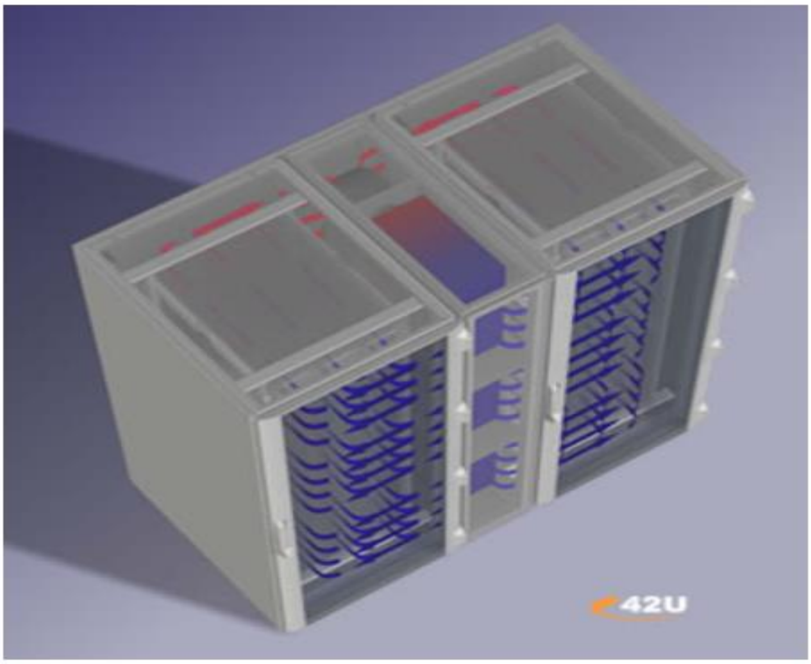

Figure 6 In-rack cooling (42U Data Center Solutions)

This type of rack and in-rack cooling system displayed in the previous figure (Figure 6) is used in this thesis. As this type of rack environment is designed so that the cold air only flows inside the rack it was seen that its suites as a good target for the evaluation of the created security solution. Detecting some changes in the temperature values should be easy when this cooling containment is broken by opening the rack door.

4.2 Environmental sensors

It is safe to say that every modern equipment has environmental sensor installations. These sensors’ main function is to protect the equipment from overheating. When the temperature rises too high, the management chips monitoring these sensors will give instructions for equipment to shut down. This shut down process reduces the probability that the equipment will be damaged.

Figure 7 illustrates the different temperature sensor which are installed in a hp-server-blade. Different sensors are distributed around the server because different components may heat up in different ways depending on how the specific equipment is used. The figure 7 also illustrates the location of the ambient zone sensor. The information collected from the ambient zone insight lights out (ILO) sensor #1 acts as the main information source used in this thesis. This sensor measure the level of ambient air temperature coming to these servers from the rack cooling systems. The location of these sensors is seen as the most suitable for information gathering as they are not near any other heat source. Heat sources such as CPU and other components may affect the temperature information in a way that it would have a negative impact on the physical access detection method.

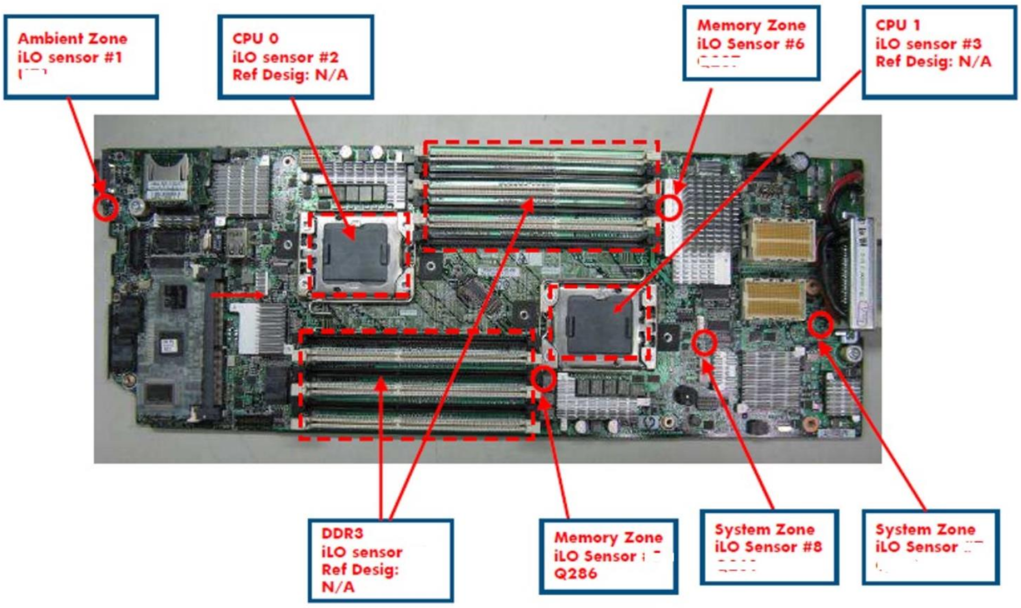

Figure 7 Temperature sensor locations on BL460c server (HPE)

4.3 The hypothesis

Now that the basics of the different components related to the detection method have been discussed, a closer look can be taken at the proposed hypothesis. The hypothesis predicts that the different states that a rack door has can be detected based on the inlet ambient sensor temperature values. As the cold and hot air is contained inside these high-density types of racks, the opening of the rack door should at least in theory break or change the containment, and in this way, also, affect temperature values on the equipment sensors.

A measurement taken from the server inlet temperature sensors located physically nearest to the cold-aisle (cold-air containment) should give sufficient information to get a trusted indication that the rack door has been opened. If the detection works as the hypothesis predicts, there will be a change in the temperature values. If the prediction proves to be true, the information can be analyzed and used as an indicator of physical change in the state of the rack door. This indicator of a change in the physical domain could then be relied to a security monitoring station or even used to launch possible countermeasures to further mitigate the possible security incident. That said, launching countermeasures would need a combination of multiple different security control information and this kind of study is outside of the scope of this thesis.

The prediction of the hypothesis relies heavily on the fact that there should be some temperature differences between the environmental measurement points. If the devices inside the rack and computer room environments have the same temperature values, it is hard or even impossible to find any anomalies in the data. This should not be a problem in the rack environment that is used in this thesis as the measurement taken from the server inlet temperature sensors which are located physically nearest to the cold-aisle (cold-air containment) should give good enough metrics to get a trusted indication that the physical security layer, in this case the rack door is opened.

5 Test environment

Testing of the predictions made on the hypothesis were conducted with real enterprise hardware. A real data center environment was used so that the solution could be implemented from evaluation to production with little or no extra work.

5.1 Rack and equipment installation

The first figure (Figure 8) in this chapter illustrates the basic structure of the data center where the test was conducted. Data collected from the rack cabinet and other data sources including the UPS system and external temperature sensor. Another value of the test environment was the temperature inside the computer room which was around 22 degrees Celsius. The temperature difference was about five degrees between the room and the cooled air flowing inside the test rack environment's cold-aisle. This difference in temperatures between the data center and rack environment was seen crucial for the hypothesis to work as intended.

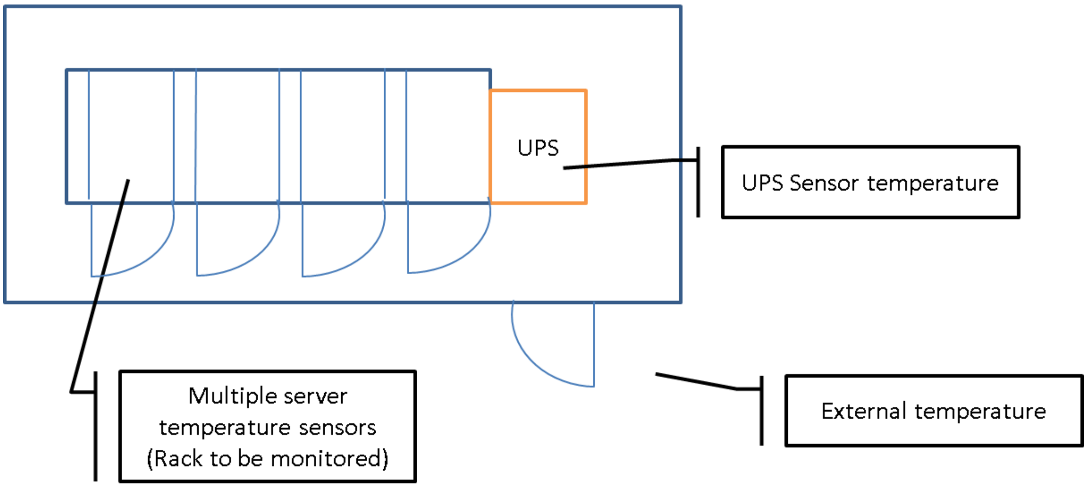

Figure 8 Server Room Illustration

The servers installed in the testing environment consisted of four generation 9 hp-blade-servers and 22 generation one hp-blade-servers. In case of the servers, the number of sensors differ between types, generation and the model of server. In the test cases, one ambient inlet temperature sensor from every server was used to make the detection process possible.

Figure 9 shows the overall illustration of the test environment, the rack to be monitored had two HP C7000 rack enclosures which include a total of 31 servers. Liquid cooling packs (LCP) provided the cooling for the rack cabinet. Six LCP units were present in the test environment.

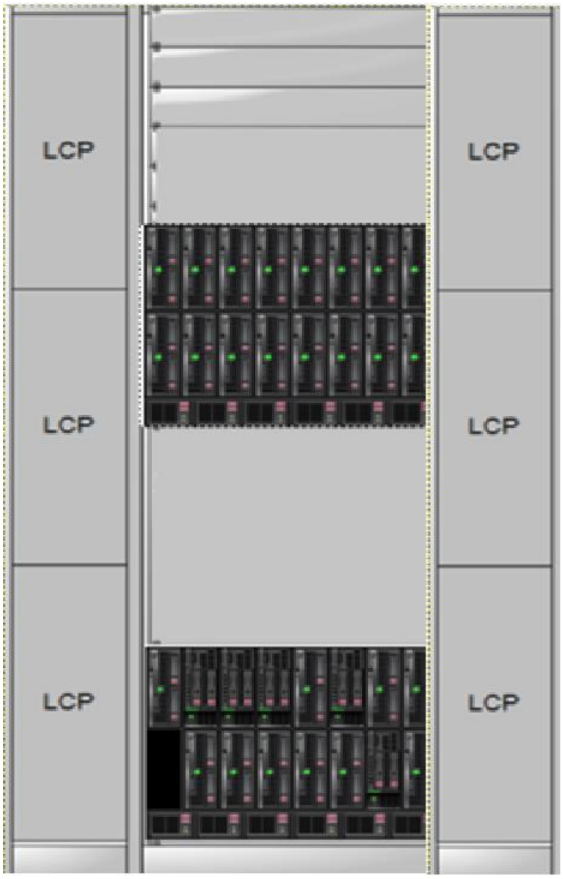

Figure 9 Front side view of the test rack.

5.2 Software architecture and data collection process

Testing the hypothesis and evaluation of the detection solution required collection of the data and the introduction of controlled test cases that would be detectable trough analysis of temperature data.

Multiple various technologies were used. The main analytics database and collection was conducted with Elastic Stack. The temperature information from the onboard administrator (OA) cards was extracted with Linux shell scripts. The function of these cards as explained by HPE (2018), these management cards can provide the information of local and remote status reporting of system health, power usage and cooling status HPE (2018). The OA management card handles this information collection from the management processors of the blade servers. This information was further transformed and uploaded to the Elastic database.

While the Elastic Stack has its own tool Kibana for data query and visualization the main time series analysis was conducted by Grafana. The change from Kibana to Grafana was mainly due to the alerting capabilities of Grafana, as this were one of the goals that this study should meet. The next flowchart, Figure 10 illustrates the whole information data flow and where the tools are situated along the data flow.

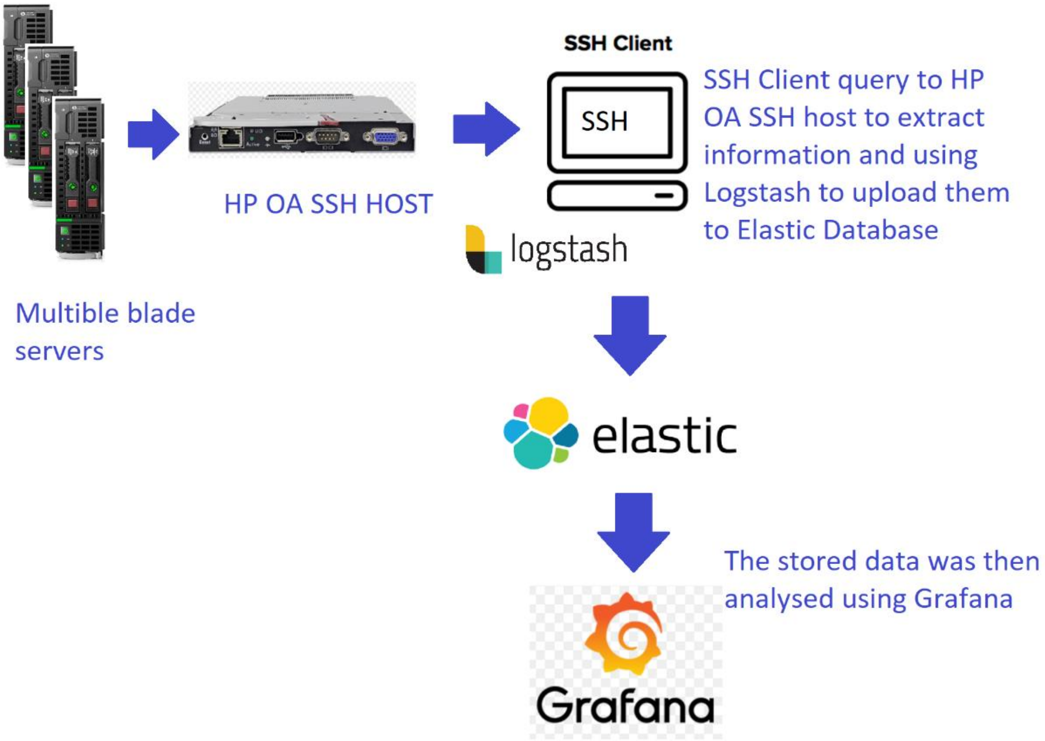
SSH Client
SSH
SSH Client query to HP OA SSH host to extract information and using Logstash to upload them to Elastic Database
HP OA SSH HOST
logstash
Multible blade servers
elastic
The stored data was then analysed using Grafana
Grafana

Figure 10 Information dataflow

Not included in the previous figure was the collection of temperature information from the UPS system, which was carried out using the same SSH client that used SNMP query to extract environmental information from the UPS system.

As previously mentioned, for the data extraction a combination of different technologies was used. The data was gathered using a combination of Linux shell scripts, SNMP queries and Logstash component. The shell scripts were responsible for data query and data transformation alongside with Logstash component, which also applied the mappings of field names and oversaw the loading of the transformed data to the Elasticsearch database.

While it was not clear that the hypothesis would work and would create usable results, the implementation of the testing environment was still decided to be created so that these results could be implemented as part of a working monitoring system. Elastic database can be easily queried to other monitoring systems. The component Logstash responsible for the data loading phase can be configured so that it can write the collected information to many different target databases or even store the collected information to a file. This then supports the information loading, so it can be implemented to any other security information and event management (SIEM) or intrusion detection system (IDS) system. Elastic stack with its different components combined with Grafana provided the most suitable platform for the whole testing and evaluation process. Microsoft Excel was used mainly for offline analysis of the collected data.

Example of one elastic document is shown in Figure 11.

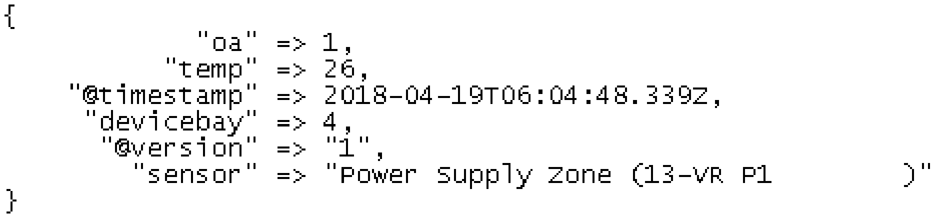

```
{
            "oa" => 1,
          "temp" => 26,
    "@timestamp" => 2018-04-19T06:04:48.339Z,
     "devicebay" => 4,
      "@version" => "1",
        "sensor" => "Power Supply Zone (13-VR P1         )"
}
```

Figure 11. Elastic document

5.3 The field tests

The design of the field test was planned so that a controlled event would be introduced to the monitored environment. The events were: opening the computer room door, the rack’s back and front doors and keeping a log of the times and when these tests were conducted. The main goal was to create points in time that should be detectable through the analysis of the environmental data. Marking time with one-minute intervals was used to indicate when the different states of the rack door were open or closed. Without knowing how long it would take to make a possible detection, it was decided that the doors should be opened at least for a few minutes. In the end, the field tests then ranged from three to ten-minutes.

Table 1 shows an example of the logbook of the test performed on 13.4.2018. The value one indicates when a door was in the opened state. Different parts of the rack and computer room were tested. At this time, data was also gathered from the system; however the data analysis phases were conducted after the field-tests had been completed.

Table 1 Field logbook example

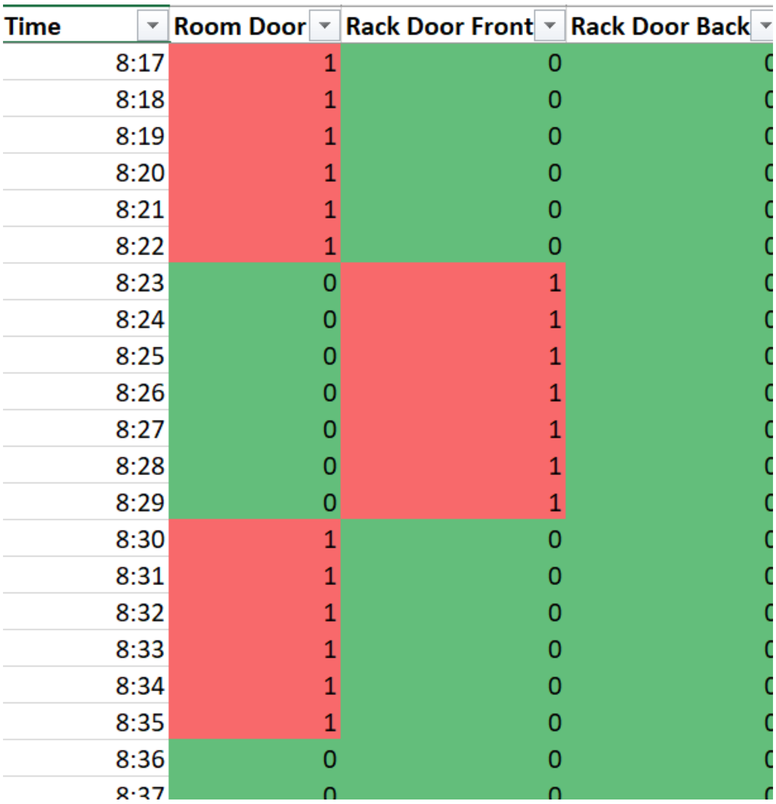

Time	Room Door	Rack Door Front	Rack Door Back
8:17	1	0	0
8:18	1	0	0
8:19	1	0	0
8:20	1	0	0
8:21	1	0	0
8:22	1	0	0
8:23	0	1	0
8:24	0	1	0
8:25	0	1	0
8:26	0	1	0
8:27	0	1	0
8:28	0	1	0
8:29	0	1	0
8:30	1	0	0
8:31	1	0	0
8:32	1	0	0
8:33	1	0	0
8:34	1	0	0
8:35	1	0	0
8:36	0	0	0
8:37	0	0	0

The field-testing phases were carried out in four different dates. These dates were April 11, April 13, April 19 and April 30. The system was in production use on all the chosen test-days. There was no negative impact from the information gathering on the production.

5.4 Excluded test case

One of the first test cases was to see if the UPS system temperature sensor could be used to detect a possible access to the computer room. The first analysis focusing only on the UPS environmental data. The conclusion from the analysis for this specific test was that opening the computer room door and monitoring the UPS environmental information did not show any promising changes in temperature values. This information then would not be used be used in physical access detection purposes. The temperature values outside and inside the computer room were too close to each other. This finding let to the conclusion that as there would not be enough variance in the temperature values between these two datapoints. Trying to detect an field-test introduced event based on this low variance between these datapoints would not be possible. As the UPS system monitoring did not show any promise the later tests only focused on the rack from door and the possibility to detect the different states in temperatures inside the rack environment.

Chapter 5 summary

Along with the designing and execution of the field tests plenty of software had to be installed to create the base for the whole solution. Transformation of collected data had to be done before the whole data collection process was complete. The different software components had to be installed and configured correctly, cleaning up the data with Logstash and creating the correct configuration files for it. SSH authentication keys had to be installed to the onboard administrator cards so that queries without passwords would be possible. The Elastic database had to be installed and the workings of this analysis software figured out; how it could be used against the collected data was the biggest part of the whole work. Plenty of work had to be done before the main data analysis could be conducted.

6 Research results and detected anomalies

The results from the different test are described in this chapter with what was found and how it was found as well as the detected temperature baselines and identifying possible anomalies from this baseline. More understanding was gained how different mathematical approaches could be used to detect these anomalies. It is also discussed here what was incorrect in the initial prediction how temperature values may change in the even that a door was opened.

6.1 Data vs prediction

The first analysis focused on how the predictions made by the hypothesis would correlate with the collected data of the first field-tests. The hunt for possible anomalies started by analyzing the data with simple excel spreadsheets. As there was around seven-degree Celsius temperature difference between the rack environment and server room, this difference in temperatures led one to believe that opening the rack door would lead to a temperature raise on every ambient inlet sensor. The collected information did not support this prediction.

The first finding showed some increasing and decreasing temperature values inside the rack environment; however, not only an increasing trend in these values. In conclusion, the prediction made by the initial hypothesis would not be true. This led to the situation where new ways of analysis of the data had to be considered. Some new mathematical methods were used to get to the desired point where an anomaly would be detected from the data sets.

Table 2 illustrates the way that the ambient inlet sensor data was affected when rack door was opened. The rack open state is named in the following chart as LK1 Front. As seen in the table, there was a combination of increasing temperatures and some temperature degreases in the situation where the rack door was open. No concurrent rise of every ambient sensor information was detected.

Table 2 Excel Chart of test case 1

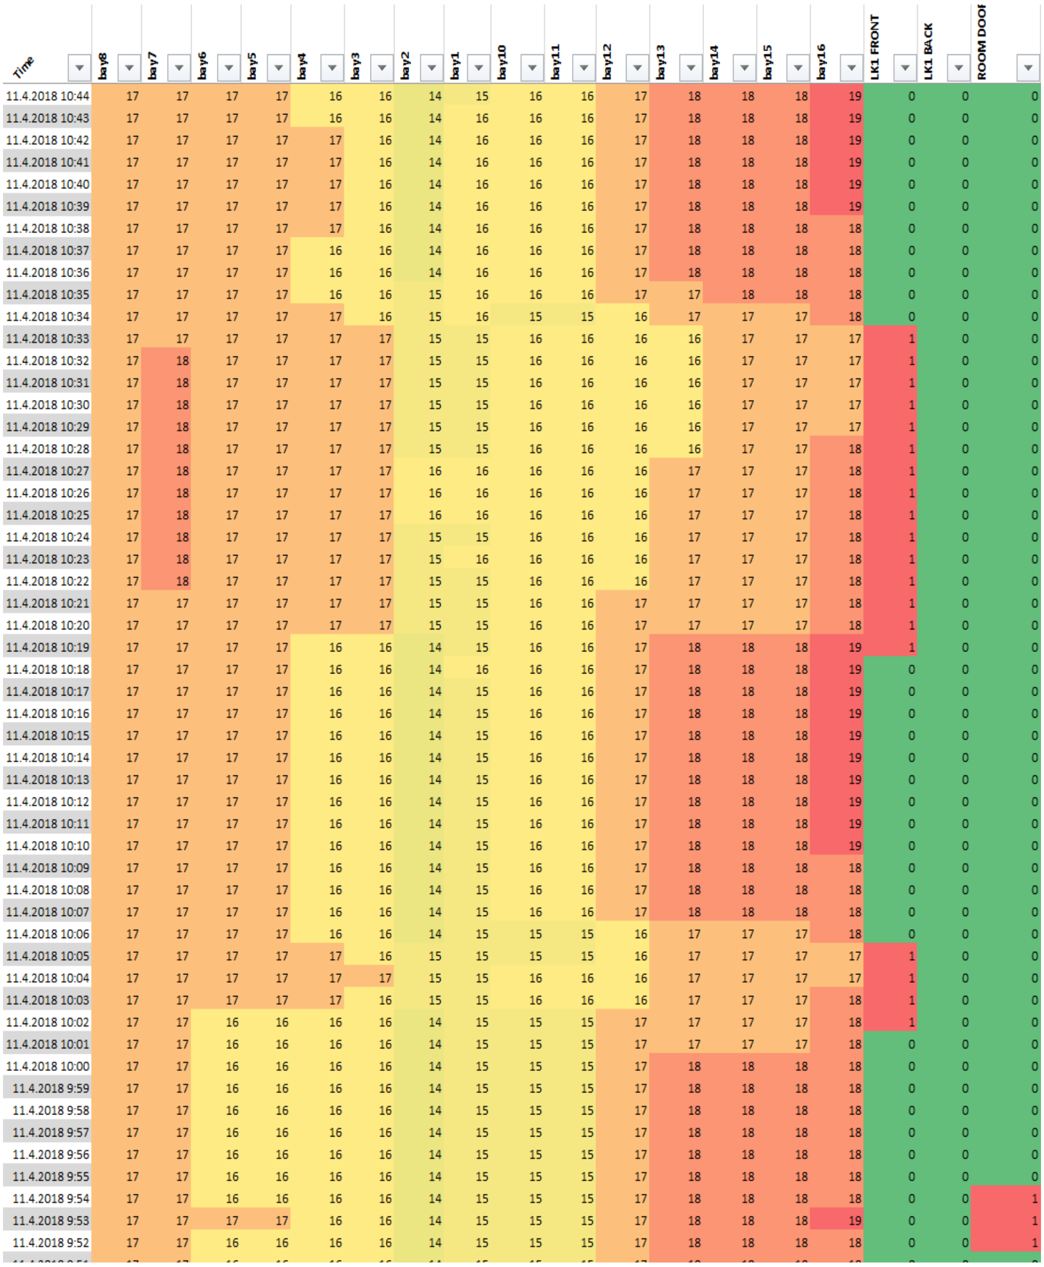

Time	bay8	bay7	bay6	bay5	bay4	bay3	bay2	bay1	bay10	bay11	bay12	bay13	bay14	bay15	bay16	LK1 FRONT	LK1 BACK	ROOM DOOF
11.4.2018 10:44	17	17	17	17	16	16	14	15	16	16	17	18	18	18	19	0	0	0
11.4.2018 10:43	17	17	17	17	16	16	14	16	16	16	17	18	18	18	19	0	0	0
11.4.2018 10:42	17	17	17	17	17	16	14	16	16	16	17	18	18	18	19	0	0	0
11.4.2018 10:41	17	17	17	17	17	16	14	16	16	16	17	18	18	18	19	0	0	0
11.4.2018 10:40	17	17	17	17	17	16	14	16	16	16	17	18	18	18	19	0	0	0
11.4.2018 10:39	17	17	17	17	17	16	14	16	16	16	17	18	18	18	19	0	0	0
11.4.2018 10:38	17	17	17	17	17	16	14	16	16	16	17	18	18	18	18	0	0	0
11.4.2018 10:37	17	17	17	17	16	16	14	16	16	16	17	18	18	18	18	0	0	0
11.4.2018 10:36	17	17	17	17	16	16	14	16	16	16	17	18	18	18	18	0	0	0
11.4.2018 10:35	17	17	17	17	16	16	15	16	16	16	17	17	18	18	18	0	0	0
11.4.2018 10:34	17	17	17	17	17	16	15	16	15	15	16	17	17	17	18	0	0	0
11.4.2018 10:33	17	17	17	17	17	17	15	15	16	16	16	16	17	17	17	1	0	0
11.4.2018 10:32	17	18	17	17	17	17	15	15	16	16	16	16	17	17	17	1	0	0
11.4.2018 10:31	17	18	17	17	17	17	15	15	16	16	16	16	17	17	17	1	0	0
11.4.2018 10:30	17	18	17	17	17	17	15	15	16	16	16	16	17	17	17	1	0	0
11.4.2018 10:29	17	18	17	17	17	17	15	15	16	16	16	16	17	17	17	1	0	0
11.4.2018 10:28	17	18	17	17	17	17	15	15	16	16	16	16	17	17	18	1	0	0
11.4.2018 10:27	17	18	17	17	17	17	16	16	16	16	16	17	17	17	18	1	0	0
11.4.2018 10:26	17	18	17	17	17	17	16	16	16	16	16	17	17	17	18	1	0	0
11.4.2018 10:25	17	18	17	17	17	17	16	16	16	16	16	17	17	17	18	1	0	0
11.4.2018 10:24	17	18	17	17	17	17	15	15	16	16	16	17	17	17	18	1	0	0
11.4.2018 10:23	17	18	17	17	17	17	15	16	16	16	16	17	17	17	18	1	0	0
11.4.2018 10:22	17	18	17	17	17	17	15	15	16	16	16	17	17	17	18	1	0	0
11.4.2018 10:21	17	17	17	17	17	17	15	15	16	16	17	17	17	17	18	1	0	0
11.4.2018 10:20	17	17	17	17	17	17	15	15	16	16	17	17	17	17	18	1	0	0
11.4.2018 10:19	17	17	17	17	16	16	14	15	16	16	17	18	18	18	19	1	0	0
11.4.2018 10:18	17	17	17	17	16	16	14	16	16	16	17	18	18	18	19	0	0	0
11.4.2018 10:17	17	17	17	17	16	16	14	15	16	16	17	18	18	18	19	0	0	0
11.4.2018 10:16	17	17	17	17	16	16	14	15	16	16	17	18	18	18	19	0	0	0
11.4.2018 10:15	17	17	17	17	16	16	14	15	16	16	17	18	18	18	19	0	0	0
11.4.2018 10:14	17	17	17	17	16	16	14	15	16	16	17	18	18	18	19	0	0	0
11.4.2018 10:13	17	17	17	17	16	16	14	15	16	16	17	18	18	18	19	0	0	0
11.4.2018 10:12	17	17	17	17	16	16	14	15	16	16	17	18	18	18	19	0	0	0
11.4.2018 10:11	17	17	17	17	16	16	14	15	16	16	17	18	18	18	19	0	0	0
11.4.2018 10:10	17	17	17	17	16	16	14	15	16	16	17	18	18	18	19	0	0	0
11.4.2018 10:09	17	17	17	17	16	16	14	15	16	16	17	18	18	18	18	0	0	0
11.4.2018 10:08	17	17	17	17	16	16	14	15	16	16	17	18	18	18	18	0	0	0
11.4.2018 10:07	17	17	17	17	16	16	14	15	16	16	17	18	18	18	18	0	0	0
11.4.2018 10:06	17	17	17	17	16	16	14	15	15	15	16	17	17	17	18	0	0	0
11.4.2018 10:05	17	17	17	17	17	16	15	15	15	15	16	17	17	17	17	1	0	0
11.4.2018 10:04	17	17	17	17	17	17	15	15	16	16	16	17	17	17	17	1	0	0
11.4.2018 10:03	17	17	17	17	17	16	15	15	16	16	16	17	17	17	18	1	0	0
11.4.2018 10:02	17	17	16	16	16	16	14	15	15	15	17	17	17	17	18	1	0	0
11.4.2018 10:01	17	17	16	16	16	16	14	15	15	15	17	17	17	17	18	0	0	0
11.4.2018 10:00	17	17	16	16	16	16	14	15	15	15	17	18	18	18	18	0	0	0
11.4.2018 9:59	17	17	16	16	16	16	14	15	15	15	17	18	18	18	18	0	0	0
11.4.2018 9:58	17	17	16	16	16	16	14	15	15	15	17	18	18	18	18	0	0	0
11.4.2018 9:57	17	17	16	16	16	16	14	15	15	15	17	18	18	18	18	0	0	0
11.4.2018 9:56	17	17	16	16	16	16	14	15	15	15	17	18	18	18	18	0	0	0
11.4.2018 9:55	17	17	16	16	16	16	14	15	15	15	17	18	18	18	18	0	0	0
11.4.2018 9:54	17	17	16	16	16	16	14	15	15	15	17	18	18	18	18	0	0	1
11.4.2018 9:53	17	17	17	17	16	16	14	15	15	15	17	18	18	18	19	0	0	1
11.4.2018 9:52	17	17	16	16	16	16	14	15	15	15	17	18	18	18	18	0	0	1

The example table (Table 2) shows the temperature changes on each of the server device bays on the C7000 blade chassis located on the bottom of the rack. The coloring schema from green to red was applied to highlight different values. Higher temperatures across device-bays are shown with deeper reds. The field-test cases and their events are on the right most columns and are colored independent of the temperature values.

As previously mentioned, the table shows when the rack door was open (LK1 Front). The table also shows the data center room door (Room Door) and, the rack back door (LK1 Back). The values on the three right most columns have the following meaning where value one indicates the door was in an opened state and zero that door was closed in the controlled field-test with other values at the time of this specific test case ambient temperature of the server room was 22 degrees Celsius. The LCP units were introducing cold-air to the inlet ambient zone. This airflow temperature from the first LCP ranged from 15 to 17 degrees Celsius. The second LCP units was operating at temperatures ranging from 17 to 18 Degrees Celsius. Figure 12 illustrates how the actual directions of the airflow inside the test-rack

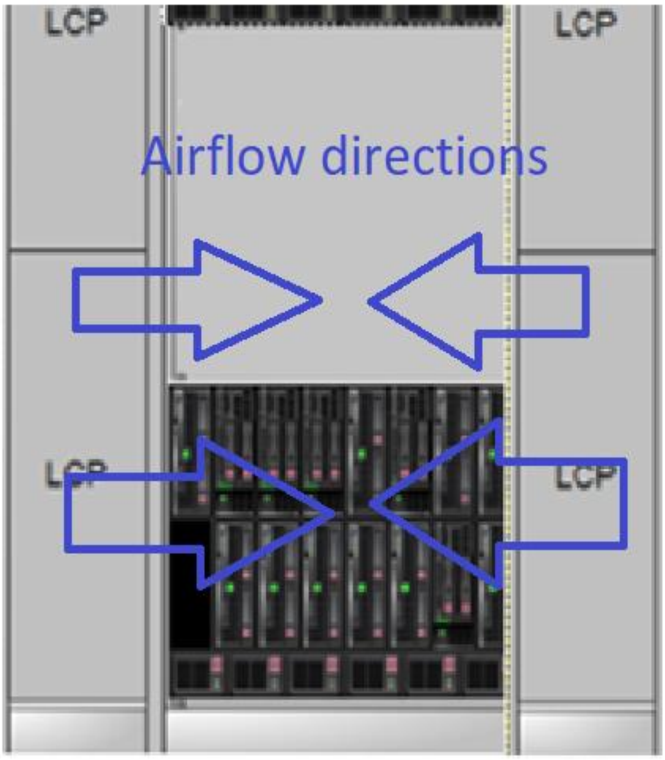

Figure 12. Airflow direction

What was not known in the first prediction aws how the airflow from LCP units would affect the system when the rack door was open. The CRAC unit does not stop providing cooled air to the system and this will lead to the situation that the temperatures will not raise as fast as predicted. The CRAC also dynamically adjusts the cooling power when a raising temperature is detected with the CRAC sensors. The dynamic operation will keep the temperature in the rack at a certain level even if the door is open.

The sensors nearest to the left LCP units such as bays 1 or 10 will be affected by cooler air than the middle and right parts of the rack. An example table (Table 3) the average server inlet ambient temperatures of the blade servers in a period of 8 of April to 13 of April. The average values are calculated from 7,200 datapoints.

Table 3. Average Temperature 7,200 data points

bay1	bay2/bay10	bay3/bay11	bay4/bay12	bay5/bay13	bay6/bay14	bay7/bay15	bay8/bay16
15	16	16	17	18	18	18	18
	16	16	17	18	18	18	18

6.2 Average temperature in anomaly detection purposes

The conclusion of the first analysis of the information was that the temperature may rise but this can also happen in the situation that the CRAC unit starts to a warmer air-flow. When using average maximum temperature as a threshold for indication of an anomaly the rate of false detections would be too high. The next figure (Figure 13) taken from the time that test case one was conducted shows the average temperature of the server ambient temperature sensors. The red line displays the state of the rack door ranging from zero to 16 where the numeric value of 16 indicates when the rack door was open.

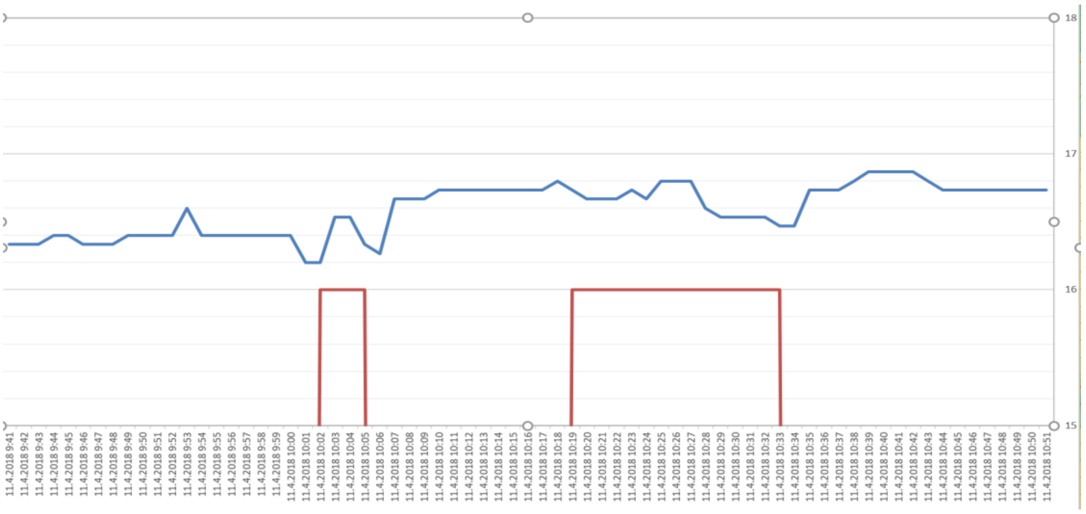

Figure 13. Average ambient sensor temperatures

There is a slight change in the temperature values when the controlled test was done; however, not enough for a reliable anomaly to be detected. The average temperatures would not work in the anomaly detection purposes.

As stated by Easttom (2014), it is often quite difficult to establish proper threshold values or proper time frames at which to check those threshold values. This statement proves to be true when using the average temperature values.

6.3 Standard deviation

Why does the standard deviation work better than average values? As there is a total of 31 servers, there is the same amount of temperature value samples every minute. As can be seen in Table 2, the inlet temperature values started to move closer to each other and this led to a new prediction that this could be the case in every situation where the rack door was open.

Standard deviation explained by Weaver, a low value for the standard deviation indicates that the data points tend to be close to the mean or the expected value of the set, while a high value indicates that the data points are spread out over a wider range. The area under each of the three curves in the diagram below are the same, representing the same population. The values could be millimeters or days depending on what’s being measured: where the SD is small, at 0.5, everything is clustered close to the mean, some 98% of all the values will be in the range of +2 to -2. Where the SD is large, at 2.0, a much larger number of measurements are outside of the +2 to -2 range (Weaver 2016). The Figure 14 illustrates the statement made by Weaver.

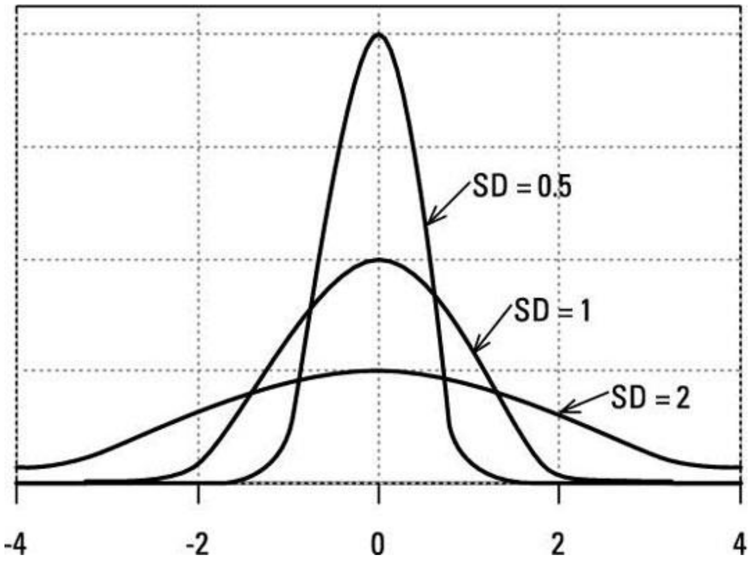

Figure 14 Standard deviation (Waver 2016)

From the same time interval, the mathematical model was changed from average values to calculation of standard deviation using the server ambient inlet sensors as the sample population. The states of the door indicated in the next figure (Figure 15) are value one as an open rack door and zero as a closed rack door. The same test-case one was used.

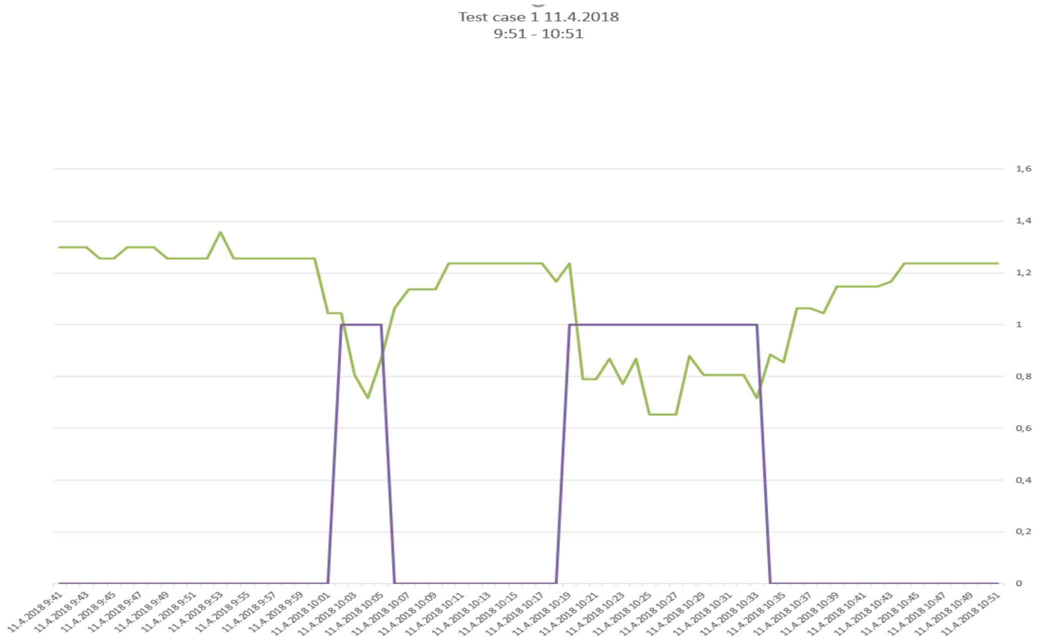

Figure 15 Standard Deviation Test 1

From the test-case one analysis of effect on the standard deviation result one can see how low the standard deviation value dropped, and this possible detection threshold was also used for other detections. As stated by Cornelissen (2013, 445), a threshold can occur in several ways and the method of reacting to them varies accordingly to the way design the monitor. The monitor will be designed so that when the standard deviation drops to certain threshold a detection has been made.

6.3.1 Sample Standard Deviation

Using all the available ambient temperature sensor information as data-points while gives us good results in detection perspective also creates the situation that the number of sensors used in the test was relatively high. If using all the servers that were installed in the rack environment the total would be 31 server sensors. Minimizing the number of sensors will in some cases make the proposed solution more usable. If the detection relies on all the data-points to be available, this leads to the situations where there are not enough data-points available and the detection would fail. Dropping down the number of sensors used as data point had to be considered. The beauty of standard deviation is a part of the population can be taken and a standard deviation calculation done against this population and one can still reach good results. With the same standard deviation model, the temperature value count is dropped to half of the total. Selecting only the bottom rack chassis and collecting server sensor temperature information from these servers. Now the total monitored sensor is dropped to 15 and the detections illustrated in the following Figure 16 with detections highlighted with yellow.

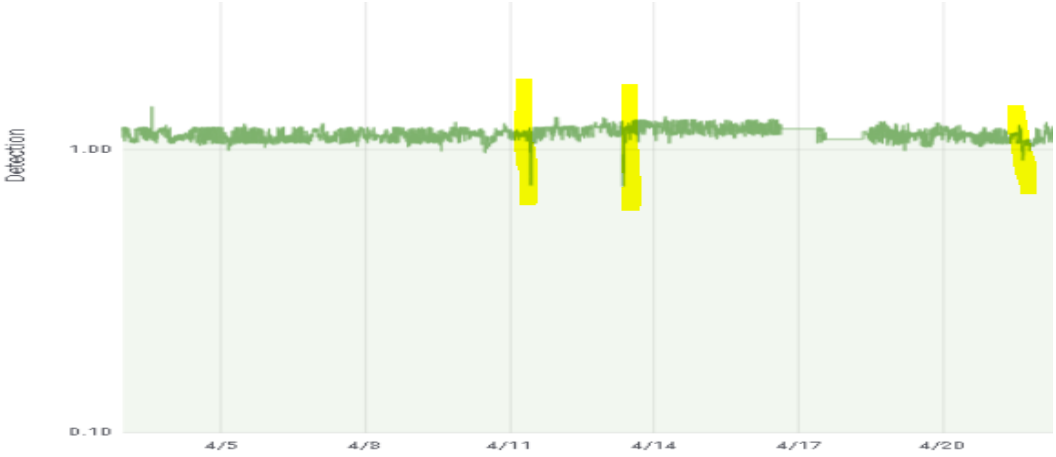

Figure 16. STD 15 Sensors

If the standard deviation threshold is below 0.85, this indicate that an open-door state may be detected. This value would be adjusted if too many false positive results are detected. Running the of 0.85 standard deviation against a dataset with 7202-minutes of information and show only rows below this 0.85 standard deviation threshold. The following table is created, note that the information presented in the next chart only include test case one and two. The naming for the next table is as follows.

The first column of Table 4 indicates the time and the second column the calculated value for the standard deviation for all the server inlet sensor value population. The value 1 in columns LK1 Front2, LK1 Back3 and Room_Door4 in this table indicate whan a controlled open-door test has been made in these specific elements. The last column shows the number 1 if standard deviation is lower than 0.85. This last column was used to filter only the rows that would be under the detected threshold.

Table 4 7202-minute dataset and standard deviation threshold detection

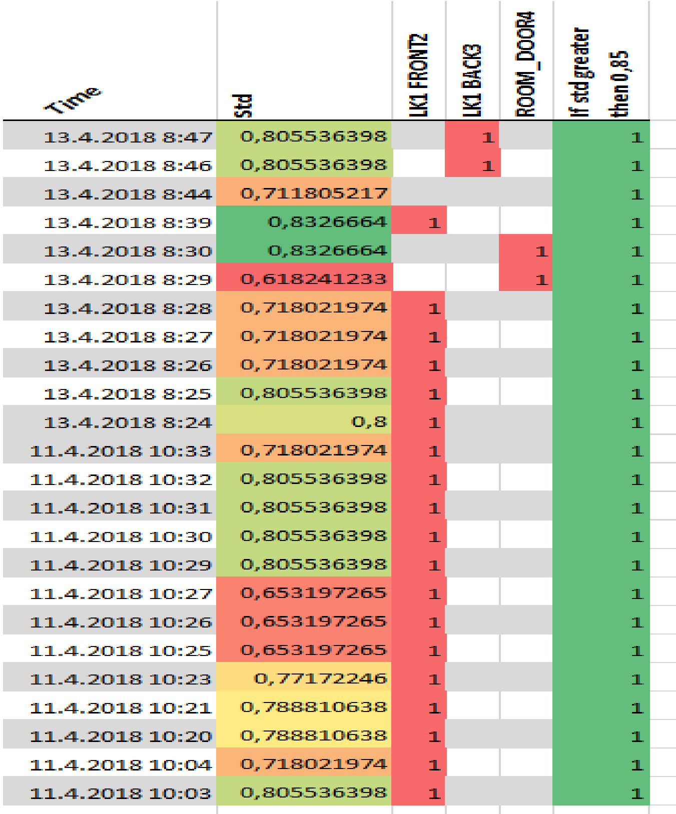

Time	Std	LK1 FRONT2	LK1 BACK3	ROOM_DOOR4	If std greater then 0,85
13.4.2018 8:47	0,805536398		1		1
13.4.2018 8:46	0,805536398		1		1
13.4.2018 8:44	0,711805217				1
13.4.2018 8:39	0,8326664	1			1
13.4.2018 8:30	0,8326664			1	1
13.4.2018 8:29	0,618241233			1	1
13.4.2018 8:28	0,718021974	1			1
13.4.2018 8:27	0,718021974	1			1
13.4.2018 8:26	0,718021974	1			1
13.4.2018 8:25	0,805536398	1			1
13.4.2018 8:24	0,8	1			1
11.4.2018 10:33	0,718021974	1			1
11.4.2018 10:32	0,805536398	1			1
11.4.2018 10:31	0,805536398	1			1
11.4.2018 10:30	0,805536398	1			1
11.4.2018 10:29	0,805536398	1			1
11.4.2018 10:27	0,653197265	1			1
11.4.2018 10:26	0,653197265	1			1
11.4.2018 10:25	0,653197265	1			1
11.4.2018 10:23	0,77172246	1			1
11.4.2018 10:21	0,788810638	1			1
11.4.2018 10:20	0,788810638	1			1
11.4.2018 10:04	0,718021974	1			1
11.4.2018 10:03	0,805536398	1			1

The threshold of 0.85 standard deviation seems to show a great amount in the first analysis. The missing information in the time period of test case one and two was treated as the value zero, so the standard deviation would also be 0 and this would be treated as a detection. No situation was seen where all the server sensor information values would be the same on a given minute. If the standard deviation was 0 these missing values were not included in this first analysis.

6.3.2 Document count affecting the detection.

Through the testing different elastic document counts where used. Table 4 shows how with only one population of 17 temperature values the detection can be possible. A measurement was made to verify this to be true.

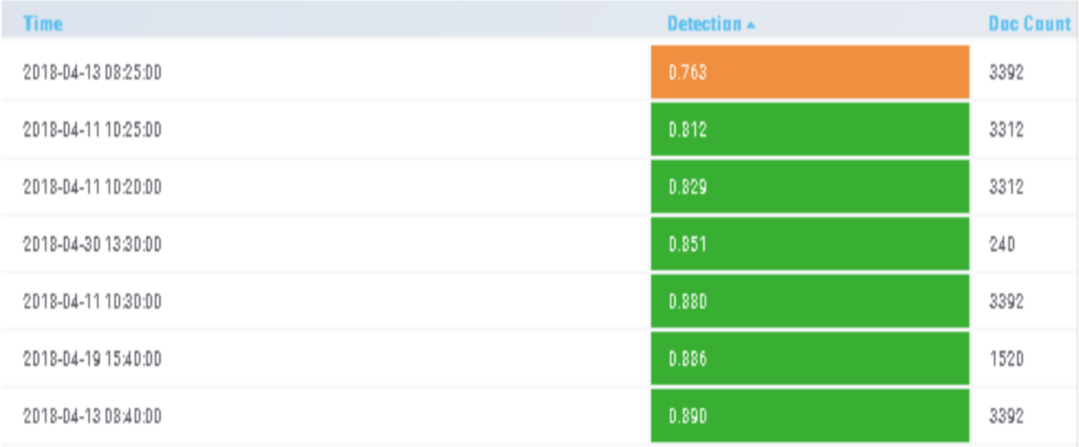

Time	Detection ▲	Doc Count
2018-04-13 08:25:00	0.763	3392
2018-04-11 10:25:00	0.812	3312
2018-04-11 10:20:00	0.829	3312
2018-04-30 13:30:00	0.851	240
2018-04-11 10:30:00	0.880	3392
2018-04-19 15:40:00	0.886	1520
2018-04-13 08:40:00	0.890	3392

Figure 17. Data rates and detections

The first two tests were conducted with over 3,000 documents per a 5-minute interval, which makes the document count per-minute on one sensor over 10. The third test was conducted with half of that; thus the test was conducted with 5 documents per sensor per minute. The fourth and final controlled test was conducted with one to two per-minute per server sensor document rate. Figure 17 shows all the detections even with these different document data rates. In conclusion, the document data-rates did not affect the detection in a positive way even if the document data rates were high.

The collection of data could be conducted with this lower amount of document per minute safely. This will increase the total amount of information gathered to the system without any decrease in the speed of data analysis or fear that the collection will create high network load.

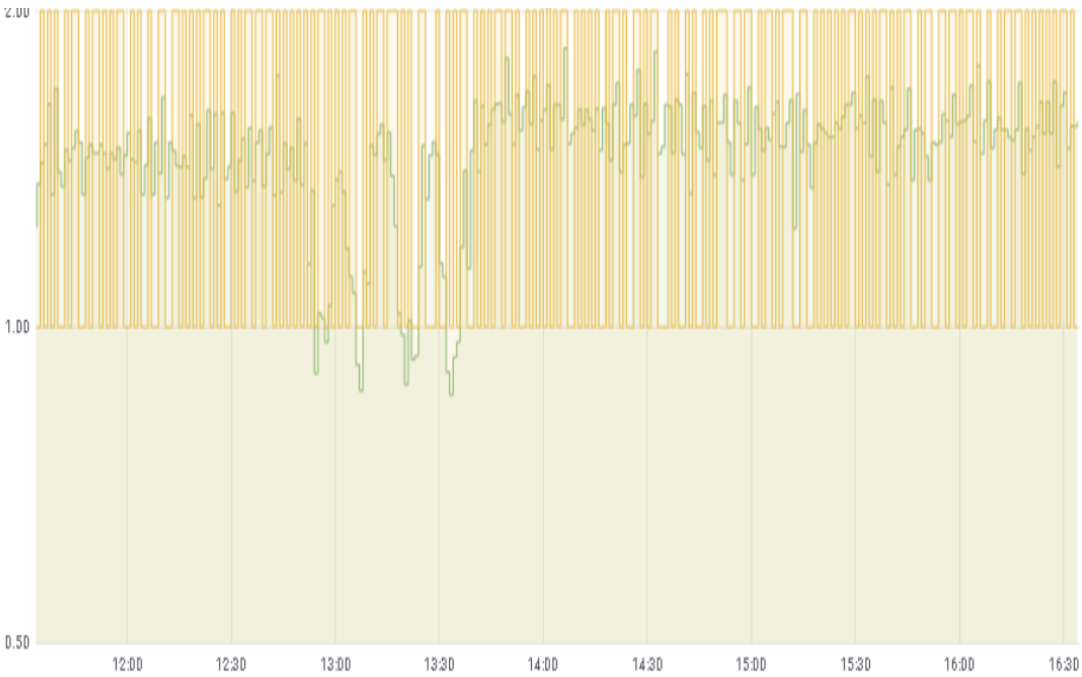

Figure 18. Example STD and Doc Count Test case 4

7 Test cases

In this chapter, all conducted controlled test-cases are discussed with an examination of every test case: what was done, when the test was conducted by the minute, and with an analysis of every test case result. For most cases time graphs are used to visualize the findings.

7.1 Test case 1: 11 April 2018 Results

As the first analysis indicated the standard deviation of inlet ambient sensor temperature information could be used as a way on how an open-door state could be detected. In the first test an open-door state was introduces twice. The first test was conducted at 10:03 to 10:05 and the second 10:23 to 10:33. The idea was to create a short 3-minute test where the rack door was opened and a longer one that lasted for 10 minutes. There was an 18-minute break between tests.

The standard deviation of information is illustrated in Figure 20.

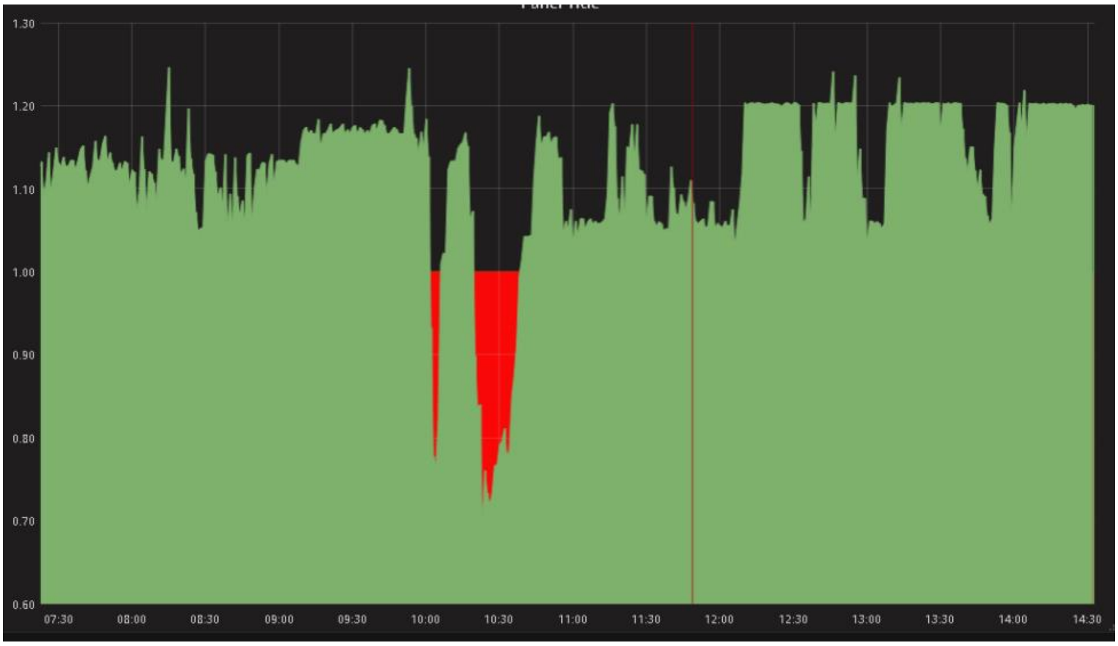

Figure 19. Test case 1: 11 April 2018

The analysis shows clearly when the rack door was opened. The standard deviation value dropped nicely below the set threshold and a detection was possible. The previous figure is self-explanatory.

7.2 Test case 2: 13 April 2018 Results

The next controlled test was conducted two days later from 8:23 to 9:00. Four test cases were conducted at specific times of 8:23 to 8:28 (6 minutes), 8:38 to 8:43 (6 minutes), third one at 8:51 to 8:53 and a fourth one at 8:58 (1 minute).

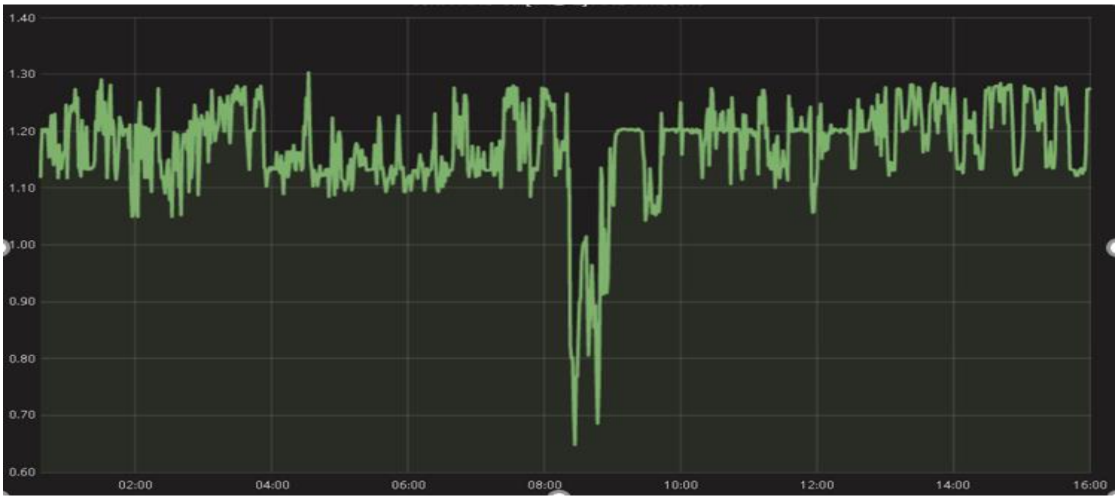

Figure 20. Test case 2 13.4.2018

Once again, a good result is gained. There was doubt where and when the threshold value has been reached. A trusted a detection would be made, illustrated by the Figure 21. The third and fourth test were not detected.

7.3 Test case 3: 19 April 2018

Test case three was conducted on 19 April 2018 from 15:29 to 15:40 where the rack door was opened at 15:29 and closed at 15:31 (open 3 minutes). The second test was from 15:32 to 15:34 (open for 3 minutes) and the third test from 15:42 to 15:43 (open for 2 minutes). Once again, the standard deviation of the measured ambient inlet temperatures dropped below the set threshold of 0.85; hance, it was still working as expected. The third test did not reach the threshold. Figure 22 illustrates the test case.

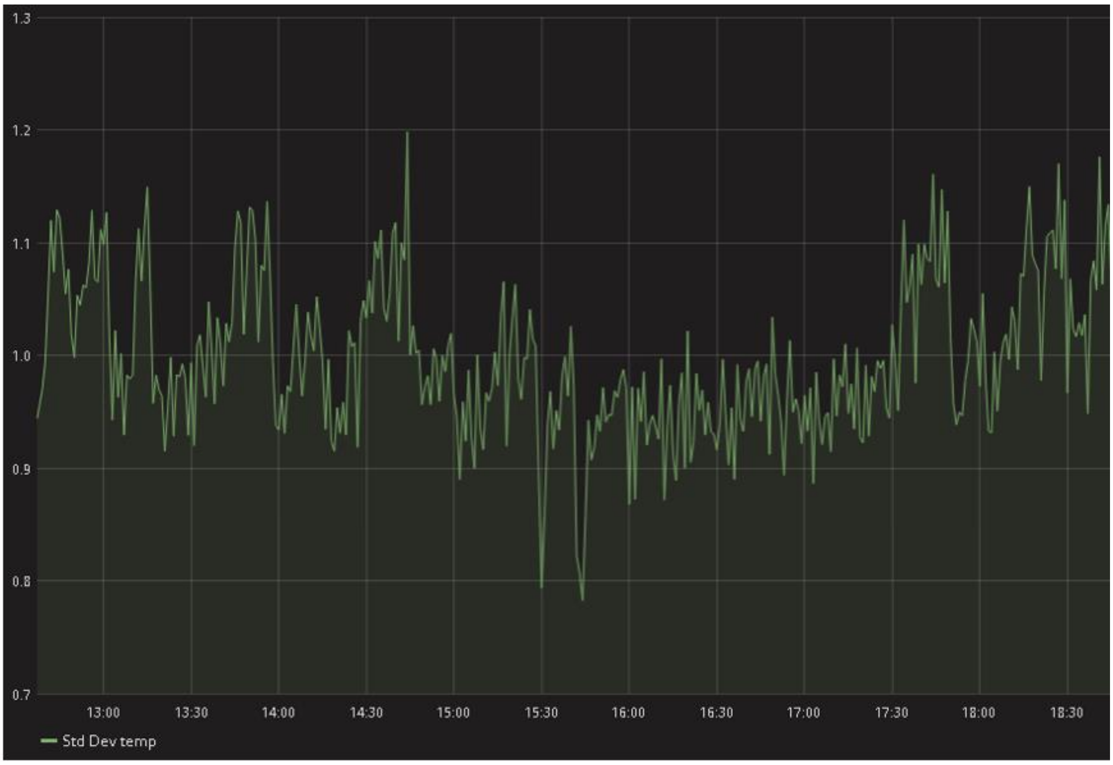

Figure 21. Test case 3 19 April 2018

7.4 Test case 4: 30 April 2018

Test case 4 was conducted with four five-minute tests between 12:52 to 13:35. The results showed the same drop of standard deviation values than as in the previous tests. The First five-minute test was conducted on 12:52 to 12:57, the second from 13:02 to 13:07, third 13:18 to 13:23 and the last one 13:30 to 13:35. The anomaly or detection is clearly visible when the rack door was open.

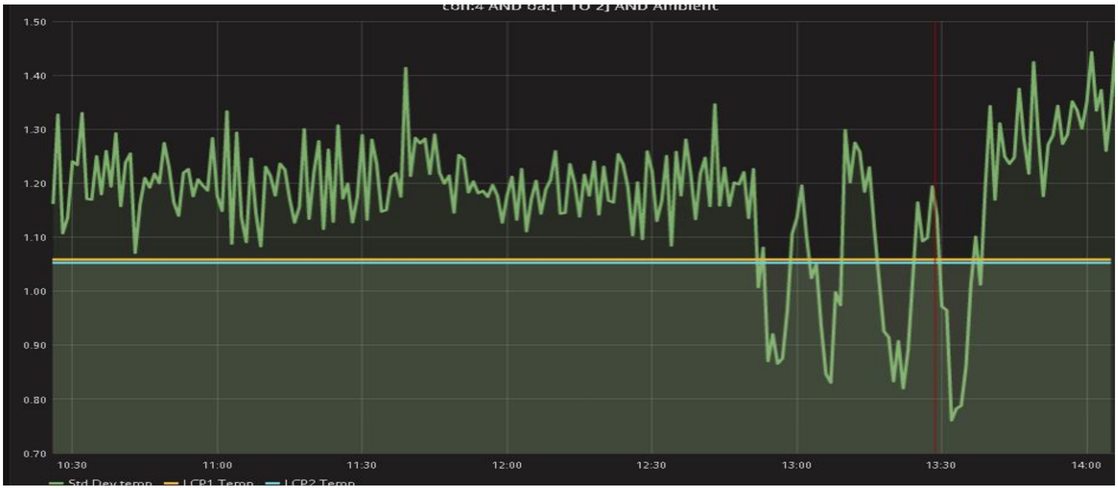

Figure 22. Test case 4 30.4.2018

As figure 23 clearly indicates using standard deviation of data points only hits our field test cases.

7.5 Cleaning up missing data and reducing false positives

These tests were conducted in a real-life situation and there were some problems with collecting the temperature data, which led to the situation where some missing was present in the data-set. As this missing data treated as zero, will mean that a specific point in time all values will be zero, this then will lead to the situation that for this specific time the STD is 0, and under the threshold of 0.85 this is an unwanted situation.

Missing data can occur for many reasons but none of the missing data-points did hit our controlled test-case time intervals. Figure 24 shows the way missing data is displayed without any filtering. In the following, the example average temperature values are displayed and used to demonstrate how the filtering works. The same filtering would also apply to standard deviation data.

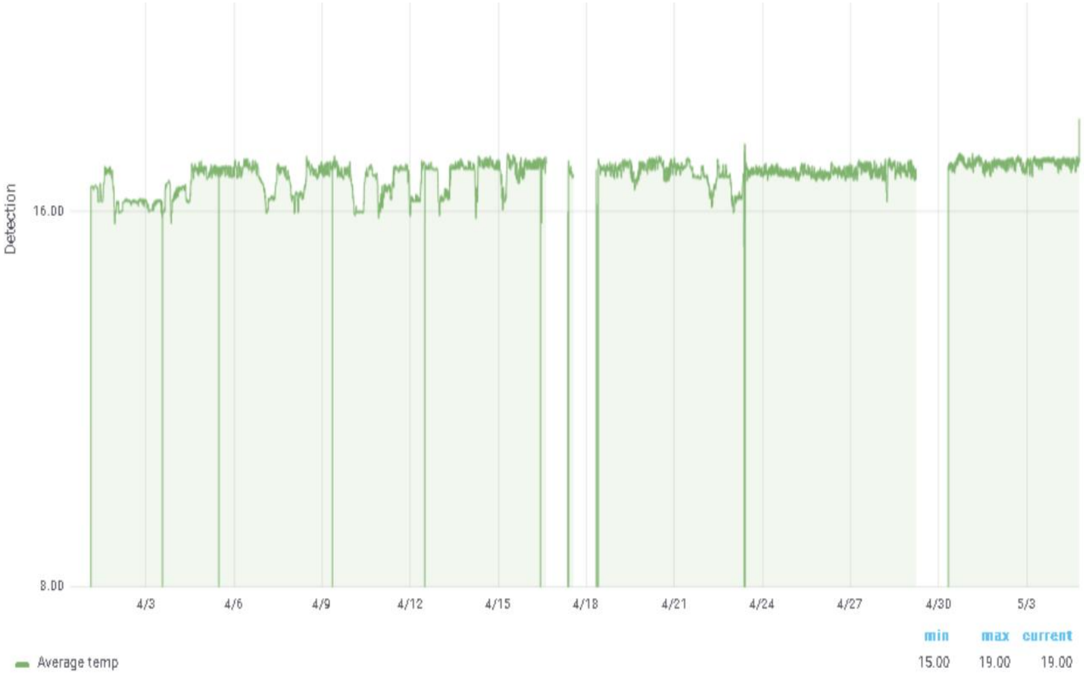

Figure 23. Dirty Graph

Next, a filter was applied where Grafana was instructed to only collect and show empty buckets where minimum document count is less than 55 documents in a 5-

minute interval. This simplified means that Grafana will not show any data sets collected over 5 minutes which do not have over 55 elastic documents. This is much better as one does not want to see the value zero in our graph. As stated before, situation where zero is present would lead to a false detection because zero is under our 0.85 threshold. Figure 24 illustrates this.

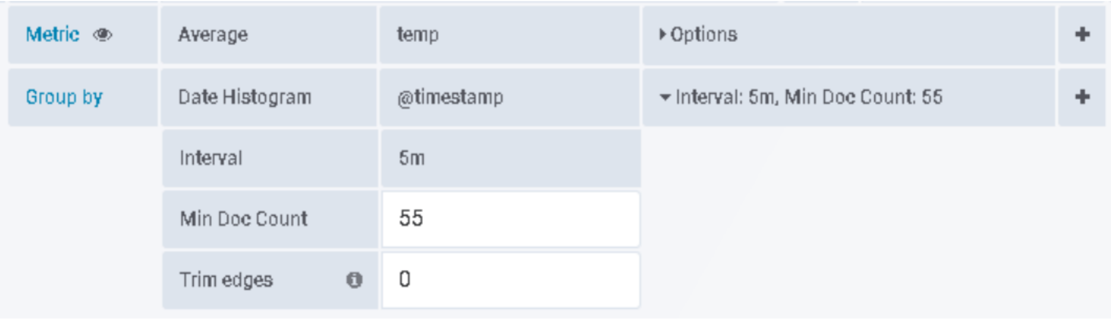

Metric 👁	Average	temp	▸ Options	+
Group by	Date Histogram	@timestamp	▾ Interval: 5m, Min Doc Count: 55	+
	Interval	5m		
	Min Doc Count	55		
	Trim edges ℹ	0		

Figure 24. Min Doc Count

When applying this filtering of the minimum document count the graph will look as illustrated in Figure 25.

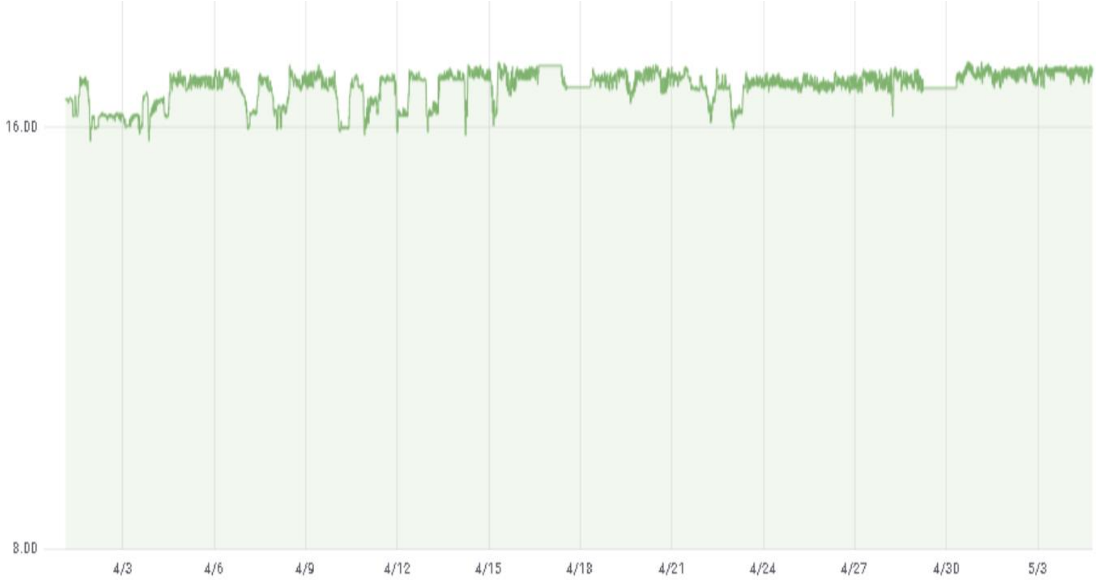

Figure 25. Cleaned Graph

The graph has no more missing values showing and will not create false positive detections. The missing information is still absent; however, it will not be shown in the graph and will not be detected.

7.6 The detection speeds and time it takes to detect

Even though the speed was not an issue when looking at the test cases and the data with one-minute intervals we can see that opening the rack door and the first detection took place during the same minute. In the example Table 5 shown as example of this.

Table .5 Detection speeds

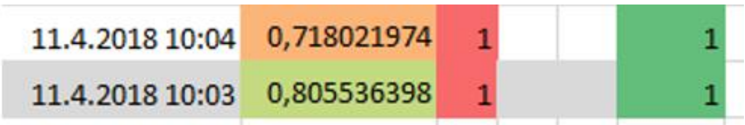

11.4.2018 10:04	0,718021974	1			1
11.4.2018 10:03	0,805536398	1			1

The threshold value of standard deviation was reached at the same minute as the rack door was opened, which was a surprisingly fast detection. The detection speed stayed the same in all test cases.

The time it takes to do the detection was around 3 minutes. The detection could be faster; yet for usability reasons 5-minute intervals were used because the calculation of 1-minute data over longer periods of time is expensive for the Elastic search engine. This 3-minute detection speed was seen still as adequate for detection purposes.

7.7 Different visuals and back in time a bit more

Four test days were used to create thirteen controlled test cases of which every over two-minute test was detected. Sometimes it can be useful to look back in time with different tools such as a carpet plot. Using carpet plot to visualize this information identifies the controlled test cases on an hourly basis. This will come in handy when looking back a long period in time, e.g. in the case of post analysis of an indecent. Figure 27 was created from the data over one month. The settings are so that the lowest standard deviation in a hour period is measured and displayed with a different color depending on the value.

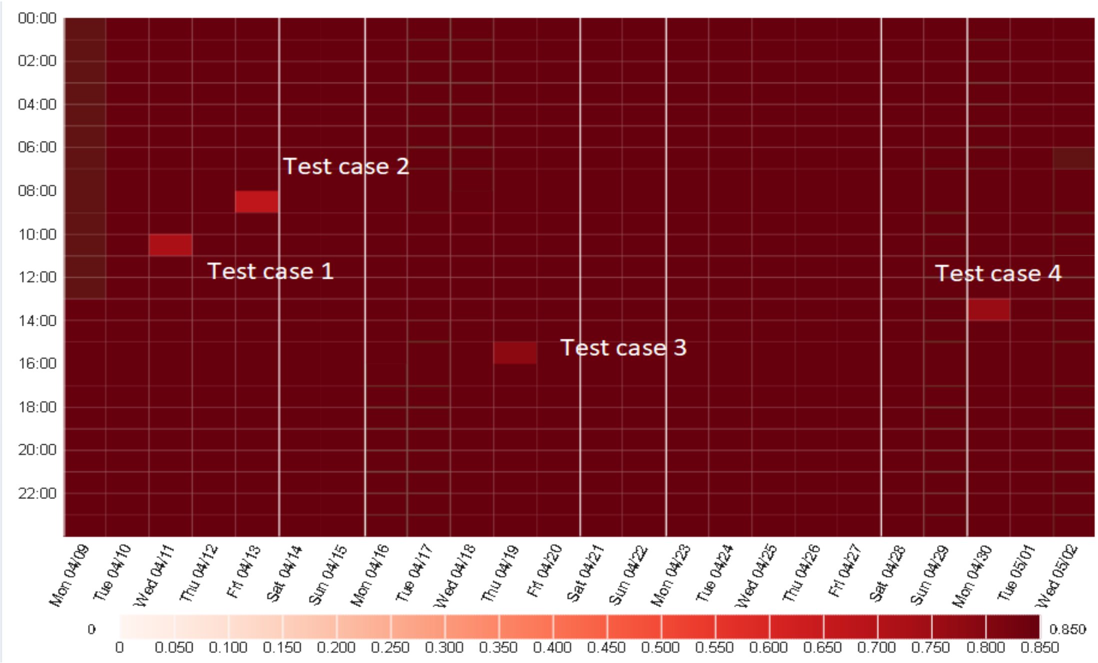

Figure 26. One-month carpet plot

All controlled test cases are clearly displayed. No data hits false positives and the missing zero values are cleaned out by the filter minimum document count of 55.

Looking over an even longer period with the same carpet plot Figure 28 can be created. This illustration has plenty of missing data for a couple of months from March to the start of May, the start of June to the start of July and from mid-July to the start of October.

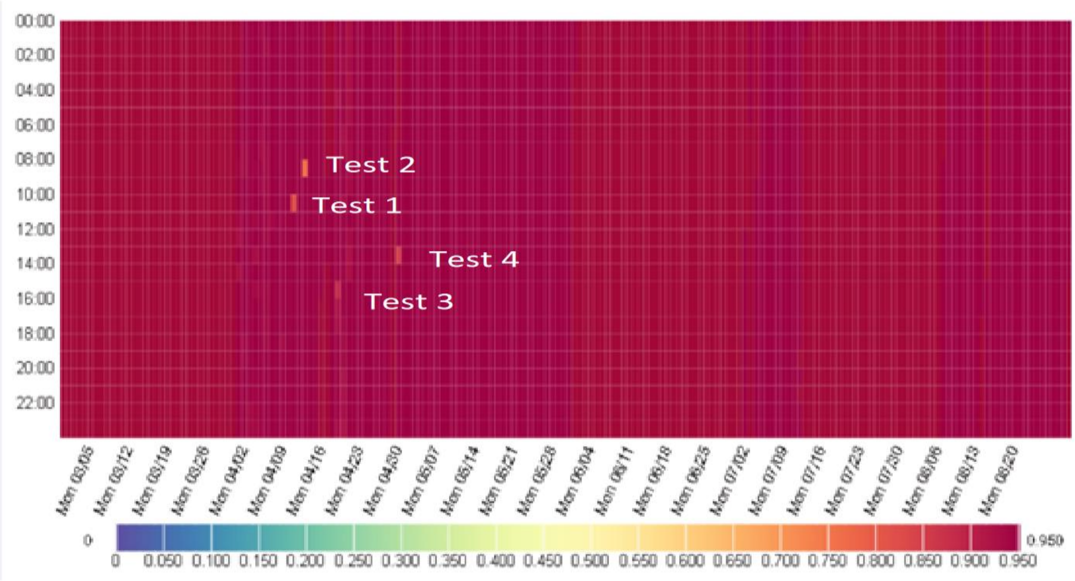

Figure 27. Carpet plot over five months

7.8 Blind testing detection solution on similar rack environments

For verifying that this method would suite other rack environments, a blind testing test set was conducted. The same method was applied to three other similar rack systems without the knowledge of when they had been accessed. The information was gathered from the timeframe ranging 8 April 2018 to 8 July 2018. Standard deviation with the same threshold 0.85 was counted with five-minute intervals from temperature values collected from the ambient inlet sensors again using 15 server inlet ambient sensors.

Filtering the data and querying for values under standard deviation of 0.85 provided the following table on the first of three of the new test rack environments. Characters U and X are used to indicate if the detection was verified where X means verified and U means unverified. Table 6 illustrates this.

Table 6. Another rack, rack number 1

Time 5 minute intervals		RACK 1
22.4.2018 13:10	U	0.84
14.5.2018 1:35	U	0.84
14.5.2018 5:15	U	0.84
14.5.2018 6:30	U	0.84
14.5.2018 7:25	U	0.84
14.5.2018 7:50	U	0.84
14.5.2018 9:50	U	0.84
23.5.2018 15:10	X	0.63
23.5.2018 15:15	X	0.84

The detection was successful with some unverified detections. After the test was carried out, it was verified that at the specific time and date of 23 May 2018 there was an access made to this specific rack cabinet. False detections are visible but when lowering the detection threshold, we could only see one day with the value 0.63, which was verified. With little tweaking of the detection threshold value one could only see the day on which the access had been made.

The same test against rack 2 provided the following information (Table 7). Where the days 15 May and 16 May 2018 were verified: that an access had been made to the rack environment.

Table 7. Another rack, rack number 2

Time 5 minute intervals		RACK 2
30.5.2018 10:15	U	0.76
16.5.2018 18:20	X	0.78
16.5.2018 18:25	X	0.81
16.5.2018 15:30	X	0.82
16.5.2018 16:40	X	0.82
16.5.2018 18:15	X	0.82
15.5.2018 17:50	X	0.83
16.5.2018 15:25	X	0.83
16.5.2018 15:35	X	0.83
16.5.2018 18:35	X	0.83
15.5.2018 17:55	X	0.84
16.5.2018 16:45	X	0.84
16.5.2018 18:30	X	0.84
16.5.2018 18:40	X	0.84
16.5.2018 18:45	X	0.84
6.7.2018 22:30	U	0.84
11.4.2018 9:50	U	0.85
2.5.2018 10:05	U	0.85
15.5.2018 14:35	X	0.85
15.5.2018 17:45	X	0.85
15.5.2018 18:05	X	0.85
7.7.2018 19:10	U	0.85

Rack 3 had a different standard deviation threshold than the previous racks environments as the value range was higher for the threshold as no detections were made by the threshold value of 0.85. It was then decided to select the seven lowest standard deviation values and the following table was created using this information. Once again, with a little tweaking of the standard deviation threshold value detection would be possible. When rising the detection threshold to 1.30 only the verified days would be present. (Table 8)

Table 8 Another rack, rack number 3

Time 5 minute intervals		RACK 3
21.4.2018 17:15	U	1.33
6.5.2018 15:20	X	0.94

6.5.2018 15:25	X	1.08
6.5.2018 15:30	X	1.24
11.5.2018 11:45	U	1.32
22.5.2018 13:50	X	0.92
22.5.2018 13:55	X	1.29

Concluded from the blind test cases it became clear that even if some false positives are possible every access made in this time period on these different rack environments was detected. This follows the same pattern as the controlled-field test and the concluded analysis of these test. The method can and could be used for physical access detection.

7.9 Alarms

Alarms in Grafana are possible for different notification types. These types are email, Slack, PagerDuty, webhook, DingDing, Kafka (Grafana). Sending alarms with Grafana is a simple process. The process only consists of setting a desired threshold for detection and then selecting the notification channel that is used for the alert to be sent.

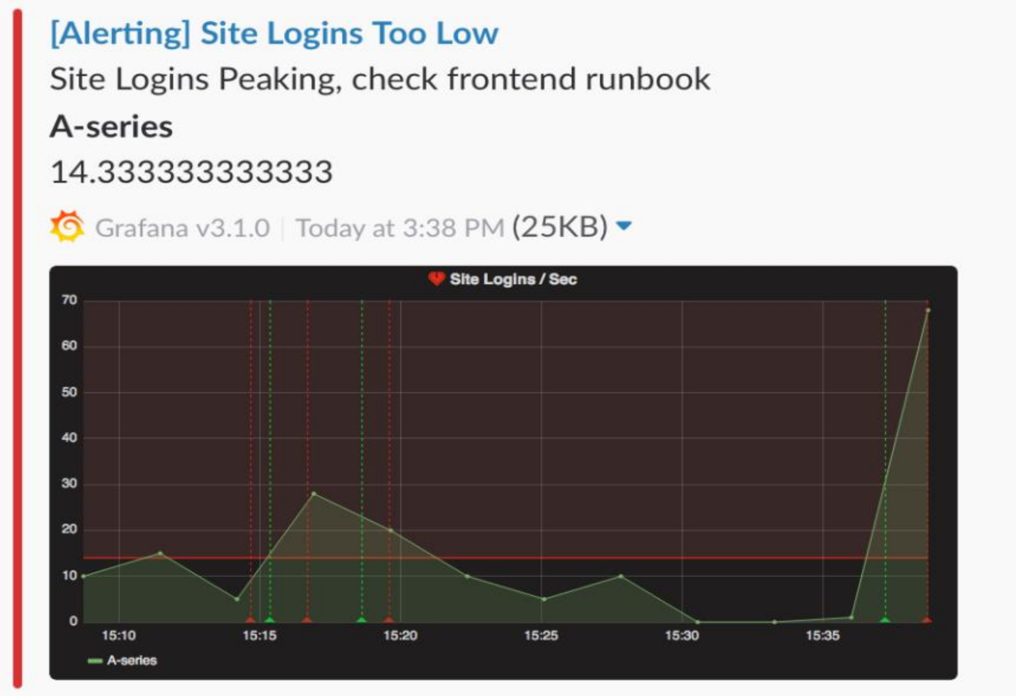

[Alerting] Site Logins Too Low

Site Logins Peaking, check frontend runbook

A-series

14.333333333333

Grafana v3.1.0 | Today at 3:38 PM (25KB)

Figure 28. Grafana alerting

As seen in the example Figure 28. it would be possible to implement the alarms based on the detections. But at the time of writing this thesis Grafana did not support Elasticsearch alerting. This feature is desired by many and is scheduled in the 5.2 version release of Grafana.

7.10 The objectives and how did the solution meet them?

The solution can be used to detect physical access through monitoring of the server ambient temperature values. Now in the current Grafana version there is alerting capabilities that could be used to generate alarms based on the detections made by the method. The amount of false detections is low, and the threshold values can be adjusted to minimize the amount even more. Even if it was impossible to detect the test cases where the door was open for two minutes or less it was seen that the objective was met.

Even if the objectives for this specific case were succeeded in a good way a more comprehensive study should be conducted to verify that the proposed method would work in different kinds of rack environments with different amount of equipment and cooling solutions. This study strongly focused only on a enclosed rack environment with in-rack cooling system and the method may not suit these different environments.

8 Conclusions and discussion

We asked some questions what this thesis should answer. The first question was what the limitations are of the currently implemented traditional physical security controls that may affect the effectiveness of real time physical access detection? The second question asked if there could be limitations on the current security implementations that could pose a risk to the information assets?

Traditional controls such as locks can only protect the assets from physical access. They cannot generate real time information that could be used to determinate when or where one of the locks had been opened. The information can be obtained from the key management process if handled correctly. All of locks should be separate and no same key should fit in all of them. Only this kind of implementation and key management arrangement could inform where someone has gained access. The same goes for CCTV cameras as they cannot be used to mitigate the risk in real-time; there will always be some sort of delay from detection to mitigation of the incident. The information that CCTV cameras generate can still be useful but only when used to conduct post analysis on a possible security incident.

The third question asked about the possibility to mitigate the residual risks left by e.g. traditional locking mechanisms using existing IT equipment ambient inlet temperature sensor information.

As this study shows with mathematical methods such as calculation of standard deviation on the population of environmental sensor values, can be used to do physical access detection. When the existing software such as Grafana is supported by a database such as Elastic, it can handle plenty of information and these tools can carry out fast calculations of incoming data; hence it is possible to conduct anomaly detection using these tools. The software suit very well in different kinds of data analysis in many shapes and forms. Every environment will differ in some way from others but when using analysis software like Grafana the data can be dynamically inspected in different ways and forms. While the tests were limited and there was a low amount of different environments targeted by the solution the same method could be used in other types of rack environments also.

Finally, based on a simple hypothesis we were able to create a working solution for access detection. The same solution could be modified so that it could be used to monitor our rack cabinets and give out alarms for e.g. overheating. Extending the reach of the detection system can be extended by adding more targets to the Logstash query component. Creating everything from scratch was a real eye-opening experience. Even if this solution works like a fire alarm you never want to hear but when you need it you will be happy that it was there.

Even though the study shows that detection can be possible, a more extensive study could be conducted with a wider range of different types of racks, devices and cooling solutions.

References

42u Data Center Solutions 2015. *High Density In-Rack Cooling Solutions for Server Racks, Computer Rooms & Data Centers.* Accessed on 30 April 2018. Retrieved from http://archive.42u.com/cooling/in-rack-cooling/in-rack-cooling.htm

Messina, G. 2018. *CISSP: Perimeter Defenses.* Accessed on 30 April 2018. Retrieved from http://resources.infosecinstitute.com/category/certifications-training/cissp/domains/security-engineering/perimeter-defenses-and-the-cissp/

Cornelissen, B., Greene K.,Hadzhiyski I., Keely, P., Allen, S. Telmo Sampaio, T., 2012 Mastering System Center 2012 Operations Manager. Cybex.

Easttom, C. 2015. Computer Security Fundamentals Second Edition. Indianapolis: Pearson.

Kim, D., & Solomon, M,G. 2014. Fundamentals of Information Systems Security Second Edition. Burlington: Jones & Bartlett Learning.

Grafana Web 2018. *Alert Notifications.* Accessed on 30 April 2018. Retrieved from http://docs.grafana.org/alerting/notifications/

Hewlett Packard Enterprice 2018. *HPE Onboard Administrator*. Accessed on 30 April 2018. Retrieved from https://www.hpe.com/us/en/product-catalog/servers/bladesystem-enclosures/pip.hpe-onboard-administrator.3188465.html

Hewlett Packard Enterprice Community 2014. *Temperature Sensor information of WS460c / BL460c Gen8.* Accessed on 30 April 2018. Retrieved from https://community.hpe.com/t5/BladeSystem-Server-Blades/Temperature-Sensor-information-of-WS460c-BL460c-Gen8/td-p/6675226

IBM Security 2015. *IBM 2015 Cyber Security Intelligence Index*
Accessed on 30 April 2018. Retrieved from https://essextec.com/wp-content/uploads/2015/09/IBM-2015-Cyber-Security-Intelligence-Index_FULL-REPORT.pdf

InfoSec Handbook 2018. *Information Security Frameworks and Information Security Architecture.* Accessed on 30 October 2018. Retrieved from https://ebrary.net/26641/computer_science/information_security_frameworks_information_security_architecture

Oriyano, S-P,. 2014 Hacker techniques, Tools, and incident Handling Second Edition. Burlington: Jones & Bartlett Learning.

Williamson, S. 2017 *COUNTERING THE THREAT OF PHYSICAL SECURITY BREACHES*
Accessed on 30 October 2018. Retrieved from http://www.datacenterjournal.com/countering-threat-physical-security-breaches/

The Server Rack FAQ 2018. *Information about rack issues and racking servers.* Accessed on 30 October 2018. Retrieved from https://www.server-racks.com/rack-cabinet.html

Weaver, P. The project manager.com.au 2016. *Can you use standard deviation in project management?.* Accessed on 30 October 2018. Retrieved from http://projectmanager.com.au/can-you-use-standard-deviation-in-project-management/