

Cyber-risks in the Industrial Internet of Things (IIoT): towards a method for continuous assessment.

Carolina Adaros Boye, Paul Kearney, and Mark Josephs

Birmingham City University, Birmingham, B4 7XG, UK
carolina.adarosboye@mail.bcu.ac.uk,
paul.kearney@bcu.ac.uk
mark.josephs@bcu.ac.uk

Abstract. Continuous risk monitoring is considered in the context of cybersecurity management for the Industrial Internet-of-Thing. Cyber-risk management best practice is for security controls to be deployed and configured in order to bring down risk exposure to an acceptable level. However, threats and known vulnerabilities are subject to change, and estimates of risk are subject to many uncertainties, so it is important to review risk assessments and update controls when required. Risks are typically reviewed periodically (e.g. once per month), but the accelerating pace of change means that this approach is not sustainable, and there is a requirement for continuous monitoring of cybersecurity risks. The method described in this paper aims to alert security staff of significant changes or trends in estimated risk exposure to facilitate rational and timely decisions. Additionally, it helps predict the success and impact of a nascent security breach allowing better prioritisation of threats and selection of appropriate responses. The method is illustrated using a scenario based on environmental control in a data centre.

Keywords: Internet of Things, Industrial IoT, Industrial Control Systems, Cyber-security, Control Systems, Risk Analysis

1 Introduction

The US National Institute for Standards and Technology (NIST) defines risk monitoring as “maintaining ongoing awareness of an organisations risk environment, risk management program, and associated activities to support risk decisions” [6]. Nevertheless, it is not unusual for risk monitoring to be done as a discrete activity, once over a period of one, or even several, months with a low level of integration with the operational processes.

The Industrial Internet of Things (IIoT) present important concerns regarding cyber security including risks where consequences go beyond the realm of information systems to interact with the physical world. Therefore, it is advantageous to have timely information about the possible development of a threat scenario.

Suspicious events that can be detected during operation occur frequently and can overload Security Operations Centre (SOC) personnel with data. Introducing a risk-based approach to automated threat detection would allow prioritising security resources and improve decision making. Many methods fail to do this by focusing only on the threat and on the direct consequences, detaching the analysis from the operational impacts, and from the business context. The solution proposed in this research contributes to improving IoT cyber security by monitoring risks continuously. The main idea is to provide a holistic view of the potential impacts of an attack considering how the consequences at an operational level can affect business processes and strategic objectives. The aim of developing this method is to provide relevant, accurate and timely information about cyber-security risks in IoT systems.

While the focus of the method is to adjust risk indicators in near real time, it is necessary to have a good level of understanding about the variables that will be used in the calculations.

The method considers a variety of inputs divided in two groups: dynamic and static. Dynamic inputs will provide (near) real time information about the state of the system to a risk calculation engine that will update the key risk indicators. This should shorten response times for allocation and adjustment of security controls. Continuous updates to the risk treatment plan will procure a better integration between operational processes, risk management, and security processes.

As one of the main “blind-spots” in IoT security is the physical layer, it is proposed to use anomaly detection techniques to monitor variables that can be correlated with possible security issues. Examples are electricity consumption, server performance, and other side-channel information. Establishing direct correlations between an anomaly and its root cause will be challenging and in many cases it will require the involvement of an expert. Also, it may be more difficult to obtain these data compared with other dynamic inputs such as software and network monitoring, because not all of them will be necessarily provided by already available detection tools. Addressing this is among the main challenges that this research project will face.

At the current stage of this research, a conceptual model has been developed with the potential to be adapted to different sorts of IoT systems. Section 2 of this paper explains the problem and current gaps that are addressed. Section 3 gives an example of a use case to provide a setting for explaining the method. Section 4 gives a the general description of the method. Section 5 provides a threat scenario based on the case described in Section 3 and explains how the method would work in this case. Section 6 mentions relevant considerations and future challenges of this project and Section 7 provides the conclusions.

2 Outline of the problem

Although there are many expectations about introducing new technologies in the industrial industrial control system domain there are also many legacy systems

that cannot be easily replaced. These systems are still widely used and deployed and will need to coexist with the concept of Industry 4.0. One important concern is that their original design did not consider security sufficiently for the current levels of connectivity. Industrial Control Systems (ICS) is closely related to IoT in the sense that they both fit within the definition of an "ecosystem of interconnected devices and services that exchange and process data" [9]. Throughout this paper, the term IoT will be used under the understanding that Industrial Control Systems fit among this definition. Some authors will refer to these systems as Industrial IoT (IIoT). Examples of use cases in IoT can be found but are not restricted to the following industries[21]:

- Transportation
- Health-care
- Government
- Public safety and military
- Retail and hospitality
- Food and farming
- Manufacturing and heavy industry
- Entertainment and sports
- Energy and utilities
- Finance and banking
- Education
- Information and communications technology

In most of these industries, performance, time to market and cost pressures have been a priority over cyber-security [21]. The lack of standards and regulations, and poor security awareness of manufacturers and users has not helped to improve this situation. In the past (and in some cases, still in the present) electro-mechanical or cyber-physical systems based their security mostly on isolation and perimeter security. The circumstances have changed and the vulnerability of these systems has increased. Even critical systems that are isolated from public networks present risks. For example, the malware Stuxnet, discovered in 2010, was allegedly infiltrated to an Iranian nuclear plant through an USB drive connected to one of the computer terminals. This terminal was connected to the control system and was used as foothold to spread the malware to the Siemens PLCs of the plant. This is an example of the "air gap myth" which proves that isolation by itself is insufficient.

Although attacks on IoT systems are nothing new, the amount of connected IoT systems currently exceeds the human population [10][14], giving more opportunities to attackers. An industrial report released this year based on the study of different attack vectors in industries reveals that in 82% of the cases an internal attacker could have penetrated the industrial system from the corporate network. Significant flaws in network segmentation and separation of privilege were also found, among many other vulnerabilities [25]. Attacks on Symantec's IoT honeypots almost doubled in less than a year [29]. According to Cisco, no industry vertical is safe from cyber-attacks [3] and it is believed that IoT devices

“are becoming the attack infrastructure of the future” [1]. Cyber-crime has become to be known as a profitable business and cyber-weapons also have started to be commonly used by nations for surveillance and national intelligence. smart TV’s have already been known to be part of plans to develop tools for espionage, and successful cases of sabotage of national critical infrastructure have been attributed to nation states.

Because IoT systems are based, in part, on computer and network systems, they inherits all their security issues, as well as presenting additional cyber-risks. Their complex architecture increases their attack-surface [16] [21] by the addition of devices that interact with the physical world. The variety of hardware involved will have distinctive requirements and constraints which makes security more challenging. In many cases, typical security mechanisms could be not feasible or be insufficient [26]. Limitations in memory and processing capabilities, as well as real time response requirements present constraints to encryption and authentication processes. Also, special attention regarding physical security is required as often the systems have components distributed in a wide area. The use of wireless communications has an inherent risk enhanced by the variety of protocols and enabling technologies. There are fewer standards, regulations, and overall less experience in IoT security [21] and manufacturers tend to have less knowledge in the matter than professionals from the software development world [4]. Whereas in information systems the main concern usually is related to confidentiality, in IoT systems implications of a cyber attack go beyond information theft. Risks can include also damage of physical assets, and even threat to human life [26]. For example by compromising the integrity or availability of critical systems such as life support equipment or systems working in safety-critical environments.

Currently, an important amount of available literature proposes solutions for particular aspects of IoT security such as authentication, secure communications, and attack modelling. These solutions are relevant, but not sufficient by themselves to provide acceptable levels of security. Security issues of one layer of an IoT system cannot be solved in another [16]. This means that different solutions need to be integrated and effectiveness monitored in the context of the overall system.

Security should be implemented as a combination of processes, technology, and people [6] so it is important to consider all these factors in the equation. Automated tools can help to deal with big data and recognise patterns of behaviour, but these patterns need to be put in the context of the operational and business processes and their objectives. The input of experts is essential to achieve this.

Several methods have been developed for IoT cyber security risk assessments based on existing techniques, including game theory [28], fuzzy logic [17], and Bayesian Networks [31]. Some of these methods are general and others focus on specific type of system. A review of 24 existing cyber security risk assessment methods applied for SCADA systems was done in [2] where the main opportunities of improvement that were found are the following:

1. Addressing context establishment
2. Overcoming attack or failure orientation ¹
3. Accounting for the human factor
4. Capture and formalisation of expert opinion
5. Improvement of the reliability of probabilistic data
6. Evaluation and validation
7. Tool support

Risk analysis continues to be understood as a discrete activity, often done using spreadsheets or other tools which are not integrated with operations and are fed manually with data. Nevertheless, the NIST recommends transitioning to near real time risk management [7]. With new threats and vulnerabilities been discovered on a regular basis, it is likely that many of the data used in a risk assessment would expire in a short period of time. This would make the results irrelevant. Very little academic work has been done related to real time or continuous risk evaluation.

A limited amount of research proposing dynamic or real time cyber security risk evaluation methods has been published [12] [13] but most of them are not specific for IoT or IIoT. Other models for real time risk assessment reviewed were mainly focused on threat and anomaly detection and did not consider the impacts. Anomaly detection can be useful to detect threats in an IoT system by comparing variables with a model of their expected behaviour. However, the picture will be incomplete if this information is detached from its context. Several publications about anomaly detection in IoT, Industrial Control, and SCADA systems propose techniques such as machine learning [11][20][17], data mining [24], statistical analysis [8][32], and hybrid methods [19]. The work published by Zhang et al [33] on incident prediction and risk assessment for industrial control systems considers both real time processing and asset valuation, but, it only provides proof of concept through simulation experiments on a single type of system. In conclusion, there is a lack of risk assessment methods for IoT that are both holistic and dynamic and that have been tested in different scenarios.

3 Example of a risk scenario in IoT system

A simplified temperature control system will be used to illustrate the method. Temperature control is use case that can be found in domestic, commercial, and industrial environments. Nevertheless, in different domains the system will typically present different characteristics, types of technology, and architectures. This research, rather than in IoT domestic or consumer devices, is focused on Industrial Systems. Temperature control can have different purposes in industries. For example, avoiding products such as food and chemicals to decompose or degrade, or allowing different process to perform in optimal conditions. A temperature control malfunction will have different consequences depending on the

¹ This means basing the analysis only on known attack mechanisms and failure modes

business processes involved. Understanding these consequences is crucial when assessing risks, and also a key part of finding possible signals of compromise.

The scenario developed in this example corresponds to the temperature control system for a Data Centre, and is shown in Figure 1. The scenario was validated with an engineer that works in a consultancy in Chile who has over ten years of work experience in configuration, installation, and maintenance of Industrial Control Systems.

A data centre, ideally should operate in an environment with a temperature between 24 and 27 degrees Celsius [23]. This includes a margin of error, as servers typically can tolerate up to 30C. At higher temperatures, servers do not achieve their best performance, and their fans will need to spin at their maximum rate, increasing power consumption. To avoid the temperature surpassing an established limit, Direct Digital Control (DDC) devices are used which are connected directly to temperature sensors and to the control valves for the heating and cooling systems. The controllers communicate with a Building Management System (BMS) that runs in an application server. A local PC located in the same premises is in charge to run the control and monitoring software interface. The BMS sends alerts via email and SMS messages when an event requires attention.

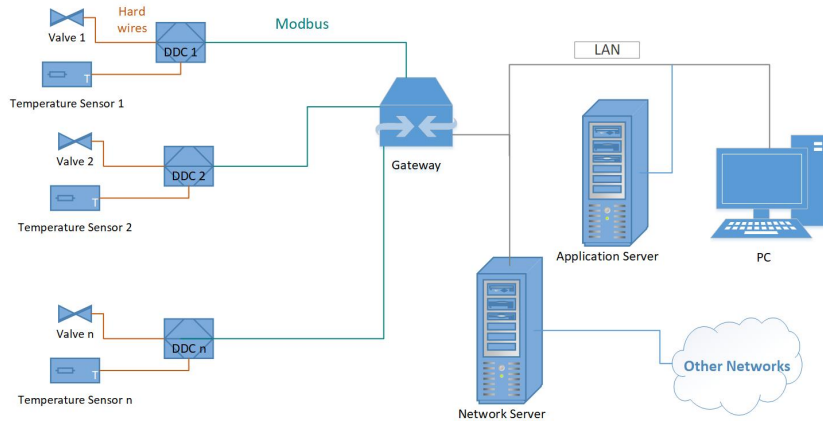


Fig. 1. Diagram of Temperature Control System

Figure 2 shows a diagram from the Industrial Internet Reference Architecture (IIRA) [18] that will be used to describe the system based on the definition of three tiers. The edge tier collects data from the real world through the proximity network where sensors and actuators are connected. The platform tier exchanges, consolidates, and processes data from the other tiers. This can include commands generated in the enterprise tier to control variables in the edge tier. The enterprise tier implements domain-specific applications and provides user interfaces.

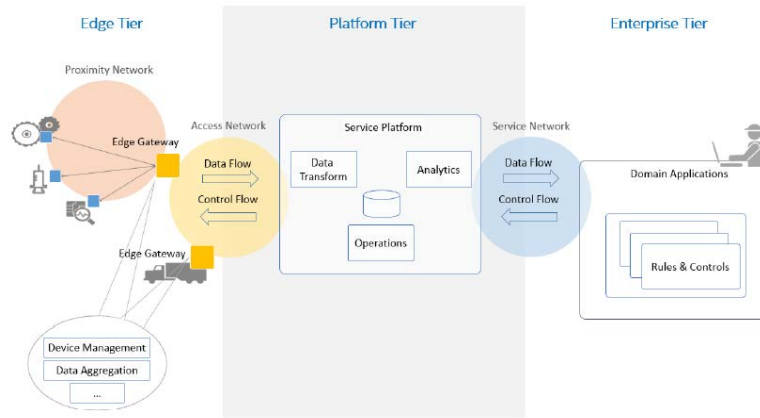


Fig. 2. Three tier architecture pattern from the IIRA Implementation Viewpoint

The following specification represents the system “as is”, before applying any risk treatment:

Edge Tier: comprises sensors, actuators and DDCs. The sensors and actuators are hard wired to the controllers providing inherent trust. They use electrical signals to communicate. Thus, they are not “smart”. The DDCs have a keyboard allowing authentication through a 4 digit security code. In the perimeter network, the protocol used by the controllers is Modbus (other protocols commonly used in these systems are Bacnet and Lonworks). The controllers are connected to a gateway that converts the signals to a standard internet protocol (TCP/IP) and connects to a Local Area Network (LAN).

Platform Tier: comprises the BMS software installed in an application server that processes operational data. The access network which connects the Platform tier to the Edge Tier is the same as the service network, corresponding to the LAN. The system is insulated from other networks for security reasons, except for the connection to an email server that allows sending alerts to operators in case of certain events and the connection to a service that sends SMS alerts. There are no firewalls or any network monitoring and detection mechanisms in place. The network server has separate cards for the LAN and other networks.

Enterprise Tier: comprises a monitoring and control software running in a PC terminal. Authentication is done through user and password without enforcing a secure credentials. There are no defences against brute force attacks in place. Privilege separation options include user, administrator and engineer roles. There is no remote connection, therefore the software only can be accessed within the perimeter. Remote monitoring is based only on the alerts sent by the BMS system.

Physical Security: authorised personnel is authenticated through an ID card, a 4 digit password, and their digital print. Special authorisation is required

for visitors and contractors which need to register. Although they should be accompanied by an authorised member of staff at all times, some contractors might be left alone for small periods of time, as sometimes they require to work there for several hours. The hardware of the control systems has often physical ports open. Personnel only visit the data centre when it is necessary, but there is nobody permanently in the area.

Cyber-security policies and practices: before the risk assessment, the BMS was considered in the cyber-security policies of the data centre. Some isolated cyber-security controls were in place, such as some degree of authentication, and the control for physical access described. There is not a clear differentiation of roles and privileges and most users just share credentials. This includes contractors. Network security is based only on isolation, and for this reason malware detection is not considered important. At the enterprise platform level the software registers and stores event logs but they are not monitored. Regarding the configuration of the temperature control settings, there is no registration of any changes or events and there are no configuration management policies in place. Backups of the system are done every six months, but there are no assurance processes to audit this or any other cyber-security practice.

4 Description of the method proposed

This project aims to make use of different sources of data to analyse cyber-risks in a continuous basis, integrating this activity with the operational process. The objective of the method is to generate useful and meaningful information for decision makers. A “decision maker” is any actor that is in position to make a decision that can affect security. These decisions can be related to business operations that can cause collateral effects in security or to security management itself. Figure 3 shows a general view of the method. The Security Operations Centre (SOC) which is the area that monitors and deals with security issues on an organisation will be provided with a more comprehensive view of attack vectors, by including IoT and operation technologies (industrial systems) in their scope. They will also be able to establish priorities for alerts regarding to the level of risk involved. The risk analysts will be allowed to monitor risks continuously, evaluating the effectiveness of security measures and control, and providing up-to-date inputs to decision making processes that involve or affect security.



Fig. 3. Illustration of the method

While the whole purpose of risk management is to improve decision making, the landscape changes too quickly to have a picture of the situation without expecting it to vary in a short time. Different internal and external factors, will continuously shape the degree of risk. New information can modify the levels of uncertainty regarding occurrence of an event, and also internal and external changes can affect risks factors. This means that security plans based on previous evaluations may become quickly obsolete. In the case of IoT and IIoT, there are more attack vectors and less visibility of the system from end-to-end in comparison with IT systems. Thus, analysing, monitoring and managing risks is critical.

Figure 4 shows a conceptual model describing the main building blocks of the method proposed. The idea is that the results should generate decisions that affect the risk treatment plan of the organisation, modifying the situation of the security levels of the IIoT system. The inputs for the risk calculations will come from three main sources: detection tools, systems variables, and a knowledge base. The first two type of sources are categorised as dynamic inputs and would be transmitted in a continuous stream. The data that is stored in the knowledge base is categorised as static inputs which either remains unchanged or is subject to eventual updates. The risk calculation engine will process the information about threats, vulnerabilities and impacts and issue alerts in the event of any condition that might change the risk scores.

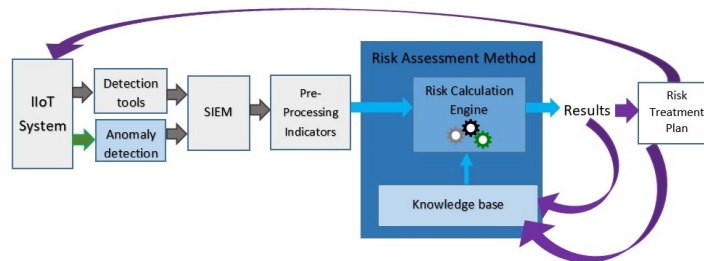


Fig. 4. Building blocks of the method

4.1 Dynamic inputs

These inputs are captured in run time from tools for malware detection, intrusion detection, network traffic analysis, and logs monitoring, and threat intelligence sharing. An anomaly detection tool is in charge of analysing data that can reveal signs of threats, including operational data. A module based on a SIEM software will be used to process the dynamic inputs that will feed the risk calculation engine. The anomaly detection tool will help detecting issues related to the edge tier.

Currently, there are existent commercial tools for anomaly detection based on machine learning. While normal operation conditions of the system and business rules could, in theory, eventually be learned by artificial intelligence (AI), this will need a training period and stable conditions over time. Therefore, known business rules and thresholds should be previously set. Also it is necessary to include expert's knowledge regarding other variables like dependencies between processes, impacts at different levels, regulations and strategic decisions. Therefore, while machine learning and AI techniques will contribute to anomaly detection, it is proposed to combine it with other methods. It is not aimed in this project to work with tools that work following a 100% unsupervised dynamic. The human factor and experts knowledge are elements that require to be acknowledged in risk management [2].

4.2 Static inputs

Static inputs correspond to the data that cannot be collected online. The reasons could be that there are no tools available to collect this information on real time or, that the data just does not change continuously. For example, the valuation of an asset, quantification of impacts, and risk acceptance criteria. These inputs will be provided during the set up process of the tool. The data will be stored in a knowledge base and updated in a periodic basis or after events. An example of such events is a change in the risk appetite of the organisation. The provenance of this data will be diverse, some inputs will be loaded from tools or data bases and others manually. Examples of inputs stored in the knowledge base are: asset inventory, asset and impact valuation, initial threat quantification, traceability between threats and assets, and between processes and business objectives, and definition of normal and abnormal states of operation.

4.3 Risk Calculation Engine

The risk calculation engine would be in charge to process the inputs and generate the results. This engine will be implemented in a software tool. The risk calculation engine would be based on different standards to quantify vulnerabilities, threats, and impacts, the three variables that define a cyber-risk. The FAIR method (Factor analysis of information Risks) [30] is a useful start point to understand the different factors that are involved in a cyber-risk. This method defines the vulnerability as a combination of the "threat capability", the resourcefulness of the threat agents to act against an asset, and the "control strengths", the probability that the current controls resist the attack. While this definition is conceptually useful, the quantification of these two factors, as defined by this method can present problems. Both are defined according to their position within a probabilistic distribution of the threat population, meaning that it is necessary to be able to make plausible assumptions about the possible threat agents. The Common Vulnerability Scoring System (CVSS) [5], provides a score from zero to ten depending on eight variables which are related to the vulnerability. These

variables are: attack vector, attack complexity, privileges required, user interaction, scope, confidentiality, integrity, and availability. Additionally, temporal and environmental metrics help giving a score are used for more accuracy within an specific context.

Different methods were reviewed for quantification of threats. Some methods base the probability of a threat event in the frequency of occurrence [30][15]. This approach can be plausible in environments that maintain similar conditions over time. Nevertheless, in cyber-security assuming that threats will behave in the future following the same trend they have done in the past is dangerous. It has to be considered that often new attack mechanisms and zero-day vulnerabilities appear. The model proposed by SANDIA [22] analyses and scores threats by building a profile according to seven different attributes: intensity, stealth, time, technical personnel, cyber-knowledge, kinetic knowledge, and access. Different combinations of this attributes are used to describe 8 different threat profiles, where 1 represents the highest level of threat and 8 the lowest.

To calculate the impact, it is common transforming every consequence into monetary values, because it is an useful way to add up and compare impacts of diverse nature, such as time loss and reputation damage, among others. As a mean of normalising this value, it will be suggested within this method to use a ratio between the total impact of a risk and a referential budget that the organisation will define according to its risk appetite.

The quantification of risks done in the initial assessment would be subject to continuous updates during operation mode. This updates will be related to the threat analysis, which is the factor that presents the higher levels of uncertainty. The threat value then, is the risk component that will be subject to change dynamically according to the information provided by the dynamic inputs. When events that imply possible threats are detected, they will have an effect of modifying the quantification of the corresponding threat values, and therefore, the risk scores.

4.4 Results

Continuous re-calculations of Key Risk Indicators (KRI) will be performed in run time for monitoring purposes, and stored in a data base for ex-post analysis. The risk analyst and SOC operator will have different views of the KRI according to their roles. The risk analyst will be more interested in monitoring the behaviour of the risks and evaluating the effectiveness of current controls, as a mean to make better informed security-related decisions. The SOC operator will be more focused on alerts and any indicator of an attack developing in any of its stages. This is explained through an example in section 5. Risks are based on uncertainty. Thus, in the cases of an imminent attack, this is not considered a risk but an issue. In most occasions a cyber-attack will not take place in a single instance but it will follow a sequence of stages. Detection of an issue such as an unauthorised access or malware presence, can help to avoid the risk of an attack progressing into further stages, like privilege escalation, maintaining foothold, and establishing command and control capabilities.

4.5 Initial risk assessment and continuous monitoring dynamic

The method will consider two stages: the initial risk assessment and the continuous risk assessment. In the first stage the initial KRI are calculated and the risk monitoring tool is configured according to the context. This activity will condition the success of the continuous risk assessment. Therefore, it is crucial to develop a good understanding of the likelihood and impacts of a breach. Different forms and questionnaires based on standards (e.g. ISO 27005) will be developed to capture expert's opinion and guide the set-up process of the tool in a way that cyber risks can be mapped with their impacts at all levels of the organisation, including the business point of view. The second stage is the continuous risk assessment which consists in the recalculation of risk scores according to the information provided by the detection tools.

5 Demonstration of the method through the example

To demonstrate how the continuous risk assessment method would work, a threat scenario was built using the example of section 3. The initial risk assessment and tool configuration stage will consist in evaluating the system "as is". After this, a risk treatment plan is developed, incorporating controls to mitigate risks and residual risks are formally accepted. Then, the continuous risk assessment process starts.

5.1 Initial risk assessment and tool configuration.

In this stage risks are identified, quantified and evaluated. It is expected that after this assessment a risk treatment plan will be developed, incorporating controls. The next stage will provide the means to monitor the effectiveness of these measures in a continuous basis. The risk management approach will follow the process described by the ISO27005 standard shown in figure 5.

Context establishment. In this stage it is defined the scope, including assets that need protection. To illustrate the method some of the attack vectors related to the temperature control system will be reviewed. The assets considered in the assessment belong to both the data centre domain and the company business domain, because the temperature control system can be used as a bridge to other systems. Other control systems in the data centre such as CCTV, fire alarm system, electrical supply and UPS, as well as servers and network equipment are included in the monitoring. Their dependencies with the temperature control system need to be established at this point. Operational and business processes, as well as support processes such as IT, finance, and human resources should be considered. To quantify impacts, a series of possible consequences were listed including damages to assets implying replacement or reparation, effort spent in recovery, downtime, fines, and compensations to customers and third parties, and damage in brand reputation. This last can be reflected in loss of revenue or need to expend in marketing strategies to recover the trust of the customers. All

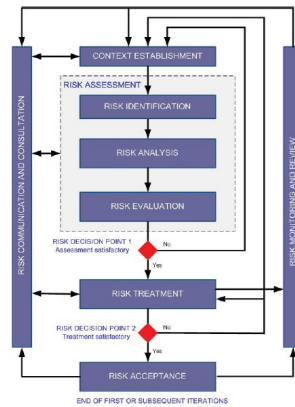


Fig. 5. Overview of the risk management process, ISO27005

the impacts are quantified in monetary values. To put this value in context according to its significance to the organisation, an impact indicator is calculated to evaluate the impact in terms of the monthly budget that the company assigns for cybersecurity.

The risk acceptance criteria is established to define a risk frontier depending on the organisations risk appetite. Once established this frontier, any risk that is out of limits is reviewed and included in the risk treatment plan. The decision of retaining any risk outside of these limits needs to be formally approved by senior management. We will suppose this company defined that any risk that is rated medium or higher or whose potential impact exceeds a predefined value should be reviewed. The roles involved in this risk assessment were: risk analyst, GRC manager, SOC analyst, Security manager, data centre manager and operators, the asset owners of all assets identified in the assessment, and senior management (CSO, CIO, CEO).

Risk assessment. In this stage, risks are identified, analysed, and evaluated. The analysis process consist in understanding possible attack paths, and consequences of a breach, and making plausible assumptions that allow identifying and quantifying threats, vulnerabilities, and potential impacts. Threats and vulnerabilities are linked with possible assets compromised and impacts at different levels. The likelihood is estimated based on the threat and vulnerability levels, and the total impact is quantified considering the effects on the operational processes as well as in the business.

From the analysis of the system “as is”, it was found possible that internal personnel or a contractor could download malicious code in the computer terminal using a USB drive. Thus, bridging the air-gap. The motivations of the attacker which could be many, for example, been bribed by a competitor to sabotage the servers. The malicious code could have different purposes, which means that there are different risks. The present example will focus on two risks: Risk 1 will correspond to the manipulation of the temperature to increase failure

rates of the servers and Risk 2 is related to the use of the BMS as a bridge to access the business networks. Each of the two risks is analysed independently.

Risk 1 was analysed obtaining a low level of vulnerability and a low level of threat. The impact was also considered low, since the system could be easily reset and reconfigured in the case of been compromised. As a result, this risk was rated as low. Risk 2 was rated as medium, because, although the threat and vulnerability were also considered low, the impact was considered high as critical processes and data could be compromised. Both risks were evaluated by comparing them with the risk acceptance criteria (risk frontier) concluding that only Risk 2 needed to be included in the risk treatment plan.

Risk treatment. After the risk assessment, a risk treatment plan is developed, to address all the risks that exceeded the risk frontier choosing from a set of three possible actions: reduce, avoid, or share. A fourth possible action is retaining the risk, which can be consider under senior management approval. In the case of Risk 2, using the BMS to access other networks, it was decided to reduce the risk by introducing a firewall in the local network.

Tool configuration. Residual risks are analysed and evaluated after the risk treatment plan is put in action. The information is loaded in the knowledge base, including the scores of all the risks identified. As Risk 2 was reduced, it is re-evaluated and defined as low which means that it is within acceptable limits. At this point, all the risk scores should be below the risk frontier. There could be two reasons why some risks might not meet this condition. One is that the control defined in the risk treatment plan has not been fully implemented yet, and the other is that there is a formal authorisation of senior management to retain certain risk. Each risk is linked with the events that might change the current scores in order to start the continuous monitoring process.

5.2 Continuous risk assessment.

In the previous stage, an initial iteration of the information security risk management process was done. The following step consists in monitoring the risks and perform subsequent iterations of the whole process. Through the continuous risk assessment, it is intended that these iterations will be repeated in short intervals of time and whenever a certain development of events requires it, rather than in a periodic basis.

For the current example, we will imagine that a contractor introduces malicious code through an USB drive in the computer terminal to change the temperature control settings. The malware will work in a similar way as Stuxnet, changing temperature setting of the DDCs and disguising this action by displaying the original set values of temperature. Therefore, there will not be any condition that triggers an alert directly related to overheating. For example, the set value is 25 degrees Celsius, and all the instances of the system display 25 degrees, but the temperature will really be 50 degrees Celsius. The malicious actions are scheduled to take place at hours where is more unlikely to be staff

on site to notice this, allowing persistence of the attack and increasing the potential damage. But, while the malware might remain unnoticed, the anomaly detection system will detect an unusual behaviour in other processes which are linked to this. The fans of the servers will spin faster, increasing the electrical power consumption and the servers might not perform in their best capacity. Considerable outliers in the energy consumption levels and in the performance of the servers will trigger an alert by the anomaly detection system. The risk calculation engine will process this information and modify the risk indicators of all the risks that can be related to this event, including Risk 1. An alert will be sent to the SOC operator to investigate the situation and call for action. In parallel, another alert will be sent to the risk analyst indicating that Risk 1 has now been rated as high. The risk analyst will also have the information about the processes that the affected servers are running and how they impact the business.

The previous example shows how a risk that was initially considered low changes to a higher value dynamically, through the development of events. The threat score is amplified when suspicious events are detected resulting in a higher risk value. It has to be noticed that this would lead to two courses of action. First, the SOC operator can generate an immediate response regarding remediation and recovery actions. Second, the risk analyst will generate all the necessary actions to develop a risk treatment plan to establish controls to avoid this risk becoming again an issue in the future. Examples of these actions are blocking all unused physical ports by default, restricting the privileges to download software, and adding malware detection and stricter regulations regarding not leaving any third party unattended in the perimeter.

6 Considerations and challenges

The development of this method will not be exempt of challenges. There are still unsolved issues in this project which need to be tackled in future stages of this project for the method and tool to have a practical application.

Because IoT and IIoT systems are very heterogeneous, the method only can be tested in a limited amount of systems. Tailoring guidelines can be provided for adapting the method to different use cases, and it would be a matter of further research to confirm its applicability in different contexts and scenarios. Big data issues including processing, storage and retrieval of information will also be a challenge, as well as the development of interfaces between tools and normalisation of the data.

As much as there might be a lot of ground in common with regular IT systems, this project aims to tackle the particular requirements of IIoT. One of the challenges of this is that the amount of processes, stakeholders, dependencies with other systems, and assets involved is bigger. Also, there might be more expectation for these systems to have automated security controls. Nevertheless, it must be recognised that the autonomy of any system will always be within certain limits. Establishing these limits, mapping all the processes affected, as

well as providing appropriate rules and training mechanisms to the anomaly detection system is part of the challenge.

False positives is a known problem of detection tools which would affect, as well this method. It is necessary to find solutions that do not undermine the ability of the method to alert when there is a real threat. The user will have the mission of calibrating the tool by identifying and giving feedback about any misinterpretation of the data. A case management system would be a possible alternative to support a continuous improvement mechanism for the method. Overall, it is important to understand the data in order to avoid providing misleading results. An example of this is the huge amount of noise that failed attack attempts can cause which may lead to think that there is a developing threat when actually according to [22] one attribute that increases the threat is, precisely, stealth. Attacks that are easy to detect and stop might not be a threat at all!

Another aspect is to distinguish cyber-attacks from other issues such as physical attacks or malfunctions of the system. It is considered on the best interest of an organisation to know about any potential threat even if it is not caused by a cyber-attack. Therefore, the detection of an issue whose causes end to be from a different nature, rather than been dismissed, should be reported to the relevant stakeholders. Although the scope of the method is to monitor cyber-risks, from the risk management point of view, other types of risk can also be of interest. For example, in [27] it is argued that physical attacks and cyber-attacks should not be treated separately proposing 4 types of attacks: physical-only, cyber only, cyber-enabled physical and physical-enabled cyber.

7 Conclusions

The presence of IoT in several industries and the increasing amount of cyber threats predicts a growth in the demand for cyber security solutions. Developing methods to maintain cyber-situational awareness through a continuous risk monitoring process can support rational and well informed decisions. The approach proposed takes into account the context of the system, as well as the business objectives and priorities. By linking the potential threats with the impacts and vulnerabilities it is possible to do a better prioritisation of security resources. Shorter iterations for risk assessments will make it possible to react in a more timely manner to changes in the environment. The underlying principle is that risk management should not be detached from the system's operations, it should be integrated, since both processes serve as input to each other.

Currently there are not widely used and tested solutions to evaluate IIoT cyber security risks in run time that include an holistic perspective of the system. The present paper gives a general description of a solution that has the potential of addressing several gaps of existent risk assessment approaches. Considering the context establishment and capturing expert's opinion is addressed on the initial assessment and tool configuration, and subject to updates. The feedback loop of the system requires experts to be involved in the process as well as to give

input to calibrate the method. The method also goes beyond attack or failure orientation because it is not limited to known attack mechanisms. By including anomaly detection and other tools it allows issuing alerts under any event that diverts the system's behaviour from what is considered normal. This is relevant, because it is not feasible to analyse all the possible attack mechanisms. The combination of different tools to support this method, as well as the development of the risk calculation engine as an automated tool, will allow the method to be implemented in a practical and effective way. If decision makers are well informed of cyber-security risks, this will allow better application of policies and control mechanisms, improving the overall security of the system.

References

1. Boddy, S., Shattuck, J.: Threat analysis report. the hunt for iot. the growth and evolution of thingbots ensures chaos (2018)
2. Cherdantseva, Y., Burnap, P., Blyth, A., Eden, P., Jones, K., Soulsby, H., Stoddart, K.: A review of cyber security risk assessment methods for scada systems. *computers & security* 56, 1–27 (2016)
3. Cisco: Cisco 2017 annual security report. Tech. rep. (2017)
4. Cook, E., Kearney, P.: Security challenges and cybercrime. *Journal of the Institute of Telecommunications Professionals* 9, 22–25 (2015)
5. Common vulnerability scoring system sig. <https://www.first.org/cvss/>, Accessed on: 2018-04-9
6. Dempsey, K., Chawla, N.S., Johnson, A., Johnston, R., Jones, A.C., Orebaugh, A., Scholl, M., Stine, K.: Information security continuous monitoring (iscm) for federal information systems and organizations: National institute of standards and technology special publication 800-137 (2012)
7. Dempsey, K., Ross, R., Stine, K.: Supplemental guidance on ongoing authorization (2014)
8. Desnitsky, V., Kotenko, I., Nogin, S.: Detection of anomalies in data for monitoring of security components in the internet of things. In: *Soft Computing and Measurements (SCM), 2015 XVIII International Conference on*. pp. 189–192. IEEE (2015)
9. ENISA: Security Recommendations for IoT in the context of Critical Information Infrastructures. <https://www.enisa.europa.eu/publications/baseline-security-recommendations-for-iot> (2017), Accessed on: 2018-04-09
10. Gartner newsroom. <http://www.gartner.com/newsroom/id/3598917> (2017), Accessed on: 2017-07-30
11. Greensmith, J.: Securing the internet of things with responsive artificial immune systems. In: *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*. pp. 113–120. ACM (2015)
12. Henrie, M.: Cyber security risk management in the scada critical infrastructure environment. *Engineering Management Journal* 25(2), 38–45 (2013)
13. Huang, H., Xie, D.: Real-time network risk evaluation paradigm-inspired by immune. In: *Natural Computation (ICNC), 2015 11th International Conference on*. pp. 786–790. IEEE (2015)
14. IBM Institute for Business Value: Internet of threats. securing the internet of things for industrial and utility companies (2018)

15. ISO/IEC: Iso/iec 27005:2011. information security risk management (2011)
16. Jing, Q., Vasilakos, A.V., Wan, J., Lu, J., Qiu, D.: Security of the internet of things: perspectives and challenges. *Wireless Networks* 20(8), 2481–2501 (2014)
17. Kotenko, I., Saenko, I., Ageev, S.: Countermeasure security risks management in the internet of things based on fuzzy logic inference. In: *Trust-com/BigDataSE/ISPA, 2015 IEEE*. vol. 1, pp. 654–659. IEEE (2015)
18. Lin, S.W., Miller, B., Durand, J., Joshi, R., Didier, P., Chigani, A., Torenbeek, R., Duggal, D., Martin, R., Bleakley, G., et al.: Industrial internet reference architecture. Industrial Internet Consortium (IIC), Tech. Rep (2015)
19. Linda, O., Manic, M., Vollmer, T.: Improving cyber-security of smart grid systems via anomaly detection and linguistic domain knowledge. In: *Resilient Control Systems (ISRCs), 2012 5th International Symposium on*. pp. 48–54. IEEE (2012)
20. Liu, C., Zhang, Y., Zeng, J., Peng, L., Chen, R.: Research on dynamical security risk assessment for the internet of things inspired by immunology. In: *Natural Computation (ICNC), 2012 Eighth International Conference on*. pp. 874–878. IEEE (2012)
21. Macaulay, T.: *RIoT Control: Understanding and Managing Risks and the Internet of Things*. Morgan Kaufmann (2016)
22. Mateski, M., Trevino, C.M., Veitch, C.K., Michalski, J., Harris, J.M., Maruoka, S., Frye, J.: *Cyber threat metrics*. Sandia National Laboratories (2012)
23. Moss, D.L.: *Data center operating temperature: The sweet spot* (2011)
24. Pan, S., Morris, T., Adhikari, U.: Developing a hybrid intrusion detection system using data mining for power systems. *IEEE Transactions on Smart Grid* 6(6), 3104–3113 (2015)
25. Positive Technologies: *Industrial companies attack vectors*. <https://www.ptsecurity.com/upload/corporate/ww-en/analytics/ICS-attacks-2018-eng.pdf> (2018), Retrieved on: 2018-06-25
26. Sadeghi, A.R., Wachsmann, C., Waidner, M.: Security and privacy challenges in industrial internet of things. In: *Design Automation Conference (DAC), 2015 52nd ACM/EDAC/IEEE*. pp. 1–6. IEEE (2015)
27. Smith, B.J., Sholander, P.E., Phelan, J.M., Wyss, G.D., Varnado, G.B., Depoy, J.M.: *Risk assessment for physical and cyber attacks on critical infrastructures*. Tech. rep., Sandia National Laboratories (2005)
28. Spyridopoulos, T., Maraslis, K., Tryfonas, T., Oikonomou, G., Li, S.: Managing cyber security risks in industrial control systems with game theory and viable system modelling. In: *System of Systems Engineering (SOSE), 2014 9th International Conference on*. pp. 266–271. IEEE (2014)
29. Symantec: *Istr- internet security threat report*. Tech. rep. (2017)
30. The Open Group: *Fair iso/iec 27005 cookbook* (2010)
31. Wang, J., Fan, K., Mo, W., Xu, D.: A method for information security risk assessment based on the dynamic bayesian network. In: *Networking and Network Applications (NaNA), 2016 International Conference on*. pp. 279–283. IEEE (2016)
32. Yang, Y., McLaughlin, K., Sezer, S., Littler, T., Im, E.G., Pranggono, B., Wang, H.: Multiattribute scada-specific intrusion detection system for power networks. *IEEE Transactions on Power Delivery* 29(3), 1092–1102 (2014)
33. Zhang, Q., Zhou, C., Tian, Y.C., Xiong, N., Qin, Y., Hu, B.: A fuzzy probability bayesian network approach for dynamic cybersecurity risk assessment in industrial control systems. *IEEE Transactions on Industrial Informatics* (2017)