Case Study: Conducting a Risk Assessment for an Electrical Utility

Edwin Covert, CISSP-ISSAP, CISM, CRISC, SCF, PMP
Colorado State University Global, Aurora, Colorado, USA

edwin@edwincovert.com

Abstract: Risk management is an important part of effective cybersecurity. This paper presents a hypothetical risk assessment as a case study for one of the largest electricity providers in the southern California region using the approach outlined in the National Institute of Standards and Technology Special Publication 800 series on cybersecurity.

Keywords: Risk Management, Risk Assessment, Case Study

1. Introduction

Electrical generation is the lifeblood of the US and the global economy, according to Webb (2016). One firm dominates the creation of electricity in southern California. This paper will discuss various aspects of cybersecurity risk management relevant to a local utility company's (referenced as either company or utility in henceforth) management, including how to document specific risks it might face. Next, this paper will address high-risk concerns through the application of proven mitigation techniques. As a business, the company has standard cybersecurity concerns with corporate networks, such as access control, data loss prevention and business continuity. However, what makes this firm unique is that it manages a large electrical grid. The company recognizes that it possesses critical information technology systems, sensitive customer and employee data, network infrastructure, and information vital to its business.

2. The Company's Organizational Risks

Electrical generation requires Industrial Control Systems (ICS) to create and distribute power where it needs to go. Within the utility sector, risks to ICS remain high (Khodabakhsh et al., 2020;). ICS systems are those that are "used to control industrial processes such as manufacturing, product handling, production, and distribution" (National Institutes of Standards and Technology, 2019, para. 1). According to Khodabakhsh et al. (2020), they comprise "assets such as supervisory control systems and data acquisition (SCADA), distributed control systems (DCS), and human-machine interfaces (HMI) that are commonly used for monitoring and control of their critical infrastructure" (p. 1). Vendors increasingly developed them to use internet-based communication protocols (Ovaz Akpinar & Ozcelik, 2018).

The Cybersecurity and Infrastructure Security Agency (CISA) (2021) has released detailed advisories surrounding ICS prompted by several examples of ICS attacks that made headlines. One of the first occurred in 2000 when Vitek Boden used stolen radio equipment and drove to 46 radio-controlled sewage control systems in Queensland, Australia (Abrams & Weiss, 2008). In each case, he used the stolen radio equipment to shut off industrial controls remotely that the maintenance company installed for remote access. This led to 800,000 gallons of untreated sewage released into the water supply. Citing the Australian Environmental Protection Agency, Abrams and Weiss (2008) note, "Marine life died, the creek water turned black, and the stench was unbearable for residents" (p. 1).

A second example involves Programmable Logic Controllers (PLC) and a type of malware called a worm. PLCs are another type of ICS. According to Kaspersky (2021), threat actors design worms with an eye towards self-replication: "[w]orms do not require activation—or any human intervention—to execute or spread their code" (para. 1). A worm is a self-replicating program designed to disrupt operations. Spenneberg et al. (2016) documented as a proof-of-concept a worm in a popular PLC that would consume PLC resources and significantly affect operations. Researchers presented this scenario at a renowned cybersecurity conference.

A final example of an in-the-wild risk to ICS is WIN32/Industroyer. This malware targets ICS used in electrical substations. According to Cherepanov (2017), WIN32/Industroyer targets four distinct ICS protocols, showing the threat actors had advanced knowledge of ICS design. The malware installs a separate backdoor to additional attacks, wipes data, and contains additional payloads designed to control ICS (Cherepanov, 2017).

Threats to ICS are increasing (Radanliev et al., 2018). Therefore, organizations must determine their ICS risk; for this, there is a standard formulation: risk equals the likelihood of a threat exploiting a vulnerability and causing an impact on operations (National Institute of Standards and Technology, 2012). Figure 1 shows this graphically. This paper will explore each of these as they apply to the company.

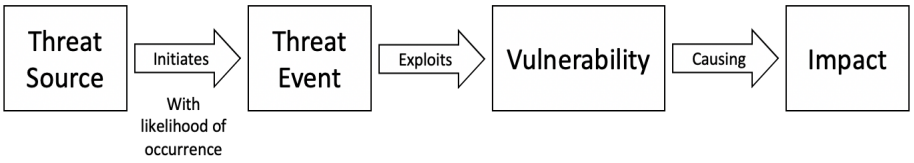

Figure 1: Determining risks

Note: Adapted from "Guide for Conducting Risk Assessments (SP 800-30, rev. 1)" by National Institute of Standards and Technology, 2012.

2.1 Threats

The National Institute of Standards and Technology (NIST) (2012) defines a threat as "any circumstance or event with the potential to harm organizational operations and assets, individuals, other organizations, or the Nation through an information system via unauthorized access, destruction, disclosure, or modification of information, and/or denial of service" (p. 8). Applied to this organization, threats come from a variety of sources. The US Intelligence Community (IC) states cybersecurity attacks from "nation states and their surrogates will remain acute" (Office of the Director of National Intelligence, 2021, p. 20). The IC notes that cybersecurity operations are increasingly a tool that countries apply to increase their national power; these operations will affect civilians and non-military targets. The Annual Threat Assessment from the IC also calls out threat actors who have attacked "software and [information technology] service supply chains" (p. 21).

While nation state threat actors are important to counter, threats come in a variety of flavors, both intentional and accidental. Additional intentional threats include criminals and activists who might target the company; the techniques might be the same but the motivations are different. Nation states seek to cause disruptions to rivals; criminals seek financial gain; activists look to increase the visibility of their cause. The company must account for accidental threats, including natural disasters and errors caused by users as well.

An important way to understand what threats the utility faces is via a threat model. A threat model creates a picture of the overall attack surface it presents to the outside world. The attack surface is what a threat actor sees at the perimeter of a system where they will try to enter (National Institutes of Standards and Technology, 2019). Understanding the attack surface is knowing what parts of the company's information systems are vulnerable and need to be tested, according to OWASP (2021).

Developing a threat model involves decomposing the information system into its individual components, identifying the risks presented (because of design or software elements used) and developing countermeasures to each of the risks identified (Shostack, 2014). Analysts use the threat modeling process to "analyze potential attacks or threats, and can also be supported by threat libraries or attack taxonomies" (Xiong & Lagerström, 2019, p. 56). Analysts then collate each of these attacks into a sense of situational awareness, i.e. the attack surface of the systems in question.

Using the STRIDE model allows organizations to consider the range of threats it faces. Each letter in the mnemonic stands for a distinct threat tactic: *S* for spoofing, *T* for tampering, *R* for repudiation, *I* for information disclosure, *D* for denial of service, and *E* for elevation of privileges. Table 2 defines each of these terms. Each intentional threat actor in this section can and will use one or more of these threat tactics to reach their aim. MITRE (2020) has created a detailed list of attack techniques threat actors have employed that the company can use to inform its threat model.

Table 2: Definitions for each letter in the STRIDE threat model method

Term	Definition
Spoofing	Pretending to be something the user or system is not
Tampering	Modifying something in an unauthorized manner
Repudiation	Falsely claiming a user or system did or did not do something
Information Disclosure	Revealing information meant to be protected
Denial of Service	Preventing a service that suppose to run from happening
Elevation of Privileges	Doing things a user or system should not be able to do

Note: Adapted from "Threat Modeling: Designing for Security" by A. Shostack, 2014. Copyright 2014 by Wiley.

2.2 Vulnerabilities

Cybersecurity professionals define a vulnerability as a "weakness in an information system, system security procedures, internal controls, or implementation that could be exploited by a threat source" (National Institute of Standards and Technology, 2012, p. 9). There are a multitude of vulnerabilities in ICS the company must know. NIST published its *Guide to Industrial Control Systems (ICS) Security* in 2015. Appendix C of this NIST document breaks down vulnerabilities into five categories: architecture and design, configuration and maintenance, physical access, software development, and communications and network (Stouffer et al., 2015).

2.2.1 Architecture and Design

Architecture and design vulnerabilities are flaws built into how an organization might create an ICS. For example, Stouffer et al. (2015) describe a particular vulnerability where the ICS does not have a defined security perimeter. This lack of perimeter means the utility cannot ensure they have properly deployed the controls they rely on for protecting the confidentiality, availability, and integrity of the ICS. This can lead to potentially unauthorized access to the ICS and its data (information disclosure in the STRIDE process).

2.2.1 Configuration And Maintenance

Configuration and maintenance vulnerabilities occur when organizations like this one do not properly care for their ICS. Here, improperly configured ICS creates openings to the utility from unnecessary services, functions, ports, and protocols (Stouffer et al., 2015). Leaving default configurations from an installation will expose weaknesses and services to attack, leading to tampering or elevation of privileges under a STRIDE model.

2.2.1 Physical Access

Stouffer et al. (2015) cite allowing unauthorized personnel to have physical access to an ICS as a physical vulnerability the company must remediate. Failure to do so can lead to theft or destruction of data or hardware and unauthorized changes to the ICS through a variety of means. This is potentially a denial of service within STRIDE's framework.

2.2.1 Software Development

Software development vulnerabilities occur when developers use bad coding practices, such as improper data validation techniques on user inputs and received data to the ICS. Such a vulnerability will create many "vulnerabilities including buffer overflows, command injections, cross-site scripting, and path traversals" (Stouffer et al., 2015, p. C-9). Such attacks will lead to tampering or elevation of privileges under a STRIDE model.

2.2.1 Communications And Network

Communications and network vulnerabilities are those that happen when part of the ICS is communicating with each other. For example, if two ICS components communicate over a wireless network, Stouffer et al. (2015) note there needs to be sufficient data protection between the two points. These ICS elements transmit sensitive information on behalf of the utility and should have strong encryption in place to secure the communications link. A lack of encryption on these communications can lead to tampering noted in the STRIDE rubric.

2.3 Likelihood

Returning to NIST's (2012) *Guide for Conducting Risk Assessments*, the document defines likelihood as a "weighted risk factor based on an analysis of the probability that a threat can exploit a vulnerability (or set of vulnerabilities)" (p. 10). Unfortunately, determining the likelihood of a cyber attack is challenging. While the Cyber Risk Task Force of the American Academy of Actuaries (2021) notes such attacks are real in the modern world, using past data to determine future probability is hard because data sharing about past attacks is notoriously difficult to come by.

CISA (n.d.) within the US government requests the business community to share cyber threat data with them to better inform the larger cybersecurity community, but is not a requirement. Recently, the US Senate removed a requirement that would have mandated such reporting to CISA from the National Defense Authorization Act, dooming it to possibly passing at a later date (Starks, 2021). However, each critical infrastructure sector maintains an Information Sharing and Analysis Center (ISAC) and the utilities sector is no different.

The Electricity Information Sharing and Analysis Center (E-ISAC) (2020) serves to lessen the risk of cyber and physical security risks for its members; principally, the E-ISAC does this through its Cybersecurity Risk Information Sharing Program (CRISP) program. The company is presumably a member of E-ISAC based on its

geographical coverage and importance to the southern California region; E-ISAC does not publish its membership roster.

Irrespective of the determining actual probabilities that a threat actor could exploit a vulnerability, the utility needs only look at recent zero-day attacks announced to see there is always a high likelihood of attempted penetration. For example, researchers discovered a new vulnerability for Log4j (a key part of many Unix logging systems) and issued an alert on December 10, 2021 (National Institute of Standards and Technology, 2021). Four days later, attackers launched over 840,000 attacks on Log4j components (Jeffrey, 2021).

2.4 Impact

Impact is the amount of damage expected from a particular threat exploiting vulnerability for the company (National Institute of Standards and Technology, 2012). In more practical terms, impacts are usually costs associated with an ICS component failure and the need to replace it, paying to address reputational damage from angry customers, or for addressing new regulatory requirements on the heels of an attack it might have failed to address.

3. Calculating The Organization's Risk Exposure

NIST (2012) describes a method for document specific risks. The first step in the process is to determine the threat sources or events. Next, the company would determine the vulnerabilities it faces within its attack surface from each threat source. Third, the utility should determine the likelihood a threat would exploit each identified vulnerability and then determine what impact that exploitation would have on its operations. Figure 2 below outlines the overall risk assessment process.

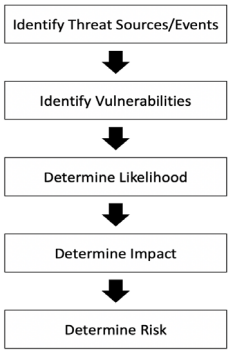

Figure 2: Risk Assessment Process

Note: Adapted from "Guide for Conducting Risk Assessments (SP 800-30, rev. 1)" by National Institute of Standards and Technology, 2012.

3.1 Threat-Vulnerability Pairs

A complete cybersecurity risk assessment for this company is beyond this paper. However, using the threats and vulnerabilities already presented, examples become obvious using the process described in Figure 2. For brevity's sake, this paper will only use the intentionally malicious threats described earlier: nation-state attackers, financially motivated criminals, and activists seeking to promote their cause. This paper randomly pairs each of these three threat sources with one of the five vulnerabilities discussed: lack of security perimeter, improper system configurations, unauthorized physical access to company facilities, improper data validation, and a lack of data protection between nodes. As noted previously, in a full risk assessment, the utility should map each threat source to each vulnerability.

This creates five threat-vulnerability pairs (P). Table 3 below provides them along with a brief narrative description of each. Each description also notes the associated MITRE (2020) attack technique in parentheses relevant to each P value.

Table 3: Organizational Threat-Vulnerability Pairs with Associated MITRE Attack Techniques

Identifier	Threat-Vulnerability Pair	Description
P1	Nation state - Perimeter	A nation state might attack an ICS device Edison International believes is protected by its perimeter but is not (*T0829*)
P2	Crime - System Configuration	A cyber criminal could attack an ICS with default system configurations (*T0885*)
P3	Activity - Physical	An activist group might gain access to an Edison electrical substation and destroy it (*T0879*)
P4	Nation state - Data Validation	A nation state could attack a database that stores ICS information using SQL injection techniques (*T0811*)
P5	Crime - Comms Protection	A cyber criminal might change vital communications from between ICS nodes if there is no encryption on the link (*T0830*)

Note: Adapted from "Guide for Conducting Risk Assessments (SP 800-30, rev. 1)" by National Institute of Standards and Technology, 2012 and "ICS Attack Techniques" by MITRE, 2020. Copyright 2020 by MITRE.

3.2 Risk Matrix

With qualitative values for likelihood and impact, the company can build a table that shows their intersection. The company should define these qualitative values (high, moderate, and low) in terms of their business needs. High likelihood might mean there is an openly available proof-of-concept piece of software for a known vulnerability on the internet, while a low impact could be that the affected system would not disrupt electrical generation and distribution. Table 4 below presents such a matrix combining likelihood and impact. This is only an example matrix; the company would need to tailor its values to its risk tolerance levels.

Table 4: Likelihood-impact matrix to show risk levels

		Impact		
		Low	Moderate	High
Likelihood	Low	Low	Moderate	Moderate
	Moderate	Moderate	Moderate	High
	High	Moderate	High	High

Note: Adapted from "Guide for Conducting Risk Assessments (SP 800-30, rev. 1)" by National Institute of Standards and Technology, 2012.

Table 4 implements a high-water mark process. Whatever the highest value is for either likelihood or impact becomes the risk level associated with that combination. The exception to this is the two combinations of low and high likelihood and impact. For those two items (*low likelihood, high impact* and *high likelihood, low impact*) this matrix splits the difference and marks both as moderate. For each threat-vulnerability pair in table 3, the company can enter the business appropriate likelihood and impact and then calculate the risks each pair presents to its business. Table 5 presents examples of what these results might look like using the matrix in Table 4**.**

Table 5: Likelihood, impact, and associated risk for each identified threat-vulnerability pair

Identifier	Likelihood	Impact	Associated Risk
P1	Low	High	Moderate
P2	High	High	High
P3	Low	Moiderate	Moderate
P4	High	High	High
P5	Low	Low	Low

Note: Adapted from "Guide for Conducting Risk Assessments (SP 800-30, rev. 1)" by National Institute of Standards and Technology, 2012.

Like most organizations, this utility does not possess unlimited resources and must prioritize where to get the greatest return on its investment in controls. Reviewing the information in Table 5, the company should focus its resources on resolving those concerns with the highest risks: TV2 (*a cyber criminal could attack an ICS with default system configurations*) and TV4 (*a nation state could attack a database that stores ICS information using SQL injection techniques).*

4. Mitigation Techniques for the Organization's High-Risk Pairs

How might the company address the two high-risk items? MITRE (2020) recommends specific countermeasures for each attack technique associated with the threat-vulnerability pairs.

4.1 P2 - Crime - System Configuration

To mitigate against ICS attack technique *T0885 - Commonly Used Port*, MITRE (2021b) recommends several changes. First, each controller deployed by the utility in the field should "require users to authenticate for all remote or local management sessions" (para. 3). MITRE's (2021b) second recommended mitigation is to disable all unnecessary ports and protocols; only allow what is absolutely necessary to function on the ICS component. The company can confirm this standard by using a commercially available Security Content Automation Protocol (SCAP) enabled vulnerability scanners to seek open ports and protocols. SCAP is "a suite of specifications for exchanging security automation content used to assess configuration compliance and to detect vulnerable versions of software" (National Institute of Standards and Technology, 2016, para. 1).

Third, it should install a network-based intrusion prevention system. Such a system analyzes specific information from the network attached ICS components to seek the characteristics that show an attack is occurring and prevent the attacking traffic from reaching its target (Wang & Xu, 2020). The last recommendation MITRE (2021b) makes is to segment the ICS components into their own distinct network space. Experts call this network segmentation. Specifically, the company should "[c]onfigure internal and external firewalls to block traffic using common ports that associate to network protocols that may be unnecessary for that particular network segment" (para. 3).

4.2 P4 - Nation state - Data Validation

To address the concerns in *T0811 - Data from Information Repositorie*s, MITRE (2021a) similarly makes a series of recommendations. First, the company must encrypt the information in its databases. There are many options for such an encryption scheme; the company should find the one that provides the correct balance between key management and flexibility (Ocenas et al., 2020). Second, the company should make use of a Privileged Account Management (PAM) solution that will "minimize permissions and access for service accounts to limit the information that may be exposed or collected by malicious users or software" (MITRE, 2021a, para.1). There are several available for purchase and implementation.

Third, MITRE (2021a) recommends the company restrict permissions to files and directories on the database server. This will prevent threat actors from interacting with and collecting data from the database server. This coincides with the fourth recommendation: ensuring users and groups have only the permissions via a properly configured Identity and Access Management (IdAM) program (MITRE, 2021a). Finally, the company should "develop an auditing mechanism to conduct periodic reviews of accounts and privileges for critical and sensitive repositories" (para. 3)

4.3 Beyond the High-Risk Findings

Once the utility addresses these two high-risk threat-vulnerabilities pairs, they can then focus on the moderate-risk items (P1 and P3) before moving on to the low-risk item (P5) presented in this example risk assessment. For each of the techniques listed in Table 3 and the calculated risks in Table 5, MITRE (2020) provides specific mitigation tactics it should implement.

5. Conclusion

This utility provides energy for an important part of the world. Understanding the threats and vulnerabilities companies face will improve cybersecurity across the country's critical infrastructure. Within the utility sector, risks to Industrial Control Systems (ICS) remain high. Cybersecurity and Infrastructure Security Agency (CISA) (2021) has released detailed advisories surrounding ICS.

Tried-and-true processes exist for assessing risks. While there are many processes for conducting a risk assessment, the US government's process is straight-forward. This particular utility company should find the process that best meets their business needs. In the process described previously, the company must identify its threats and understand its vulnerabilities. From there, they should estimate, with key management officials, the likelihood of those threat-vulnerability pairs affecting operations. The company can rank each threat-vulnerability pair, and its associated likelihood and impact, using a standard risk matrix. This will allow it to focus their resources on those pairs most likely to cause significant damage to its operations and business.

References

Abrams, M. & Weiss, J., 2008. Malicious control system cyber security attack case study–Maroochy Water Services, Australia.

Cart, J., 2021. California's 2020 Fire Siege: Wildfires by the numbers. CalMatters. Available at: https://calmatters.org/environment/2021/07/california-fires-2020/ [Accessed January 15, 2022].

Census Bureau, n.d.. U.S. Census Bureau quickfacts: Los Angeles County. QuickFacts. Available at: https://www.census.gov/quickfacts/fact/table/losangelescountycalifornia,CA/PST045219 [Accessed December 10, 2021].

Cherepanov, A., 2017. WIN32/Industroyer: a new threat for industrial control systems.

Cyber Risk Task Force, Casualty Practice Council, 2021. Cyber risk toolkit.

Cybersecurity and Infrastructure Security Agency, n.d.. Systemic cyber risk reduction. Cybersecurity and Infrastructure Security Agency. Available at: https://www.cisa.gov/systemic-cyber-risk-reduction [Accessed January 9, 2022].

Electricity Information Sharing and Analysis Center, 2020. E-ISAC long-term strategic plan update.

Fontinelle, A., 2021. Getting a grip on holding companies. Investopedia. Available at: https://www.investopedia.com/terms/h/holdingcompany.asp [Accessed December 10, 2021].

Hughes, R.A., 2020. If California were a country. Bull Oak Capital. Available at: https://bulloakcapital.com/blog/if-california-were-a-country/ [Accessed December 10, 2021].

Jeffrey, C., 2021. LOG4J flaw turns into pandemic with over 840,000 attacks initiated within 72 hours. TechSpot. Available at: https://www.techspot.com/news/92633-hackers-launch-over-840000-attacks-through-log4j-flaw.html [Accessed January 9, 2022].

Kaspersky, 2021. What's the difference between a virus and a worm? www.kaspersky.com. Available at: https://www.kaspersky.com/resource-center/threats/computer-viruses-vs-worms [Accessed December 16, 2021].

Khodabakhsh, A. et al., 2020. Cyber-risk identification for a digital substation. Proceedings of the 15th International Conference on Availability, Reliability and Security.

LinkedIn, 2021. Adam Tuzzolino. LinkedIn.com. Available at: https://www.linkedin.com/in/adam-tuzzolino-40a7596 [Accessed December 10, 2021].

MITRE, 2020. ICS attack techniques. Techniques - Attacks ICS. Available at: https://collaborate.mitre.org/attackics/index.php/All_Techniques [Accessed January 8, 2022].

MITRE, 2021. T0811 - data from information repositories. Data from Information Repositories - attack ICS. Available at: https://collaborate.mitre.org/attackics/index.php/Technique/T0811 [Accessed January 9, 2022].

MITRE, 2021. T0885 - Commonly used port. Commonly Used Port - attack ICS. Available at: https://collaborate.mitre.org/attackics/index.php/Technique/T0885 [Accessed January 9, 2022].

National Institute of Standards and Technology, 2012. Gaithersburg, MD: US Department of Commerce.

National Institute of Standards and Technology, 2016. Security content automation protocol: CSRC. CSRC. Available at: https://csrc.nist.gov/Projects/Security-Content-Automation-Protocol/faqs [Accessed October 14, 2021].

National Institute of Standards and Technology, 2021. CVE-2021-44228. NVD. Available at: https://nvd.nist.gov/vuln/detail/CVE-2021-44228 [Accessed January 9, 2022].

National Institute of Standards and Technology, n.d.. Baseline configuration - glossary. CSRC. Available at: https://csrc.nist.gov/glossary/term/baseline_configuration [Accessed November 24, 2021].

Ocenas, M. et al., 2020. Security and encryption at modern databases. Proceedings of the 2020 4th International Conference on Cryptography, Security and Privacy, pp.19–23.

Office of the Director of National Intelligence, 2021. Washington, DC: US Government.

Ovaz Akpinar, K. & Ozcelik, I., 2018. Development of the ECAT preprocessor with the Trust Communication Approach. Security and Communication Networks, 2018, pp.1–16.

OWASP, 2021. Attack surface analysis cheat sheet. Attack Surface Analysis - OWASP Cheat Sheet Series. Available at: https://cheatsheetseries.owasp.org/cheatsheets/Attack_Surface_Analysis_Cheat_Sheet.html [Accessed October 14, 2021].

Radanliev, P. et al., 2018. Future developments in cyber risk assessment for the internet of things. Computers in Industry, 102, pp.14–22.

Shostack, A., 2014. Threat modeling: designing for security, Wiley.

Spenneberg, R., Bruggemann, M. & Schwartke, H., 2016. PLC-Blaster: a worm living solely in the PLC. Blackhat. Available at: https://www.blackhat.com/docs/asia-16/materials/asia-16-Spenneberg-PLC-Blaster-A-Worm-Living-Solely-In-The-PLC-wp.pdf [Accessed January 3, 2022].

Starks, T., 2021. Cyber incident reporting mandates suffer another Congressional setback. CyberScoop. Available at: https://www.cyberscoop.com/cyber-incident-reporting-ransomware-payments-congress-ndaa/ [Accessed January 9, 2022].

Stouffer, K. et al., 2015. Guide to industrial control systems (ICS) security. NIST Special Publication 800-82, pp.1–247.

Wang, D. & Xu, G., 2020. Research on the detection of network intrusion prevention with SVM based optimization algorithm. Informatica, 44(2).

Webb, E.L., 2016. The internet of things: cybersecurity, insurance, and the national power grid. Natural Resources & Environment, 30(4), pp.35–39.

Xiong, W. & Lagerström, R., 2019. Threat modeling – A systematic literature review. Computers & Security, 84, pp.53–69.