Designing an Industrial Cybersecurity Program for an Operational Technology Group

Matthew Sell
mrsell@uw.edu
University of Washington
Bothell, Washington, USA

Marc Dupuis
marcjd@uw.edu
University of Washington
Bothell, Washington, USA

ABSTRACT

The design of a cybersecurity program for an Information Technology ("IT") group is well documented by a variety of international standards, such as those provided by the U.S. National Institute of Standards and Technology ("NIST") 800-series Special Publications ("SP 800"). However, for those wishing to apply standard information security practices in an Operational Technology ("OT") environment that supports industrial control and support systems, guidance is seemingly sparse. This project expands on the abstract concepts described in textbooks by documenting the implementation of an industrial cybersecurity program for a local manufacturing firm. The project started with hardware and software asset inventories, followed by a risk assessment and gap analysis, and then implemented mitigating controls using a combination of manual and automated procedures. Security posture of the OT group was constantly evaluated against corporate security goals, the project-generated risk assessment, and NIST SP 800-171 requirements. Improvements in security posture and compliance to corporate requirements were achieved in part through alignment with existing policies and procedures developed by the organization's IT group, with the balance implemented and documented by the author of this project. The materials generated by this project may be used to assist other organizations starting their journey towards securing their industrial control assets.

CCS CONCEPTS

• **Security and privacy → Human and societal aspects of security and privacy**; • **Applied computing → Command and control**; **Enterprise computing**.

KEYWORDS

industrial control, cybersecurity program, operational technology

ACM Reference Format:
Matthew Sell and Marc Dupuis. 2023. Designing an Industrial Cybersecurity Program for an Operational Technology Group. In *The 24th Annual Conference on Information Technology Education (SIGITE '23), October 11–14, 2023, Marietta, GA, USA.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3585059.3611438

SIGITE '23, October 11–14, 2023, Marietta, GA, USA

ACM ISBN 979-8-4007-0130-6/23/10.
https://doi.org/10.1145/3585059.3611438

1 INTRODUCTION

1.1 Problem Definition

For those manufacturing enterprises developing or expanding an information security program, the need to secure network-connected Industrial Control Systems (ICS) presents unique and challenging problems. A combination of legacy protocols, insecure design practices, and expensive equipment with long lifetimes and slow replacement schedules all conspire to complicate the creation of security policies and procedures [12, 18]. Frequently, special considerations must be made in the selection of products and tools when securing industrial control systems [1], resulting in increased costs from the purchase, training, and deployment of specialized applications, hardware, and network configurations.

Frequently, manufacturing enterprises delegate the task of supporting different computing environments into two major groups: information technology (IT) and operations technology (OT) [1]. While the traditional IT group manages equipment and tools utilized in the office environment, the OT group focuses on equipment used in the manufacturing environment. Both groups are tasked with the challenge of securing their assets, although the focus of information security efforts between the two are slightly different. With respect to the cybersecurity "triad" (confidentiality, integrity, and availability), IT groups are generally focused more on the confidentiality aspect, while OT groups will usually focus more on availability aspects of information security [7].

Guidance for establishing an information security program that supports OT assets seemingly is focused on "top-down" approaches, where the IT organization can leverage existing tools, personnel, and support from vendors to extend into the industrial control domain [7]. Planning and deployment guides are helpfully provided by ICS product vendors such as Cisco Systems and Rockwell Automation to aid engineers and security personnel in the herculean task of architecting secure control networks [8].

For those OT groups developing an information security program from the "bottom-up", guidance on the subject is much sparser. The deployment of policies and tools must be performed in such a manner as to not interfere with existing operations or negatively affect legacy equipment and applications. The option to replace older equipment for the purposes of removing insecure communication protocols is typically not available in the bottom-up approach [16]. Additionally, legacy software applications and execution platforms must also be included in the information security program, as they must interact with existing equipment and processes [14].

1.2 Goals

This project aimed to develop a "bottom-up" information security program for an operations technology group supporting a large electronics manufacturing organization. This approach was selected after consultation with operations management since the bottom-up approach appeared to provide the lowest risk of disruption to production systems while the project was being implemented. While assistance was made available from the corporate information security group, the author's OT group ("Factory Information Systems", or FIS) was ultimately responsible for improving the security posture of assets under its control.

The primary goal of the project was to improve the security posture of the operations technology group by developing policies that complemented and extended the existing IT-oriented information security program. Progress was measured by the mitigation of vulnerabilities found early in the project through the performance of a security risk assessment and gap analysis. Additional metrics relating to compliance with regulations and vulnerability scoring were measured and evaluated as well. Where practical, software automation tools were employed to ease the burden of applying baseline controls, security patches, and compliance auditing. A training program was developed to ensure that personnel employed by the operations technology group were able to maintain the security controls developed during this project.

A secondary goal of this project was to develop a process by which other operational technology groups within the company can develop their own information security program, utilizing policies, procedures, and automation tools developed during this project to reduce their development time.

1.3 Scope

Within any information security program, there are many aspects that must be examined and considered. This includes the human factor and how organizations educate, train, and bring awareness to cybersecurity issues. The human factor of cybersecurity may be incorporated into an overall risk management program through awareness posters, gamification, computer-based training (CBT) [10], the use of fear or shame to engender policy compliance [19, 20], among other approaches. While all of these aspects are important to organizational risk management, they were not a focus of the current project.

2 RELATED WORK

2.1 Asset Inventory

The National Institute of Standards and Technology ("NIST") recently published the details of a reference implementation (SP 1800-5, September 2018) for managing information technology assets, both hardware and software-based [2]. Although this implementation was geared towards large financial institutions, it embodies best practices documented in other NIST Special Publication ("SP") 800-series [3] publications, such as those that specifically cover manufacturing operations [7].

NIST SP 1800-5 effectively describes the complete asset lifecycle from determining which assets to procure, through use and maintenance, and finally to decommissioning and secure disposal. While this reference implementation provides a useful roadmap for managing a large population of assets, organizations that have smaller numbers in their inventory may find the proposed framework itself difficult to deploy and manage. Some individual aspects of the framework, such as the use of "Splunk" [6] for analyzing log messages and sensor data, may be used independently or as part of other asset management tools more suited for the organization.

2.2 Risk Assessment

In "Critical Infrastructure Risk Assessment", a comprehensive strategy for finding and documenting vulnerabilities in manufacturing facilities is outlined [11]. The author takes the reader through an example consultation with a critical infrastructure facility and points out the numerous areas in which such a facility may be vulnerable to bad actors. Through a methodical process of conducting interviews, reviewing documentation, and making observations, the book describes how to analyze findings and assess risks found in the facility. Finally, Hayden provides an example of a completed risk assessment report that would be provided to responsible personnel at the facility. The framework and methodologies presented in this book were used in the preparation of the risk assessment report for this project.

2.3 Gap Analysis

Another major aspect of this project is to drive compliance with NIST SP 800-171, which aims to make United States Government contractors more resilient to cybersecurity attacks. To help determine which aspects of cybersecurity this project will focus upon requires an examination of current state against this standard; a gap analysis was prepared that compared existing policies and procedures against the requirements of NIST SP 800-171. To assist organizations in complying with SP 800-171, NIST provides a checklist [23] consisting of numerous questions; answers to the questions provide a clear picture as to the state of compliance. Completion of the questionnaire formed a significant portion of the gap analysis conducted for this project.

2.4 Policies and Procedures

In "Industrial Cybersecurity – Efficiently Secure Critical Infrastructure Systems", a multi-layer model ("defense in depth") for securing industrial control systems is proposed [7]. The practice of employing defense in depth measures is recommended by many in both academic and industrial environments [8, 14]. The defense in depth model relies on multiple layers of security measures, often overlapping, that make it substantially more difficult for bad actors to penetrate through to critical systems. Recommendations presented by this textbook for the securing of industrial control systems devices, such as the use of network segmentation and demilitarized zones ("DMZ"), were adopted by this project.

2.5 Automation

"Ansible" is a popular software tool used by network and system administrators to automate tasks such as system configuration, operating system hardening, and application of security controls [4]. Using tools such as Ansible, baseline configurations can be implemented, and auditing of security controls can be performed. In fact,

some service providers provide subscription-based tools built using Ansible to implement consistent baseline configuration levels for a variety of operating systems and software service platforms. This project implements many security controls using Ansible to ensure consistency across the population of assets. The use of Ansible by this project is limited by comparison; for example, a commercial offering of Ansible provides an automation controller ("Ansible Controller"), which claims to provide a simplified interface for managing large-scale deployments within an enterprise environment [5]. While Ansible was adopted for use by this project, Ansible Controller, however, was not implemented.

2.6 Top-Down vs. Bottom-Up Approaches

In many organizations, it is quite common for direction to always be provided top-down. And while this may make sense in many respects, it is not without its drawbacks [9]. In particular, there are processes and experiences that those in non-management positions may be intimately more familiar with than upper management. If these are not taken into account, then it is possible both inefficiencies and security issues will persist. Part of what is needed in any organization is a culture that not only allows feedback from those not in upper management, but actively encourages it. It is from this bottom-up approach that often provides a more robust risk management perspective. Additionally, it provides for a more collaborative environment from multiple levels of employees, which may encourage greater buy-in when there are specific mandates developed that must be followed.

2.7 Systems Thinking Approaches

Finally, one critical aspect of security is that it is rare for a single entity to operate in isolation from the whole. More commonly, security is often best understood when the context in which it exists is fully understood and delineated from other aspects of the environment. Systems thinking approaches are important in general given the explanatory power it provides [15], but also instrumental in helping security researchers and practitioners better understand the implications a change in one aspect of a larger system may have on other aspects of the same system. Beyond that, it also helps make systems more usable and effective [17]. Thus, while a full systems thinking approach was not the goal of this project, it was nonetheless something that we kept at the forefront of our efforts.

3 METHODS

3.1 Introduction

The primary goal of improving the security posture of this group (FIS) through the establishment of an information security program was itself broken down into primary and secondary tasks. The primary task was to mitigate high risk vulnerabilities after performing a risk assessment for the assets owned and maintained by this group. The secondary task was to establish policies and procedures for this employer to achieve compliance with NIST SP 800-171 ("Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations") [21].

The process this project followed was iterative in nature, and followed recommendations given by Ackerman [7] and Hayden [11]. First, hardware and software assets managed by this group were cataloged into two separate inventories, which included details such as links to source code repositories, defect trackers, vulnerability database queries, and network addresses. Secondly, a Risk Assessment was performed that documented risks to the information systems and components identified in the inventories, with emphasis placed on risks unique to the nature of industrial control systems and equipment. Next, a Gap Analysis was performed that highlighted nonconformances and provided basic guidance on how to achieve compliance.

For each item listed in the Risk Assessment and Gap Analysis, one or more policies were developed to provide guidance on achieving compliance in that area and associated security family. Existing corporate policies, NIST recommendations [3], and freely available security policy templates [2] were used as references for the creation of policies for this project. Ansible playbooks were developed to implement policies as baseline system configurations and to audit compliance by checking software package installation status, configuration parameters, and service execution status (such as firewalls).

3.2 Requirements

NIST SP 800-171 is a series of guidelines published by the United States National Institute of Standards and Technologies [3]. NIST is one of the primary agencies tasked by the United States Federal Government to establish standards for the implementation of information security ("cybersecurity") policies, programs, and methods by government agencies and contractors. As part of a company that is a U.S. Government contractor, this group (FIS) is responsible for implementing the requirements of SP 800-171.

The company and its parent organization have established policies for information security, and this group is also responsible for compliance to those policies in addition to those specified in SP 800-171. This project included requirements from both SP 800-171 and corporate policies in in addition to vulnerabilities identified as part of the Risk Assessment. This implementation phase of this project consisted of the development of a series of policies, procedures, and IT automation scripts. These three components were then used to mitigate vulnerabilities identified in the risk assessment and to work towards compliance with NIST SP 800-171.

3.3 Asset Inventories

After establishment of requirements, the next major task was to identify which "assets" were to be covered by the new information security program. Assets in this project comprised physical hardware devices and software applications owned and managed by this group; hardware and software assets already covered by corporate information security and managed by our information technologies group were not considered along with any devices without a network connection.

A hardware inventory was taken that included network-connected traditional compute devices (such as servers and workstations), virtual machines, and industrial control devices.

For the software inventory, applications covered by this project were cataloged and categorized as follows:

- Custom applications developed by FIS personnel ("Custom")

- Custom applications not developed by FIS personnel, but transferred to FIS for ongoing maintenance and support ("Assigned")
- Acquired software applications, commercial and open source ("COTS")

3.4 Risk Assessment

Upon completing the inventories of assets covered by the new information security program, it was then possible to determine risk exposure for each asset. It was necessary to inspect each asset's physical and network presence and its operational environment. Guidance from Ernie Hayden's book, "Critical Infrastructure Risk Assessment" [11], was invaluable for performing this phase of the project in part for its detailed methodology for performing a risk assessment and for providing guidance related specifically to industrial control systems.

For each asset, observations were made considering the following topics:

- Applicable policies and procedures
- Credential management
- Environmental conditions
- Hardware interfaces
- Immediate and surrounding physical environment
- Network interfaces and traffic flows
- Physical security
- Power sources
- Previously reported alerts and vulnerability notices for installed software
- Support available from the vendor
- User management, including identity and access controls

When asset observations were being made, considerations of the following methods, tools, and information sources helped to build a complete picture of the environment the asset operates within:

- Documentation review
- Examinations of applicable policies and procedures
- Inspection of vulnerability scanner reports, port scans, and other networking tools
- Interviews with key support personnel
- On-site observations
- Source code and configuration file inspections

While asset observations were being made, threats to the asset from many sources were considered. The risk assessment for this project considered the following threat sources:

- Cyber attacks o Denial of Service (DoS) attacks o Man-in-the-middle (MITM) attacks or eavesdroppers o Privilege escalation o Packet replay and modifications
- Human-caused events o Insider (malicious and non-malicious intent) o Trusted insider (malicious and non-malicious intent) o Outsider (malicious intent)
- Natural hazards o Fire, flood o Power disruption
- Organizational o Internal, local o Internal, remote o External, business competitor

Once the list of vulnerabilities was compiled, it was necessary to determine the actual risk posed to the organization should the vulnerability be exploited. There are several methods available to calculate severity; this project adopted a "qualitative" approach, where each vulnerability was assigned a "likeliness of occurrence" and "impact" [13]. As an operational technology group, with availability being the next highly valued criteria after safety, it was determined to measure risk in terms of the impact on production should a vulnerability be exploited.

Vulnerabilities were also assigned a likeliness of occurrence, which represents a "best guess" of the probability that the vulnerability might be exploited. Consistent with the use of a qualitative rather than quantitative approach, a rating system was devised based on an assumption of a vulnerability being exploited by a disgruntled employee or accident (the most likely scenarios).

Once an impact rating and likeliness of occurrence ratings were assigned, the vulnerabilities were assigned a final "risk rating" based upon a risk matrix. To obtain the final risk rating, the convergence of likeliness of occurrence (rows) and impact (columns) determined the rating. This risk rating represented the priority for application of mitigations.

3.5 NIST SP 800-171 Gap Analysis

As a secondary goal, this project attempted to meet requirements of the NIST SP 800-171 ("Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations") standard [21], which is applicable to the author's particular manufacturing organization. To begin the effort to comply with the standard, a gap analysis report was prepared immediately after completion of the risk assessment. The gap analysis report highlighted several areas of non-compliance that were not identified in the risk assessment; these items were addressed during the mitigation phase of the project.

3.6 Vulnerability Scanner

Shortly after the hardware inventory was compiled, an arrangement was made with the corporate information security group to include these assets in a weekly network-based vulnerability scan. This network-based vulnerability scanner, "Nexpose", had been performing weekly scans of all corporate devices but this group was not receiving reports until after this list was submitted.

The Nexpose vulnerability scanner works by executing a series of scripts or procedures against a network host and recording the response to various probing attempts [39]. Responses are checked to see if the host responds in a manner that indicates it is vulnerable to a particular attack. Nexpose contains a large database of vulnerabilities and is updated frequently to include signatures for newly discovered attack patterns.

Once Nexpose detects a vulnerability, it assigns a score based on metrics defined by our corporate information security group. The corporate target for assets is a score less than 5,000, although naturally it is desirable to achieve a score as close to zero as possible; this indicates a very low level of risk to the organization by the host. The goal is to have as low of a score as possible, both total score for all assets combined and the average score per host.

3.7 Policies and Procedures

As part of this project, policies and procedures were developed to help mitigate vulnerabilities exposed by the risk assessment, gap

analysis, and observations from Nexpose. Policies were developed to operate in conjunction with already established policies provided by corporate information security; policies developed during this project cover security aspects unique to industrial control systems and manufacturing operations in general.

Policy documents established guidelines that help achieve compliance by standardizing the practices this group uses to configure servers, develops applications, integrate software components, and performs many other similar tasks that involve network-connected devices and applications. Policies are organized by "security families", which are defined by NIST SP 800-53a [22] and SP 800-171 [21].

Procedures were written as necessary to implement these policies. For example, one policy ("FISSEC-SC-001") required specific encryption algorithms and key lengths to be used for certificate generation; a procedure was written that described the steps necessary to create certificates using an approved algorithm and key length. Procedures that were not practical to implement using automation scripts were developed as documents and checklists that system administrators followed to implement security controls.

3.8 Task Creation and Progress Tracking

With the risk analysis, gap analysis, and policy creation tasks complete, remaining work was primarily focused on implementing mitigations for the vulnerabilities discovered. To keep track of the work required, "issues" were created in a tracking system used by this group for task management. These issues were assigned to projects and personnel that would be responsible for implementing the mitigation.

3.9 Automation

"Ansible" is a software tool commonly used by network and system administrators to perform configuration and auditing tasks against network-connected devices and software [4]. Ansible executes scripts, known as "playbooks", to perform tasks in a consistent manner across a population of devices. As part of this project, Ansible playbooks were created to implement policies and to perform audits against security controls to verify compliance. In addition, Ansible playbooks were also created to assist with the deployment of mitigating controls (such as configuration changes and firewall rules), which greatly reduced the time required to implement controls among the population of hardware assets.

3.10 Metrics

During this project, the following metrics were established and continuously monitored to gauge progress towards compliance goals:

- Risk Assessment Mitigation Progress: The primary measure of success for this project was the number of vulnerabilities fully mitigated or in-progress towards completion.
- NIST SP 800-171 Gap Analysis Mitigation Progress: A secondary measure of project success was the number of recommendations implemented or deficiencies mitigated as determined by the Gap Analysis Report.
- Network-Based Vulnerability Scans: Weekly reports from the Nexpose vulnerability scanner were monitored during this project, with total risk and average risk per host scores plotted on a chart to observe progress made in implementing mitigations.

4 CONCLUSION AND FUTURE WORK

4.1 Conclusion

In this project, we have implemented an operational technology information security program from the "bottom-up" for the purposes of securing, without disruption, the critical infrastructure of a manufacturing operations group. This program, acting as an extension of an existing business-oriented information security program, focused on the personnel, systems, equipment, and data unique to the manufacturing organization and aimed to complement existing corporate security controls. By focusing on high priority vulnerabilities identified by the risk assessment conducted as part of this project, this team was able to focus on mitigating vulnerabilities even as the program was in development.

The primary goal of this project was to improve the security posture of this group, which was primarily accomplished by identifying, prioritizing, and mitigating vulnerabilities exposed by the risk assessment. The secondary goal of the project was to develop a "roadmap" that other similar organizations could use to start their own information security program, which was accomplished through the development of policies, procedures, and Ansible playbooks.

Work performed during this project has increased the security posture of this group. Through the application of policies, procedures, and Ansible playbooks, we have been able to fully mitigate all critical risks, 9% of high-level risks, 14% of medium level risks, and 26% of low-level risks. A remaining 16% of risks (all levels) have in-progress mitigation work ongoing.

Significant progress with respect to NIST SP 800-171 compliance was made during this project. Initially, performing the gap analysis indicated that only 9% of 363 checklist items were fully in compliance with requirements while 29% were partially compliant. At the time this report was created, however, 27% of checklist items were fully compliant and 40% partially compliant.

While developing this information security program, a corporate network-based vulnerability scanner gave measures of progress in the form of "scores", with lower scores representing fewer host vulnerabilities. At the beginning of this project the average score (risk) was 3,597 per host, and upon conclusion was 2,753. This represents a decrease of 24%, and always remained below the corporate benchmark average score of 5,000 per host. Total risk exposure of all FIS assets, as measured by Nexpose, dropped from 150,351 to 66,076, representing a decrease in risk to the organization by 56%.

4.2 Reflection

Considering the number of vulnerabilities found during the risk assessment phase of the project, it probably would have been sufficient to focus on those issues first and then perform the NIST SP 800-171 gap analysis after critical, high, and medium-level vulnerabilities were mitigated. Delaying work on mitigating known vulnerabilities exposed the organization to risks that could have been mitigated sooner. Going forwards, the gap analysis phase will

be conducted after vulnerabilities identified in the risk assessment have been significantly mitigated.

Another challenge was coordinating the different groups that would need to provide resources to mitigate vulnerabilities discovered during the risk assessment. In the future, it will be necessary to involve affected groups earlier in the project for the purposes of application of security controls, insights into technical solutions, and for general coordination of activities to minimize disruptions to operations.

A valuable insight was in the management and prioritization of vulnerabilities; prior to starting this project, many risks that this group felt were of high priority to mitigate were overshadowed by other risks that were found to be more critical after using a consistent formula for calculating risk exposure.

Through a deeper investigation of the risks posed by our use of legacy industrial control systems and devices, this group has come to appreciate the necessity for adding layers of additional controls to protect sensitive and critical operational technology equipment. Although this project did not have the opportunity to deploy security controls such as network segmentation or industrial control DMZs, the resulting discussions with our networking group has triggered the beginnings of a formal project to create these layers. It is recommended that any operational technology groups following the template outlined by this project should begin such work immediately after performing the hardware inventory and identifying affected ICS devices.

Finally, it became clear that it is not always necessary to use sophisticated or expensive tools to fully mitigate risks; for example, after evaluating the risk posed by remote administration facilities on obsolete servers, it was determined that it would be far more effective to simply replace the old servers rather than spend time and money on new security tools, additional network segments, and more detailed monitoring strategies.

4.3 Future Work

Moving forward, this program will expand to cover other groups within manufacturing operations and the policies, procedures, and playbooks developed by this project will serve as a model for expansion to other geographic regions. Yearly assessments will uncover new risks, and existing policies and procedures will be augmented as necessary to mitigate these vulnerabilities. Work to mitigate the remaining vulnerabilities is already in progress, and the group expects to have a majority of risks mitigated before the next risk assessment is performed in 2023.

REFERENCES

[1] 2016. *Deploying Industrial Firewalls within a Converged Plantwide Ethernet Architecture.* 113 pages. https://literature.rockwellautomation.com/idc/groups/literature/documents/td/enet-td002_-en-p.pdf

[2] 2020. *NIST Cybersecurity Framework Policy Template Guide.* 13 pages. https://www.cisecurity.org/wp-content/uploads/2020/07/NIST-CSF-Policy-Template-Guide-2020-0720-1.pdf

[3] 2023. Current Publications, SP 800 Series. https://csrc.nist.gov/publications/sp800

[4] 2023. How it works. https://www.ansible.com/overview/how-ansible-works

[5] 2023. Red Hat Ansible - Automation controller. https://www.ansible.com/products/controller

[6] 2023. Splunk | The Key to Enterprise Resilience. https://www.splunk.com/

[7] Pascal Ackerman. 2017. *Industrial Cybersecurity: Efficiently secure critical infrastructure systems.* Packt Publishing Ltd.

[8] Paul Didier, Fernando Macias, James Harstad, Rick Antholine, Scott A. Johnston, Sabina Piyevsky, Mark Schillace, Gregory Wilcox, Dan Zaniewski, and Steve Zuponcic. 2011. Converged plantwide ethernet (cpwe) design and implementation guide. *Cisco Systems and Rockwell Automation* (2011).

[9] Marc Dupuis, Barbara Endicott-Popovsky, Hank Wang, Ilanko Subramaniam, and Yuejin Du. 2011. Top-Down Mandates and the Need for Organizational Governance, Risk Management, and Compliance in China: A Discussion. *China-USA Business Review* (2011), 319.

[10] Marc Dupuis and Collin Gordon. 2018. Evaluating Prevalence, Perceptions, and Effectiveness of Cyber Security and Privacy Education, Training, and Awareness Programs. In *The Colloquium for Information Systems Security Education.* New Orleans, Louisiana, USA.

[11] MIPM Ernie Hayden and CEH CISSP. 2020. *Critical Infrastructure Risk Assessment: The Definitive Threat Identification and Threat Reduction Handbook.* Rothstein Publishing.

[12] Alexander Giehl and Sven Plaga. 2018. Implementing a performant security control for industrial ethernet. In *2018 International Conference on Signal Processing and Information Security (ICSPIS).* IEEE, 1–4.

[13] David Kim and Michael G. Solomon. 2016. *Fundamentals of Information Systems Security: Print Bundle.* Jones & Bartlett Learning.

[14] Irina Mashkina and Ildar Garipov. 2018. Threats modeling and quantitative risk analysis in industrial control systems. In *2018 International Russian Automation Conference (RusAutoCon).* IEEE, 1–5.

[15] George E. Mobus and Michael C. Kalton. 2015. *Principles of systems science.* Vol. 519. Springer.

[16] Rohit Negi, Aneet Dutta, Anand Handa, Ujjwal Ayyangar, and Sandeep K. Shukla. 2020. Intrusion Detection & Prevention in Programmable Logic Controllers: A Model-driven Approach. In *2020 IEEE Conference on Industrial Cyberphysical Systems (ICPS)*, Vol. 1. IEEE, 215–222.

[17] Jessica Nguyen and Marc Dupuis. 2019. Closing the Feedback Loop Between UX Design, Software Development, Security Engineering, and Operations. In *Proceedings of The 20th Annual Conference on Information Technology Education (SIGITE '19).* ACM, Tacoma, WA, USA, 93–98. https://doi.org/10.1145/3349266.3351420

[18] Karl-Heinz Niemann. 2019. IT security extensions for PROFINET. In *2019 IEEE 17th International Conference on Industrial Informatics (INDIN)*, Vol. 1. IEEE, 407–412.

[19] Karen Renaud and Marc Dupuis. 2019. Cyber Security Fear Appeals: Unexpectedly Complicated. In *New Security Paradigms Workshop.* San Carlos, Costa Rica, 42–56. https://doi.org/10.1145/3368860.3368864

[20] Karen Renaud, Rosalind Searle, and Marc Dupuis. 2021. Shame in Cyber Security: Effective Behavior Modification Tool or Counterproductive Foil?. In *New Security Paradigms Workshop.* ACM, Virtual Event USA, 70–87. https://doi.org/10.1145/3498891.3498896

[21] Ron Ross, Victoria Pillitteri, Kelley Dempsey, Mark Riddle, and Gary Guissanie. 2019. *Protecting controlled unclassified information in nonfederal systems and organizations.*

[22] Ronald S. Ross. 2014. Assessing security and privacy controls in federal information systems and organizations: building effective assessment plans. (2014).

[23] Patricia Toth. 2017. Self-Assessment Handbook For Assessing NIST SP 800-171 Security Requirements in Response to DFARS Cybersecurity Requirements. (2017).