PhD school: KIDS, Kestrel Based Intrusion Detection System for Industrial Control Systems

Nowshaba Jeelani Wani
PhD (Second Year), School of Computer Science and IT, University College Cork
Cork, Ireland
nowshabawani@umail.ucc.ie

Abstract

Security of industrial systems is hugely impacted by the convergence of Information Technology (IT) and Operational Technology (OT). While the focus has been largely on improved security of IT systems, less effort has been spent on securing the OT side of industrial processes, in particular Industrial Control Systems (ICS). This integration of IT and OT has introduced a significant gap between Intrusion Detection Systems (IDS) used for IT and those required for OT. The majority of tools to improve OT security ignore its process based nature and only concentrate on management components, which are essentially IT systems. Specifically, IDS for IT are applied to OT. This approach has serious limits as an attack on the industrial process is often invisible. The aim of this PhD research is to close this gap by developing a unified approach to IDS that addresses the specific needs and constraints of OT while also leveraging the strengths of IT based security mechanisms. The current goal is to design an IDS called "Kestrel Intrusion Detection System (KIDS)", which is a threat hunting framework based on Kestrel that aims at bringing IT and OT security closer together to improve IDS for industrial environments. The IDS proposed represents a foundational query-based design that functions with the ICS components irrespective of vendor or implementation complexities. The flexibility also allows its quick adaptation to different manufacturing processes. Once developed and tested, KIDS will be compared to state-of-art IDS used in OT and improved gradually. This work will contribute to both theoretical and practical understanding of threat detection in converged IT, OT infrastructures.

CCS Concepts

• **Security and privacy** → *Intrusion detection systems.*

Keywords

Industrial control systems, Intrusion detection systems, Operational technologies, Threat hunting, Security

1 Introduction

Cyber attacks are ever increasing in volume, complexity, and variety of attack methods. Specifically, Industrial Control Systems (ICS) as part of the Operational Technology (OT) side of industrial infrastructures is increasingly the focus of attackers. The Stuxnet malware [3], the Ukrainian power plant attack, and other similar attacks have shown the vulnerability of ICS to cyber attacks [2]. Such attacks often lead to a loss of trust, reputation, physical damage, and organisational security crises, raising concerns regarding the safety and security of ICS. The drive to closely integrate IT and OT within industry, i.e. the merger of enterprise information networks and automation system networks, amplifies this threat. This integration leads to exposing OT devices, from sensor endpoints to Programmable Logic Controllers (PLCs), as part of an Industrial Internet of Things (IIoT), to the Internet [5, 7]. Due to Internet connectivity, attackers have a potential path to access and modify the control of physical processes. However, most ICS are ill equipped to defend against attacks as they were originally created to operate as standalone systems without any interaction with the outside world as Internet connectivity is a recent addition to the ICS world.

IDS play a critical role in cybersecurity by monitoring and analyzing network traffic for suspicious activity or policy violations. Upon detecting malicious activity, IDS are designed to take countermeasures such as alerting the security team or taking predefined actions to protect against the detected intrusion [8], [4]. IDS in industrial environments are usually provided by teams in charge of IT, leading to IDS treating OT components as if they were IT elements. IDS are used to monitor communication with the management interface of ICSs and related components such as Human Machine Interface (HMI) or Historians. However, the details of the industrial processes are not exposed to the IDS. From an organisational perspective, control engineers designing process automation systems typically have no input to the design of an IDS used to protect OT. However, there are usually key parameters in an industrial system that can provide insight into operations. For example, the filling rate of a tank in an industrial process may be a key indicator. Such information can be combined with information from the IT side to construct an improved IDS with a holistic view of a production environment. For example, a login to an IT system at an unusual time might not be considered suspicious. Similarly, a significant increase in tank filling rates may also be part of normal operation. However, if both events coincide an alarm may be raised to investigate the situation.

In this work we introduce KIDS, a Kestrel based IDS that aims to bring together IT and OT to improve IDS in industrial environments. Kestrel is an open-source language for threat hunting [6] that enables enhanced behavioral analysis, granular visibility, real-time querying and zero-day attack detection, allowing for custom rule creation and dynamic threat identification. It also supports cross-layer and cross-system analysis, customizable responses, anomaly detection, and monitoring of unusual patterns. We leverage Kestrel to formulate so called *hunting scripts* to include OT components such as Programmable Logic Controllers (PLCs) and analyse heterogeneous data across IT and OT. Once developed, a hunting script can execute periodically in an automated fashion to construct an IDS. We decided to base our work on Kestrel as it already provides a framework for the IT domain and offers a high degree of flexibility. We believe this flexibility is key when integrating the OT domain

as industrial processes vary and adaptation for specific processes is necessary.

2 Kestrel Based Intrusion Detection System

We propose KIDS, a Kestrel query based IDS for OT, implementing a multi-layered approach that provides a comprehensive, cost-effective, and flexible intrusion detection solution. It facilitates intrusion detection at all ICS levels (0 to 5 in the Purdue Model) by abstracting ICS specifics (e.g. ICS vendor specifics, protocols and automation process details) [9]. To integrate OT into Kestrel, software modules need to be embedded within ICS components such as PLCs, Historian, HMI, and Field Programming Devices. These modules are used to expose data points to Kestrel. KIDS modules can be created from standard Kestrel components by adding ICS specifics. Core to integrating ICS in Kestrel is the addition of a PLC module, which we discuss in the next paragraphs. We leave the integration of other ICS components for future work.

2.1 PLC module for Kestrel

A Kestrel module for a PLC should provide security relevant information on three layers of abstraction: *OS Layer*, *PLC Layer*, *Process Layer*.

- *OS Layer:* On this level information related to the operating system of the PLC module is made available. This includes information such as past logins, current connections to the PLC, system uptime or error codes.
- *PLC Layer:* On this level, information is made available to Kestrel that is generic to specific PLCs and do not require understanding of the specific industrial process. Kestrel is able to query process independent information, such as general PLC configurations, the number of sub-processes installed, specific operating conditions such as state of sub-processes, memory mapping and access.
- *Process Layer:* On this level critical information related to the specific industrial process (and sub-processes) are made available to Kestrel. For example, a process controlling pumps and valves on a tank exposes maximum filling levels or a flow rate as critical parameters. To expose critical information to Kestrel, the process engineer designing the process would need to perform an annotation. We describe this aspect in section 2.2.

A Kestrel module for a PLC may be implemented directly on the PLC management module. However, as vendors are reluctant to open their platform for third party modules a proxy approach may be used. The Kestrel PLC module can be installed on a dedicated system which queries the PLC using the usual vendor specific protocol. The general approach to developing KIDS, also involves configuring data profiles that allow Kestrel to connect, authenticate and query numerous data sources for information. It also involves writing a collection of hunting steps in form of a hunt book. These steps articulate the threat hypotheses and corresponding data queries and analytics to validate them. A hunt book in KIDS is capable of performing simple queries against the data source as well as looking for more complex patterns and behaviours that are characteristic of cyber threats in ICS environments. Additionally, KIDS provides extensibility through the creation of new plugins including custom data source interfaces to integrate with ICS specific tools and data formats. Once suitable queries are determined, they can be automated and added continuously to the IDS to detect any suspicious behaviours before an irreversible loss or damage is caused by an adversary.

2.2 Process Variable Annotation

With the PLC model, we also introduce process variable annotation. This feature involves labeling critical process variables, determined by control engineers during the initial definition phase of processes, such as within PLC logic programming. Thus, bridging the gap between automation and security efforts by simplifying the extraction of critical process information. This integration simplifies the extraction of vital process information for intrusion detection purposes, enabling focused reasoning. Consequently, it augments security analytics by granting direct access to crucial operational data, thereby improving the clarity, speed, and accuracy of threat detection within ICS.

3 Future Work

Future work focuses on developing a prototype of KIDS and evaluating its effectiveness in detecting intrusions that operate across IT and OT. These next steps are briefly outlined below:

(1) *Developing a Prototype*: After identifying unique characteristics and security challenges of ICS and defining specific objectives of integration for the framework, it has to be developed into a working prototype. This involves identifying data sources(logs, network traffic, PLC, Supervisory Control and Data Acquisition (SCADA) systems) and creating unified data pipelines where both IT and OT systems feed data into KIDS. Threat hunting workflows can be based on threat intelligence from resources like MITRE ATT&CK [1]. In order to achieve this, multiple additions and modifications to the Kestrel language are required for script execution and data normalization to ensure interoperability. The prototype will be used to create custom scripts to query systems across both environments. Once suitable scripts are identified, they can be automated to run periodic checks across IT and OT systems.

(2) *Implement the Prototype in a Testbed*: The next step involves setting up a test bed to validate the working of our developed prototype for KIDS. The test bed will ideally simulate a typical ICS environment which includes components like PLCs, HMI, SCADA, historians, servers, engineering workstations and so on. Attack scenarios will be simulated and specific scripts targeting potential attack vectors in OT will be customized to identify such attacks using KIDS. After a successful implementation, an initial evaluation will be done on the effectiveness of KIDS to interoperate between ITand OT and to identify various attack scenarios.

(3) *Implement the Prototype in an Industrial Setting*: In order to evaluate the effectiveness and feasibility of KIDS in a practical scenario, it will be implemented in an industrial setting. This step is crucial for identifying real world attack scenarios and collecting critical threat intelligence by monitoring the real time data. It will also help us assess the

detection accuracy, response time and ability to adapt to OT-specific constraints like low latency.

(4) *Evaluate and optimize the solution*: We will gather performance matrices and compare KIDS to other state of art solutions available. Based on the results, KIDS can be optimized and made scalable in order to meet industry requirements. KIDS will be continuously evaluated and optimized to achieve the set goals.

While Kestrel was chosen for this framework because of its open-source nature, ability to create custom, high-level threat-hunting workflows, as well as its support for complex queries, automation of repetitive tasks, and its compatibility with different frameworks, other security tools could also be used to achieve similar goals. The methodology outlined emphasizes a generalized approach to IT/OT integration and is not tied to a single tool.

4 Open Challenges

Research in the field of cybersecurity for OT presents numerous challenges that make it a complex and evolving domain. One of the most significant ones is the prevalence of legacy systems. These outdated technologies, often deeply embedded in critical infrastructure, were not designed with modern cybersecurity threats in mind, making it difficult to implement effective security measures without disrupting operations. In addition to these legacy systems, OT environments are governed by strict protocols and regulations, which often impose constraints on the flexibility and creativity of such solutions.

Another significant challenge is the unforeseen limitations that frequently arise during the research process. These may include compatibility issues, performance trade-offs, or security mechanisms that prove impractical for real-world implementation. This gap between theoretical research and practical application can be frustrating, especially when solutions that look promising in a controlled environment fail to meet the requirements of industrial-scale systems.

Moreover, OT cybersecurity demands expertise in both IT and OT, two domains that have historically operated in silos. The lack of cross-disciplinary knowledge can delay progress, as one often needs to become well-versed in the intricacies of both fields to fully understand the challenges and develop comprehensive solutions. This knowledge gap is further widened by the rapid pace at which both fields evolve, requiring constant learning and adaptation.

In addition to technical challenges, the process of networking with professionals developing similar solutions is crucial, yet difficult. Building these connections is essential to stay up to date with the latest advancements and exchange ideas. Research in OT cybersecurity is a niche field, and finding the right network of researchers and industry professionals to collaborate with or learn from can be a daunting and a time-consuming task.

Communicating the importance and complexity of one's research to a broader audience is another challenge. Explaining a specialized research topic to a general audience, whether it is colleagues from different fields, industry professionals, or the public, requires translating technical concepts into more relatable terms without losing the essence of the work. This is a vital skill, especially when seeking funding, support, or awareness, but many new researchers struggle with this including myself.

Also, PhD students often face the additional burden of managing expectations from different stakeholders, industry partners, and even themselves. It can be difficult to balance ambitious research goals with the practical realities of what can be achieved within the time frame and resources available. Setting realistic expectations, communicating progress effectively, and staying focused on the most impactful aspects of the research can be hard at times.

Balancing these technical and non-technical challenges with the demanding journey of a PhD adds another layer of complexity. The path to a PhD requires continuous refinement of research questions, experimentation, and evaluation, often in the face of obstacles that can slow progress. In this context, the importance of regular feedback from academic advisors and industry experts cannot be overstated. Their guidance helps refine ideas, address blind spots, and ensure that the research remains both relevant and feasible.

Acknowledgments

This publication has emanated from research conducted with the financial support of Science Foundation Ireland under Grant number 18/CRT/6222. For the purpose of Open Access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

References

[1] MITRE ATT&CK. 2015-2024. MITRE ATT&CK for ICS. https://attack.mitre.org/techniques/ics/

[2] Defense Use Case. 2016. Analysis of the cyber attack on the Ukrainian power grid. *Electricity Information Sharing and Analysis Center (E-ISAC)* 388, 1-29 (2016), 3.

[3] James P Farwell and Rafal Rohozinski. 2011. Stuxnet and the future of cyber war. *Survival* 53, 1 (2011), 23–40.

[4] Hung-Jen Liao, Chun-Hung Richard Lin, Ying-Chih Lin, and Kuang-Yuan Tung. 2013. Intrusion detection system: A comprehensive review. *Journal of Network and Computer Applications* 36, 1 (2013), 16–24.

[5] Md Noor-A-Rahim, Jobish John, Fadhil Firyaguna, Hafiz Husnain Raza Sherazi, Sergii Kushch, Aswathi Vijayan, Eoin O'Connell, Dirk Pesch, Brendan O'Flynn, William O'Brien, et al. 2022. Wireless communications for smart manufacturing and industrial IoT: Existing technologies, 5G and beyond. *Sensors* 23, 1 (2022), 73.

[6] IBM Research. 2021. The thrill of cyber threat hunting with Kestrel. https://research.ibm.com/blog/kestrel-cyber-threat-hunting. [Accessed: 05-Oct-2023].

[7] Emiliano Sisinni, Abusayeed Saifullah, Song Han, Ulf Jennehag, and Mikael Gidlund. 2018. Industrial internet of things: Challenges, opportunities, and directions. *IEEE transactions on industrial informatics* 14, 11 (2018), 4724–4734.

[8] Aurobindo Sundaram. 1996. An introduction to intrusion detection. *Crossroads* 2, 4 (1996), 3–7.

[9] Theodore J Williams. 1994. The Purdue enterprise reference architecture. *Computers in industry* 24, 2-3 (1994), 141–158.