

Poster: Towards Selecting Threat Appropriate Industrial Intrusion Detection Systems

Stefan Lenz
lenz@spice.rwth-aachen.de
RWTH Aachen University
Aachen, Germany

Johannes Weidmann
weidmann@spice.rwth-aachen.de
RWTH Aachen University
Aachen, Germany

Martin Henze*
henze@spice.rwth-aachen.de
RWTH Aachen University
Aachen, Germany

Abstract

As the threat landscape against industrial control systems is dynamic, effective security requires detection strategies capable of timely reactions to these evolving threat situations. To address this problem, we propose the idea of a counter-threat intelligence sharing based mechanism to select appropriate detectors for current circumstances. To highlight the potential of this mechanism, we conduct attack-level performance evaluations of various intrusion detection systems. Results show the variance of intrusion detection performance depending on the attack scenario, emphasizing the benefits of such a mechanism for industrial control system security.

CCS Concepts

• Security and privacy → Intrusion detection systems; • Networks → Cyber-physical networks.

Keywords

Security, industrial control systems, intrusion detection

ACM Reference Format:

Stefan Lenz, Johannes Weidmann, and Martin Henze. 2026. Poster: Towards Selecting Threat Appropriate Industrial Intrusion Detection Systems. In *Proceedings of the 2026 ACM SIGSAC Conference on Computer and Communications Security (CCS '26)*, November 15–19, 2026, The Hague, Netherlands. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3830454.3846437>

1 Motivation

Industrial control systems (ICSs) become evermore interconnected and present high-value targets for adversaries facing increasingly sophisticated attacks (e.g., Stuxnet). Since ICS are famously hard to secure, e.g., due to decade-long device-lifetimes [7], retro-fittable security measures such as intrusion detection became popular for ICS [17]. Consequently, a large body of research on novel, and well-performing detection approaches has been published.

As ICS exhibit predictable and deterministic behavior, attack detection for such systems is typically anomaly-based. Thus, detectors leverage features such as network timings (e.g., [11]) or the state of the physical process (e.g., [16]) to detect deviations from known, benign behavior. To test and develop novel detection mechanisms researchers typically rely on benchmarking datasets [3]

and use metrics such as recall or F1-Score [9]. However, as these metrics present average performance across the complete dataset and modern datasets typically contain diverse behavior and attack scenarios [3], such “dataset-level” metrics average-out performance in specific threat situations. Furthermore, as different detection approaches leverage different features, individual detectors might be suited to detect certain attacks better. For example, a timing-based detector could detect denial of service (DoS) attacks (e.g., flooding) quite well, while a detector using process states may not.

As effective defense strategies benefit from a holistic approach towards intrusion detection [10], related work attempts to solve this issue by combining multiple detectors with different strengths into so-called ensembles [5, 15]. However, these again try to improve dataset-level detection, not optimizing according to the currently observed threat landscape. In principle however, ICS operators could leverage cyber threat intelligence (CTI) sharing to make such threat specific adaptations. For example, if an ICS operator learns from other operators, that DoS attacks are currently prevalent, they could select the most suitable detector for this threat and thus improve their security.

Therefore, in this paper we present our idea of such CTI-sharing based mechanism for dynamic intrusion detection system (IDS) selection and report our initial results. To this end, we present the idea behind our selection mechanism (§2) including initial results showing the potential of the approach (§3). Then, we discuss potential challenges of realizing the selection mechanism (§4)

2 Threat Appropriate IDS Selection

Intrusion detection research and development currently focusses on dataset-level performance [3, 9] obfuscating scenario-specific detection performance. However, through sharing information about ongoing threats, organizations could optimize their detection by selecting the most suitable IDS for the current situation. To this end, we present our idea for such a CTI-sharing based selection mechanism for (industrial) intrusion detection systems. Broadly, this mechanism consists of three parts (shown by Fig. 1): (i) CTI analysis and sharing, (ii) an IDS performance database for individual attacker techniques, and (iii) IDS selection by cross-referencing CTI reports with the database. In the following, we motivate the potential benefits of such a mechanism in detail.

Step 1 - CTI Sharing: By increasing situational awareness, CTI sharing is widely considered to benefit (cyber-)security operations in IT systems and ICS [8]. To this end, organizations share knowledge of adversary tactics, technique, and procedures (TTPs) in standardized formats as provided by, e.g., MITRE ATT&CK® [13]. This knowledge about currently prevalent TTPs can also serve as the decision basis for an informed adaptation of defense strategies.

*Also with Fraunhofer FKIE.



This work is licensed under a Creative Commons Attribution 4.0 International License. *CCS '26, The Hague, Netherlands*

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2871-6/2026/11

<https://doi.org/10.1145/3830454.3846437>




Figure 1: Sharing CTI reports among ICS operators allows selecting appropriate IDS for the current threat landscape.

In our case, enabling ICS operators to select the “best” detector for the current situation. Therefore, the first step of our IDS selection mechanism consists of such a CTI sharing mechanism, where ICS operators share knowledge of detected TTPs with each other. Then, assuming that similar threats might target their system as well, the administrators can select a fitting detection strategy (Fig. 1-1).

Step 2 - IDS Performance Lookup: The second step of our selection mechanism comprises a “database look-up” telling the administrator which IDS works best for the current situation, i.e., currently observed TTPs (Fig. 1-2). To this end, such a database must be created, by mapping attack descriptions of available benchmarking datasets to TTPs *and* cross-reference performance evaluations of multiple detectors on these TTPs. Additionally, maintaining and growing such a database allows for increasingly informed decisions and better IDS selection.

Step 3 - IDS Selection: The last step of the mechanism contains the actual IDS selection (Fig. 1-3). After cross-referencing CTI reports with the IDS database, the administrator chooses and deploys the “most appropriate” detection mechanism. Moreover, this step also may include a feedback-loop, by confirming (or denying) an increase in detection of previously reported TTPs.

Although, such a mechanism does not exist (yet) in an end-to-end manner, research shows the benefits of CTI sharing [8], which can also benefit intrusion detection. To show how such an approach can be beneficial for intrusion detection, we evaluate TTP-level intrusion detection performance in the following.

3 Proof of Concept

To highlight the potential of a threat appropriate IDS selection mechanism, we show the variance in IDS detection performance across different TTPs. To this end, we map each attack of the WDT [4]




Figure 2: Labelling attacks of WDT [4] with TTPs (top row) shows how detectors alert (black boxes = true alerts - red boxes = false alerts) differently based on the current TTP. Especially, InterArrivalTime performs particularly well in DoS (salmon) scenarios, but struggles with other TTPs.

dataset to a MITRE ATT&CK® TTP for ICS [13] and evaluate each detector on these scenarios.

Dataset Selection & Mapping Procedure: As a basis for our experiments, we select the WDT dataset by Faramondi et al. [4], since it contains well-labelled attacks against ICS communication (e.g., flooding attacks) and against the physical process (e.g., physical manipulation or adversary-in-the-middle) enabling the comparison of *fundamentally different* detection mechanisms. Furthermore, to discuss detector performance in a generalizable manner (see §2), we map each attack to a TTP from the MITRE ATT&CK® Matrix for ICS [13] using official CISA mapping guidelines [2].

Comparison & IDS Selection: To conduct the comparison, we utilize the IPAL framework [17] and select four fundamentally different IDS approaches: two monitoring communication behavior and two monitoring the state of the physical process. (i) The InterArrivalTime approach by Lin et al. [11], leveraging consistent packet timings, (ii) Kitsune [12], monitoring various communication features, (iii) a comparatively simple detector utilizing support vector machines OneClassSVM [14], and (iv) the process state invariant-based detector GeCo [16]. We choose these particular detectors not only because they promise high detection performance, but also because they utilize different mechanisms *and* features of the ICS for anomaly detection.

Results: First, mapping the attacks from WDT [4] to TTPs reveals a non-trivial challenge: Mapping attack descriptions to TTPs does not result in clear 1-to-1 assignments. Instead, most attack descriptions fit to multiple TTPs each describing distinct aspects of the attack. For example, the description “Physical breakdown of pump X” fits to the TTPs Denial of Control and Loss of View (brown and indigo in Fig. 2). This known challenge for attack labelling [1], complicates *consistent* TTP assignment. Ultimately, we assign six TTPs to three attack scenarios (Fig. 2 - top row) covering both process-oriented TTPs (e.g., Manipulation of View), as well as TTPs against network components (e.g., Denial of Service).

To gain a high-level understanding of the behavior of each IDS across the different attack scenarios, we visualize the alerts of each detector (Fig. 2). Firstly, we observe that the communication-based detectors (i.e., InterArrivalTime [11] and Kitsune [12]) provide

Table 1: The best performing detector (in green) differs between attack scenarios and dataset-level evaluations do not fully capture IDS capabilities.

IDS	WDT [4]	TTP					
		AitM	DoC	DoS	LoV	MoV	UM-RM
InterArrivalTime [11]	0.20	0.38	0.69	0.87	0.12	0.64	0.71
Kitsune [12]	0.07	0.36	0.36	0.64	0.09	0.32	0.25
OneClassSVM [14]	0.37	0.61	0.32	0.32	0.32	0.53	0.61
GeCo [16]	0.62	0.61	0.77	0.55	0.77	0.71	0.62

All scores in eTaR [6] - higher is better. Acronyms: AitM=Adversary in the Middle; DoC=Denial of Control; DoS=Denial of Service; LoV=Loss of View; MoV=Manipulation of View; UM-RM=Unauthorized Message: Reporting Message.

persistent alerts in the two DoS scenarios (salmon in Fig. 2). However, the process-aware detectors (i.e., OneClassSVM and GeCo) do not provide reliable alerts for this TTP. Conversely, OneClassSVM and GeCo alert consistently for scenarios targeting the process state, e.g., Adversary in the Middle or Manipulation of View (gold and magenta in Fig. 2), while the communication-based IDS do not.

To solidify this understanding, we analyze the detection performance on the complete dataset and each individual TTP (Tb. 1). To this end, we use eTaR [6], a recall metric for time-series data, showing how well a detector recognizes each TTP. On the complete dataset GeCo is the best performing IDS with a score of 0.62. However, when considering individual TTPs, e.g., InterArrivalTime performs substantially better in DoS scenarios achieving an eTaR of 0.87 vs. 0.55 of GeCo. Similarly, although Kitsune achieves a score of 0.09 for Loss of View, its performance on DoS increases to 0.64 eTaR, a relative difference of more than 700%.

Overall, our results show the variance in the performance in intrusion detection based on the current attack scenario. Furthermore, we show that utilizing TTPs for IDS evaluation enables the “decoupling” of performance from datasets to specific scenarios. Ultimately, these results highlight the potential of a threat specific selection mechanism for more fine-grained intrusion detection.

4 Future Challenges

Realizing such a selection mechanism also presents technical and non-technical challenges. To consider these and avoid pitfalls down the road, we discuss potential problems in the following.

CTI Sharing Incentives and Privacy Concerns: First, we identify the challenge of incentivizing organizations to participate in CTI sharing and not adopting the mechanism out of privacy concerns. As any CTI-sharing approach, such problems cannot be avoided and must be addressed by providing trust of participants and guaranteeing no negative impact from sharing CTI reports [8]. For the ICS sector, comprising a small amount of participants that are sometimes government controlled, utilizing government agencies as a trusted third party could alleviate these problems.

Consistent TTP Labels: As a technical challenge w.r.t. to creating, maintaining, and expanding the IDS database, we identify potential problems in inconsistent and unclear labelling of attacks to TTPs. Such inconsistencies could cause a warped picture of IDS capabilities in the database leading to wrong decisions in the IDS selection phase. However, these challenges can be minimized by utilizing automated mapping tools based on, e.g., large-language models, which show promising results for this task [1].

Timely IDS Selection and Switching: At last, we identify the problem of determining when to “switch” to another IDS for the current situation. For example, switching as soon as new CTI reports come in might not be the best approach in all situations, since previous threats might still be relevant. Unfortunately, simply using *all* available detectors in parallel also creates uncertainty on when to report individual IDS alerts to operators [17]. Similarly, completely switching to an IDS specialized on single TTPs might also hinder detection of other TTPs and thus potentially reducing overall detection rates. To address such a problem, using IDS selected based on current TTPs for ensemble deployments could reduce these problems. For example, by giving IDS for current TTPs more weight in a majority vote while still considering others to keep TTP-coverage high.

5 Conclusion

As threat landscapes against ICS are rapidly evolving and intrusion detectors leverage various features for anomaly detection, their performance across attack scenarios differs, too. Leveraging this fact for better detection rates and improved ICS security, we present the idea of a CTI-sharing based IDS selection mechanism. This mechanism cross-references the current threat landscape with an attack-scenario specific performance database and selects appropriate detectors. Initial experiments show the variance of detection mechanisms between different attacks, highlighting the potential of precisely selecting a well-suited detector for the current threat landscape. Ultimately, this mechanism aims to not only improve attack detection in ICS, but also simplify and demystify detection capabilities of IDS for practical application.

References

- [1] Marvin Büchel et al. 2025. SoK: Automated TTP Extraction from CTI Reports – Are We There Yet?. In *USENIX '25*.
- [2] CISA. 2023. *Best Practices for MITRE ATT&CK Mapping*. Technical Report. <https://www.cisa.gov/sites/default/files/2023-01/Best%20Practices%20for%20MITRE%20ATT&CK%20Mapping.pdf> accessed July 2026.
- [3] Mauro Conti et al. 2021. A Survey on Industrial Control System Testbeds and Datasets for Security Research. *IEEE Communications Surveys & Tutorials* 23, 4.
- [4] Luca Faramondi et al. 2021. A hardware-in-the-loop water distribution testbed dataset for cyber-physical security testing. *IEEE Access* 9.
- [5] Jun Gao et al. 2021. Omni SCADA Intrusion Detection Using Deep Learning Algorithms. *IEEE Internet of Things Journal* 8, 2.
- [6] Won-Seok Hwang et al. 2022. “Do you know existing accuracy metrics overrate time-series anomaly detections?”. In *SAC '22*.
- [7] Eric D. Knapp. 2024. Industrial Cybersecurity History and Trends. In *Industrial Network Security*.
- [8] Rubin Krief et al. 2026. An evidence-driven analysis of threat information sharing challenges for industrial control systems. *Computers & Security* 169.
- [9] Olav Lamberts et al. 2023. SoK: Evaluations in Industrial Intrusion Detection Research. *Journal of Systems Research* 3, 1.
- [10] Stefan Lenz et al. 2026. On the Challenges of Holistic Intrusion Detection in ICS. In *SPRING Graduate Workshop '26*.
- [11] Chih-Yuan Lin et al. 2018. Timing-Based Anomaly Detection in SCADA Networks. In *CRITIS '21*.
- [12] Yisroel Mirsky et al. 2018. Kitsune: An Ensemble of Autoencoders for Online Network Intrusion Detection. In *NDSS '18*.
- [13] MITRE Org. 2026. MITRE ICS Matrix. online. <https://attack.mitre.org/versions/v19/matrices/ics/> accessed July 2026.
- [14] Takashi and Kiuchi. 2012. Analysis of Intrusion Detection in Control System Communication Based on Outlier Detection with One-Class Classifiers. In *ICONIP*.
- [15] Konrad Wolsing et al. 2023. One IDS is not Enough! Exploring Ensemble Learning for Industrial Intrusion Detection. In *ESORICS '23*.
- [16] Konrad Wolsing et al. 2025. GeCos Replacing Experts: Generalizable and Comprehensible Industrial Intrusion Detection. In *USENIX '25*.
- [17] Konrad Wolsing et al. 2022. IPAL: Breaking up Silos of Protocol-dependent and Domain-specific Industrial Intrusion Detection Systems. In *RAID '22*.