

CYBER RESILIENCE ASSESSMENT OF INDUSTRIAL SYSTEMS USING MITRE CALDERA ADVERSARY EMULATION AND SIEM DETECTION: A PRELIMINARY STUDY

C. M. CIOBANU, National University of Science and Technology POLITEHNIC, Bucharest, ROMANIA, ciobanucorinamaria@gmail.com

O. R. CHIVU, National University of Science and Technology POLITEHNIC, Bucharest, ROMANIA

M. HELSTERN, National University of Science and Technology POLITEHNIC, Bucharest, ROMANIA

Abstract: *Industrial organizations often rely on default security configurations, assuming they provide a reliable safety net. This paper presents a preliminary empirical study evaluating the detection of Microsoft Sentinel against six Defense Evasion techniques from the MITRE ATT&CK framework, simulated in a controlled laboratory environment using MITRE CALDERA. The results were definitive: without manual intervention, the system remained silent. Detection was only achieved through the development of custom analytics rules in Kusto Query Language (KQL). These findings highlight a critical gap between expected and actual security, demonstrating that active detection testing is essential for meeting the high security standards of the ISA/IEC 62443 framework.*

Key words: cyber resilience, adversary emulation, defense evasion, SIEM detection, MITRE CALDERA

1. INTRODUCTION

Reliability engineering has long provided industrial organizations solid tools for predicting and preventing system failures. These frameworks share a common assumption: that failures happen by accident.

The increasing presence of cyber threats in industrial environments challenges this view. Unlike mechanical faults, cyberattacks are intentional, carefully planned, and often designed to go unnoticed for long periods of time, bringing a completely new type of risk that traditional reliability models were never built to handle [1].

For decades, the resilience of industrial systems was measured through well-known metrics such as Mean Time Between Failures (MTBF) or Failure Mode and Effects Analysis (FMEA). These methods work well for physical degradation, but they were not designed with an adversary in mind. Historically, industrial control systems (ICS) were considered safe precisely because they were installed in isolated or air-gapped networks. As control system architectures are increasingly connected to traditional enterprise IT networks, the risk that ICS can be compromised by the same type of attacks targeting conventional IT environments has grown substantially. [2]

This shift has been carefully studied and documented. The MITRE Corporation developed the ATT&CK for ICS knowledge base specifically to capture the tactics, techniques, and procedures used by attackers targeting industrial environments. [3] The framework gives security teams a structured way to understand how adversaries behave inside industrial

networks and to build defenses that reflect real attacks patterns rather than theoretical assumptions.

Having a frame is not the same as having reliable detection. Penetration testing and adversary emulation are still relatively new practices in industrial setting, and the question of whether existing security monitoring tools can actually detect ICS-relevant attacks in practice remains largely unanswered in the academic literature. [4] This is the gap the present study aims to address.

In order to evaluate contemporary defensive efficacy, this study selects six distinct techniques categorized under the Defense Evasion tactic of the MITRE ATT&CK framework.

These techniques are executed within a controlled laboratory environment using the MITRE CALDERA adversary emulation platform,[5] facilitating a rigorous assessment of Microsoft Sentinel (a cloud native SIEM detection capability). Where built-in detection rules prove insufficient, custom detection logic is developed to ensure complete coverage of the techniques under study.

The findings are discussed in relation to cyber resilience and industrial risk management, with reference to IEC 62443, [6] the international standard series developed by the International Electrotechnical Commission to support the security of industrial automation and control systems. The goal is not only to report detection gaps, but to highlight that actively testing and validating detection coverage should become a standard part of how industrial organizations manage cybersecurity risk.

2. FROM ATTACK TAXONOMY TO DETECTION COVERAGE IN INDUSTRIAL CYBERSECURITY

Cyber threats targeting industrial environment have grown steadily in both frequency and sophistication over the past decade. Documented vulnerabilities within industrial control systems experienced a 52% escalation in 2021. The trend intensified throughout 2022, characterized by an 87% increase in ransomware campaigns directed at industrial entities. Same, the adversarial landscape witnessed a 35% proliferation of specialized ransomware collectives focusing their operations on operational technology (OT) infrastructures. [7]

These numbers reflect a broader shift in attacker behavior. Industrial systems are no longer considered peripheral targets. They are primary objectives.

What makes this threat category particularly difficult to manage is not just its scale, but its nature. Threat actors are increasingly pivoting toward “living-off-the-land” methodologies, which prioritize the exploitation of authorized engineering utilities, valid credentials, and inherent system functionalities over conventional malicious payloads. [8] Using standard system tools for malicious purposes makes detection much harder, as these actions look exactly like regular engineering tasks. This similarity allows attacks to hide within normal daily operations. Automated security frameworks face a critical obstacle when attackers hide their presence using normal system functions and the absence of traditional malicious signatures allows these unauthorized actions to cross detection thresholds without being caught.

Recent threats such as FrostyGoop, a malware strain identified in 2024, [9] demonstrate how attackers can now directly manipulate industrial process commands while remaining invisible to antivirus software.

Such attacks highlight that the risks to industrial systems transcend simple data loss, often resulting in physical equipment damage, and serious safety hazards. More than one in five organizations reported an ICS or OT cyber incident in the past year, many resulting in operational disruption and prolonged recovery. [10]

The MITRE ATT&CK for ICS framework was developed to provide a clear structure to the modern threat landscape. It acts as a detailed catalog of adversary behaviors observed in actual industrial attacks. Rather than just listing threats, it connects them to actionable defense strategies, allowing manufacturers to make more informed security decisions.

The framework is organized around the goals an attacker pursues at each stage of an intrusion, from gaining initial access and establishing persistence, to evading detection and ultimately disrupting or manipulating the industrial process. [11] This practical approach makes the framework a valuable tool for both security experts and plant engineers, providing a shared language for conducting laboratory-based security assessments.

While a taxonomy of attack techniques provides a vital foundation, its utility depends on an organization's ability to verify if their monitoring systems can detect these threats in practice. Adversary emulation platforms serve as the essential bridge between theoretical knowledge and operational validation. MITRE CALDERA, an open-source platform, allows security teams to construct specific threat profiles and launch them within a network to pinpoint vulnerabilities. The introduction of CALDERA in OT in late 2023 marks a pivotal advancement for industrial security. Developed through a partnership between MITRE and federal agencies like CISA, [12] this extension enables automated exercises tailored specifically to operational technology. By building directly on the ATT&CK for ICS framework, the tool allows organizations to move beyond "assumed security" and rigorously evaluate their OT defense perform as intended against real-world adversaries.

While SIEM platform like Microsoft Sentinel offer powerful real-time event correlation and alerting capabilities, their effectiveness in industrial settings is not guaranteed. Sentinel's flexibility allows for sophisticated custom rules via KQL, yet the platform's performance is heavily dependent on the quality of the incoming data. [13] Without rigorous data collection, security breaches may go undetected, leading to significant operational harm. A critical limitation is the commercial SIEM products typically arrive with out of the box rules designed for corporate IT.

These default settings often lack the necessary context to recognize the specific adversarial techniques used against operational technology (OT) and industrial control systems. [14].

3. FROM EMULATION TO EVIDENCE: THE EXPERIMENTAL APPROACH

The approach adopted in this study follows a structured adversary emulation methodology, designed to move beyond theoretical threat modeling and produce concrete, verifiable evidence about detection capability. Rather than assuming that a SIEM platform covers a given set of attack techniques, each technique is actively simulated in a controlled environment and the resulting detection output is systematically evaluated.

The laboratory environment consists of a set of virtual machines running both Windows and Linux operating systems, configured to replicate a simplified but realistic enterprise network. This setup provides a safe and isolated space in which attack techniques can be executed without any risk to production systems, while still generating the types of system

events and network activity that would be observed in a real industrial environment. The virtual nature of the environment also allows for rapid and repetition of tests, which is essential for consistent and comparable results across different techniques.

MITRE CALDERA serves as the adversary emulation platform throughout this study. Rather than relying on manual attack execution, CALDERA automates the deployment of specific adversarial abilities mapped directly to ATT&CK techniques, ensuring that each test is executed in a consistent and repeatable manner. For each technique under evaluation, the corresponding CALDERA ability is launched against the target virtual machines, and the resulting system activity is captured by Microsoft Sentinel.

Microsoft Sentinel works as the detection and monitoring layer in this setup. Sentinel, as a cloud-native SIEM, ingests log data from the virtual machines in real time, applying both built-in detection rules and custom analytics to identify suspicious activity. The evaluation of detection coverage is carried out in two stages. In the first stage, each technique is executed and the built-in detection rules of Sentinel are assessed for their ability to generate a relevant alert. In the second stage, for techniques where built-in rules prove insufficient, custom detection logic is developed using the Kusto Query Language (KQL) and validated against the observed activity.

The criteria used to verify the successful detection of a technique remain intentionally stringent. Detection is validated only when a direct, verifiable correlation exists between the specific command's issues by the CALDERA platform and the granular log entries ingested by Microsoft Sentinel.

This verification process relies on targeted KQL queries, which are engineered based on the precise commands and parameters employed during the experimental phase. An incidental or generic alert occurring near the time of a test is deemed insufficient for confirming detection, instead, every alert must be explicitly traced back to the specific technique being evaluated. This rigorous standard ensures that the study's findings represent genuine defensive coverage rather than coincidental alerting, providing a transparent and reproducible path from adversary emulation to empirical evidence.

The six techniques selected for evaluation belong to the Defense Evasion tactic of the MITRE ATT&CK framework. This tactic was chosen deliberately, as defense evasion techniques are specifically designed to avoid triggering security monitoring systems. In an industrial context, an attacker who successfully evades detection can remain active inside a network for extended periods, observing processes, manipulating configurations, or preparing more disruptive actions without raising any alarm.

The specific techniques evaluated, together with their ATT&CK identifiers and their relevance to industrial environments, are presented in the following section.

4. EMPIRICAL EVIDENCE OF DETECTION GAPS IN SIEM-BASED MONITORING

The six techniques evaluated in this study were selected for their documented relevance to industrial environments and their shared characteristic of being specifically designed to avoid detection.

Two of the techniques, T0820 and T0851, are drawn from the ATT&CK for ICS matrix, representing adversarial behavior observed specifically in industrial control system

environments. The remaining four T1134, T1197, T1202, T1205 belongs to the ATT&CK Enterprise matrix but have direct applicability to industrial networks where Windows-based engineering workstations, human-machines interfaces, and IT/OT integrated infrastructure are common. An analysis of the study's findings reveals that none of the six selected techniques triggered an automated response from the standard detection library of Microsoft Sentinel.

Every CALDERA execution left a clear activity that was captured in the logs but did not trigger any alert. These entries remained "silent" under the platform's default alerting thresholds. For identifying these was needed the creation and deployment of custom KQL queries that targeted the specific behavioral patterns of the adversary. This finding provides concrete, reproducible evidence that commercial SIEM platforms cannot be assumed to provide comprehensive ATT&CK coverage without active validation and deliberate custom rule development.

To understand why these gaps, exist, it is important to consider the nature of the techniques themselves. Defense evasion is not an incidental category within the ATT&CK framework. Maintaining a persistent presence within a compromised environment without detection constitutes a primary adversarial objective. Techniques categorized under this tactic are created to mirror legitimate system behaviors and leverage trusted operating system utilities, deliberately avoiding the red flags that trigger standard security alerts.

When these methods are tested against a commercial SIEM platform using only default configurations, such "by-design" invisibility manifests as measurable detection gaps. This gap highlights the inherent limitation of relying on baseline security logic to identify sophisticated, stealth-oriented threats.

Each of the six techniques selected for this study reflects a different dimension of that evasion capability, from firmware-level obfuscation and token manipulation to the abuse of legitimate Windows services and covert network signaling. Their selection was guided by two criteria: documented relevance to industrial control system environments, as established in the MITRE ATT&CK for ICS and Enterprise matrices, and practical feasibility of emulation within the laboratory setup using MITRE CALDERA.

Table 1 summarizes the six techniques evaluated, their definitions as documented in the MITRE ATT&CK knowledge base, and their potential operational impact in an industrial environment. The description is drawn directly from the official MITRE sources to ensure accuracy and consistency with the broader understanding of each technique. [11]

By aligning these finding with the operational realities of modern industrial environment, the study ensures that the analyzed risks reflect actual observed behaviors rather than hypothetical models.

Table 1. MITRE ATT&CK Defense Evasion techniques evaluated in the laboratory study and their potential impact on industrial systems.

ATT&CK Technique	Description	Potential impact on Industrial Systems
T0820 – Exploitation for Evasion	By exploiting vulnerabilities within applications, services, or the operating system kernel, adversaries can effectively bypass security controls and deactivate defensive features. In this way, the attackers operate within the environment while minimizing the risk of triggering automated detection systems.	Bypassing protective mechanisms on ICS and SCADA systems, enabling undetected manipulation of control processes.

T0851 Rootkit	-	Sophisticated malware maintains persistence by compromising the integrity of operating system APIs to hide its presence. The software can suppress the visibility of specific files and components, by intercepting core system requests and successfully evading detection by distorting the system's own reporting mechanisms.	Persistent and invisible access to HMI or engineering workstations, with the ability to feed false data to control devices.
T1134 Access Token Manipulation	-	Token modification enables adversaries to shift their operational context, granting them the permission necessary to bypass security barriers without re-authentication. Acting as a trusted system or user profile, attackers can navigate restricted environment while leaving standard identity-based alerts undetected.	Elevated host privileges serve as a critical bridge for adversaries seeking to control industrial processes. This shift in authorization allows unauthorized users to move from general network access to the high-level control functions required to disrupt or manipulate operational technology.
T1197 – BITS Jobs		Using the Background Intelligent Transfer Service provides a reliable and stealthy channel for maintaining long-term access to a host. Hiding malicious command within this trusted Windows utility, adversaries can bypass security filters that typically flag unusual file transfers or unrecognized process executions.	Convert data transfer and payload execution using a trusted Windows service, evading network-level monitoring.
T1202 Indirect Command Execution	-	Executing command through trusted system components allows attackers to neutralize common defensive safeguards, such as application whitelisting. Since these native utilities are pre-authorized by the system, their malicious application does not trigger the same scrutiny as unrecognized executables or traditional script interpreters. This approach enables unauthorized activity to blend into the background of legitimate administrative tasks, effectively blinding security tools that rely on monitoring command-line behaviors.	Execution of malicious commands on industrial workstations without triggering policy-based security controls.
T1205 Traffic Signaling	-	Covert triggers allow adversaries to maintain persistent backdoors that remain invisible until a specific network signal is received. By monitoring for unique packet signatures or sequences, a compromised system can be instructed to perform high-impact actions, such as disabling security controls or initiating a reverse shell, all while appearing to handle routing network traffic.	Exploiting manipulated network traffic for the covert activation of backdoors allows adversaries to maintain a persistent, but sleeping presence within industrial networks. Using non-standard signaling to trigger remote access, attackers can bypass perimeter defenses and re-establish control over compromised OT assets without generating the continuous traffic patterns typical of standard administrative connections.

The results presented in Table 1 confirm that defense evasion technique, whether originating from the ICS or Enterprise ATT&CK matrix, consistently fall outside the detection scope of default SIEM configurations. The need for custom detection rules across all six cases is not an exception but a pattern, one that has direct implications for how industrial

organizations approach cybersecurity monitoring. These findings set the foundation for the discussion that follows.

5. IMPLICATIONS FOR CYBER RESILIENCE IN INDUSTRIAL ENVIRONMENTS

The findings of this study carry implications that extend beyond the technical domain of SIEM configuration. They reveal a gap in how industrial organizations conceptualize and operationalize resilience. A system that functions reliably under normal operating condition, but whose security monitoring fails to detect deliberate adversarial activity, cannot be considered fully resilient. [1] Cyber resilience, understood as the ability of a system to anticipate, withstand, recover from, and adapt to adverse cyber events, must be treated as an integral dimension of industrial system reliability, not as a separate or secondary concern.

This framing has direct consequences for industrial risk management. A detection gap of the kind identified in this study represents an unquantified risk: one that does not appear in standard risk registers, is not addressed by routine maintenance procedures, and cannot be mitigated through mechanical redundancy alone. The ISA/IEC 62443 series of standards explicitly recognizes this challenge, defining risk assessment processes as critical to protecting industrial control systems and requiring that cybersecurity be addressed throughout the full lifecycle of industrial automation and control systems. [15] The results of this study suggest that organizations may satisfy formal requirements while still leaving significant gaps in their practical ability to detect sophisticated, evasion-oriented attacks.

Evidence from the evolving ISA/IEC 62443 standards suggests that modern OT security depends on the active validation of detection systems. [15] Because technical safeguards are insufficient without proof of their effectiveness, organizations must move toward a model of continuous, evidence-based testing. Utilizing adversary emulation to identify monitoring gaps represents a proactive strategy for strengthening industrial defenses. As demonstrated in this study, the integration of CALDERA and Sentinel provides a reliable pathway for organizations to verify their detection coverage and ensure their response capabilities are grounded in operational reality.

The findings presented here should be interpreted within the specific parameters of the experimental design. Notably, the simulated network environment provides a foundational model but lacks the particular behaviors found in live industrial systems. Because the current assessment utilized a finite number of techniques specifically tailored for Microsoft Sentinel, the results represent a platform-specific snapshot of detection efficacy. Also, replicating the unique challenges of real-world OT environments, specifically the integration of legacy hardware and specialized industrial protocols, will be critical for validating these findings across the broader industrial sectors.

6. CONCLUSIONS

Through the lens of adversary emulation, this study assessed whether a commercial SIEM platform could autonomously detect various defense evasion techniques. The evidence gathered across all test cases indicates that standard detection libraries lack the necessary granularity to identify these threats. Since identification was only made possible through the development of tailored detection analytics, the results underscore a critical security reality: comprehensive

coverage against evasive maneuvers is not a default feature of SIEM platforms but rather a product of active, continuous rule development and validation.

The broader significance of this finding extends beyond the technical domain. Detection gaps of this nature represent unmanaged risks in environments where operational continuity, safety, and physical process integrity cannot be addressed without first being made visible through active validation.

Future work will extend this assessment to a broader set of ATT&CK techniques and tactic categories, working toward a comprehensive detection validation model applicable to real industrial control system environments.

7. REFERENCES

- [1]. **Hopkin, P.** *Fundamentals of Risk Management: Understanding, Evaluating and Implementing Effective Risk Management.* Kogan Page, 2018.
- [2]. **Stouffer, K., Pillitteri, V., Lightman, S., Abrams, M., and Hahn, A.** *Guide to Industrial Control Systems (ICS) Security.* NIST SP800-82 Rev.2. National Institute of Standards and Technology, 2015
- [3]. MITRE Corporation. MITRE ATT&CK for ICS Knowledge Base. MITRE Corporation 2020.
- [4]. **Jiang, Y., Meng, Q., Shang, F., Oo, N., Le T.H.M., Lim, H.W., and Sikdar, B.** *MITRE ATT&CK Applications in Cybersecurity and The Way Forward.* arXiv preprint arXiv:2502.1085, 2025
- [5]. MITRE Corporation. CALDERA: Automated Adversary Emulation Platform. MITRE Corporation, 2023
- [6]. International Electrotechnical Commission. IEC 62443: Security for Industrial Automation and Control Systems. IEC, 2021
- [7]. Dragos, Inc. ICS/OT Cybersecurity Year in Review 2022. Dragos, Inc., 2023
- [8]. Dragos, Inc. Living Off the Land in ICS/OT Cybersecurity. Dragos, Inc., 2019
- [9]. O'Meara, K., Graham, M., and Ahlers, C. Impact of FrostyGoop ICS Malware on Connected OT Systems. Dragos Intelligence Brief, Dragos, Inc., 2024
- [10]. **Christopher, J.D.** The 2024 State of ICS/OT Cybersecurity. SANS Institute Survey Report, SANS Institute, 2024.
- [11]. **Strom, B.E., Applebaum, A., Miller D.P., Nickels, K.C., Pennington, A.G., and Thomas, C.B.** MITRE ATT&CK: Design and Philosophy. MITRE Corporation Technical Report, 2020.
- [12]. **Portase, R.M., Colesa, A., and Sebestyen, G.** *SpecRep: Adversary Emulation Based on Attack Objective Specification in Heterogeneous Infrastructures.* Sensors, Vol.24, No.17, 2024.
- [13]. **Shirazi, P., and Padyab, A.** *Discerning Challenges of Security Information and Event Management (SIEM) Systems in Large Organizations.* In: Clarke, N., Furnell, S. (eds) Human Aspects of Information Security and Assurance. IFIP Advances in Information and Communication Technology, Vol. 721. Springer, 2025.
- [14]. **Manzoor, J., Waleed, A., Jamali, A.F., and Masood, A.** *Cybersecurity on a Budget: Evaluation Security and Performance of Open-Source SIEM Solutions for SMEs.* PLOS ONE. Vol 19, No.3, 2024
- [15]. International Society of Automation. ANSI/ISA-62443-2-1-2024: Security Program Requirements for IACS Asset Owners. ISA, 2024.