Attacking Operational Technology Without Specialized Knowledge: The Unspecialized OT Threat Actor Profile

Stash Kempinski*†, Savio Sciancalepore*, Emmanuele Zambon*, and Luca Allodi*
* *Eindhoven University of Technology*
Eindhoven, The Netherlands
{s.p.kempinski, s.sciancalepore, e.zambon.n.mazzocato, l.allodi}@tue.nl
† *Secura*
Eindhoven, The Netherlands
stash.kempinski@secura.com

***Abstract*—Due to the unique characteristics of Operational Technology (OT), i.e., technology centered around cyber-physical activities, performing OT-related cyber-attacks is traditionally thought to require both specialized- and generic IT-related knowledge. However, in recent years, the need for specialized knowledge decreased, and OT-related cyber-attacks became increasingly easier to perform. In this paper, we profile a new threat actor, referred to as the *unspecialized OT attacker*, who performs targeted, OT-related cyber-attacks with at most basic generic knowledge. We show the relevance of this threat actor by identifying past OT-related cyber-attacks that match this threat actor profile's capabilities; we do so by mapping the types of tools used during these cyber-attacks and the knowledge required to use them. To further substantiate our analysis, we investigate readily-available tools that can assist threat actors in performing OT-related cyber-attacks. The combination of our findings highlights the present-day lowered entry level requirements to attack OT environments while limiting the scope of current assumptions.**

1. Introduction

Cyber-attacks on Operational Technology (OT) / Industrial Control Systems (ICSs) date at least as far back as 1988 [1]. Due to the unique characteristics of OT environments, such as the use of proprietary protocols and types of assets not found in IT environments, these cyber-attacks used to require specialized knowledge and/or physical access. Moreover, they used to be performed either by insiders, people without actual malicious intent, or in an untargeted manner by worms in (accidentally) Internet-connected environments [2]. This trend changed around 2010, when a shift in (OT-related) cyber-attacks can be observed [3]. Examples such as Stuxnet and Triton show that attacks became more targeted, malicious, and destructive [4]. Remarkably, asides from few high-profile attacks, this shift did not imply an increase in skilled threat actors (characterized by specialized knowledge). On the contrary, due to the IT/OT convergence [5], accessibility to OT environments became easier [4], and the need for specialized knowledge decreased [6].

Today, many readily-available tools exist that can be used to perform OT-related cyber-attacks. These tools can be found in underground markets and the clearnet, including both special-made ones that target OT(-components), and ones used for regular everyday professional usage. As discussed in Section 7, a significant number of these special-made tools are not created with malicious intent, but they can be (and have been) abused with little to no effort. Moreover, IT-related cyber-attack tools are also increasingly proving their usefulness in OT environments [7]. The present day availability of these tools further contributes to less-skilled threat actors being able to perform cyber-attacks that previously required specialized knowledge [8].

Current literature acknowledges this reduction in skill set requirement for performing OT-related cyber-attacks. However, to the best of our knowledge, there is no dedicated study on the topic. Moreover, although OT-related threat actor profiles are discussed in literature (see Section 2), the threat actor profile highlighted above is always missing. In this work, we aim to contribute to the awareness that this reduction in skill set is a credible threat to OT, and show how these threat actors can perform OT-related cyber-attacks without specialized knowledge.

Contribution. First, we formally define a new OT threat actor profile, which we name the "unspecialized OT attacker", who exclusively uses readily-available tools to carry out OT-related cyber-attacks without any specialized (OT) knowledge and at most basic generic (IT) knowledge. Second, we show that this threat actor profile is a credible threat to OT by performing an extensive analysis of the tools used in OT-related cyber-attacks between 1988 and 2022. To do so, we introduce a methodology to classify the tools used during these cyber-attacks and to map them to the well-known MITRE ATT&CK® for ICS framework. This mapping allows us to systematically determine what types of tools were used during the different steps of each cyber-attack, in turn identifying those that match the unspecialized OT attacker profile. Additionally, we investigate the availability of relevant tools in underground markets and on the clearnet. We map these tools to the ATT&CK for ICS framework, showing that at least one readily-available tool exists for each step of an OT-related cyber-attack. Through this mapping, we demonstrate the feasibility of performing OT-related cyber-attacks using only readily-available tools and the threat that the unspecialized OT attacker poses to OT today.

The rest of the paper is structured as follows. Section 2 presents related work; Section 3 provides preliminary

notions; Section 4 gives a detailed definition of the unspecialized OT attacker, including examples; Section 5 focuses on the analysis of the tools used in past OT-related cyber-attacks; Section 6 presents the research on readily available tools on the Internet; Section 7 discusses the results of both parts of our research; Lastly, Section 8 concludes the paper.

2. Related Work

Our work is not the first that profiles threat actors that target OT and Cyber-Physical Systems (CPSs). Multiple profiles have already been formalized in literature, which range from cybercriminals (with various intentions) to nation state actors. These studies have been summarized by Rocchetto and Tippenhauer in [9]. We compare the unspecialized OT attacker to the profiles defined in [9] in Table 1, showing a distinct difference in skill set and knowledge with the other threat actor profiles. Note that this is not the complete scheme that Rocchetto and Tippenhauer use to compare threat actor profiles. We omitted a set of comparison dimensions, such as "financial support", "camouflage", and "manpower", as we define the unspecialized OT attacker using only their skill set and knowledge, as will be explained in more detail in Section 4. Furthermore, we omitted the terrorist profile as its intentions are not cyber-attack related. We also do not consider hacktivists to be a different profile from cybercriminals, as we do not consider motivation. Furthermore, the skill sets characterizing hacktivists can vary significantly, e.g., some have the specialized (OT) knowledge to make custom tools [10], making such profile unsuitable to map using these characteristics. We kept both profiles in Table 1 as Rocchetto and Tippenhauer consider them as different profiles with different characteristics. Lastly, although we do not consider background, we did include the insider profile to show the difference between insiders that have OT-specific knowledge (e.g., engineers), opposed to those who do not, i.e., those fitting the unspecialized OT attacker profile, such as a system administrator.

Even though this threat actor has not been formally profiled before, threat actors matching this profile have been identified by Mandiant [11]. They recognize that there is an increase of low-sophisticated cyber-attacks on OT, describing multiple cyber-attacks from 2020 onwards. Our research goes a step further by formalizing the threat actor characteristics and taking a structured approach to identifying cyber-attacks compatible with the unspecialized OT attacker profile. Note that Mandiant's research includes non-public cyber-attacks as well, which make up the largest part of their findings, whereas we only make use of public information sources.

Mandiant also recognizes the increasing availability of OT-related readily-available tools [8] that help threat actors performing these cyber-attacks today. Through their research they identify a wide set of tools and categorize them based on their "functionality", e.g., hardware-based or network discovery tools. They also observe that most development of OT-related tools started around 2010, but their first observation dates back to 2004.

Table 1. THREAT ACTOR PROFILES AND CHARACTERISTICS AS DEFINED BY ROCCHETTO AND TIPPENHAUER, COMPARED TO THE UNSPECIALIZED OT ATTACKER PROFILE.

	Knowledge	Offensive	Physical	Network	Software	System	Source code	Protocols	Resources	Tools	Aim-Physical
B	○	○	○	○	○	○	○	○	○	○	○
UOA	○	○	○	◐	◐	○	○	○	○	○	◐
C	◐	◐	○	●	◐	○	○	○	◐	●	○
H	◐	◐	○	●	●	○	○	○	◐	◐	○
I	◐	○	○	○	◐	●	●	●	◐	◐	◐
N	◐	●	●	●	●	○	○	○	●	●	●

B = Basic user, UOA = Unspecialized OT Attacker, C = Cybercriminal, H = Hacktivist, I = Insider, N = Nation State. Skill level/knowledge/intention [*basic* < *intermediate* < *advanced*] expressed as [○ < ◐ < ●].

3. Background

To analyze the capabilities of threat actors performing OT-related cyber-attacks, it is important to classify the nature and purpose of the tools they use. In Section 3.1, we introduce cyber-attack tools and their features relevant for this work. In Section 3.2, we describe the ATT&CK for ICS framework.

3.1. Tools and Their Characteristics

In the context of cyber-attacks, we define tools as instruments contributing to the successful exploitation of system vulnerabilities [12]. Thus, cyber-attack tools are not limited to software, e.g., malware or dual-use tools (not maliciously intended, but still abusable, software) [13]. Cyber-attack tools also include general-purpose information, e.g., access credentials and other relevant intelligence, and hardware, e.g., RF equipment and lock-picking kits.

In this paper, we distinguish tools through the following characteristics: nature (benign/malicious), skill set required for effective/correct usage, and availability. These characteristics can vary significantly between tools, and, in turn, require different capabilities for attackers to be able to (ab)use them during cyber-attacks. Table 2 summarizes the three classes of tools we consider in this paper: *auxiliary*, *commodity*, and *non-commodity*. Note that we choose to explicitly avoid the dual-use tool categorization as it does not encompass the nature and skill set required to effectively use the specific tool.

"Auxiliary tools" are inherent to everyday (professional) system usage. They are created for benign purposes and are either available to any interested party (freely or commercially) or only distributed to professionals who intend to use them in their daily work. The distinguishing feature of these tools is that they are not created to exploit any vulnerabilities, but are prone to having their features exploited [14]. Hence, they are interesting for threat actors to (illegally) acquire and abuse them. The act of abusing these tools is often labeled as a "living-off-the-land" attack [13], and these tools are often categorized as dual-use tools (by definition). Examples of auxiliary tools are email clients, used to spread malware, telnet/SSH clients, used to execute malicious commands,

or the software controlling an ICS, which can be used for multiple malicious purposes [15].

"Commodity tools" are both readily available and malicious in nature, i.e., they are created to identify or exploit system vulnerabilities, regardless of their user's nature. In other words, this tool classification does not consider the context in which these tools are being used, e.g., (maliciously) by a threat actor or (benign) by a penetration tester. The tools themselves are malicious in nature as they are created with the intention to contribute to the exploitation of systems. Typically, these tools can be acquired by almost any interested party through either underground markets, commercial parties that sell them as assessment tools, or open source repositories. Examples of commodity tools are Agent Tesla, a Remote Access Tool (RAT) sold on underground markets [16], and CobaltStrike, a well-known penetration testing tool that is also being sold illegally on underground markets and abused during cyber-attacks [7]. Note that the latter is also labeled as a dual-use tool by industry [17], even though CobaltStrike and auxiliary tools differ in nature (and potentially in skill set as discussed below). Hence, the need to distinguish auxiliary from commodity tools as the dual-use categorization does not provide this required granularity.

"Non-commodity tools" are malicious in nature, but not readily available to any interested party. Thus, they share the same definition as commodity tools, but their distribution is limited or highly controlled by, typically, their creator(s). We refer to these tools as "non-commodity tools", but in literature they are also referred to as "bespoke tools" [14] or "custom tools". Examples of such tools are Stuxnet [18] and INCONTROLLER [19].

Characteristic variations unique to OT. When reasoning about these tool classes from an OT threat actor perspective, it is important to discuss the characteristic variations unique to OT. These variations stem from the inherent differences between IT and OT environments, such as the heterogeneity of OT vendors, who primarily create their own proprietary hardware and software, and the vast differences in sector-specific processes. As a result, the skill set potentially required for OT-related cyber-attacks is wider than for IT-related ones. Table 2 summarizes the differences in characteristics between these tools from an OT perspective. Both commodity and non-commodity tools generally require generic (sector-agnostic) knowledge for effective usage, e.g., IT-related knowledge for software-based tools. For example, threat actors must know when and why to perform a port scan and how to weaponize its outputs. However, they do not need to know how/why a malicious script influences a Programmable Logic Controller (PLC), only how to execute it successfully. Conversely, auxiliary tools may require specialized knowledge, e.g., OT-specific tools usually require training or having (sector-specific) process knowledge. For example, an engineer making changes to a PLC needs to know how to use the programming software correctly. They do not require knowledge related to the inner workings of the underlying network, e.g., on what ports the relevant protocols operate. Furthermore, they need to have sector-specific knowledge to ensure that the intended changes are correctly implemented. Thus, an engineer specialized in process automation in the oil and gas sector is not necessarily able to make meaningful changes to PLCs used in the food manufacturing sector.

Table 2. TOOL CLASSES AND THEIR CHARACTERISTICS FROM AN USER'S PERSPECTIVE.

	Tool Class		
Characteristic	**Auxiliary**	**Commodity**	**Non-commodity**
Tool-literacy	✓	✓	✓
Generic knowledge		✓	✓
Specialized knowledge	(✓)		(✓)
Readily available	(✓)	✓	
Nature	Benign	Malicious	Malicious

✓= required, (✓) = requirement differs per tool.

Specialized knowledge is generally required for creating any tool capable of interacting with OT environments [8]. Hence, when discussing threat actor capabilities, it is important to differentiate between those who only have the ability to use tools and those who are able to create them. Consequently, threat actors that are unable to create tools themselves must acquire them through other means.

By definition, auxiliary and non-commodity tools cannot always be acquired by any interested party. Additionally, even if an attacker would have access to a tool that in principle could serve a desired purpose does not mean that the tool can be used in the target's environment. For example, radio equipment used to configure a Remote Terminal Unit (RTU) can likely only be obtained legitimately through its vendor, which only sells said equipment to accredited customers. Moreover, this equipment can likely only be used to configure RTUs from this specific vendor, not others. Note that this does not mean that threat actors are unable to acquire any auxiliary- or non-commodity tools at all; these tools are just not always readily available.

Summarizing, OT-targeting threat actors who are unable to create tools can still successfully perform OT-related cyber-attacks using the tools available to them. However, their tools may not be suitable for the target's environment, and require the threat actors to have the knowledge to use them effectively.

3.2. ATT&CK for ICS

ATT&CK for ICS is an actively maintained knowledge base[1] developed by MITRE and describes, among others, adversarial behavior [20]. It provides a categorization of goals, named *tactics*, that a threat actor could want to achieve during a cyber-attack. For example, if a threat actor wanted to *impair process control* in an OT environment, they would have to first achieve *initial access* to the network, and possibly perform *lateral movement* to reach the OT assets within the network. In this paper, we use these *tactics* to systematically describe the different steps a threat actor can perform to execute their cyber-attack.

4. The Unspecialized OT attacker

We define the *unspecialized OT attacker* as a threat actor that performs targeted, OT-related cyber-attacks with,

1. During our research we used version 13.

at most, a basic generic skill set. They know how, why, and when to use tools widely used by both hackers and vulnerability testers (e.g., nmap and Metasploit). However, they are unable to develop such tools themselves or make significant adjustments to them, nor do they have any industrial or process knowledge. The former is a key differentiating factor between unspecialized OT attackers and cybercriminals, as characterized in Table 1. In terms of tools, the unspecialized OT attacker is limited to using commodity and readily-available auxiliary tools that do not require OT-specific (specialized) knowledge. In terms of technical knowledge, the unspecialized OT attacker is somewhat comparable to a script kiddie [21]. However, there are three main features distinguishing them from a script kiddie. First and foremost, script kiddies are defined as having little to no IT skills, whereas unspecialized OT attackers do have a generic (but basic) skill set. The script kiddie would be equivalent to, at most, basic user level of knowledge and skills as characterized in Table 1. Second, script kiddies are often associated with performing untargeted cyber-attacks, whereas unspecialized OT attackers perform targeted attacks. Lastly, opposed to script kiddies, unspecialized OT attackers are well aware of the consequences of their actions.

When characterizing an unspecialized OT attacker we focus only on their technical skills. We explicitly do not consider their intentions or other skills (e.g., social skills). Thus, an unspecialized OT attacker could be very proficient in convincing victims to click links or open malicious files, just not in creating them. Furthermore, we do not profile the unspecialized OT attacker through their working activities. Hence, it is possible for them to bypass their employer's external cyber-attack mitigation strategies, e.g., by using their own login credentials, without necessarily having the technical skills to bypass them otherwise. Although this is also an important aspect of disgruntled employees (insider in Table 1); the insider profile has, as required for their work activities (characterized in Table 1), industrial- or process knowledge. This knowledge is the key differentiating factor between an insider and a unspecialized OT attacker. Namely, having this knowledge allows engineers to perform malicious activities using only auxiliary tools, such as PLC programming software, whereas supporting staff may have to use commodity tools to achieve the same outcome.

Examples. An example of a unspecialized OT attacker is Jesse Willian McGraw, a janitor who installed malware on computers of a hospital where he was employed. Among the affected systems there was a HVAC system [22]. According to reports, McGraw used his employee privileges to physically access these computers, bypassing any external mitigations. Then, he used a commodity tool named "OphCrack" to bypass the computers security measures. Finally, he installed the (remote access) auxiliary tool "LogMeIn" to gain persistence on the HVAC system. During his trial, he admitted that he was aware of the potential consequences, namely affecting temperature-sensitive patient treatments and supplies. This classifies him as a unspecialized OT attacker rather than a script kiddie, due to him being aware of the consequences and explicitly targeting these systems, or a disgruntled employee, due to him not having the specialized knowledge to handle HVAC systems or enter them through legitimate means, i.e., valid credentials.

By contrast, someone who is not a unspecialized OT attacker, but a disgruntled employee, is Vitek Boden, who caused the Maroochy Shire sewage spill [23]. Boden was a contractor in charge of installing radio-controlled sewage equipment in Maroochy Shire who used stolen radio equipment (an auxiliary tool) to issue malicious commands to the sewage system causing raw sewage spill. He could issue these commands by tuning into the frequencies of the radio-controlled equipment, spoofing the sewers' control system and spoofing one of the assets in this system.

Although both examples abused their (ex-)employee privileges and used auxiliary tools during the attacks, Boden had the industrial and process knowledge to use the specialized tool that caused the spill.

5. Tool Usage in Past OT Cyber-Attacks

We determine that the unspecialized OT attacker is a credible threat by reasoning about the tools and capabilities needed to carry out past OT-related cyber-attacks, and verifying if they are compatible with this attacker profile. To do so, we first present a methodology to analyze our cyber-attack data set and identify those that could possibly have been performed by an unspecialized OT attacker. It is important to note that our intention is not to identify cyber-attacks that have been performed by unspecialized OT attackers, but to identify those that are compatible with their capabilities. The reason for doing so is that more advanced threat actors can also solely rely on skills and tools matching those of an unspecialized OT attacker [3], thus making it impossible to determine their profile only from the tools used. However, through our methodology, we can argue that those cyber-attacks *could* have been performed by an unspecialized OT attacker, regardless of the threat actor's actual capabilities.

5.1. Tool Mapping & Identification Methodology

We aim to identify tool usage in OT-related cyber-attacks in relation to our unspecialized OT attacker definition as unambiguously as possible. To do so, we created the following tool mapping scheme, which maps, per cyber-attack, the ATT&CK for ICS *tactics* to one of the ten classification options presented in this section. The goal of this scheme is to identify the type of tools and required skill set at each step of the cyber-attacks considered in our study. We chose to use *tactics* as cyber-attack step representation as they depict the distinct objectives a threat actor could want to achieve (with the help of tools), and they are widely accepted in the community. This representation method is more suitable for our research than, e.g., the ICS cyber kill chain [24], as that representation focuses more on the preparation of a cyber-attack and actions taken to ensure successful execution, rather than the execution itself.

We derived the methodology classification options from the tool classes described in Section 3.1, and we extended it with a "*non-cyber*" class, which we use to indicate that a step was performed by physically accessing and exploiting the victim asset(s). To achieve the required granularity for this research, we split the *auxiliary* and

Table 3. TOOL CLASSIFICATION OPTIONS

Category	Description
Commodity	Commodity tool(s) used.
Non-commodity	Custom or not seen before tool(s) used.
Auxiliary conforming	Auxiliary tool(s) used in a way that did not require process- or industrial knowledge.
Auxiliary non-conforming	Auxiliary tool(s) used that required process- or industrial knowledge.
Auxiliary unknown	Auxiliary tool(s) used, but no further classifiable information is available.
Non-cyber conforming	Directly abused the victim asset in a way that did not require process- or industrial knowledge.
Non-cyber non-conforming	Directly abused the victim asset in a way that required process- or industrial knowledge.
Non-cyber unknown	Directly abused the victim asset, but no further classifiable information is available.
Unknown	Tactic was performed by adversaries, but no further classifiable information is available.
Not performed	Tactic not performed by adversaries.

non-cyber classes into three classifications: "*conforming*", "*non-conforming*", and "*unknown*", to model to what extent the use of these tools required knowledge compatible with the unspecialized OT attacker profile, i.e., if specialized knowledge was required, or if there is not enough information available in our cyber-attack sources to determine this. We provide the complete list of the classification options in Table 3 and a more extensive description in Appendix A.

During the mapping process we relied on the written information in the cyber-attack sources as much as possible. However, because most sources do not provide a detailed description of each step (i.e., they do not mention what exact tool is used and how), we inferred most of the mappable information. To remain as objective as possible, the level of inference was kept at a minimum and only used when the required information could be inferred beyond reasonable doubt. For example, consider the *execution* of the cyber-attack wherein Alisha Sult, who was employed as a lift operator, tampered with the Gondola Transit System in Colorado, causing multiple shutdowns [25]. We can only infer that a *non-cyber* method was used; however, further inferring what happened is unfeasible due to her non-technical job description (leaning towards *confirming*) combined with the amount of work required to determine and fix the root cause (leaning towards *non-confirming*). Hence, *execution* is mapped to *non-cyber unknown* for this cyber-attack.

We used multiple procedures that allow us to compensate for such sparse mappable information, and to capture the nuances of tool usage during cyber-attacks, while still maintaining a realistic and credible mapping. Namely, a) how to classify *tactics* when multiple tools of varying classes are used, b) how to map the use of login credentials (i.e., information-based auxiliary tools) in *initial access* cases, c) how to map the access to auxiliary tools obtained through previous steps during a cyber-attack, d) how to consider the possible commoditization of non-commodity tools over time, e) how to classify *tactics* for which mappable information is almost never available. Note that the last point led us to exclude the *privilege escalation* from the mapped *tactics*. We provide in Appendix B the codebook allowing one to reproduce our mapping.

Table 4. UNSPECIALIZED OT ATTACKER MATCHING CLASSIFICATION OPTIONS.

Tactic	Options
Initial access	*All classification options*
Execution	*Commodity, Auxiliary conforming, Non-cyber conforming, Not performed*
Persistence	
Evasion	
Discovery	
Lateral movement	
Collection	
Command and control	
Inhibit response function	
Impair process control	

5.1.1. Unspecialized OT Attacker Identification Scheme. Table 4 summarizes the classification options that match the unspecialized OT attacker profile per *tactic*. We allow all *initial access* classifications. We include *non-cyber non-conforming* and *auxiliary non-conforming* to fit the profile of a unspecialized OT attacker possibly having insider access. Through this reasoning, we also include the *auxiliary*- and *non-cyber*-based *unknown*-classifications, as all their "known" counterparts are part of the profile. We include *non-commodity* because we argue that if the remaining *tactics* match the profile, the cyber-attack could have been performed by a profile-matching insider as well, regardless of how access was acquired. Consequently, *unknown* can be included, as the profile matches in all other cases.

The remaining tactics should be mapped to either *commodity*, *auxiliary conforming*, *non-cyber conforming*, or *not performed*, as by definition they match the unspecialized OT attacker profile. Conversely, *non-commodity* and *non-conforming* classifications are excluded by definition. Indeed, the latter is in contrast with unspecialized OT attackers possibly having insider access. However, this exclusion allows us to distinguish between insiders with specialized knowledge and insiders with supporting roles within an organization, i.e., those without that knowledge. We further exclude all *unknown*-related categories due to the uncertainty and possible misclassifications that they introduce. This exclusion allows us to be conservative in our identification by not erroneously identifying cyber-attacks as unspecialized OT attacker-compatible.

5.2. Data Collection

Most of the cyber-attacks analyzed during this research come from the Operational Technology Cyber-Attack Database (OTCAD) [2], which consists of OT-related cyber-attacks occurred between 1988 and 2020. We extended this data set with cyber-attacks from 2021 and 2022 through ICSSTRIVE [26] and Mandiant's research [11], using OTCAD's principles, i.e., we only consider cyber-attacks related to the OT environment of an organization. The resulting data set contains 190 cyber-attacks. [2] However, not all attacks in this data set are suitable for this research, either due to their nature (as will be explained below) or the lack of (publicly available) information needed to evaluate them. After filtering the

2. This set includes the Oldsmar Water Treatment attack despite the conflicting opinions about its factuality.

cyber-attacks based on these criteria, we obtain a data set of 47 cyber-attacks.

First, we removed cyber-attacks without publicly available specifics of their lateral movement. We made this decision because we believe that the complexity of this step heavily depends on the victim's implemented cybersecurity measures, which can vastly differ between organizations, and thus can require a vastly varying skill level to perform successfully. Hence, if the lateral movement specifics are unknown, we cannot reliably reason about the capabilities required for these cyber-attacks, and thus cannot determine if they match the unspecialized OT attacker profile. For example, consider the LockerGoga ransomware attack on Norsk Hydro [27]. It is known that this ransomware was deployed via Norsk Hydro's domain controller and that LockerGoga does not have any intrinsic propagating capabilities [28], thus lateral movement must have happened using a different tool. Depending on Norsk Hydro's cybersecurity measures, this might have required a non-commodity tool. As this information is unknown to us, we cannot reasonably classify which tool was used. In turn, we are unable to identify the required capabilities for this cyber-attack, making this and similar cyber-attacks unsuitable for our research. This criterion also served as a preprocessing optimization to our methodology as it removed 101 cyber-attacks from the data set, including all but two ransomware cases, limiting the amount of incidents for which a full mapping was needed.

Second, we removed all cyber-attacks characterized by a worm-type malware, as the untargeted nature of worms makes them incompatible with the unspecialized OT attacker profile. Furthermore, as worms are released by their creator with the intention to spread as much as possible, we cannot establish if a specific incident has happened due to the original or following releases of the worm. For example, the Conficker worm has been active between 2008 and (at least) 2017 [29]. Hence, it is impossible to determine if it was released again by someone other than its creator, and with what intentions.

Lastly, we excluded all the cyber-attacks whose only impact was the theft of non-OT related information from OT-centered organizations as, by definition, they do not fit the OT-targeting nature of the unspecialized OT attacker. For example, the Night Dragon attacks [3]. Even though these cyber-attacks specifically targeted oil, energy, and petrochemical organizations, the threat actors were collecting operations and project-financing information. This information is not related to the OT subsystem of these organizations, nor do they require specialized knowledge. Hence, these cyber-attacks are essentially generic IT-based ones, thus not matching the scope of our research.

5.3. Results

Figure 1 shows the number of unspecialized OT attacker profile matches, as defined in Section 5.1.1, compared to the number of cyber-attacks in our filtered data set per year. In total, we identify 18 cyber-attacks that could have been performed by an unspecialized OT attacker out of the 47 considered, i.e., $\approx$ 38.3%. We see that, from 2009 onward, cyber-attacks matching this profile become regular occurrences. Such a finding is in line with our hypothesis, i.e., that unspecialized OT attackers appeared

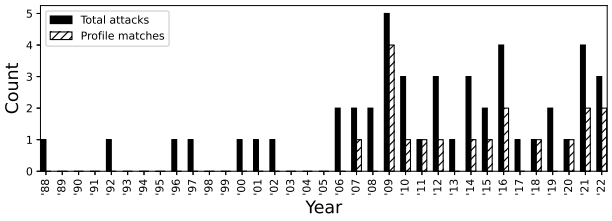

Figure 1. Total attacks and unspecialized OT attacker profile matching attacks per year.

around 2010, when a shift in OT-related cyber-attacks occurred. The single instance in 2007 originates from an attack where a disgruntled system administrator of a manufacturing company deleted server boot files, causing a downtime of several days in both the operational and enterprise aspects of the company [30]. Although this is an IT-centered attack, it caused a significant loss in productivity, hence it is included in the data set. In our opinion, the first distinct occurrence of an unspecialized OT attacker-matching cyber-attack, used as an example in Section 4, was in 2009. The threat actor abused his employee status to gain physical access to, among others, the HVAC system of a hospital, used commodity tools to bypass authentication mechanisms of this system, and an auxiliary tool to gain persistence to the system. Although this is also an IT-centered attack, it compromised the integrity of the HVAC system and, in turn, could have impacted the operations of the hospital.

These two examples do not mean that unspecialized OT attackers can only perform IT-centered cyber-attacks with an OT impact, i.e., where no specialized knowledge is required in any case. We also observe OT-centered cyber-attacks that could have been carried out by an unspecialized OT attacker, including those that would require specialized knowledge if not for readily-available tools. For example, in 2022, the SiegedSec group attacked systems using two OT-exclusive protocols, namely IEC 104 and Ethernet/IP, by using Metasploit modules [10]. It is unknown, but improbable, the attackers had the specialized knowledge to perform these cyber-attacks without these tools. However, the existence of these tools enables everybody with generic knowledge (i.e., how to use Metasploit) to attempt these attacks, leading to our threat actor profile.

Figure 2 reports the classified *tactics* used in profile-matching cyber-attacks. Note that in our methodology, *evasion* is mapped to *not performed* if the respective information sources do not report about this step. Furthermore, *unknown* and *auxiliary non-conforming* are only valid for *initial access*, because of the unspecialized OT attacker profile definition (see Section 5.1.1).

As can be derived from the high number of *not performed* classifications, during most cyber-attacks one or more *tactics* were not performed. This correlates with the noticeably low complexity of these profile-matching cyber-attacks (discussed further in Section 7). For example, consider the GhostSec group who attacked multiple Berghof PLCs, which were seemingly directly accessible from the Internet [31]. These cyber-attacks consisted of accessing the web interface of the PLCs and stopping the

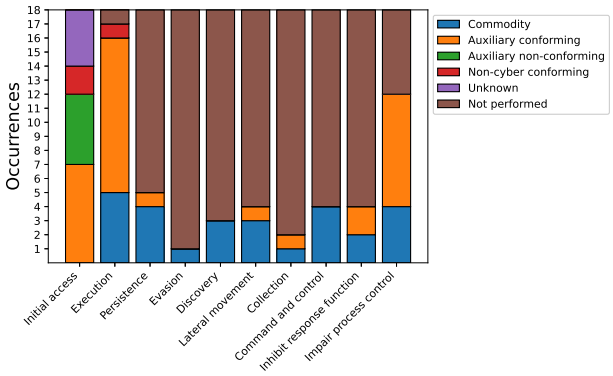

Figure 2. Stacked bar graph showing the obtained classifications per *tactic* of the cyber-attacks matching the unspecialized OT attacker profile.

execution of the user program. They were not required to pivot through an organization's enterprise network to access the PLCs (*lateral movement*) or maintain a foothold there (*persistence*).

Considering the five cyber-attacks that have *auxiliary non-conforming* mapped to *initial access*, i.e., cyber-attacks performed by insiders, none of the threat actors were employed in a role that suggested that they had any industrial or process knowledge. Next to these five, another profile-matching cyber-attack can be attributed to an insider, namely the example given in Section 4. Notably, this is the only of the six insider cases where the threat actor did not have a technical background, in the other cases they all had an IT-related job description. This indicates that all insiders should be considered a potential threat to the OT environment of an organization, not just those with specialized knowledge or those related to the IT environment of an organization. These six insider attacks all happened between 2007 and 2011, and make up all but one of the profile-matching attacks in that period.

Table 5 shows a timeline of the number of performed *tactics* per year. This shows that the arguably more advanced *tactics*, i.e., those in the middle of the cyber-attack steps, only started being used more recently. The same holds for the usage of commodity tools, as opposed to auxiliary tools. The outliers in 2009 are from the previously mentioned non-IT background insider, who relied largely on commodity tools. This correlates with our reasoning that insiders with relevant knowledge can perform malicious actions using the available auxiliary tools at their disposal, as opposed to needing malicious tools.

We derive two more observations from Table 5. First, in 2009, 2012, 2015, and 2021, not all cyber-attacks had a mapping to either *inhibit response function* or *impair process control*, meaning that there was no (noteworthy) OT-related impact. Still, we consider these cyber-attacks in our data set as the threat actors in question gained access to an OT environment and compromised its integrity. If they had other intentions, these cyber-attacks could have had an OT-related impact. Second, the threat actor of the cyber-attack in 2014 only performed *initial access*, without any *execution* or other *tactics*. As described in [32], the threat actor simply gained access to an OT-related system; they did not attempt to perform any other activities. Although we do not know the threat actor's intentions, we can argue that due to no other activities being tried, the threat actor might have been aware of potential consequences and decided not to pursue the attack further. This matches the unspecialized OT attacker's characteristics.

Lastly, we identified the industries in which the profile-matching cyber-attacks took place to determine if attacks compatible with the unspecialized OT attacker are inherent to a specific domain. As can be seen in Table 6, this is not the case. The threat the unspecialized OT attacker poses should thus be considered by all industries.

6. Readily-Available Tools on The Internet

Next to identifying past cyber-attacks that could have been performed by the unspecialized OT attacker, we carry out an empirical study on commodity tools currently available in underground markets and on the clearnet. Our goal is to identify tools that can be used by an unspecialized OT attacker during OT-related cyber-attacks, confirming that it is possible to perform them solely using readily-available tools today. Specifically, we looked for software, exploits, and information, such as tutorials or access to organizations. Note that our intention is only to show the feasibility of OT-related cyber-attacks by an unspecialized OT attacker, not to create an exhaustive list of the commodity tools available to perform them. To show that these tools exist for each cyber-attack step, we again leverage the ATT&CK for ICS *tactics*.

During our tool search, we used two sets of keywords, one set for underground markets and one set for the clearnet, due to the differences in format and the type of tools that they advertise. The studied clearnet sources consist exclusively of exploit listings, without any meaningful form of interaction possible between the creator and (potential) user. Hence, we created a set of keywords for the clearnet that consists solely of OT vendor names. Conversely, the studied underground markets are forum-based and include, next to the marketplace, also boards with discussions about tools. To be able to identify OT-related discussions on these forums, we included general OT-related terminology, common OT-using industries, and popular protocols in this set of keywords.

6.1. Underground Markets

We focused on three hacking-centered markets to which we have access and deemed relevant for this study: two Russian-speaking and one English-speaking market. We do not refer to these markets by name in this paper, as this might undermine our future monitoring capabilities of these markets.[3] Note that, due to our inability to read Russian, we asked a Russian-speaking cybersecurity professional to translate our keywords to ensure that the context of our keywords did not get lost, which might happen when using an online translator.

We used the evaluation criteria set in [33] to determine the credibility of these markets, namely a) enforcement of market regulation mechanisms, b) evidence of trade, and c) presence of prominent attack tools reported by the industry. We apply this strategy to ensure that we are

3. We will share the names of the studied markets in private at request.

Table 5. NUMBER OF PERFORMED TACTICS IN PROFILE-MATCHING CYBER-ATTACKS PER YEAR.

Tactic	2007	2009	2010	2011	2012	2014	2015	2016	2018	2020	2021	2022
Initial access	1	4	1	1	1	1	1	2	1	1	2	2
Execution	1	**4**	1	1	1		1	2	1	1	2	2
Persistence		**1**						2	1	1		
Evasion								**1**				
Discovery								**1**	1	1		
Lateral movement								**2**	1	1		
Collection		**1**										1
Command and control		**1**						1	1	1		
Inhibit response function	1			1				**1**	1			
Impair process control	1	2	1	1				**2**	1	1	1	2

x = first year a commodity tool classification occurred.

Table 6. UNSPECIALIZED OT ATTACKER PROFILE MATCHES PER INDUSTRY.

Industry	Count
General manufacturing	1
Transportation	3
Healthcare	2
Power and utilities	1
Petroleum	2
Automotive	1
Water/waste water	4
Unknown	4

considering relevant markets [34]. The considered markets all pass the evaluation. The details of this evaluation can be found in Appendix D.

To search for threads containing our set of keywords, rather than using scrapers (as usually done in underground market research such as [34], [35]), we used the built-in search functionalities of the markets. This strategy allowed us to quickly go through multiple years of forum posts. Our keywords had a low hit rate, as further discussed in Section 6.3, which enabled us to manually process the search results. This method had the added benefit of avoiding being banned from these forums, as the Russian-speaking markets explicitly forbid scrapers. We searched for keywords in full threads, but only processed threads that had potentially relevant titles.

6.2. Clearnet

We also looked into three clearnet-accessible exploit listings: namely *0day.today*, *exploit-db.com*, and *Metasploit*. We name these sources here because it does not interfere with our anonymity: the former does not require any account for the part we scraped and the latter two are publicly accessible by nature. Note that we discuss *0day.today* in this section as it vastly differs from the other underground markets, i.e., the lack of discussions and no account required to browse the listings.

We only scraped these listings for keywords consisting of OT vendor names because of the way exploits are listed: vendor name, software / embedded system name, and usually version number. As a result, there was no reason to search for other OT-related terms. We created the list of vendors from those mentioned in the ICS Shodan protocol list [36], combined with those listed in [37]. Note that, due to the vast heterogeneity of OT vendors [38] and the nature of this research, we only searched for a limited set of vendors. However, to widen our search, we also included all Metasploit exploits categorized as “SCADA".

As the three sources overlap in exploits that they advertise, we manually removed all duplicates based on their vulnerability disclosure identifier (whenever possible), software name, version, and type of exploit combination. In uncertain cases, we compared the exploits, e.g., for two exploits where one only included the major version number and another the complete version number. Furthermore, we ensured that we only included relevant exploits, i.e., those of software and embedded systems exclusively used in OT environments. This extra step had to be performed as some of the vendors also produce non-OT-related assets (e.g., Toshiba). For example, we removed the exploits related to the Toshiba e-Studio, which is an office printer.

We categorized the exploits by their target (Windows-based *software* or *embedded system*) and their purpose. We make this distinction between targets because of the generic skill set that is usually associated with software exploit development, as opposed to the specialized knowledge usually required for developing embedded system exploits. Furthermore, through these classifications, we show that exploits exist for all assets found in OT environments, not just for those also found in traditional IT environments (such as Windows-based computers). We use the following categories: *arbitrary execution*, *denial of service*, *web-based exploits*, *exposure of (potentially) sensitive information*, *remote file Create, Read, Update, and Delete (CRUD)*, *authentication and authorization bypass*, and *protocol-native commands*. More details about these categories, such as their definitions and our reasons for grouping them can be found in Appendix C. We use this categorization as it enables us to concretely show how these exploits can be used in the context of *tactics*.

6.3. Results

During our underground market study, we did not find any noteworthy OT-tailored tools, except for access to OT-centered organizations. It is important to note that, as argued by Campobasso *et al.* in [34], our lack of findings does not mean that there are no OT-tailored tools available in underground markets, only that they are not being discussed or sold on those we studied. The tools being advertised on these markets included general-purpose information stealers, crypto-related malware, access into organizations, and bulletproof hosting. Furthermore, we

Table 7. MAPPING BETWEEN EXPLOIT CATEGORIES AND ATT&CK FOR ICS TACTICS.

Exploit category	Tactics
Arbitrary execution	*Inhibit response function, Impair process control, Execution, Privilege escalation*
Denial of service	*Inhibit response function, Impair process control*
Web-based exploits	*Initial access*
Exposure of (potentially) sensitive information	*Initial access, Lateral movement, Collection*
Remote file CRUD	*Persistence, Collection*
Authentication or authorization bypass	*Initial access, Lateral movement, Persistence*

Table 8. NUMBER OF IDENTIFIED EXPLOITS PER EXPLOIT CATEGORY AND TARGET.

Exploit category	Embedded System (protocol-native)	Software
Arbitrary execution	28 (10)	79
Denial of service	13 (1)	11
Web-based exploits	14	4
Exposure of (potentially) sensitive information	17 (9)	13
Remote file CRUD	4 (1)	10
Authentication or authorization bypass	15	20

also found advertisements of cracked CobaltStrike versions and other commodity tools intended for legitimate penetration testing. These generic tools are relevant as they can (and have) help(ed) threat actors during OT-related cyber-attacks [7]. For example, CobaltStrike assists in performing nearly all *tactics*. Therefore, these underground markets provide threat actors with tools to perform most *tactics*, except for OT-specific ones, namely *inhibit response function* and *impair process control*.

Interestingly, tools supporting these two *tactics* can be easily acquired via the clearnet. Table 7 shows our mapping from exploit categories to *tactics*, the most relevant being *arbitrary execution* and *denial of service*, which are both mapped to these *tactics*. As summarized in Table 8, we identified in total 228 exploits published between 2008 and 2022. The majority of these exploits are categorized either as *arbitrary execution* or *denial of service*. Indeed, these exploits are not necessarily abusable by a basic skilled threat actor such as the unspecialized OT attacker; for example, they could require the threat actor to write their own shell code. However, when combined with tools available via underground markets such as CobaltStrike, which provides the relevant shell code, lower skilled threat actors might still be able to use them.

Furthermore, a sizeable number of exploits (91) target embedded systems, i.e., specialized hardware (such as PLCs). Although their development usually requires specialized knowledge, this is not necessarily required for using them. For example, the protocol-native exploits we identified are paired with the code that executes the exploits for its user, only requiring the target's necessary information, such as its IP address. In turn, when commoditized, these exploits can be used by any generic skilled threat actor without them ever needing specialized knowledge. For example, the SiegedSec case described in Section 5.3 uses one of these exploits.

7. Discussion

The Unspecialized OT Attacker Prevalence. The results presented in Sections 5.3 and 6.3 show that the unspecialized OT attacker is a realistic threat to OT and that this threat actor could successfully perform OT-related cyber-attacks today. In particular, we see that a significant percentage (38.3%) of the cyber-attacks in our data set could have been performed by an unspecialized OT attacker. However, there are nuances that we must consider when discussing this percentage, which can potentially influence the magnitude of our results and claims. First, note that when filtering our data set from the complete set of 190 cyber-attacks, we tried to be as conservative as possible w.r.t. the criteria for unspecialized OT attacker-profile matching. For instance, we removed cyber-attacks where *lateral movement* information is not available. Due to this choice, a set of 58 ransomware attacks were not included in our dataset. In all likelihood, a (significant) subset of these ransomware attacks is compatible with the capabilities of the unspecialized OT attacker, especially considering the state of cybersecurity measures in most OT environments. Furthermore, we can only reason about publicly reported OT-related cyber-attacks. Reports of OT-specialized cybersecurity organizations show that there are considerably more OT-related cyber-attacks every year than made public. For example, the low-sophistication cyber-attacks studied by Mandiant [11] consist primarily of "non-public" attacks, indicating that the unspecialized OT attacker might be more prevalent than we can observe. These nuances, and recognizing that more advanced threat actors can also solely use unspecialized OT attacker-matching tools for their attacks, make it unfeasible to reason quantitatively about the prevalence of this threat actor profile. We can however conclude that the unspecialized OT attacker is a realistic threat.

Low Complexity of Profile-Matching Cyber-Attacks. As shown in Figure 2, the *tactics* in profile-matching cyber-attacks are largely mapped to *not performed*. Through these *not performed* mappings, we can infer the following aspects about these cyber-attacks. The *not performed persistence* mappings indicate that the threat actors did not require any repeated access to their victims' systems. Furthermore, the *not performed discovery* and *lateral movement* mappings indicate that the threat actors did not need to pivot through a network, meaning that the victims' systems were directly accessible over the Internet. For example, the GhostSec cyber-attacks described in Section 5.3 did not require any of these three *tactics*.

This does not mean that an unspecialized OT attacker is unable to perform these *tactics*. Conversely, every *tactic* is mapped at least once, meaning that unspecialized OT attackers are able to attempt them if they have to. Note that their success is not a given, but depends on the implemented cybersecurity measures by the targeted organization. However, during most profile-matching cyber-attacks it was not required to perform these *tactics* in the first place, indicating that the victim OT environments allowed for such low-complexity cyber-attacks. In turn, this shows that the lack of cybersecurity measures taken by victim organizations also contributes to the feasibility of the unspecialized OT attacker to perform these cyber-attacks.

Relevance of OT-specific Exploits. We identified a sizeable list (228) of exploits for embedded systems and software used exclusively in OT environments. However, as mentioned in Section 6, the OT vendor market is quite fragmented, and most use their own proprietary software, protocols, and embedded device hardware. Hence, the coverage of the exploits in terms of number of targetable OT systems is probably smaller than what this number might suggest. We also need to consider that most exploits only work for a specific firmware and software version, so older ones have less chance of working when patches are available. Finally, we could not confirm that all identified exploits work as advertised, as we could not systematically test them with the diverse set of hardware, software and versions they target. However, note that old exploits are still very relevant: threat actors are known to use Metasploit modules from 2012 [10], and vulnerabilities from 2016 are known to be still exploited in 2023 [39].

Threat Implications. The emergence of the unspecialized OT attacker shows that it is becoming increasingly easier to carry out OT-related cyber-attacks. As shown by our research, this can be partly attributed to the increasing accessibility to tools allowing non-specialized threat actors to attack OT environments. Furthermore, the IT/OT convergence causes IT-centered cyber-attacks to impact OT environments more likely as well, even unintentionally.

Remarkably, current OT-related cybersecurity research focuses on the characteristics that make OT unique, such as its uncommon and usually proprietary protocols, types of assets unique to OT environments, and their peculiarities. In turn, research in OT mostly targets cyber-attack scenarios that involve threat actors with specialized knowledge, i.e., highly skilled actors that develop the tools themselves to perform the cyber-attacks in these scenarios. Offensive security research may also benefit unspecialized OT attackers and other generic skilled threat actors, as most exploits identified during our clearnet study were originally published by researchers. Note that we are not advocating for security-by-obscurity here or that such research must not be made public. However, we recognize that cyber-attacks described in [10] have made use of research-intended tools published in 2012 for malicious purposes in 2022. Our study indicates that it is also necessary to study effective defenses against low-complexity cyber-attacks on OT, particularly targeting the intersection between IT and OT. For instance, studying how traditional (IT) cybersecurity measures could be effectively applied to OT, such as the timely implementation of patches, which could help improve protection against such lower sophistication cyber-attacks. According to our study, even with the most conservative approach, i.e., not matching any removed cyber-attacks from the data set, 9.5% of the 190 cyber-attacks could have been performed by a non-specialized threat actor such as the unspecialized OT attacker. We believe the efforts of the community in this direction should follow a multi-disciplinary approach involving both technical countermeasures and organizational aspects, such as how to effectively employ these countermeasures in real-world scenarios.

8. Conclusion

In this paper, we identified and formalized a new threat actor profile, i.e., the unspecialized OT attacker, that successfully performs OT-related cyber-attacks without the specialized knowledge often thought to be needed for such cyber-attacks. At the time of writing, although informally acknowledged by some previous research [11], this threat actor has not been formally profiled in literature, nor has the prevalence of cyber-attacks compatible with this profile been studied. We profiled the unspecialized OT attacker by means of the tools available to them and their skill set, i.e., the lack of specialized (OT) knowledge. We showed that this threat actor profile is not just theoretical, by identifying 18 past cyber-attacks that match the unspecialized OT attacker profile, out of 47 we deemed suitable for our research. Finally, we showed that today threat actors can perform OT-impacting cyber-attacks through tools available on the Internet. We identified 228 exploits for embedded systems and (Windows-based) software used exclusively in OT environments that can help threat actors to perform cyber-attacks to OT without specialized knowledge.

As future work, we would like to explore the likelihood of the identified exploits being useful in a given OT-related cyber-attack. Furthermore, we would like to research the effectiveness of using solely IT-intended cyber-attack tools in OT environments, to see what IT mitigation strategies should be considered in such environments.

Data Availability

The unprocessed results of this paper can be found in [40]. These results consist of the filtered data set of mapped cyber-attacks used in Section 5 and the categorized list of exploits used in Section 6. Furthermore, the keyword sets used in Section 6.1 and Section 6.2 can be found in [40].

Acknowledgements

We would like to thank Michele Campobasso for his insights in underground markets and help shape the unspecialized OT attacker profile.

This work has been supported by the INTERSECT project, Grant No. NWA.1162.18.301, funded by the Netherlands Organisation for Scientific Research (NWO). Any opinions, findings, conclusions, or recommendations expressed in this work are those of the author(s) and do not necessarily reflect the views of NWO.

References

[1] "PLC Password Change," *RISI Online Incident Database*. [Online]. Available: https://www.risidata.com/Database/Detail/plc-password-change

[2] S. Kempinski, "OTCAD: Operational Technology Cyber Attack Database," *Secura*, 2021. [Online]. Available: https://www.secura.com/whitepapers/otcad

[3] "Global Energy Cyberattacks: Night Dragon," *McAfee*, 2011.

[4] A. Lakhani, "The Evolution of OT Cyberattacks from 2010 to Present," *Fortinet*, 2020. [Online]. Available: https://www.fortinet.com/blog/industry-trends/evolution-of-cyber-threats-in-ot-environments

[5] "Rapid Digitization is Transforming Industries Worldwide Including OT," *Otorio*, 2023. [Online]. Available: https://www.otorio.com/blog/the-ot-security-evolution/

[6] G. Murray, M. N. Johnstone *et al.*, "The convergence of IT and OT in critical infrastructure," *Australian Information Security Management Conference*, 2017.

[7] "Cyber Threat Bulletin: The Cyber Threat to Operational Technology," *Canadian Centre for Cyber Security*, 2021.

[8] J. Ashcraft, D. K. Zafra *et al.*, "Monitoring ICS Cyber Operation Tools and Software Exploit Modules To Anticipate Future Threats," *Mandiant*, 2020. [Online]. Available: https://www.mandiant.com/resources/blog/monitoring-ics-cyber-operation-tools-and-software-exploit-modules

[9] M. Rocchetto and N. O. Tippenhauer, "On attacker models and profiles for cyber-physical systems," *ESORICS*, 2016.

[10] D. K. Zafra, K. Lunden *et al.*, "We (Did!) Start the Fire: Hacktivists Increasingly Claim Targeting of OT Systems," *Mandiant*, 2023. [Online]. Available: https://www.mandiant.com/resources/blog/hacktivists-targeting-ot-systems

[11] ——, "Crimes of Opportunity: Increasing Frequency of Low Sophistication Operational Technology Compromises," *Mandiant*, 2021. [Online]. Available: https://www.mandiant.com/resources/blog/increasing-low-sophistication-operational-technology-compromises

[12] B. Arief, M. A. B. Adzmi *et al.*, "Understanding cybercrime from its stakeholders' perspectives: Part 1–attackers," *IEEE Security & Privacy*, 2015.

[13] C. Wueest and H. Anand, "Living off the land and fileless attack techniques," *Symantec*, 2017. [Online]. Available: https://docs.broadcom.com/doc/istr-living-off-the-land-and-fileless-attack-techniques-en

[14] "Common cyber attacks: Reducing the impact," *National Cyber Security Centre*, 2016.

[15] K. Proska, J. Wolfram *et al.*, "Sandworm Disrupts Power in Ukraine Using a Novel Attack Against Operational Technology," *Mandiant*, 2023. [Online]. Available: https://www.mandiant.com/resources/blog/sandworm-disrupts-power-ukraine-operational-technology

[16] "2023 cyber security report," *Checkpoint*. [Online]. Available: https://go.checkpoint.com/2023-cyber-security-report/chapter-09.php

[17] J. Agcaoili and E. Earnshaw, "Locked, loaded, and in the wrong hands: Legitimate tools weaponized for ransomware in 2021," *Trend Micro*, 2021. [Online]. Available: https://www.trendmicro.com/vinfo/us/security/news/cybercrime-and-digital-threats/locked-loaded-and-in-the-wrong-hands-legitimate-tools-weaponized-for-ransomware-in-2021

[18] N. Falliere, L. O. Murchu *et al.*, "W32.Stuxnet Dossier," *Symantec*, 2011.

[19] N. Brubaker, K. Lunden *et al.*, "INCONTROLLER: New State-Sponsored Cyber Attack Tools Target Multiple Industrial Control Systems," *Mandiant*, 2022. [Online]. Available: https://www.mandiant.com/resources/blog/incontroller-state-sponsored-ics-tool

[20] "ATT&CK for ICS," *MITRE*. [Online]. Available: https://attack.mitre.org/

[21] N. R. Mead, E. Hough *et al.*, "Security Quality Requirements Engineering (SQUARE) Methodology," 2005.

[22] "Former Security Guard Who Hacked Into Hospital's Computer System Sentenced to 110 Months in Federal Prison," *U.S. Attorney's Office*, 2011, https://archives.fbi.gov/archives/dallas/press-releases/2011/dl031811.htm.

[23] M. Abrams and J. Weiss, "Malicious Control System Cyber Security Attack Case Study - Maroochy Water Services, Australia," *MITRE*, 2008.

[24] M. J. Assante and R. M. Lee, "The Industrial Control System Cyber Kill Chain," *SANS Institute*, 2021.

[25] A. Taylor, "Telluride liftie facing jail time," *Pique News Magazine*, 2003. [Online]. Available: https://www.piquenewsmagazine.com/whistler-news/telluride-liftie-facing-jail-time-2463460

[26] ICSSTRIVE. https://icsstrive.com/.

[27] "Hackers hit norsk hydro with ransomware. the company responded with transparency," *Microsoft*. [Online]. Available: https://news.microsoft.com/source/features/digital-transformation/hackers-hit-norsk-hydro-ransomware-company-responded-transparency/

[28] "What you need to know about the lockergoga ransomware," *Trend Micro*. [Online]. Available: https://www.trendmicro.com/vinfo/mx/security/news/cyber-attacks/what-you-need-to-know-about-the-lockergoga-ransomware

[29] "Examining CONFICKER/DOWNAD's Impact on Legacy Systems," *Trend Micro*, 2017. [Online]. Available: https://trendmicro.com/en_us/research/17/l/conficker-downad-9-years-examining-impact-legacy-systems.html

[30] "Former Worker Sentenced for Wrecking Servers," *CSO*, 2008. [Online]. Available: https://www.csoonline.com/article/522576/access-control-former-worker-sentenced-for-wrecking-servers.html

[31] D. Krivobokov, "Pro-Palestinian Hacking Group Compromises Berghof PLCs in Israel," *Otorio*, 2022. [Online]. Available: https://www.otorio.com/blog/pro-palestinian-hacking-group-compromises-berghof-plcs-in-israel/

[32] "ICS-CERT Monitor January - April," *U.S. Department Of Homeland Security*, 2014.

[33] L. Allodi, "Economic Factors of Vulnerability Trade and Exploitation," *CCS*, 2017.

[34] M. Campobasso, R. Rădulescu *et al.*, "You Can Tell a Cybercriminal by the Company they Keep: A Framework to Infer the Relevance of Underground Communities to the Threat Landscape," *WEIS*, 2023.

[35] D. Georgoulias, R. Yaben *et al.*, "Cheaper than you thought? a dive into the darkweb market of cyber-crime products," *ARES*, 2023.

[36] "Industrial Control Systems," *Shodan*. [Online]. Available: https://www.shodan.io/explore/category/industrial-control-systems

[37] "Top 20 programmable logic controller manufacturers," *Robotics & Automation News*, 2020. [Online]. Available: https://roboticsandautomationnews.com/2020/07/15/top-20-programmable-logic-controller-manufacturers/33153/

[38] "DefenderSphere – ICS Vendors," *Industrial Defender*. [Online]. Available: https://www.industrialdefender.com/defendersphere-ics-vendors

[39] "Annual ICS Advisory Summary," *ICS Advisory Project*, 2024. [Online]. Available: https://drive.google.com/file/d/1HYDE_rD1dvJb30r7CuoJuXEYSPNs7MxU/view

[40] S. Kempinski, S. Sciancalepore, E. Zambon, and L. Allodi. [Online]. Available: https://github.com/StashK/Unspecialized-OT-Attacker

[41] C. Bing, "Hackers try to contaminate florida town's water supply through computer breach," *Reuters*, 2021. [Online]. Available: https://www.reuters.com/article/us-usa-cyber-florida-idUSKBN2A82FV/

[42] "United States of America v. Adam Flanagan," *United States District Court For The Eastern District of Pennsylvania*, 2017.

[43] M. Chapman, "Teenager hacks Polish tram system," *IT News*, 2008. [Online]. Available: https://www.itnews.com.au/news/teenager-hacks-polish-tram-system-100838

[44] T. Bateman, "Police warning after drug traffickers' cyber-attack," *BBC*, 2013. [Online]. Available: https://www.bbc.com/news/world-europe-24539417

[45] "Lesson Learned: Risks Posed by Firewall Firmware Vulnerabilities," *NERC*, 2019.

[46] "Malware Shuts Down Milling Factory," *RISI Database*. [Online]. Available: https://www.risidata.com/Database/Detail/malware-shuts-down-milling-factory

[47] K. E. Hemsley and R. E. Fisher, "History of Industrial Control System Cyber Incidents," *Idaho National Laboratory*, 2018.

[48] M. Almodawah, "Shamoon 3 Malware Sample," 2018. [Online]. Available: https://github.com/m-almodawah/Shamoon-3

[49] "What is vidar malware?" *Checkpoint*. [Online]. Available: https://www.checkpoint.com/cyber-hub/threat-prevention/what-is-malware/what-is-vidar-malware/

[50] A. Marín, "Everything you need to know about the lummac2 stealer," *Outpost24*, 2023. [Online]. Available: https://outpost24.com/blog/everything-you-need-to-know-lummac2-stealer/

A. Tool Classification Options

Commodity indicates the usage of a commodity tool, i.e., a readily available and malicious-in-nature tool. For example, the tools used during a cyber-attack by Sieged-Sec were Metasploit modules created for industrial equipment [10].

Non-commodity indicates the usage of a non-commodity tool. A straightforward example of such a tool is Stuxnet, malware that used multiple 0-day vulnerabilities and was created to target nuclear facilities in Iran [18].

Auxiliary conforming indicates that an auxiliary tool was abused without the need for any specialized knowledge. For example, during the *execution* of the Oldsmar water attack [41], the threat actor altered the chemical contents through the SCADA software that was open on the machine that was accessed over TeamViewer. Although it is not self-evident that this software does not require any process knowledge, the source describes an threat actor that performed unstructured actions to change the water treatment process, indicating a lack of process knowledge.

Auxiliary non-conforming indicates that an auxiliary tool was abused in a way that required specialized knowledge, which can be either (sector-specific) industrial or (possibly organization-specific) process knowledge. For example, consider the tools that were used in the previously mentioned Maroochy Shire sewage spill [23]. During this cyber-attack, the attacker used specialized equipment to send command messages into the network, which is part of its intended usage. From the cyber-attack source, we can infer that both process- and industrial knowledge was needed to perform the attack, namely knowing what network frequency identifiers to impersonate an asset within the network and what meaningful commands to send.

Auxiliary unknown indicates that there is no clear indication of the kind of knowledge required to use the tool; only that it can be inferred that an auxiliary tool was used. For example, consider the cyber-attack where Adam Flanagan changed, among other things, the communication channel frequency of base stations, rendering them unusable for the company he installed them for in the first place [42]. The indictment states that he simply logged into the stations and adjusted the frequencies. We can infer that no malicious tools (thus auxiliary tools) were used for the *execution* and *impair process control* steps of this cyber-attack. However, it is unclear if these actions required specialized knowledge: how to adjust the frequencies of the radio equipment, or not (e.g., the tool already provided a number of valid options and the attacker could have chosen a random one from the list).

Non-cyber conforming indicates that a step consisted of physically accessing an asset and abusing it without the need for specialized knowledge. For example, in 2008 a Polish teenager was able to switch tram tracks of the Polish tram system through the use of a "TV remote control"-like unit [43]. Although it is assumed that the *execution* of this attack required specialized knowledge (i.e., creating the remote-like tool), the *initial access* consisted of him physically being near the tram tracks, i.e., something anyone can do.

Non-cyber non-conforming indicates that a step consisted of physically accessing an asset and abusing it in a way that required specialized knowledge. For example, consider the *initial access* into the systems of the Port of Antwerp, during which threat actors physically accessed and connected devices to the assets that handled the location details of containers that were able to act as key loggers and screen grabbers [44]. To be sure that these devices were connected to the right assets, process knowledge was likely required, i.e., knowing what assets have access to the container location details and where these assets were physically located in the Port of Antwerp.

Non-cyber unknown indicates that we can infer that a step consisted of physically accessing an asset and abusing it, but we cannot determine with reasonable certainty if specialized knowledge was required to perform the malicious actions. For example, consider the *execution* of the cyber-attack wherein Alisha Sult tampered with the Gondola Transit System in Colorado causing multiple shutdowns [25]. From the sources, we can only infer that a *non-cyber* method was used to tamper with the system; however, further inferring what happened is unfeasible due to Sult her job description (leaning towards *confirming*) and the amount of work required to determine and fix the root cause (leaning towards *non-confirming*).

Unknown indicates that no information is available about how a step was performed, but we can infer that it was performed during the respective cyber-attack. For example, during a cyber-attack on the U.S. power grid, the attackers used a known vulnerability to reboot firewalls [45], which in turn caused a denial of service between field devices and their control center. Even though the source states that a known vulnerability is used, it is unknown if these firewalls were directly internet-facing and if the vulnerability was exploitable through commodity tools. Furthermore, as firewall details (brand or model) are unknown, we are unable to determine if any commodity tools supporting that exploit were available at that time.

Not performed indicates that a step has not been performed. For example, during the previously mentioned cyber-attack, no *collection* was performed.

B. Mapping Procedures and Codebook

A1. If we know that multiple tools of mixed *conformity* were used during a cyber-attack, we mapped the non-profile matching categories to show that, at some point, more advanced knowledge was used to perform a step.

A2. When considering *initial access* achieved through the use of information-based (*auxiliary*) tools, e.g., credentials, we considered if any insider could have had access to those credentials, or if these credentials have been leaked in the past. Through this reasoning, we consider credentials for production systems to be process knowledge, thus mapping their usage during *initial access*

to *auxiliary non-conforming*. However, in case there is a strong indication that relevant credentials were leaked before an attack, the related tactics are mapped to *auxiliary conforming* as long as the other constraints for its classification hold. This consideration also applies to the use of default credentials, although such a case is not present in the cyber-attack set used for this research.

A3. If a previous cyber-attack step provided access to the auxiliary tool(s) used in a step that did not require any specialized knowledge, we mapped to *auxiliary confirming*, regardless of whether (any of) the previous step(s) required non-profile matching tools. This assumption covers the cases where a threat actor only gets access to the relevant tools during the cyber-attack, such as the *auxiliary conforming* example in Appendix A.

A4. We classified specific malware as *non-commodity* only for its first occurrences in the data set within a reasonable time frame, and when there is no indication that the malware has been previously used. After the first usage, it is assumed that samples of this malware eventually became available online, so becoming a commodity tool. In turn, we cannot say with certainty anymore that following cyber-attacks are launched by its creator. This assumption influences the mappings related to two malware cases, namely Stuxnet and Shamoon. Although Stuxnet was a targeted attack, the way it spread caused other organizations to be impacted by the malware as well [46]. Both this incident and the intended target of Stuxnet were impacted at around the same time (the same year in this case). Due to the complexity of the malware, both cyber-attacks are mapped to *non-commodity*. Shamoon, at the time of writing, has three cyber-attacks attributed to it, two in 2012, and one in 2016 [47]. The first two cyber-attacks are classified, where applicable, to non-commodity tool usage, due to the short time frame between the two attacks and the attribution to the same cyber-terrorist group. However, we classified the third cyber-attack as a commodity tool usage, due to the malware being available online for some years now [48].

A5. We mapped *evasion* to *not performed* unless there is information that explicitly reports anything about the use of this tactic by adversaries. We made this decision due to the little reporting usually available about this tactic due to the state of OT cybersecurity, and we assume that there was no detection mechanisms present in the victim environment, hence no evasion necessary.

A6. We excluded *privilege escalation* and *impact* from the mapping. There is too little information in the cyber-attack sources available to create a meaningful mapping for the former, while the latter describes the consequences of an attack rather than the way it is performed.

C. Exploit Categorization

Arbitrary execution: exploits such as buffer overflow and remote code execution. We grouped these exploits as they all enable threat actors to (somewhat) freely execute code on the victim machine.

Denial of service: exploits that stop systems from performing their intended functionality. We created a separate category for this type of exploit because of the directly noticeable impact this exploit type causes.

Web-based exploits: different forms of cross-site scripting, HTML injection, and cross-site request forgery. We grouped these exploits as they require victim interaction.

Exposure of (potentially) sensitive information: exploits that abuse vulnerabilities such as directory traversal, user enumeration, and exposure of cryptographic keys. We grouped these exploits because they can all potentially lead, directly or indirectly, to the exposure of sensitive information.

Remote File CRUD: exploits related to the Creation, Reading, Updating, or Deleting (CRUD) of arbitrary files. We grouped these exploits due to them specifically being able to interact in some form with the victim's file system.

Authentication and Authorization bypass: vulnerabilities that relate to the bypassing of the intended authentication and authorization mechanisms, such as privilege escalation, but also the possibility to recover passwords or expose credentials. We grouped these exploits as they allow threat actors to gain access or perform actions on systems. Note that we included under this category exploits related to the exposure of passwords and credentials, and not under the *exposure of (potentially) sensitive information* category, as they specifically allow for the bypassing of authentication or authorization.

Protocol-native commands: protocol messages that are native to the protocols embedded systems and used to communicate with each other and other relevant software. Exploits of this nature abuse the functionality provided by these commands, e.g., to start or stop PLCs or extract potentially sensitive data. We created this category to show that they are inherent to the (usually insecure-by-design) protocols embedded systems use rather than a vulnerability contained within a single piece of software / embedded system firmware.

D. Market Evaluation

In the interest of our anonymity on those markets, we refer to them as *RU1*, *RU2*, and *EN*. In particular, we found evidence of the enforcement of rules in the significant amount of banned users in all three markets, as discovered during our manual search-result processing. All three markets also use a reputation system in which users can mark other users' posts as (un)helpful, providing an indication of the user's trustworthiness. It must be noted that *EN*'s reputation system is not very representative as lots of users try to improve their reputation through low-effort posts. At the same time, *EN* does feature an elaborate pledge-based system that allows users to get into a gentleman's agreement during sales, which then can get publicly disputed, so providing a better representation of a user's reputation. *RU1* and *RU2* provide an escrow service to ensure fair exchange. In addition, *RU2* shows on users' profiles the number of usages of this service, whereas *RU1* only shows evidence of trade through positive reactions on the respective sale threads. Lastly, we found evidences of tools such as Vidar [49] and LummaC2 [50] being advertised on these markets, indicating that prominent tool creators were actively advertising on them.