Evolution of ICS Attacks and the Prospects for Future Disruptive Events

Joseph Slowik
Principal Adversary Hunter
Threat Intelligence Centre
Dragos Inc
Hanover, MD, United States
jslowik@dragos.com

Abstract:

Headlines are full of proclamations covering the latest in industrial control system (ICS) attacks and threats to critical infrastructure. But behind each prominent event lies a trendline from the 2015 Ukraine power outage through the 2017 attack on safety systems at an oil and gas facility in Saudi Arabia. When moving beyond media reporting, two clear patterns emerge in how ICS attacks have evolved: first, initial attack vectors increasingly avoid using malware and techniques that are tell-tale signs of advanced adversary activity; second, only at the final, ICS-disruptive stages of intrusions is complex malware introduced to codify ICS-specific knowledge to enable nearly any computer network operations operator to execute complex commands.

Exploration and examination of these trends reveals a definite direction in how future attacks will occur within the ICS space, as adversaries seek to satisfy the seemingly mutually-exclusive goals of evading detection while deploying increasingly advanced capabilities. By adopting and understanding a "complete kill-chain" approach to ICS attack methods, defenders – from ICS asset owners and operators to national governments to intergovernmental organizations – can begin formulating defensive plans to detect and mitigate future attacks.

To describe and defend this thesis, ICS disruptive events from the past four years will be analysed in detail to identify how these threats have evolved over time, and what complementary measures are necessary to defeat these attacks. A thorough understanding of the risk posed by ICS attacks will allow stakeholders from ICS operators to policymakers to begin identifying and implementing appropriate controls and security measures to safeguard critical infrastructure and prevent future, potentially catastrophic attacks.

Keywords: *ics, industrial control systems, cyber physical, scada, cyber warfare, cyber disruption*

Contents

1. ICS ATTACK BACKGROUND

Industrial control system (ICS) cyber events first rose to prominence with the discovery of Stuxnet in 2010. [1] [2] Since this ground-breaking, epoch-defining malware emerged, ICS asset owners and defenders discovered several additional intrusions and targeted malicious software (malware) in the following years:

Figre 1: Overview of ICS Attacks and Events

- HAVEX malware infecting various ICS-related organizations from 2011 through potentially 2015. [3] [4]
- An incident at a German steel mill in 2014 attributed by the German government to malicious cyber activity. [5] [6]
- The 2015 and 2016 Ukraine power events. [7] [8]
- The 2017 TRISIS attack on an oil and gas facility in Saudi Arabia [9]
- A series of intrusions from 2017 to 2018 targeting US, UK, and German electric utilities and grid operators. [10] [11]

Overall, as shown in Figure 1 outlining known ICS attacks, events appear to not only increase in relative frequency, but in severity. Leaving Stuxnet aside as an outlier, events have progressed from mere enumeration and data gathering (HAVEX campaigns) to active disruption of operations (Ukraine events) to potentially seeking physical destruction (TRISIS).

Firgure 2: ICS Cyber Kill Chain

While most reporting and defensive countermeasures dwell on the final impact of ICS-centric intrusions, such events rely upon a series of prior steps and required actions for success. Based on this expanded view of the overall attack progression, network defence and response operations can orient ICS events to the series of steps and pre-requisites necessary to achieve attacker objectives. In this fashion, the ICS Cyber Kill Chain (Figure 2) displays the various stages required to successfully initiate, continue, and ultimately conclude an ICS-focused network attack. [12] [13] When viewed in this fashion, ICS attacks can be analysed as a process and not a one-off, isolated event.

Based on this expanded viewpoint, ICS-centric adversaries must complete several stages to achieve a final effect. From initial survey and penetration of the enterprise IT environment (Stage 1) through crossing in to the ICS network culminating in final reconnaissance and penetration of control system resources (Stage 2), adversaries must achieve multiple, inter-linked objectives for operational success. When viewed from the perspective of the defender, this viewpoint not only serves to better illustrate the requirements for a successful ICS cyber-attack, but it also highlights the various stages where attackers must refine and implement tradecraft to ensure operational success.

When adopting this view of ICS-related intrusions and disruptive events, ICS defenders and those studying trends in ICS attacks can identify distinct attack methodologies within the same adversary or activity group corresponding to different attack stages. Specifically, one set of tactics, techniques, and procedures (TTPs) may be used for initial access and lateral movement, while a completely different set of TTPs are employed for delivering and executing an ICS disruption. More importantly, identifying commonalities in certain attack phases – such as TTP re-use common to multiple adversaries at certain stages of the kill chain – enables defenders and analysts to deploy defensive countermeasures at specific attack steps that will hold for multiple adversaries or intrusion variants. Essentially, in analysing attacks, researchers and defenders should devote more resources to areas likely to feature significant overlap or TTP reproduction – Stage 1 and initial Stage 2 activity according to the ICS Cyber Kill Chain – than on target-specific effects revealed at the attack's conclusion.

2. "Living off the Land" and ICS Attacks

Figure 3: IT-ICS Convergence

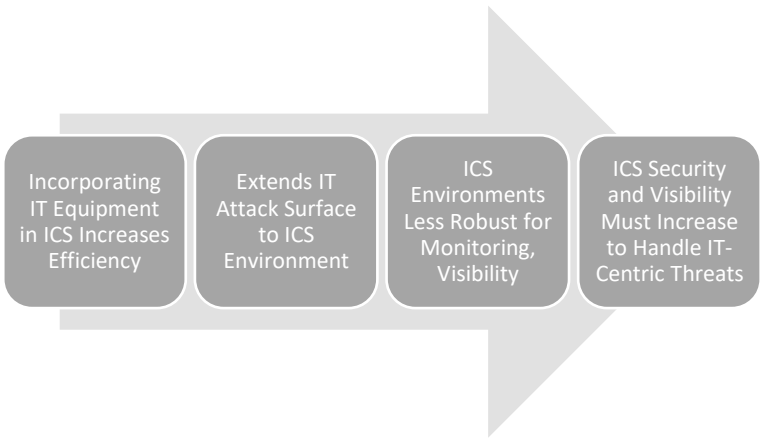

Attackers are shifting cyber operations away from custom malware and tools to an ever-greater reliance on system tools, scripts, and in-memory execution, especially for initial access operations and similar portions of the Stage 1 kill chain. [14] [15] Aside from potentially reducing development costs, such actions serve to stymie detection and defence in poor visibility environments, while also serving to frustrate attribution and actor-centric analysis given the commonality of techniques.

"Living off the land" attack methodologies have increased in popularity for several years in IT intrusions. Of note for industrial and critical infrastructure environments, what works in IT intrusions is

increasingly relevant for ICS events. There are two reasons for this: first, the "IT-ification" of ICS environments (the incorporation of ever greater numbers of Windows-based systems for various tasks, sometimes referred to as the "IT-OT Convergence") [16] [17]; and second, the need for attackers to remain reasonably well hidden (or obfuscated).

In ICS-focused attacks, we identify a significant and undeniable trend: attackers avoid custom software while focusing on "common" TTPs for initial operations (corresponding roughly to Stage 1 of the ICS Cyber Kill Chain, discussed above), and reserving development and similar resources for custom,

Figure 4: Attack Bifurcation

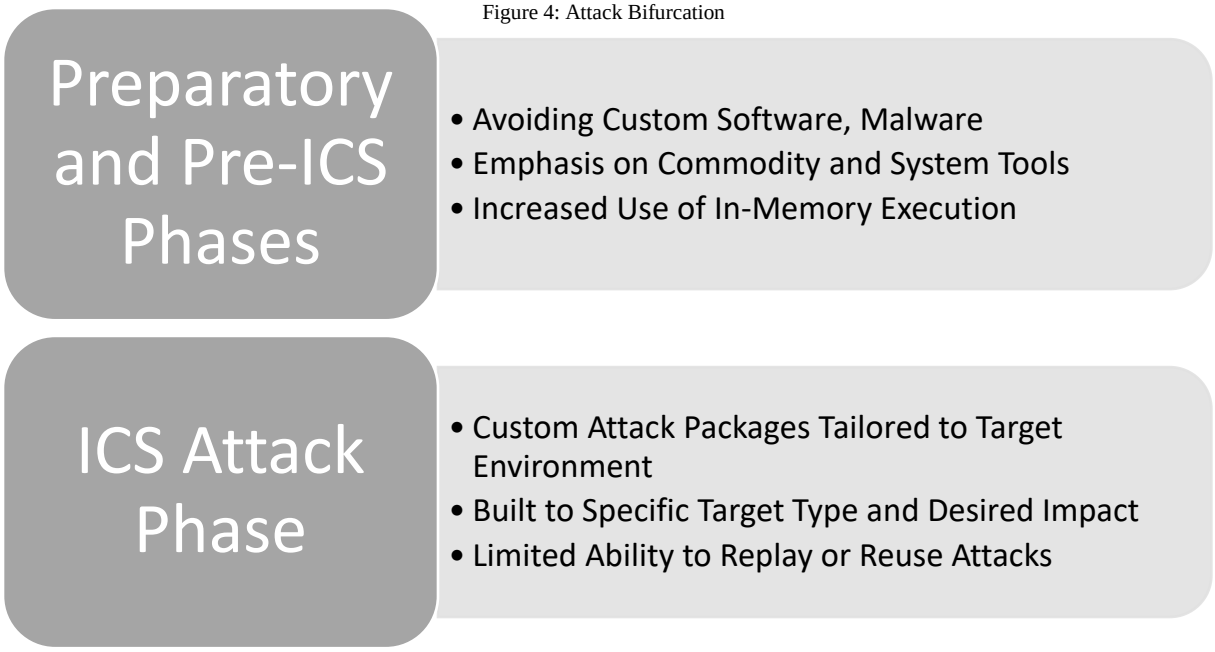

highly-targeted attacks to fulfil final-stage ICS disruptive attacks. This bifurcation of attacks, summarized in Figure 4, allows attackers to have the best of two, seemingly contradictory worlds: the ability to blend in and evade detection through built-in system tools and techniques; while deploying customized, more advanced toolkits designed to interact with industrial processes at the final attack stage.

A trend toward "living off the land" techniques until the final stages of the ICS Cyber Kill Chain presents several problems to defenders and policymakers concerned with ICS and critical infrastructure protection:

1. Ensuring appropriate telemetry and visibility to capture malicious use of otherwise benign, built-in tools in both IT and ICS environments.
2. Building defender knowledge of adversary techniques to differentiate normal or merely "strange" behaviour from malicious activity.
3. Identifying and enabling new methods for security monitoring, response, and recovery that avoid or build upon existing security solutions such as blacklisting/whitelisting, antivirus-like approaches, and anomaly detection.
4. In lieu of custom malware and similar "markers", determining alternative methods to group and attribute ICS attacks to enable policymaker responses, such as attribution to enable retribution.

An important consideration in presenting the above is to note that while attackers are migrating to IT-centric tradecraft for initial infection and lateral movement (including potentially within the ICS environment), this shift does not imply or support simply bringing IT security solutions in ICS networks. Rather, the nuances of industrial operations overall and the fundamental differences in *how* similar technology is used and deployed between IT and ICS networks mean that IT solutions cannot simply migrate over. Considerations include: prioritization of uptime and avoidance of disruption, which argue against active or preventative solutions found in Intrusion Prevention System (IPS) and Endpoint

Detection and Response (EDR) products; limited system resources that mitigate against software agent deployment; specialized systems or configurations which are not amenable (or legally allowed) to modification, such as for security software installation; and differences in how similar technologies (such as the Windows operating system) are used and employed in ICS environments compared to enterprise IT. Thus, what ICS network defenders need is not a wholesale import of IT security technologies, but a synthesis between ICS operational awareness and unique behavioural observations and IT functionality and security monitoring best practices.

Several examples best illustrate both the requirements presented previously and the limitations of both current ICS security and simply importing IT technology. Specifically, the following will be considered in depth: the CRASHOVERRIDE electric utility event; the TRISIS oil and gas attack; and persistent information gathering on US, UK, and German electric grid operators from 2017 to the present.

3. CRASHOVERRIDE IN CONTEXT

CRASHOVERRIDE, also referred to as Industroyer, represented the first known malware designed to disrupt industrial processes since Stuxnet, and the first specifically targeting electric distribution operations. [18] [19] The attack itself culminated in a brief outage through an attack on electric distribution equipment at the victim substation – essentially, opening breakers within the environment to interrupt electricity delivery. While not very sophisticated, the attack nonetheless codified knowledge of the specific victim ICS environment in software, enabling even a junior person with little understanding of electric distribution operations to successfully deliver the attack by executing a file in the "right place" at the right time. This division of labour and codification of specialized knowledge in software represent the true "breakthroughs" of CRASHOVERRIDE.

Figure 5: CRASHOVERRIDE Attack Flow

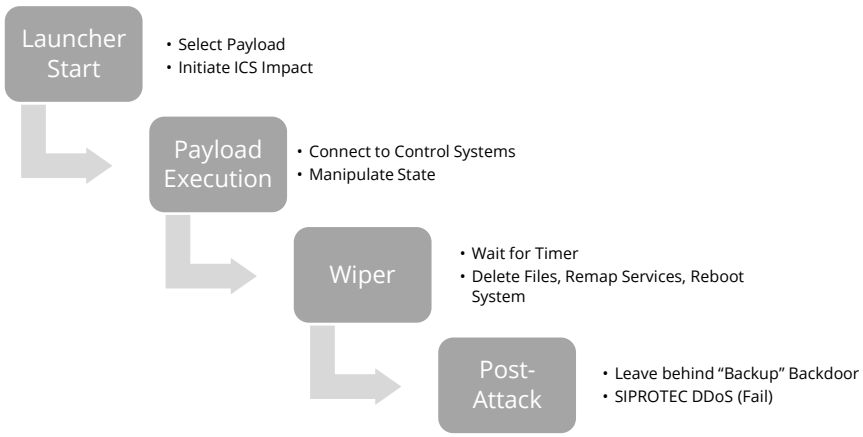

At time of initial discovery and public disclosure, much attention was lavished on *what* CRASHOVERRIDE did, but little information was available on *how* the attack came about. Subsequent details were released in October 2018, identifying an interesting division in operations: all actions prior to the delivery of the actual CRASHOVERRIDE malware relied almost exclusively on credential capture and re-use to move throughout the environment, and the use of fairly standard tools such as the Mimikatz credential theft software. [8] While the actual tradecraft deployed raises some question as to the true capability of the adversary responsible for CRASHOVERRIDE, [20] the attack nonetheless represents a two-staged event: initial access and lateral movement largely using built-in system commands fuelled through compromised credentials; followed by the deployment of custom software for the target environment just prior to actual execution.

Confounding detection and defence, the victim ICS environment simply lacked visibility into most of these techniques, or the activity in question sufficiently "blended in" with normal administrative and maintenance behaviour to fly beneath detection thresholds. Given limitations in security product

deployment, the only real way to capture the events in question would be through passive data aggregation through network capture and log forwarding, followed by correlation of observations to yield potential malicious behaviors.

While CRSHOVERRIDE itself represents a multi-staged attack, shown in Figure 5 above, the actual ICS attack "package" arrived mere hours before attack execution in the majority of observed instances from available data. Thus, from a defensive standpoint, focusing on the custom tradecraft and observables (e.g., "indicators of compromise") associated with CRASHOVERRIDE would require near-immediate response to mitigate. Instead, the victim organization needed to observe, identify, and react to operations earlier in the ICS Kill Chain when the adversary deployed common, "living off the land" techniques to secure intermediate goals leading up to final actions on objectives.

A further and final consideration would be a replay of CRASHOVERRIDE in another environment, where IOCs may be useful to prevent or quickly detect the ICS attack payload. The problem with this approach is that the IOCs are only relevant and useful in environments nearly identical to the CRASHOVERRIDE victim. The combination of communication protocols, network topology, and specific hosts as captured in both hardcoded binary values and configuration files would simply not be present (at least, not all of them) in another environment. Therefore, detecting and alerting on hashes is simply not helpful to electric utilities seeking to protect against the 'next' CRASHOVERRIDE.

4. TRISIS AND XENOTIME ACTIVITY TARGETING ICS

TRISIS, also referred to as TRITON [21] and HATMAN [22], led to an operational disruption at a Saudi Arabian oil and gas facility in 2017. Like CRASHOVERRIDE, the focus of most media and researcher attention upon initial discovery was the TRISIS malware framework: a rootkit designed to allow attacker access to a specific model of Schneider Electric's Triconex safety instrumented system (SIS). [23] TRISIS enables an attacker to control and (potentially) arbitrarily change SIS parameters, with multiple potential attack scenarios: from programming safe conditions as unsafe (causing a shutdown) to coding unsafe conditions as within safety tolerances (leading to potential physical destruction). Within this continuum of potential operations, TRISIS's perpetrators at least tacitly accepted the risk that modification of a SIS device could result in injury or even death – even if only as the result of a mistake or unintended effect of SIS modification. Much like CRASHOVERRIDE, final execution of TRISIS relies upon several interlinked stages, shown in Figure 6 below.

Firgure 6: TRISIS Attack Path

Separate from the final stage of this event, the adversary responsible for the TRISIS event, XENOTIME, [23] utilized a variety of techniques to prosecute its attack. Some custom variants of publicly-available tools such as credential capture items were present, but much like the CRASHOVERRIDE event, XENOTIME relied overwhelmingly on credential harvesting and reuse to move throughout the environment and gain sufficient privileges to execute malicious code. Through a combination of persistent credential capture and deployment multiple open- or commercial-source tools and scripts, XENOTIME successfully pivoted through the victim network and penetrated from enterprise IT into the ICS network.

Since the TRISIS event, XENOTIME remains active and has moved even further toward emphasizing a combination of "living off the land" and "in-memory" attacks to enable operations. Examples include extensive research and enumeration of target remote authentication portals and testing various publicly available PowerShell-based post-exploitation scripting frameworks. Based on observed activity,

XENOTIME's behaviour in future engagements promises to be even more difficult to detect with traditional visibility limitations, especially within the ICS environment.

From a defensive and "lessons learned" perspective, TRISIS is unique in many aspects: the malware itself simply cannot be reused except in instances targeting the same equipment with the same firmware revision as found in the original victim. While the victim is not alone in running this combination of hardware and software, this nonetheless represents a highly circumscribed set of targets with little recourse for attack replay against other organizations running either different equipment, or the same equipment running different firmware. The lesson to apply here is thus: rather than focusing on very narrowly-targeted software which will likely never again be used except in very specific circumstances, defenders and policymakers should focus on the *enabling steps* that allowed for this final attack delivery. These stages earlier in the ICS Cyber Kill Chain are far more likely to be common across multiple attack scenarios – even for different adversaries – and thus provide a more effective and useful example to follow for future engagements.

5. Electric Utility Reconnaissance Activity

Starting in early Spring 2017, electric utilities within the United States, [24] United Kingdom, [25] and Germany [26] experienced multiple, ongoing intrusions linked by several governments to Russian state interests. [27] In all observed cases, while the attacker was able to successfully penetrate the control system network of multiple organizations, including access to sensitive systems such as Human-Machine Interfaces (HMIs) and Engineering Workstations (EWs), all observed activity was limited to information gathering with no identified disruptive impact or (immediate) intent.

While very concerning from a policy and defence perspective, the activity also stood out for the technical actions enabling the intrusions. In all known cases, initial access to victim networks took place through capturing legitimate credentials – whether for the victim organization, or a trusted vendors or contractors that could be leveraged for further access – to remotely authenticate to victim environments. Further credential theft via tools such as the ubiquitous Mimikatz software enabled further migration into target environments until the attacker gathered sufficient access and privileges to penetrate into control system networks.

Aside from some publicly-available toolkits that are used by entities ranging from state-sponsored adversaries to penetration testers (such as Mimikatz), the attackers in these events, referred to as ALLANITE, completely avoided the use of custom tools or software for intrusions. To access target systems, the attacker simply "logged on" remotely, and executed built-in system commands to harvest data, capture screenshots, and move information out of the target network. In this type of an attack, antivirus, blacklisting (and even whitelisting), and similar security tools become nearly useless in monitoring – let alone blocking – intrusion activity. Furthermore, aside from access to sensitive resources (e.g., signals intelligence) tying activity back to linked perpetrators, the possibilities for technical attribution are slim, since few or no custom code, encoding schema, or other technical "tells" exist for threat researchers to fingerprint the activity.

While this activity has not at present resulted in any disruptive impact to date, thus far it represents the most extreme migration to "living off the land" techniques in ICS intrusions with the resulting impact of near complete evasion of traditional detection methodologies. Given the level of access attained and information gathered, one can reasonably anticipate that a future disruptive tool or capability linked to these intrusions would represent something like CRASHOVERRIDE: a software framework capable of interacting with control system equipment at scale (due to a modular framework with multiple attack packages) to allow for widespread impacts across multiple systems within a relatively short period of time. However, such a tool will more than likely be designed specifically for a single attack event (even if across multiple organizations, although roughly simultaneously), given the sheer diversity in ICS configurations and equipment installations. Thus any "signatures" for malicious software will likely fail to detect such a tool, while any after-the-fact detection of a malicious event will still need to deal with a disruption to electric utility operations, representing a win for the adversaries in question. To

effectively deal with such intrusions and prevent physical service disruption, these intrusions must be identified and mitigated at earlier stages – where attackers are focused on building up access within the victim networks and gathering information to enable subsequent ICS impacts.

6. Scaling Operations

The seemingly divergent trends in recent ICS-focused attacks enables adversaries to better scale and resource attacks. As indicated in the above discussions of CRASHOVERRIDE, TRISIS, and electric grid intrusion events, relatively straightforward, uncomplicated (yet quite effective) operational tradecraft against fairly typical, Windows-based environments (even in higher-level ICS applications) can be applied to enable subsequent complex, industrial-specific attacks. Once an intrusion reaches a state where target-specific, ICS-capable effects are required, tasks can be broken up into component parts to gain greater operational efficiencies: in the CRASHOVERRIDE case, specialist ICS knowledge can be encoded in software to enable a disruptive event, while in TRISIS knowledge of the inner-workings of the Schneider Electric Triconex device are incorporated into the TRISIS rootkit with which an operator can interact.

Essentially, operations bifurcate based upon position within the ICS Cyber Kill Chain. For Stage 1 and potentially initial Stage 2 operations, fairly "standard" TTPs and methodologies can be employed to gain required initial access and enable data collection. Once this stage is complete, any further operation requires a shift to ICS-specific tools and capabilities, sometimes uniquely designed for the specific victim environment. At this point, operations must be handed off to ICS subject matter experts (SMEs), or standard intrusion personnel must be provided with tools capable of automating or guiding their actions while abstracting away ICS-specific knowledge.

A division between ICS-specialist knowledge and general computer network operations tradecraft enables ICS operations to scale more widely than any previous operation since the autonomous STUXNET worm. Viewed along these lines, the Russian-linked grid intrusions from 2017 onward may very well represent the preliminary stages required to build out capabilities and conduct operational preparation of the environment (OPE) leading up to future attacks. [28] The techniques building up toward capability development remain fairly common for offensively-trained personnel, and a robust talent pool (at least among typical "Advanced Persistent Threat" adversaries) already exists to perform this initial work. Subsequent operations can then move to smaller, centralized laboratories and development shops that can process collected data and reconnaissance information to build out an attack capability for operators "in the field".

Through this approach, what were once highly specialized operations requiring some degree of knowledge for how target industrial processes function (e.g., in the 2015 Ukraine event, the attacker simply remotely logged on to an engineering workstation and manually disrupted power, requiring some basic level of knowledge for how to manipulate the device) now become commodified effects packages deployed in software.

Figure 7: Conceptual Division of Labor for ICS Attack Development

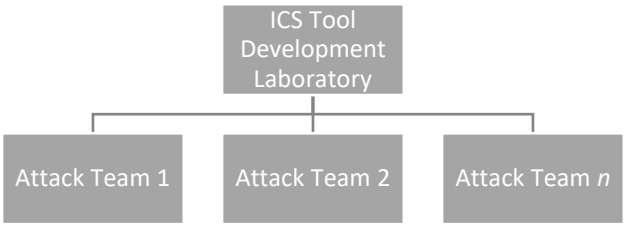

The "talent pool" of ICS-threatening personnel expands so long as specialized research and development operations exist to support offensive work. While still a non-trivial barrier to entry for ICS disruptive engagements, differentiating intrusion work from capability developments significantly increases the efficiency and scalability of ICS attacks. Essentially, a specialization of labour approach is adopted by taking advantage of SMEs to build capabilities for more general personnel spread among multiple "attack" or delivery teams, a relationship conceptually represented in Figure 7.

7. ICS-Focused Defense for ICS Networks

Much of this paper focuses on a shift by adversaries to commodified, "living off the land" techniques for intrusions. Yet in drawing attention to this real development, observers may draw the incorrect conclusion that better security as deployed in enterprise IT environments – combining EDR agents across all hosts feeding telemetry to a centralized security information and event management (SIEM) device – will be suitable to shoring up ICS-specific defence. While embracing such concepts will definitely bolster IT network defence (addressing Stage 1 intrusion scenarios), this comparison begins to break down when applied to the specific requirements and use-cases of industrial environments.

In addition to considerations previously mentioned in Section 2 of this paper, all of the above examples also highlight ICS-specific capability and understanding needed to prepare or execute attacks. While non-trivial amounts of information can be gathered on ICS operations through traditional IT intrusion means, truly accurate and in-depth information requires going "deeper". For example, HAVEX differentiated itself by understanding precisely how to interact with the OPC protocol to poll and extract ICS-specific data from control systems. [30] While some aspects of HAVEX may be detectable by traditional, IT-focused malicious binary analysis, overall this malware simply does not exhibit qualities or functionality typically found in Windows malware enabling for IT defence evasion – simply because IT defence has no idea what OPC enumeration looks like. The same applies to more concerning attacks such as CRASHOVERRIDE and TRISIS. Furthermore, preliminary operations, such as system identification and protocol enumeration, may rely on techniques as simple as an NMAP scan for initial operations, subsequent work requires an actual understanding of what traffic, exposed services, and hardware/firmware revisions mean in the context of trying to manipulate a system. Finally, the actual attack may be no more complex than simply modifying the target system via legitimate means – such as in CRASHOVERRIDE, where protocol-specific packages were deployed with no greater functionality than to change the state of a breaker between closed and open. In this sense, behavioural or binary analysis simply fails to understand not only that such software is malicious in the typical enterprise IT sense, but also fails to grasp the gravity behind the existence of such software: automating ICS manipulation in software to scale attacks.

Thus, to truly execute ICS defence, even in light of increased IT-ICS convergence, ICS-specificity and awareness are required. Simply porting enterprise IT solutions may catch preliminary activity, but should adversaries successfully evade this layer of detection they will have free reign to deliver and deploy ICS-specific capabilities without fear. While improved IT defence and visibility is required to meet the rising problem of commodified, "living off the land" attacks at scale, thus defeating initial access and data gathering operations, true ICS defence (and overall defence in depth) demands ICS-specific awareness and capabilities to fulfil the ICS defensive mission.

8. Defense and Policy Responses

Understanding the evolution of adversary tradecraft and its implications for future ICS-focused attacks is vital for preparing defence and response, and guiding policy to prevent, detect, or deter such actions. Unfortunately, as the preceding sections make clear, evolutions in initial access techniques combined with the increasing IT-ICS (or IT-OT) convergence mean that ever greater numbers of adversaries will attain the capability to at least gain entry to ICS networks. Fewer adversaries will possess the aptitude or the willingness to launch low-level, unsophisticated disruptive attacks. Fewer still will be capable of more elaborate actions – exemplified by Stuxnet and TRISIS. These higher-end actions will remain the purview of a truly elite set of adversaries capable of not only understanding industrial processes, but developing tailored capabilities to specific equipment for an attack scenario. A simple visual mapping of this relationship is provided in Figure 8, showing that the number of likely adversaries shrinks as one moves toward more sophisticated attack techniques.

Figure 8: Adversary Capability Mapping

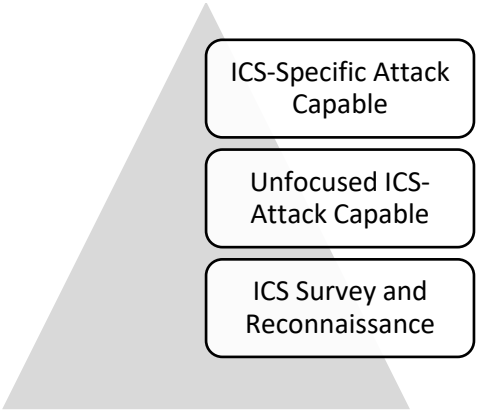

While the most sophisticated attacks will remain the purview of a few entities, industrial operators, asset owners, and defenders more generally should anticipate ever greater numbers of adversaries at least operating in this space, if only because of the powerful capabilities in terms of strategic communication afforded by critical infrastructure interference. [29] [30] The migration from standard – and unique-looking – malware and its related artefacts makes defence and detection more difficult when applied to the industrial context. While enterprise IT networks have made great strides in gaining greater endpoint or "host" visibility – through various (EDR products and freely-available frameworks such as Microsoft's Sysmon utility [31] – ICS networks lag significantly in these aspects, or may simply never be appropriate places for such technology due to other limitations. Thus ICS-centric defenders are left at a loss: importing general operational capabilities from the IT realm for efficiency and related purposes results in a more problematic threat landscape, but ICS-centric practices and restrictions prohibit more robust visibility down to the level of individual systems on the network. As a result, techniques such as abuse of native Windows tools, PowerShell scripting-based attacks, and similar items are not merely unchallenged in most industrial environments – they are invisible.

From a policy perspective, much effort is spent on compliance regimes and similar endeavours such as the North American Electric Reliability Corporation's Critical Infrastructure Protection (NERC CIP) program for electric utility operations and the European Union's Network and Information Systems (NIS) Directive. [32] [33] While these frameworks have resulted in improvements in overall security and management of industrial networks, they lag significantly behind the actual threat landscape and current attack methodologies. For example, many of these frameworks and policy responses focus on items such as software vulnerabilities in industrial equipment – when an analysis of recent attacks shows that exploitation of such vulnerabilities (whether known vulnerabilities with an available patch or true "zero day" attacks) is rare to non-existent in the past five years. By focusing efforts on improving patching cycles and vulnerability assessment, policy makers have unwittingly redirected scarce defensive resources away from current problems (the lack of visibility in industrial networks) toward issues that have thus far proven irrelevant in actual attacks.

Similarly, many government bodies and policy makers have emphasized increased information sharing to improve cyber defence. [34] [35] Unfortunately, much of the emphasis on sharing focuses on "indicators of compromise" (IOCs). [36] While the theoretical construct of an IOC emphasizes some degree of context and amplifying information, in practice IOCs are reduced to atomic, basic observables such as hash values and IP addresses. [37] [38] Unfortunately, the primary focus of most observed ICS attacks until final effect stages are *techniques* (such as command line use and credential theft and replay) instead of concrete *observables*. Even when typical IOC-like data may be present, such items may be so specific as to be useless to the majority of other potential victims, as seen in the TRISIS example above. Instead of sharing information about malware hashes or command and control addresses, sharing frameworks must instead focus on TTPs and observed adversary *behaviours*, especially when considering ICS environments. [39] A major drawback to this approach is that such items are not amenable to automated, machine-to-machine sharing desired by authorities (and most easily implemented by many stakeholders). Thus, a true data sharing framework that will be useful requires not just identifying a more robust way to share what is essentially attacker tradecraft, but also ensuring that organizations have the requisite network, host, and process visibility to actually make use of and act upon such information.

9. Articulating Attacks as Behaviors and Building ICS Defense

ICS defenders and stakeholders should view the progression of activity outlined above and the divergence between Stage 1 and Stage 2 Kill Chain activity as a call to arms to improve defence – with the obvious question of, in what way. As mentioned earlier, IT-centric security solutions may suffice to defend in Stage 1 operations, but fail utterly or are simply inapplicable for Stage 2, ICS-specific events. The current focus on IOC sharing and blacklisting protects against specific attacks, but is fundamentally backward-looking and does not take into account the necessary variation in ICS attack tools given specific victim environments.

Rather than attempting to force unworkable solutions into inappropriate environments or continually adopting a continually backward-looking approach to specific observables (IOCs), ICS defenders instead must embrace

specific, applied defensive solutions designed for industrial environments. In this case, defenders must seek solutions that are founded upon malicious behaviours or similar interactions within the ICS environment, capturing what an adversary must do from initial intrusion through ultimate effect. By capturing both the actions necessary the prerequisites for attacker methodology, defenders can build a robust posture around attacker requirements.

To illustrate this idea, typical malware identification or binary blacklisting may fail dramatically in ICS environments. But deploying mechanisms to identify ICS-like functionality within such environments – for example, code snippets or library references critical for interfacing with an industrial process – can locate binaries that are ICS attack *capable*. When combined with other observations – where that binary originated, how it moved into the network, and if it was executed – defenders can begin to build a robust set of interlocking behaviours around a new observation to disposition it as malicious or benign. In the latter case, binaries with ICS specific functionality entering sensitive portions of the network (or required areas before such sections for staging) should be explained by legitimate maintenance or operational activity. In the absence of such explanation, defenders have caught an object that should immediately be treated as possibly malicious and warranting further investigation.

10. EXPECTATIONS, THE FUTURE, AND SOLUTIONS

Adversaries are motivated to do what *works*, not what is most 'complex' or 'artful'. Therefore, so long as observed trends in ICS environments – the continued adoption of IT systems combined with a lack of sufficient host and process visibility to detect "living off the land" attack techniques for staging and ICS-specific detection for final attacks – persist, attackers have scant motivation to adjust TTPs from currently successful behaviour types. Furthermore, the totality of evidence supports a continued interest across multiple potential actors in ICS attack capabilities – whether to counter a perceived capability possessed by adversaries, [40] or simply as an end in itself. As a result, ICS-focused threats, with all of their ramifications for the orderly function of civilian societies and economies, are likely to *increase* over time rather than decreasing, while their methodology continues to move toward TTPs most environments are simply ill equipped to observe, let alone prevent.

For policy-makers and defensive planners, the resulting objective is clear, if somewhat counter to accepted wisdom within the cybersecurity space. Essentially, what is most needed to ensure viable defence of ICS environments in the face of current threats is an emphasis on *visibility:* as adversaries shift to native system tools and in-memory execution for preliminary stages prior to executing final ICS-centric attacks that are essentially invisible to IT-centric security solutions, ICS environments must ensure that such actions can first be *observed* in order to ensure that defensive personnel can then *act* on and mitigate such intrusions.

Embracing the assumption that resources are finite and that the application of scarce resources requires mapping effort to the current threat environment within ICS network security, policy-makers and others involved in attempting to ensure the security of industrial and critical infrastructure are left with a clear choice. Rather than emphasizing approaches that simply do not map to current ICS threats (imposing expensive enterprise IT security solutions on ICS, or tightly regulating vulnerability management) or actions that result in a false sense of security (encouraging and mandating IOC sharing and blacklisting), those empowered to make decisions and guide ICS owners and operators must instead drive resources towards what appears to be an obvious concept: improving visibility within ICS environments.

At present, despite the profusion of Windows- and Linux-based systems within ICS networks, the general capability to monitor, log, and centralize behaviour (including ICS-specific, process-based behaviour such as what is typically aggregated at data historians) remains minimal. Given increased adversary migration toward attacks designed specifically to take advantage of visibility gaps, standing policy and resources must shift to address the current problem. Examples include providing guidance, resources, and potentially funds to ICS-operating organizations to improve host- and industrial process-based visibility with an emphasis on security monitoring. Additionally, from a regulatory perspective, relevant bodies should investigate the possibility of guiding system and ICS equipment vendors toward improving the amount of telemetry and operational data produced by systems involved in critical infrastructure functions. Only by addressing this primary visibility gap can organizations and defenders then move on to addressing the fundamental problems posed by intrusions and gain the required

initiative to detect attacks *as they happen* – allowing for the possibility of mitigation and prevention – rather than responding to attackers *after they occur*.

Once this visibility gap is closed, more robust actions can take place to begin meeting the actual security issue. First and foremost, emphasizing the end-of-kill chain "effects" created by multi-staged intrusions essentially cedes significant initiative to adversaries. Rather than adopting an approach seeking to counter end-of-intrusion impacts, defenders and policymakers must instead seek avenues to improve defence and response across the entire chain of events that make up an intrusion scenario. In this fashion, response is enabled at earlier stages, minimizing impacts and protecting critical services.

An effective means for approaching behaviour-based ICS security lies in migrating beyond simple sharing and blacklisting of atomic IOCs – previously described as essentially irrelevant in defeating malicious *behaviours* underpinning specific attack scenarios – and toward developing detection and mitigation strategies that address the fundamental TTPs deployed by attackers to enable intrusions. [41] Examples of steps to take include typical "boiler plate" recommendations found in current policy guidance, such as implementing robust multi-factor authentication schema or ensuring secure storage of up-to-date backups to facilitate system restoration. More importantly, security recommendations and guidance must expand to include fundamental detection and monitoring strategies that can identify (or block) fundamental behaviours associated with current adversary TTPs. Examples of this include: guidance on detecting suspicious process chains (e.g., Microsoft Office launching PowerShell which in turn produces a network connection); the execution of unsigned binaries from untrusted locations; identifying new executable files entering into critical environments (such as the control system network); and thorough mapping of user logon behaviour. [42] [43] None of these are currently easy, but represent excellent areas where state-level capabilities – whether in financing, research, or deployment – can marshal resources to educate, inform, and support migration toward more effective defensive practices. Such actions will not only support ICS networks in general and critical infrastructure in particular, but form the basis for more robust IT defences as well.

11. CONCLUSION

ICS-focused attacks have shifted over time from relying almost exclusively on custom toolsets and software to a bifurcated approach embracing commodity, "living off the land" techniques for the majority of actions on objective with target-specific malware deployed only in final attack stages. While ICS-targeting adversaries have shifted their TTPs over the past five years, defenders and policymakers have not kept pace with these developments. Instead of identifying these trends and focusing resources and efforts to counter them, approaches continue to emphasize items such as patch management, raw indicator sharing, and compliance-based regulatory regimes that are at best tangentially related to the current threat environment, while at worst diverting resources from more appropriate actions.

Moving from the current threat landscape for ICS-targeting intrusions to general principles of network defence, overall defenders and institutions supporting their efforts must adopt an intelligence-driven, adversary-focused approach to ensure that defensive measures and recommendations keep pace with attacker evolution. In the current circumstances, ICS defensive recommendations and public resources largely focus on the network security problem as it existed in ICS (and IT) networks five to ten years ago: identifying malware hashes and command and control nodes, building signatures around "known bad" items, and deploying or sharing observables in the most rapid fashion possible to ensure coverage. In the meantime, adversaries evolved to evade this approach through weaponization of legitimate system tools, protocols, and functions rendering most of current defensive approaches ill-suited to many (although admittedly not all) attacks.

Thus, the recommendations in this paper come with an important caveat: the threat landscape, whether for ICS specifically or network defence more generally, is not and will not remain static. Today's recommendations and points of emphasis may prove unsuitable to addressing future developments in attacker tradecraft. Therefore, individual organizations and the national and international bodies and resources supporting defence must remain committed to an iterative, ongoing analysis of attacker

methodologies and be prepared to modify and coevolve defensive recommendations and resources to keep pace with adversaries. Only through this approach, as identified in the preceding sections related to ICS specific threat development, can network defenders at every level ensure continued, adequate coverage and response to the cyber threat environment.

Acknowledgments

As with so many endeavours, this paper represents the culmination of several efforts involving many individuals. Among others, the author wishes to thank his colleagues at Dragos for their assistance, support, and knowledge in composing this work, specifically: Selena Larson, Jimmy Wylie, Reid Wightman, Sergio Caltagirone, Dan Gunter, and Rob M. Lee.

12. REFERENCES

[1] K. Zetter, "An Unprecedented Look at Stuxnet, the World's First Digital Weapon," Wired, 03 November 2014. [Online]. Available: https://www.wired.com/2014/11/countdown-to-zero-day-stuxnet/. [Accessed 12 November 2018].

[2] D. Kushner, "The Real Story of Stuxnet," IEEE Spectrum, 26 February 2013. [Online]. Available: https://spectrum.ieee.org/telecom/security/the-real-story-of-stuxnet. [Accessed 12 November 2018].

[3] Symantec Security Response, "Dragonfly: Western Energy Companies Under Sabotage Threat," Symantec, 30 June 2014. [Online]. Available: https://www.symantec.com/connect/blogs/dragonfly-western-energy-companies-under-sabotage-threat-energetic-bear. [Accessed 12 November 2018].

[4] Kaspersky Lab Global Research and Analysis Team, "Energetic Bear - Crouching Yeti," Kaspersky, 01 July 2014. [Online]. Available: https://media.kasperskycontenthub.com/wp-content/uploads/sites/43/2018/03/08080817/EB-YetiJuly2014-Public.pdf. [Accessed 12 November 2018].

[5] R. M. Lee, M. J. Assante and T. Conway, "German Steel Mill Cyber Attack," SANS ICS, 30 December 2014. [Online]. Available: https://ics.sans.org/media/ICS-CPPE-case-Study-2-German-Steelworks_Facility.pdf. [Accessed 04 December 2018].

[6] R. A. Becker, "Cyber Attack on German Steel Mill Leads to "Massive" Real World Damage," PBS, 08 January 2015. [Online]. Available: https://www.pbs.org/wgbh/nova/article/cyber-attack-german-steel-mill-leads-massive-real-world-damage/. [Accessed 04 December 2018].

[7] SANS ICS, "Analysis of the Cyber Attack on the Ukrainian Power Grid," SANS, 18 March 2016. [Online]. Available: https://ics.sans.org/media/E-ISAC_SANS_Ukraine_DUC_5.pdf. [Accessed 12 November 2018].

[8] J. Slowik, "Anatomy of an Attack: Detecting and Defeating CRASHOVERRIDE," Dragos, 12 October 2018. [Online]. Available: https://dragos.com/media/CRASHOVERRIDE2018.pdf. [Accessed 12 November 2018].

[9] Dragos, "TRISIS Malware: Analysis of Safety System Targeted Malware," Dragos, 13 December 2017. [Online]. Available: https://dragos.com/blog/trisis/TRISIS-01.pdf. [Accessed 2018 November 2018].

[10] US-CERT, "Russian Government Cyber Activity Targeting Energy and Other Critical Infrastructure Sectors," US-CERT, 15 March 2018. [Online]. Available: https://www.us-cert.gov/ncas/alerts/TA18-074A. [Accessed 12 November 2018].

[11] Reuters, "German Intelligence sees Russia behind Hack of Energy Firms: Media Report," Reuters, 20 June 2018. [Online]. Available: https://www.reuters.com/article/us-germany-cyber-russia/german-intelligence-sees-russia-behind-hack-of-energy-firms-media-report-idUSKBN1JG2X2?il=0. [Accessed 04 December 2018].

[12] M. J. Assante and R. M. Lee, "The ICS Cyber Kill Chain," SANS, 5 October 2015. [Online]. Available: https://www.sans.org/reading-room/whitepapers/ICS/industrial-control-system-cyber-kill-chain-36297. [Accessed 14 November 2018].

[13] E. M. Hutchins, M. J. Cloppert and R. M. Amin, "Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains," *Leading Issues in Information Warfare & Security Research,* vol. 1.1, no. 80, 2011.

[14] C. Wueest, "Living off the Land and Fileless Attack Techniques," Symantec, 01 July 2017. [Online]. Available: https://www.symantec.com/content/dam/symantec/docs/security-center/white-papers/istr-living-off-the-land-and-fileless-attack-techniques-en.pdf. [Accessed 12 November 2018].

[15] Counter Threat Unit Research Team, "Living off the Land," Secureworks, 28 May 2015. [Online]. Available: https://www.secureworks.com/blog/living-off-the-land. [Accessed 14 November 2018].

[16] G. Murray, M. N. Johnstone and C. Valli, "The Convergence of IT and OT in Critical Infrastructure," in *Australian Information Security Management Conference* , Perth, 2017.

[17] Cisco, "IT/OT Convergence," Cisco, 2018. [Online]. Available: https://www.cisco.com/c/dam/en_us/solutions/industries/manufacturing/ITOT-convergence-whitepaper.pdf. [Accessed 14 November 2018].

[18] Dragos, "CRASHOVERRIDE: Analysis of the Threat to Electric Grid Operations," Dragos, 13 June 2017. [Online]. Available: https://dragos.com/blog/crashoverride/CrashOverride-01.pdf. [Accessed 14 November 2018].

[19] A. Cherepanov, "WIN32/Industroyer: A New Threat for Industrial Control Systems," ESET, 12 June 2017. [Online]. Available: https://www.welivesecurity.com/wp-content/uploads/2017/06/Win32_Industroyer.pdf. [Accessed 14 November 2018].

[20] J. Slowik, "CRASHOVERRIDE: When "Advanced" Actors Look Like Amateurs," 03 November 2018. [Online]. Available: https://pylos.co/2018/11/03/crashoverride-when-advanced-actors-look-like-amateurs/. [Accessed 14 November 2018].

[21] B. Johnson, D. Caban, M. Krotofil, D. Scali, N. Brubaker and C. Glyer, "Attackers Deplopy New ICS Attack Framework "TRITON" and Cause Operational Disruption to Critical Infrastructure," FireEye, 14 December 2017. [Online]. Available: https://www.fireeye.com/blog/threat-research/2017/12/attackers-deploy-new-ics-attack-framework-triton.html. [Accessed 14 November 20188].

[22] ICS-CERT, "MAR-17-352-01 HatMan - Safety System Targeted Malware (Update A)," US-CERT, 10 April 2018. [Online]. Available: https://ics-cert.us-cert.gov/MAR-17-352-01-HatMan-Safety-System-Targeted-Malware-Update. [Accessed 14 November 2018].

[23] Dragos, "XENOTIME Activity Group," Dragos, 24 May 2018. [Online]. Available: https://dragos.com/blog/20180524Xenotime.html. [Accessed 14 November 2018].

[24] D. Sanger, "Russian Hackers Appear to Shift Focus to US Power Grid," *The New York Times,* p. A11, 27 July 2018.

[25] National Cyber Security Centre, "Advisory: Hostile State Actors Compromising UK Organisations with Focus on Engineering and Industrial Control Companies," National Cyber Security Centre, 05 April 2018. [Online]. Available: https://www.ncsc.gov.uk/alerts/hostile-state-actors-compromising-uk-organisations-focus-engineering-and-industrial-control. [Accessed 14 November 2018].

[26] A. Braun, A. Broker and A. Hell, "BSI warnt vor Angriffen auf Stromnetze," TagesSchau, 13 June 2018. [Online]. Available: https://www.tagesschau.de/inland/stromnetze-angriffe-101.html. [Accessed 04 December 2018].

[27] US-CERT, "Russian Government Cyber Activity Targeting Energy and Other Critical Infrastructure Sectors," US-CERT, 15 March 2018. [Online]. Available: https://www.us-cert.gov/ncas/alerts/TA18-074A. [Accessed 14 November 2018].

[28] US Department of Defense Joint Chiefs of Staff, "JP 2-01.3 Joint Intelligence Preparation of the Operational Environment," 21 May 2014. [Online]. Available: http://www.jcs.mil/Doctrine/Joint-Doctrine-Pubs/2-0-Intelligence-Series/. [Accessed 04 December 2018].

[29] P. Brangetto and M. A. Veenendall, "Influence Cyber Operations: The Use of Cyberattacks in Support of Influence Operations," *2016 8th International Conference on Cyber Conflict (CyCon),* pp. 113-126, 2016.

[30] J. Slowik, "Strategic Communication and Cyber Attacks," 07 November 2018. [Online]. Available: https://pylos.co/2018/11/07/strategic-communication-and-cyber-attacks/. [Accessed 14 November 2018].

[31] M. Russinovich and T. Garnier, "Sysmon v8.0," Microsoft, 21 May 2017. [Online]. Available: https://docs.microsoft.com/en-us/sysinternals/downloads/sysmon. [Accessed 04 December 2018].

[32] North American Electric Reliability Corporation, "CIP Standards," NERC, [Online]. Available: https://www.nerc.com/pa/Stand/Pages/CIPStandards.aspx. [Accessed 14 November 2018].

[33] European Union, "Directive (EU) 2016/1148 of the European Parliament and of the Council," European Union, 06 July 2016. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016L1148&from=EN. [Accessed 04 December 2018].

[34] US Department of Homeland Security, "Automated Indicator Sharing (AIS)," US Department of Homeland Security, 21 June 2016. [Online]. Available: https://www.dhs.gov/ais. [Accessed 04 December 2018].

[35] European Union Agency for Network and Information Security, "Detect, SHARE, Protect: Solutions for Improving Threat Data Exchange among CERTs," European Union, 01 October 2013. [Online]. Available:

https://www.enisa.europa.eu/publications/detect-share-protect-solutions-for-improving-threat-data-exchange-among-certs/at_download/fullReport. [Accessed 04 December 2018].

[36] D. Kerr and W. Gibb, "OpenIOC Series: Investigating with Indicators of Compromise (IOCs)," FireEye, 16 December 2013. [Online]. Available: https://www.fireeye.com/blog/threat-research/2013/12/openioc-series-investigating-indicators-compromise-iocs.html. [Accessed 04 December 2018].

[37] D. Dittrich and K. Carpenter, "Misunderstanding Indicators of Compromise," Threatpost, 21 April 2016. [Online]. Available: https://threatpost.com/misunderstanding-indicators-of-compromise/117560/. [Accessed 04 December 2018].

[38] J. Slowik, "Indicators and Network Defense," 16 May 2018. [Online]. Available: https://pylos.co/2018/05/16/indicators-and-network-defense/. [Accessed 04 December 2018].

[39] J. Slowik, "Indicators and ICS Network Defense," Dragos, 31 May 2018. [Online]. Available: https://dragos.com/blog/20180531IndicatorsICSNetworkDefense.html. [Accessed 04 December 2018].

[40] S. Herpig, "As Germany Moves Toward a More Offensive Posture in Cyberspace, It Will Need a Vulnerability Equities Process," Council on Foreign Relations, 04 September 2018. [Online]. Available: https://www.cfr.org/blog/germany-moves-toward-more-offensive-posture-cyberspace-it-will-need-vulnerability-equities. [Accessed 04 December 2018].

[41] D. Bianco, "The Pyramid of Pain," 01 March 2013. [Online]. Available: http://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html. [Accessed 05 December 2018].

[42] P. Bhatt, E. T. Yano and P. M. Gustavsson, "Toward a Framework to Detect Multi-Stage Advanced Persistent Threats," in *2014 IEEE 8th International Symposium on Service Oriented System Engineering*, Oxford, UK, 2014.

[43] S. Caltagirone and R. M. Lee, "The Four Types of Threat Detection: With Case-Studies in Industrial Control Systems (ICS)," 13 July 2018. [Online]. Available: https://dragos.com/media/The_Four_Types%20of_Threat_Detection.pdf. [Accessed 05 December 2018].