

Using Global Honeypot Networks to Detect Targeted ICS Attacks

Michael Dodson

PhD Candidate
Department of Computer Science and
Technology
University of Cambridge
Cambridge, United Kingdom
md403@cam.ac.uk

Alastair R. Beresford

Professor of Computer Security
Department of Computer Science and
Technology
University of Cambridge
Cambridge, United Kingdom
arb33@cam.ac.uk

Mikael Vingaard

Industrial Security Researcher
Industrial Defenica
SecuriOT
Copenhagen, Denmark
info@honeypot.dk

Abstract: Defending industrial control systems (ICS) in the cyber domain is both helped and hindered by bespoke systems integrating heterogeneous devices for unique purposes. Because of this fragmentation, observed attacks against ICS have been targeted and skilled, making them difficult to identify prior to initiation. Furthermore, organisations may be hesitant to share business-sensitive details of an intrusion that would otherwise assist the security community.

In this work, we present the largest study of high-interaction ICS honeypots to date and demonstrate that a network of internet-connected honeypots can be used to identify and profile targeted ICS attacks. Our study relies on a network of 120 high-interaction honeypots in 22 countries that mimic programmable logic controllers and remote terminal units. We provide a detailed analysis of 80,000 interactions over 13 months, of which only nine made malicious use of an industrial protocol. Malicious interactions included denial of service and replay attacks that manipulated

logic, leveraged protocol implementation gaps and exploited buffer overflows. While the yield was small, the impact was high, as these were skilled, targeted exploits previously unknown to the ICS community.

By comparison with other ICS honeypot studies, we demonstrate that high-quality deception over long periods is necessary for such a honeypot network to be effective. As part of this argument, we discuss the accidental and intentional reasons why an internet-connected honeypot might be targeted. We also provide recommendations for effective, strategic use of such networks.

Keywords: *honeypot, industrial control system, ICS*

1. INTRODUCTION

Industrial Control Systems (ICS) are used to command, manage, or regulate devices or physical systems in industry (e.g., chemical processing), infrastructure (e.g., power generation), and building automation (e.g., fire suppression). Devices communicate using ICS-specific protocols, most of which are legacy point-to-point or broadcast protocols designed with the assumption that devices are connected with dedicated cabling; however, many of these protocols are now layered on top of ethernet and TCP or UDP, and devices use existing IP-based networks, including the internet, to communicate.

ICS security has not kept up with this growing digitisation and connectivity. The proprietary nature of most industrial software and the relatively low profile of industrial devices result in limited vulnerability hunting and disclosure [1] – [3]. For example, all versions of the two most popular proprietary (VxWorks) and open-source (FreeRTOS) real-time operating systems (RTOSes) have a total of 54 entries in the National Vulnerability Database (NVD) at the time of writing, compared with over 2,000 records for Windows 10 and over 800 records for Ubuntu 18.04. Further, all ‘critical’ VxWorks vulnerabilities in the NVD came from a single disclosure. Similarly, all but two of the FreeRTOS vulnerabilities came from a single disclosure. In each case, security researchers found more than 10 vulnerabilities that allowed remote code execution, data leakage, and denial of service attacks. Most were memory safety vulnerabilities and had existed in the software for more than a decade. Because these RTOSes are highly configurable, it is hard to estimate the number of affected devices; however, it is likely to exceed two billion [1], [4]. For comparison, the initial install target for Windows 10 was only one billion devices [5]. Further, when vulnerabilities

are identified, the industrial community demonstrates a strong resistance to patching, partly due to the high cost of regression testing and recertification by both the vendor and user [6]. Additionally, industrial networks have limited host-based security or logging opportunities, complicating forensic efforts. Even when forensic examination is possible, industrial network compromises are generally business-sensitive, so post-exploit forensic efforts rarely result in public disclosure of vulnerabilities, though ICS security companies often publish summary reports, such as those for Triton/Trisis [7]. Finally, few industrial protocols employ authentication or encryption; therefore, ICS devices will consider any well-formed packet to be valid, including those that request information or command changes of state [8], allowing malicious manipulation of device behaviour without actually exploiting any specific vulnerability. Together, these factors result in a vulnerable industrial environment and create unique security challenges.

Successful attacks against ICS have all targeted specific organisations and devices (e.g., Stuxnet [9], Triton/Trisis [7], CRASHOVERRIDE [10]) or have targeted vendors directly (e.g., [11]); therefore, unlike other domains where attacks are large-scale and indiscriminate, such as the Internet of Things (IoT) domain, there are limited means for researchers to gather open-source intelligence on ICS attack methods, motivations and campaigns. In domains such as IoT, honeypots have been effective tools to track and profile malicious behaviour [12], but they rely on either indiscriminate or easily deceived attackers, neither of which apply to current ICS adversaries. To date, the use of ICS honeypots for security research has been largely limited to monitoring internet-wide scanning.

Despite these challenges, we show that a geographically distributed network of high-interaction ICS honeypots can be an effective tool for identifying and profiling new, targeted attacks against ICS devices. We make the following contributions in this paper:

- A description of the largest, high-interaction ICS honeypot study to date.
- A discussion of multiple, new ICS exploits (zero days) identified by the honeypot network.
- An assessment of the growing overlap between ICS and IoT-aware scanning and botnet infections.
- An explanation of the limitations of previous ICS honeypot studies and recommendations for successful networks of ICS honeypots for security research.

2. BACKGROUND

A. Honeypots

Honeypots are computer security systems that emulate production systems and either decoy attackers away from the production system, provide warning of an intrusion, or allow attacker behaviour to be studied [12], [13]. Honeypots have been designed to emulate individual computers, such as laptops, servers, IoT and ICS devices [12], [14], and larger systems, such as electrical substations [15]. As a security device, they can be used as part of a defence-in-depth strategy alongside anti-virus software, segmented networks and firewalls. As a research tool, they are often used as stand-alone devices directly connected to the internet.

Honeypots can be characterised by their purpose and level of interaction [12]. The *purpose of interaction* refers to whether the honeypot is part of a production system, designed as part of a security solution for a given network or device, or a research device designed to attract attackers and study their behaviour [16]. The *level of interaction* refers to how well the honeypot emulates the target device, which determines how easy it is for the attacker to identify that they are interacting with a honeypot. The level of interaction is generally categorised as low, medium, or high, though these categories are not well-defined. A low-interaction honeypot may be a simple script that only emulates a login screen but no stateful device behaviour. A high-interaction honeypot may be an actual device or system, not an emulation, which is instrumented to record details of attacker behaviour on the system [17], [18].

Because honeypots have no purpose on a network except to deceive potential attackers, any interaction by an attacker with such a honeypot demonstrates that the attacker either lacks knowledge or is indiscriminate. If an attacker has sufficient knowledge and a specific target, then they can interact directly with the target device on a network and leave any honeypots untouched. If the attacker has less knowledge, but still has a specific target, they may have to scan a network to find the target device. In this case, they will interact with the honeypot and notify the defender of the attacker's presence, even if the attacker is able to avoid further interaction with the honeypot. In a less discriminate scenario, where the attacker is looking for any vulnerable device, they may go further and continue to interact with the honeypot, attempting to exploit vulnerabilities. Therefore, internet-connected, research honeypots have been effectively used to detect and monitor large-scale, indiscriminate attacks [12], but not knowledgeable, targeted attacks [19], [20].

Within the ICS community, there are several, open-source honeypots available. Conpot is a low-interaction honeypot capable of responding accurately to network scans [14]. It is easy to set up and scales well, making it a good candidate to research internet-

wide scanning [19] – [21]; however, its inability to interact with an attacker limits its utility in detecting and characterising ICS attacks, and studies using Conpot have yet to identify any new or targeted ICS attacks [19] – [21]. MiniCPS is a framework for higher-interaction honeypots and runs actual programmable logic [22]; however, it has yet to be used in a study to detect previously unknown ICS attacks, and its hardware emulation may be detectable by a capable attacker [12]. We provide a comparison of several ICS honeypot studies against our own in Section 4.

B. Targeted ICS Attacks

Most, if not all, successful attacks against ICS have been targeted, in that the attackers wish to create adverse physical effects in a specific organisation, and they knowledgeably target specific devices. Examples include Stuxnet attacks against Siemens PLCs [9]; Triton/Trisis attacks against specific models of Schneider Electric’s Triconex Safety Instrumented System [7]; and CRASHOVERRIDE attacks against the Ukrainian power grid [10]. The targeted and highly resourced nature of these attacks complicates efforts to identify and track real-world ICS exploitation, as the number of attacks is limited, and attackers have the ability and motivation to limit their exposure. As a result, ICS honeypot studies to date have not identified any attempt to maliciously modify ICS behaviour, nor have they been effectively used to disclose new ICS exploits to the community.

C. Large-scale ICS Attacks

Researchers have demonstrated scalable, proof-of-concept malware for PLCs that modifies programmable logic and automatically spreads to other devices (e.g., to create a botnet or to demand a ransom) [8], [23]. To date, no such large-scale, indiscriminate ICS malware has been observed in the wild. Furthermore, while a decade of security research has demonstrated that tens of thousands of vulnerable ICS devices are directly connected to the internet [21], [24], there has been little evidence of malicious attempts to modify the behaviour of such devices.

The lack of criminal or other large-scale malicious interest in vulnerable ICS devices can be attributed to several economic factors:

- High cost of entry: The cost of hardware for development and testing and the time to gain sufficient knowledge and experience to exploit such devices are significantly higher than in other domains (e.g., IoT).
- Fragmented population: While there may be over 100,000 internet-connected ICS devices, the population is divided amongst dozens of manufacturers running proprietary or bare-metal software on different chipsets.

- Limited resources: ICS devices have limited compute and memory resources, making them poor hosts for resource-intensive tasks such as cryptomining, and they are unlikely to store sensitive information typically used in ransomware attacks. Limited resources and proprietary software make general computing malware unlikely to succeed on ICS devices.

These economic factors are changing as industry seeks new ways to use digital technology. Industry 4.0 and Industrial IoT (IIoT) are converging with the IoT domain [25], creating a larger, more homogeneous environment of low-cost devices with general purpose compute and memory resources. In short, these changes are expected to overcome the economic factors currently inhibiting large-scale malicious interest in the ICS domain. As IIoT and IoT converge and industrial environments become increasingly attractive to cybercriminals and others looking to exploit devices at scale, ICS honeypots will be effective tools to identify and profile these attacks, as they are currently within the IoT domain.

3. SECURIOT DECEPTION TECHNOLOGY

A. ICS Honeypots

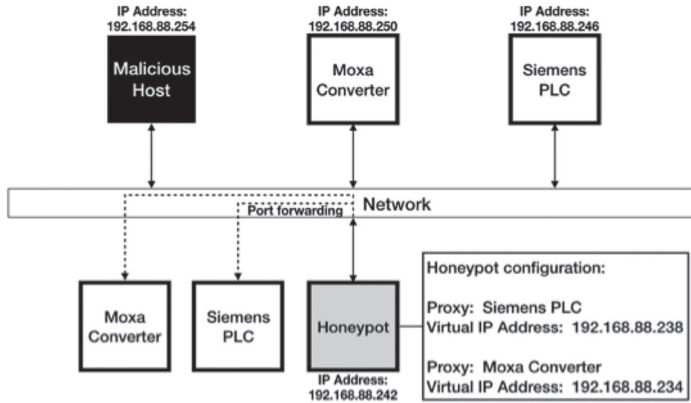
Previous ICS honeypot studies were limited in ways that reduced the likelihood of an attacker being deceived into interacting with the honeypot, such as geographic concentration, the use of cloud hosts, the use of low-interaction honeypots, and short study durations. In this paper, we demonstrate that these limitations can be overcome, showing that a sufficiently-sized, internet-connected ICS honeypot network can be effective in detecting and monitoring previously unknown, targeted attacks.

B. SecuriOT Honeypots

Low-interaction honeypots can be inexpensively deployed at scale, but they are easy to identify. Further, because they do not emulate device state, they cannot be used to profile an attacker's behaviour (e.g., attempts to modify programmable logic). High-interaction honeypots overcome these limitations, but can be expensive to develop, deploy, and maintain. To address the limitations of both low- and high-interaction honeypots, SecuriOT developed a reconfigurable device that supports multiple interaction levels with a common interface and management framework [26]. The device can be configured with templates to emulate an ICS device for low-interaction contexts, like Conpot [14], but can also act as a proxy to a production device. When acting as a proxy, the honeypot redirects traffic to a production device and acts as a man-in-the-middle between the network and the device. This proxy mode allows an adversary to exercise the full behaviour of the target device while providing the honeypot's full logging and alert functionality.

As shown in Figure 1, each physical device is capable of hosting multiple virtual IP addresses and up to three templates simultaneously, allowing a single physical device to appear as multiple devices on a given network.

FIGURE 1: SAMPLE DEPLOYMENT OF A SECURIOT HONEYPOT, SHOWING THE ABILITY TO EMULATE MULTIPLE VIRTUAL DEVICES AND ACT AS A PROXY FOR A CUSTOM DEVICE.



Each physical honeypot interfaces with a Security Information and Event Management (SIEM) system, which logs interactions and raises alerts. Since the honeypot is passive and has no production function on the network, any interaction with a virtual device is suspicious, as it implies that a host is either scanning the network segment or directly interacting with the honeypot. The SIEM is also used to manage device configurations, allowing the honeypots to maintain consistent configurations with the production devices on the network.

C. SecuriOT ICS Honeypot Network

While SecuriOT’s ICS honeypots are primarily designed for installation in production systems, the ability to act as a proxy and simultaneously support multiple virtual devices makes them a good foundation for a network of research honeypots. As part of their own intelligence-gathering operation, SecuriOT runs a network of 120 such virtual honeypots with IP addresses geolocated in over 20 countries. Each virtual IP routes traffic to a honeypot acting as a proxy to a production ICS device. Devices include PLCs, RTUs and serial-to-ethernet converters from vendors such as Siemens, Moxa, and Phoenix Contact. The virtual honeypots are supported by up to 15 production devices communicating over the following protocol/port combinations: S7comm/102, BACnet/47808, SOAP/37215, IEC-104/2404, DNP3/20000, and Modbus/502. S7comm, IEC-104, DNP3, and Modbus are used in several industrial environments, including manufacturing, automation, and power and water utilities.

BACnet is used in large-scale building automation. SOAP on port 37215 is used for configuration and management of certain routers.

The honeypots perform full packet captures and SecuriOT performs post-processing, such as fingerprinting the tool used to interact with the honeypot (e.g., NMAP [27]), identifying campaigns, and classifying packets as either reconnaissance or exploitation. The result is a dataset with the fields shown in Table I.

TABLE I: FIELDS PROVIDED PER PACKET FROM SECURIOT’S HONEYPOT NETWORK

Field	Example	Field	Example
Date	2018-03-31	Source country	Japan
Time	06:33:49	Destination country	United States
Source IP address	[REDACTED]	Source AS number	AS63949
Source port	51667	Source AS name	Linode, LLC
Destination port	102	Scanning tool	ZMAP
Protocol	S7comm	Campaign	TA-VV
Packet action	Reconnaissance		

4. DATA ANALYSIS AND DISCUSSION

Our dataset consists of 13 months of packets captured between March 2018 and April 2019 from SecuriOT’s network of 120, globally-distributed, high-interaction ICS honeypots. The dataset consists of approximately 200,000 packets, which we group into approximately 80,000 interactions. In this section, we present our analysis of the data and discuss our findings. We start with a dataset overview, including a comparison with previous, similar surveys. We then demonstrate malicious use of industrial protocols and discuss the relationships between attackers and targets. We conclude with a demonstration of large-scale attacks against non-industrial protocols recorded by the honeypot network and present early evidence that the ICS domain is affected by malicious, large-scale interest in IoT.

A. Dataset Overview

Table II provides a summary of the interactions with the SecuriOT network of industrial honeypots over the period of observation. The data demonstrates the breadth of interest in internet-connected ICS devices: thousands of individual hosts (IP addresses) are scanning industrial protocols from dozens of Autonomous Systems (ASes) in dozens of countries.

We find that a majority of these interactions originate from well-known research scanners and are expected to be benign (e.g., Censys [28], Shodan [29]), which is consistent with previous observations [19] – [21]. Both SecuriOT and the Cambridge Cybercrime Centre (CCCC) [30] maintain lists of known scanners, against which source IP addresses were compared to generate the ‘Known scanners’ percentages in Table II. Similarly, Table II shows that a vast majority of interactions are initiated by well-known scanning tools, such as NMAP [27].

Following previous studies, we classify multiple received packets from a given IP address as part of a single ‘interaction’. Comparing interactions rather than packets is preferable because the number of packets required to perform a given task can vary for different scanning tools and protocols. We define an interaction as a single scanning or exploitation event. For example, the Siemens module from the ZGrab scanner sends about 12 packets to each scanned IP address, while scanning a single port with ZMap only sends two packets (TCP SYN and RST) to each scanned IP address [21], [31]. Each of these would be considered one interaction.

TABLE II: SUMMARY OF INTERACTIONS WITH SECURIOT’S NETWORK OF INDUSTRIAL HONEYPOTS

Protocol/Port	Total packets	Related Interactions	Source IP addresses	Source ASes	Source countries	Known scanners	Known tools
Modbus/502	54,682	18,980	1,321	91	31	69.5%	99.9%
BACnet/47808	50,276	20,097	1,073	35	16	84.7%	100.0%
S7comm/102	43,203	18,422	998	85	30	49.9%	99.5%
DNP3/20000	32,534	13,283	1,040	124	42	39.1%	99.9%
SOAP/37215	12,975	7,403	337	85	29	0.0%	51.2%
IEC-104/2404	8,797	3,404	214	162	23	7.0%	99.6%

B. Comparison with Earlier Studies

Different studies use different methodologies and focus on different protocols; therefore, direct comparison is challenging. Even surveys covering the same timeframe but using different methodologies can produce different results (e.g., network telescopes versus honeypots [21]). We approach such comparisons with caution, and only draw qualitative conclusions. We selected studies for comparison for the following reasons: Mirian *et al.* is regularly used for comparison in other studies [21]; Ferretti *et al.* is a more recent study of similar size to Mirian *et al.* and has a global scope [19]; and Cabana *et al.* is the largest, low-interaction ICS honeypot study in the literature [20]. Notably, all three of these studies use low-interaction honeypots, whereas our

study uses high-interaction honeypots. There is no comparable survey in the academic literature of a large-scale, high-interaction ICS honeypot network.

Table III compares these surveys and data collection methods, showing broad agreement in the observed scanning frequency against each protocol. Table III also demonstrates that ranking based on interactions results in a different ordering than ranking based on packets, as different protocols have different packet densities.

TABLE III: COMPARISON OF RANKED POPULARITY OF SCANNED INDUSTRIAL PROTOCOLS FROM MULTIPLE SURVEYS.

‘*’ indicates raw data was not available. ‘**’ indicates an estimate based on graphical data.

Source	Method	Dataset type	Dataset size	Ranked popularity					
SecuriOT	Honeypot	Packets	202,467	Modbus	BACnet	S7comm	DNP3	IEC-104	
SecuriOT	Honeypot	Interactions	81,589	BACnet	Modbus	S7comm	DNP3	IEC-104	
Mirian et al. [21]	Telescope	Packets	2,100	Modbus	BACnet	S7comm	DNP3	Ethernet/IP	
Mirian et al. [21]	Honeypot	Interactions	5,252	S7comm	Modbus	BACnet			
Ferretti et al. [19]	Honeypot	Packets	*	Modbus	BACnet	S7comm	Ethernet/IP	IEC-104	
Ferretti et al. [19]	Honeypot	Interactions	4,986	BACnet	Modbus	Ethernet/IP	S7comm	IEC-104	
Cabana et al. [20]	Telescope	Packets	197M**	BACnet	Modbus	DNP3	S7comm	Ethernet/IP	IEC-104

While the distribution of our scanning traffic is largely consistent with previous studies, the data shows both growth and asymmetry in the DNP3 scanning traffic that has not been previously identified or evaluated. Mirian *et al.* only identified 5.1% of network telescope traffic as targeting DNP3 in 2015 [21], whereas Cabana *et al.* observed over 22% of network telescope data targeting DNP3 in 2019 [20]. Similarly, over 16% of the interactions recorded by SecuriOT honeypots targeted the DNP3 protocol. Furthermore, as shown in Table II, while the total number of DNP3 interactions is only 70% of the number of Modbus interactions (13,283 vs. 18,980), the number of IP addresses scanning for DNP3 is nearly 80% of that of Modbus (1,040 vs. 1,321), and the number of ASes from which those IP addresses originate is 136% of those for Modbus (124 vs 91). This statistic is also reflected in the number of source countries in which those IP addresses are geolocated (42 vs. 31). The asymmetry is even more pronounced when comparing DNP3 with BACnet or S7comm. Despite the challenges in quantitative comparisons between studies, there is clear evidence from multiple studies demonstrating a wider, as well as a growing, interest in DNP3 compared to other industrial protocols.

C. Targeted Attacks via Industrial Protocols

SecuriOT’s analysis concludes that only 20 of the 200,000 captured packets make use of an industrial protocol with clear malicious intent. These 20 packets can be grouped into nine attack interactions, which are summarised in Table IV. Based on feedback

from vendors and vulnerability databases, four of the nine interactions represent previously unknown attacks, or zero days, and one represents the first documentation of a previously-identified proof-of-concept attack in the wild [32]. The attack types include denial of service (DoS) and command replay attacks.

TABLE IV: ATTACKS USING INDUSTRIAL PROTOCOLS. THE PACKET COUNT DOES NOT INCLUDE TRANSPORT LAYER HANDSHAKES (E.G., INITIAL SYN PACKET FOR PROTOCOLS LAYERED ON TCP).

Date	Source country	Destination country	Protocol	Attack type	Source AS number	Number of packets
2 Apr 2018	United States	China	IEC-104	DoS	AS394828	2
17 Apr 2018	China	Poland	IEC-104	DoS	AS4134	1
20 Apr 2018	Russia	United States	S7comm	Replay	AS60307	8
27 Jun 2018	Ukraine	China	IEC-104	DoS	AS15626	4
8 Aug 2018	Vietnam	France	S7comm	DoS	AS38731	1
8 Aug 2018	Vietnam	Lithuania	S7comm	DoS	AS38731	1
9 Aug 2018	Vietnam	Poland	Modbus	DoS	AS38731	1
9 Aug 2018	Vietnam	France	Modbus	DoS	AS38731	1
19 Nov 2018	Seychelles	Czech Republic	Modbus	DoS	AS29073	1

The DoS attacks took several forms. In one case, a specially crafted packet forced a device to violate its real-time constraints, providing a low-bandwidth DoS attack on the process control. In another case, the attack targeted devices with incomplete implementations of the protocol stack; the attack provided valid, but unimplemented commands, and adversely affected the device’s process control. The attacker specifically targeted vulnerable device types, so this was not a case of accidental DoS. In a third case, a buffer overflow affected the device’s network communication capability, but did not affect the device’s process control.

Since many industrial protocols lack authentication or encryption, the receipt of any packet with a parsable command may be considered valid. In some cases, though, manufacturers have implemented protection to prevent a replay of previous commands or commands recorded in a test environment. The replay attack identified by SecuriOT was successful against a device for which the manufacturer claimed replay protection.

For most of the attacks, the source IP address was only active for the attack itself; the honeypot network had no record of other interactions from that IP address. This is not unexpected: an attacker may use one or multiple IP addresses for reconnaissance and then use a fresh IP address for the actual attack, to avoid blacklists. For three of

the attacks, however, consistent activity was observed from the source IP address. Specifically, the IP addresses used for the attacks originating in Vietnam, Ukraine, and the Seychelles performed regular scanning over the entire study duration.

These vulnerabilities and associated exploits were responsibly disclosed by SecuriOT to the device manufacturers, and public disclosure is currently being negotiated. The relationship between these vendors and SecuriOT precludes further public disclosure of the vulnerabilities at this time, but additional details may be obtained in some cases directly from SecuriOT.

D. Large-scale Attacks via Industrial Protocols

SecuriOT's honeypots also exposed a non-industrial protocol port and captured data associated with the Okiru-Satori variant of the Mirai botnet, which is indiscriminate and targets any vulnerable device across any network to which an infected device is connected.

While Okiru-Satori does not target industrial protocols, the convergence of IIoT and IoT domains may result in industrial devices being included in large-scale, non-industrial attacks. This is already the case for Windows-based industrial infrastructure. For example, the ransomware attack against the Windows-based infrastructure at Norsk Hydro in early 2019 prevented the safe and effective use of industrial devices [33]. As IIoT devices incorporate common operating systems with general purpose processing (e.g., Linux-based Azure Sphere [34]), they are more likely to become inadvertent victims of large-scale botnet or ransomware attacks targeting the IoT population. In this section, we discuss interactions with Mirai hosts and show that overlap already exists with industrial protocol scanners.

The Mirai botnet emerged in 2016 and used aggressive scanning and brute force password searches to infect hundreds of thousands of Linux-based IoT devices. At its peak, an estimated 600,000 hosts were infected [35]. At the time of writing, the CCCC [30] observes approximately 150,000 infected hosts per day scanning IP addresses in a monitored /14 network. The scanning packet used by Mirai is distinctive, allowing the CCCC to identify suspected Mirai hosts and record data such as the source and destination IP addresses and port numbers.

Many Mirai variants emerged after the public release of the Mirai source code. Variants target different device types and architectures and exploit different vulnerabilities. The Okiru-Satori variant was identified in 2017 and targeted Huawei routers on port 37215 using a previously unidentified vulnerability (CVE-2017-17215) [36]. As shown in Table V, SecuriOT's honeypots recorded 7,403 interactions from 337 IP addresses on port 37215. Of these, SecuriOT identified 222 malicious interactions, based on

attempts to brute force passwords, make use of the vulnerabilities exploited by Okiru-Satori, or modify firmware. While the malicious packets make up more than 30% of the total traffic on port 37215, only 3.0% of the total interactions are malicious, as password searches and firmware downloads necessarily require more packets than scanning.

TABLE V: SUMMARY OF INTERACTIONS ON PORT 37215.

	Packets	Interactions	Source IP Addresses	Source ASes	Dates of interaction
Overall	12,975	7,403	337	85	266
Malicious	3,919	222	13	2	15

Notably, while the scanning of port 37215 was recorded on 266 days, the honeypots were only configured as vulnerable routers over short periods in April and July 2018, resulting in only 15 days of malicious interactions. As discussed below, some of the apparently benign scanning might have transitioned to exploitation had the scanner found the honeypot in a vulnerable configuration.

To study the overlap between Mirai hosts and hosts aware of industrial protocols, we combined the CCCC database of suspected Mirai hosts [30] with SecuriOT’s honeypot data, correlating source IP addresses and interaction dates. Table VI summarises the results of this comparison from the perspective of the SecuriOT honeypots. For example, the first row should be interpreted as 792 packets received by SecuriOT honeypots on port 37215 from 26 IP addresses that the CCCC suspected to be hosting Mirai on the day of the interaction with the honeypot.

TABLE VI: SECURIOT HONEYPOT DATA CORRESPONDING TO SOURCE IP ADDRESSES AND DATES FROM THE CCCC MIRAI HOST DATASET.

Protocol/Port	Total packets	Related Interactions	Source IP addresses	Source ASes	Dates
SOAP/37215	792	789	26	11	40
DNP3/20000	116	71	4	2	4
BACnet/47808	2	2	1	1	1
Modbus/502	1	1	1	1	1

Comparing 789 SOAP/37215 interactions in Table VI with 222 malicious interactions in Table V demonstrates that the CCCC suspects many of the benign interactions with SecuriOT honeypots to have originated from Mirai hosts that simply did not find the SecuriOT honeypot to be vulnerable. This is consistent with the knowledge that the SecuriOT honeypots were only configured as vulnerable routers during limited periods.

Table VI also shows that the CCCC suspects 74 industrial protocol interactions (i.e., over DNP3, BACnet and Modbus) with SecuriOT honeypots to have originated from IP addresses hosting Mirai. As there is no known variant of Mirai that targets ICS devices, the scanning traffic by Mirai hosts against industrial protocols implies either that these scanners share an IP address with a Mirai host (e.g., a scanner behind an infected router) or that the scanner uses a similar technique to that employed by Mirai, though we are not aware of any such benign, internet-wide scanners.

This overlap between SecuriOT's honeypot data and the CCCC Mirai database, though limited, suggests that the gap between ICS-aware and IoT-aware hosts is narrowing.

5. RECOMMENDATIONS FOR HONEYPOT NETWORKS

SecuriOT's honeypot network exposed four zero-day attacks against devices running common ICS protocols, such as S7comm and Modbus. By comparing our study with previous studies that did not identify similar exploits (e.g., [19] – [21]), we provide the following recommendations for deploying networks of ICS honeypots for security research:

- Honeypot networks should be geographically dispersed. We identified nine attacks against devices in six countries, and none of the attacks originated in the same country as the target. Several honeypot studies located most or all targets in the United States [19], [21]; however, of our nine identified attacks, only one target was located in the United States.
- Honeypots should be hosted at realistic IP addresses. Several previous ICS honeypot studies used AWS or other cloud providers to host honeypots [19] – [21]. ICS devices are unlikely to be connected via a cloud service provider, so the use of AWS or similar is a red flag to an attacker.
- Honeypots should be high-interaction. Low-interaction honeypots can often be fingerprinted and generally do not allow an attacker to interact with the device beyond the initial login screen or protocol handshake. To deceive targeted attackers and understand their intentions (e.g., modifying firmware or programmable logic), high-interaction honeypots are necessary.

- Honeypot use should be systematic and continuous. This provides both authenticity and a larger window for an attacker to identify and target a given honeypot. Unlike large-scale attacks scanning for any vulnerable device, targeted attackers are looking for specific devices and may take considerable time before accidentally targeting a honeypot.

6. CONCLUSIONS

We have demonstrated that a network of high-interaction honeypots can identify and profile previously unknown, targeted ICS attacks. Specifically, we exposed four zero-day attacks against devices running common ICS protocols such as S7comm and Modbus, which were disclosed to the applicable manufacturers.

We also demonstrated that the gap between ICS-aware and IoT-aware hosts is narrowing, showing that IoT malware is co-located with ICS devices and scanners. Bridging this gap is the first major hurdle in attacking ICS devices at scale. Thus far, ICS devices have not been subjected to indiscriminate targeting, but the convergence of IIoT and IoT domains will make industrial devices more attractive targets, even if only as a vulnerable sub-population amongst the growing IoT population.

Finally, we discussed the limitations of previous ICS honeypot studies and provided recommendations for developing effective ICS honeypot networks as intelligence-gathering tools.

ACKNOWLEDGEMENTS

Michael Dodson is supported by a scholarship from the Gates Cambridge Trust; Alastair R. Beresford is partially supported by EPSRC [grant number EP/M020320/1]. The opinions, findings, and conclusions or recommendations expressed are those of the authors and do not necessarily reflect those of any of the funders. We thank Alexander Vetterl for his feedback and patient sharing of his expertise.

REFERENCES

- [1] B. Seri, G. Vishnepolsky, and D. Zusman, "URGENT/11 technical white paper," *Armis* [Online]. Available: <https://go.armis.com/urgent11>. [Accessed: 25-Nov-2019].
- [2] A. Nochvay, "Security research: CODESYS runtime, a PLC control framework," Kaspersky ICS CERT, 2019 [Online]. Available: <https://perma.cc/325P-N7AV>. [Accessed: 08-Nov-2019].
- [3] Cybersecurity and Infrastructure Security Agency, (2013) "3S CoDeSys vulnerabilities," [Online]. Available: <https://perma.cc/F8W4-7H75>. [Accessed: 28-Oct-2019].

- [4] O. Karliner, "FreeRTOS TCP/IP stack vulnerabilities," *Zimperium Mobile Security Blog*, 4 December 2018. [Online]. Available: <https://blog.zimperium.com/freertos-tcpip-stack-vulnerabilities-details/>. [Accessed: 18-Dec-2019].
- [5] C. Mihalcik, "Microsoft aims for 1 billion devices running Windows 10," *CNET*. [Online]. Available: <https://www.cnet.com/news/microsoft-aims-for-1-billion-devices-running-windows-10/>. [Accessed: 26-Feb-2020].
- [6] É. Leverett, R. Clayton, and R. Anderson, "Standardisation and certification of the 'Internet of Things'," Workshop on the Economics of Information Security (WEIS), 2017 [Online]. Available: <https://perma.cc/5Y9R-9DD3>.
- [7] Dragos Inc. "TRISIS malware: Analysis of safety system targeted malware". [Online]. Available: <https://perma.cc/K9EM-CABV>. [Accessed: 08-Nov-2019].
- [8] D. Formby, S. Durbha, and R. Beyah, "Out of control: Ransomware for industrial control systems," RSA Conference, 2017 [Online]. Available: <https://perma.cc/XD8V-LP5M>.
- [9] T. Chen and S. Abu-Nimeh, "Lessons from Stuxnet," *Computer*, vol. 44, no. 4, 2011, doi: 10.1109/MC.2011.115. [Online]. Available: <http://ieeexplore.ieee.org/document/5742014/>. [Accessed: 14-Jun-2019]
- [10] R. M. Lee, "CRASHOVERRIDE: Analyzing the malware that attacks power grids," *Dragos Inc.*, 2017 [Online]. Available: <https://dragos.com/resource/crashoverride-analyzing-the-malware-that-attacks-power-grids/>. [Accessed: 19-Jun-2019].
- [11] Cybersecurity and Infrastructure Security Agency, "Russian Government Cyber Activity Targeting Energy and Other Critical Infrastructure Sectors," [Online]. Available: <https://www.us-cert.gov/ncas/alerts/TA18-074A>. [Accessed: 26-Feb-2020].
- [12] A. Vetterl and R. Clayton, "Honeyware: A virtual honeypot framework for capturing CPE and IoT zero days," Symposium on Electronic Crime Research (eCrime), Pittsburgh, United States of America, 2019.
- [13] W. Martin, "Honey pots and honey nets - security through deception," *SANS Institute*, 2003 [Online]. Available: <https://www.sans.org/reading-room/whitepapers/attacking/honey-pots-honey-nets-security-deception-41>. [Accessed: 08-Apr-2019].
- [14] L. Rift, J. Vastergaard, D. Haslinger, A. Pasquale, and J. Smith, CONPOT ICS/SCADA honeypot.. [Online]. Available: <http://conpot.org>. [Accessed: 08-Apr-2019].
- [15] R. Rustici and I. Barak, "ICS threat broadens: Nation-state hackers are no longer the only game in town," *Cybereason*. [Online]. Available: <https://www.cybereason.com/blog/industrial-control-system-specialized-hackers>. [Accessed: 09-Dec-2019].
- [16] L. Spitzner, "The value of honeypots, part one: Definitions and values of honeypots," *Symantec*, 2001. [Online]. Available: <https://www.symantec.com/connect/articles/value-honeypots-part-one-definitions-and-values-honeypots>. [Accessed: 09-Dec-2019].
- [17] N. Provos and T. Holz, *Virtual honeypots: From botnet tracking to intrusion detection*. Pearson Education, 2007.
- [18] W. Fan, Z. Du, D. Fernández, and V. A. Villagrà, "Enabling an anatomic view to investigate honeypot systems: A survey," *IEEE Systems Journal*, vol. 12, no. 4, 2018, doi: 10.1109/JSYST.2017.2762161.
- [19] P. Ferretti, M. Pogliani, and S. Zanero, "Characterizing background noise in ICS traffic through a set of low interaction honeypots," *ACM Workshop on Cyber-Physical Systems Security & Privacy (CPS-SPC)*, 2019, doi: 10.1145/3338499.3357361.
- [20] O. Cabana, A. M. Youssef, M. Debbabi, B. Lebel, M. Kassouf, and B. L. Agba, "Detecting, fingerprinting and tracking reconnaissance campaigns targeting industrial control systems," International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment, 2019, doi: 10.1007/978-3-030-22038-9_5. [Online]. Available: http://link.springer.com/10.1007/978-3-030-22038-9_5. [Accessed: 09-Aug-2019].
- [21] A. Mirian *et al.*, "An Internet-wide view of ICS devices," *Conference on Privacy, Security and Trust (PST)*, Auckland, New Zealand, 2016, doi: 10.1109/PST.2016.7906943.
- [22] D. Antonioli, A. Agrawal, and N. O. Tippenhauer, "Towards high-interaction virtual ICS honeypots-in-a-box." presented at ACM Workshop on Cyber-Physical Systems Security and Privacy (CPS-SPC) Vienna, Austria, 2016, doi: 10.1145/2994487.2994493. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2994487.2994493>. [Accessed: 20-Dec-2019].
- [23] R. Spenneberg, M. Brüggemann, and H. Schwartke, "PLC-Blaster: A worm living solely in the PLC," presented at Black Hat Asia Marina Bay Sands, Singapore, 2016 [Online]. Available: <https://perma.cc/XWU5-TZ7L>. [Accessed: 28-Oct-2019].
- [24] É. P. Leverett, "Quantitatively assessing and visualising industrial system attack surfaces," *University of Cambridge MPhil Thesis*, 2011 [Online]. Available: <https://perma.cc/83Z9-Q5J9>. [Accessed: 26-Feb-2019].

- [25] M. Wollschlaeger, T. Sauter, and J. Jasperneite, "The future of industrial communication: Automation networks in the era of the internet of things and Industry 4.0," *IEEE Industrial Electronics Magazine*, 2017, doi: 10.1109/MIE.2017.2649104. [Online]. Available: <http://ieeexplore.ieee.org/document/7883994/>. [Accessed: 28-Oct-2019].
- [26] SecuriOT, "SecuriOT honeypot: Powered by Industrial Defenica," *SecuriOT Honeypot - Powered by Industrial Defenica*. [Online]. Available: <https://www.honeypot.dk>. [Accessed: 28-Dec-2019].
- [27] NMap Project, "Nmap: the network mapper," *NMap Project*. [Online]. Available: <https://nmap.org/>. [Accessed: 02-Dec-2019].
- [28] "Censys," *Censys*. [Online]. Available: <https://censys.io/>. [Accessed: 28-Dec-2019].
- [29] "Shodan," *Shodan*. [Online]. Available: <https://www.shodan.io/>. [Accessed: 28-Dec-2019].
- [30] Cambridge Cybercrime Centre, "Computer Laboratory: Cambridge Cybercrime Centre: Description of available datasets," Cambridge Cybercrime Centre [Online]. Available: <https://www.cambridgecybercrime.uk/datasets.html>. [Accessed: 01-May-2019].
- [31] "The ZMap project," *ZMap Project*. [Online]. Available: <https://zmap.io/>. [Accessed: 02-Dec-2019].
- [32] Cybersecurity and Infrastructure Security Agency, "PLC cycle time influences (Update A)". [Online]. Available: <https://www.us-cert.gov/ics/advisories/ICSA-19-106-03>. [Accessed: 29-Feb-2020].
- [33] Norsk Hydro, "Cyber-attack on Hydro," Norsk Hydro ASA, Oslo, Norway, 2019[Online]. Available: <https://www.hydro.com/en/media/on-the-agenda/cyber-attack/>. [Accessed: 14-Dec-2019].
- [34] Microsoft Azure, "Azure Sphere," Microsoft Corporation, Redmond, WA., [Online]. Available: <https://azure.microsoft.com/en-us/services/azure-sphere/>. [Accessed: 13-Jun-2019].
- [35] M. Antonakakis *et al.*, "Understanding the Mirai botnet," USENIX Security Symposium (USENIX Security), Vancouver, Canada, 2017 [Online]. <https://www.usenix.org/system/files/conference/usenixsecurity17/sec17-antonakakis.pdf>. [Accessed: 10-Jun-2019].
- [36] Check Point Research, "Huawei home routers in botnet recruitment," Check Point Software Technologies Inc., San Carlos, CA, Dec. 2017 [Online]. Available: <https://research.checkpoint.com/2017/good-zero-day-skiddie/>. [Accessed: 30-Nov-2019].