Honeypots for Automatic Network-Level Industrial Control System Security

Sam Maesschalck, Vasileios Giotsas, Benjamin Green, Nicholas Race

Lancaster University

Lancaster, United Kingdom

{s.maesschalck,v.giotsas,b.green2,n.race}@lancaster.ac.uk

Abstract

The proposed doctoral work investigates a new approach to implement, deploy and manage honeypots for Industrial Control Systems (ICS). Our goal is to address unique challenges of ICS security in terms of interactivity, resource utilization, timeliness of detection and uninterrupted operation, which are muchh stricter compared to traditional systems, making the existing approaches inefficient. Our proposal combines different levels of interactivity and coupling of the honeypots with the ICS network to satisfy trade-offs of detection accuracy and risk, and integrates the honeypot detection feeds with an SDN framework to enable autonomic reconfiguration.

1 Introduction

The increasing adoption of an Internet of Everything (IoE), encapsulates industrial technologies categorised under the umbrella term of Industrial Control Systems (ICS) [5]. Communication over IP for ICSs promises to improve functionality, manageability and ease of access. However, these systems were not designed with Internet connectivity in mind [1], made evident by the lack of security features in their associated network protocols [16]. These systems are often implemented as part of a nations critical infrastructure (e.g. water and electricity distribution [10]). Therefore, the security of ICS devices is paramount to the safety and economic prosperity of a nation. Although manufacturers are provided patches for known vulnerabilities, implementation times can be significantly higher when compared with traditional IT system, leaving them exposed for an extended period [8, 14]. The main reason for delayed patching can be traced to requirements aligned to continued operation (i.e. limited downtime).

EURODW 2020, April 27, 2020, Heraklion, GR

The three principle of safety, reliability, and availability, are historically favoured when compared to security [12], meaning ICS operators may prefer to leave systems unpatched than undergo downtime. However, where vulnerabilities exist, their exploitation has the potential to cause impact across the aforementioned principles. Moreover, when security features are available for ICS specific protocols, they are usually 'bolt-on' and can include vulnerabilities of their own [6].

The severity of operating unpatched ICS devices over the Internet is amplified by a growing number of known vulnerabilities, publicly available exploits, and their increasing level of sophistication [15]. In response to this, new methods to detect and mitigate attacks targeting ICSs need to be developed. It is no longer sufficient to trust traditional firewalls and anti-virus software, as they require updates in order to detect and block new forms of malicious traffic [4]. Consequently, zero-day exploits, which are not known to traditional security systems, are able to penetrate the network and infect systems. Through the introduction of "bring your own device" practices, and the prevalence of social-engineering [7], conventional perimeter defences my fail to adequately prevent initial access into a system [23].

2 Proposed System

Our proposed system aims to address the shortcomings of current industrial honeypot implementations, by combining a network of honeypots distributed across the Internet, with honeypots situated within the ICS network. Distributed honeypots are used to gather threat intelligence on commonly applied offensive security tactics and techniques, and their source. Together with blacklists, the data captured by these honeypots is used to calibrate firewalls and IDSs, and to ensure no malicious data captured by the honeypots can propagate inside the network. The internal honeypots are divided between two zones, an independent honeypot network, and the core operational network. The independent honeypot network hosts several high-interaction honeypots, including ICS devices (e.g. PLCs, HMIs, etc.) and conventional IT systems (e.g. domain controllers, web and email servers, and clients). These honeypots are designed to lure attackers, acting as a deception technique to divert attention

away from operational systems [21], and should therefore be configured in a realistic manner. Data gathered from these honeypots provides intelligence on current forms of malware and exploitation techniques [20]. Honeypots residing within the core operational network (i.e. in parallel to operational systems), will be used to capture live threats, affording administrators with the ability to not only respond according, but gain valuable insight into security deficiencies. We suggest these honeypots to be low-interaction to limit risks introduced through their use, since high-interaction honeypots can be taken over by malicious attackers and used to target operational systems [22]. The collected data will be automatically fed to a analytical server for further analysis using state-of-the-art machine learning techniques [24]. Finally, the analysed data can then be fed into an SDN controller to inform/reconfigure internal network structures to mitigate identified risk in an autonomic manner.

3 Work to be done

We are currently conducting research into the efficiency of honeypots to mimic ICS devices, and their ability to capture relevant threat intelligence. We have identified several honeypots that could be implemented in our framework, including Conpot and Dionaea. Conpot honeypots would be leveraged to gain intelligence on attacks that specifically target ICSs. Dionaea honeypots are designed to capture malware, supporting our aim to gather information on zero-day exploits and malware in the wild. As previously noted, we also implement high-interaction honeypots in order to gain a better understanding of attacker tactics and techniques. The implementation of high-interactive honeypots will not be limited to the honeypots themselves, but will also include a broader network for attackers to traverse, and allow for an enhanced level of interaction.

Data sets captured through the implementation of these honeypots will be used as inputs in a machine learning algorithm, designed to analyse and identify malicious behaviour. The machine learning algorithm will be trained with a comprehensive data set to achieve a low rate of false positives. The output of this algorithm will be used in conjunction with software-defined networking to automatically reconfigure internal operational networks to actively mitigate identified threats. Because of this, false positives could result in legitimate data being prohibited and interfere with operational processes, and due to the nature of ICS environments could lead to safety incidents. A balance must therefore be established between defensive actions and operational requirements.

Key challenges within this work include the appropriate deployment and configuration of the honeypots, making them a more attractive target compared to live operational systems. We must determine where in our network and on the internet can the most relevant data be captured. For example, what benefits are drawn through the use of company owned IP address-space over the address space of a third party. Furthermore identifying a ML algorithm that can deliver accurate outputs, with minimal false positives, is key. Related to this, we will be required to establish fundamental data types to feed into our selected algorithm from the honeypots. As the output of this algorithm will be fed into SDN to restructure network properties proportional to the risk posed, one approach that we will analyse is the use of intent-based networking.

4 Related Work

Several academic and research efforts have focused on network intrusion detection using honeypots [13, 18] However, most of these efforts usually rely on simulations, synthetic datasets, or non-industrial deployments, with little to no industry stakeholder based evaluations, and consequently, their traction within ICS has been limited. While there exist commercial security services for ICS deployments, such as Kapersky's Industrial CyberSecurity [9], these services have the risk of providing access to sensitive infrastructure and data to third parties, which may introduce additional vulnerabilities and contradict best practices on privacy and confidentiality. In addition, there have been concerns surrounding the use of such companies due to their ties with foreign governments [17].

Alata et al. [2] compared the malicious traffic attracted by a high-interaction honeypot to the traffic observed across distributed low-interaction honeypots. They found complementarity between these honeypot types suggesting their parallel use as the best security strategy. Serbanescu et al [19] analyzed traffic received in an ICS-specific honeypot in correlation with the honeypot visibility in the Shodan IP scanner. The volume of attacks increased considerably when the honeypot got indexed by Shodan, allowing to identify attackers targeting generic ICS devices. Internet traffic classification approaches have been evaluated by Kim et al. [11]. Notable results lie in the performance of supervised machine learning algorithms. SVM achieved an overall accuracy of 94.2% over ten traces, however where single-trace trained was applied, it achieved an accuracy of only 49.8% - 83.4%.

The closest work to ours is by Antonioli et al. [3], who proposed a high-interaction virtual ICS honeypot to satisfy the unique requirements of ICS honeypots compared to traditional networks, namely maintainability, low cost, timely detection and determinism. Their honeypot is connected to an SDN controller to enhance data analytics and detection. The performance of the honeypot was evaluated by six red teams during a Capture-The-Flag (CTF) competition. We plan to employ a similar evaluation to conduct a realistic assessment of the performance of our system.

References

[1] Irfan Ahmed, Sebastian Obermeier, Sneha Sudhakaran, and Vassil Roussev. 2017. Programmable Logic Controller Forensics. *IEEE Security and Privacy* 15, 6 (2017), 18–24. https://doi.org/10.1109/MSP.2017.4251102

[2] E. Alata, V. Nicomette, M. Kaâniche, M. Dacier, and M. Herrb. 2006. Lessons learned from the deployment of a high-interaction honeypot. *Proceedings - Sixth European Dependable Computing Conference, EDCC 2006* 22 (2006), 39–44. https://doi.org/10.1109/EDCC.2006.17

[3] Daniele Antonioli, Anand Agrawal, and Nils Ole Tippenhauer. 2016. Towards high-interaction virtual ICS honeypots-in-a-box. In *Proceedings of the 2nd ACM Workshop on Cyber-Physical Systems Security and Privacy*. 13–22.

[4] Leyla Bilge and Tudor Dumitras. 2012. Before we knew it: an empirical study of zero-day attacks in the real world. In *Proceedings of the 2012 ACM conference on Computer and communications security*. ACM, Raleigh, North Carolina, USA, 833–844. https://doi.org/10.1145/2382196.2382284

[5] Roland C Bodenheim. 2014. Impact of the Shodan Computer Search Engine on Internet-facing Industrial Control System Devices. (2014), 121. http://www.dtic.mil/cgi-bin/GetTRDoc?Location=U2{&}doc=GetTRDoc.pdf{&}AD=ADA601219

[6] J. Adam Crain and Sergey Bratus. 2015. Bolt-On Security Extensions for Industrial Control System Protocols: A Case Study of DNP3 SAv5. *IEEE Security & Privacy* 13, 3 (2015), 74–79. https://doi.org/10.1109/msp.2015.47

[7] Richard Derbyshire, Benjamin Green, Daniel Prince, Andreas Mauthe, and David Hutchison. 2018. An Analysis of Cyber Security Attack Taxonomies. *Proceedings - 3rd IEEE European Symposium on Security and Privacy Workshops, EURO S and PW 2018* (2018), 153–161. https://doi.org/10.1109/EuroSPW.2018.00028

[8] Debabrata Dey, Atanu Lahiri, and Guoying Zhang. 2015. Optimal policies for security patch management. *INFORMS Journal on Computing* 27, 3 (2015), 462–477. https://doi.org/10.1287/ijoc.2014.0638

[9] European Union Agency for Cybersecurity (ENISA). 2017. *Baseline Security Recommendations for IoT in the context of Critical Information Infrastructures*. Number November. https://doi.org/10.2824/03228

[10] Benjamin Green, Marina Krotofil, and Ali Abbasi. 2017. On the significance of process comprehension for conducting targeted ICS attacks. *CPS-SPC 2017 - Proceedings of the 2017 Workshop on Cyber-Physical Systems Security and PrivaCy, co-located with CCS 2017* (2017), 57–68. https://doi.org/10.1145/3140241.3140254

[11] Hyunchul Kim, K. C. Claffy, Marina Fomenkov, Dhiman Barman, Michalis Faloutsos, and Ki Young Lee. 2008. Internet traffic classification demystified: Myths, caveats, and the best practices. *Proceedings of 2008 ACM CoNEXT Conference - 4th International Conference on Emerging Networking EXperiments and Technologies, CoNEXT '08* (2008). https://doi.org/10.1145/1544012.1544023

[12] Leandros A. Maglaras, Ki Hyung Kim, Helge Janicke, Mohamed Amine Ferrag, Stylianos Rallis, Pavlina Fragkou, Athanasios Maglaras, and Tiago J. Cruz. 2018. Cyber security of critical infrastructures. *ICT Express* 4, 1 (2018), 42–45. https://doi.org/10.1016/j.icte.2018.02.001

[13] Abhishek Mairh, Debabrat Barik, Kanchan Verma, and Debasish Jena. 2011. Honeypot in network security: a survey. In *Proceedings of the 2011 international conference on communication, computing & security*. 600–605.

[14] Angelos K. Marnerides, Vasileios Giotsas, and Troy Mursch. 2019. Identifying infected energy systems in the wild. *e-Energy 2019 - Proceedings of the 10th ACM International Conference on Future Energy Systems* (2019), 263–267. https://doi.org/10.1145/3307772.3328305

[15] Angelos K. Marnerides, Daniel Prince, John Couzins, Ryan Mills, Vasileios Giotsas, Paul McEvatt, David Markham, and Catherine Irvine. 2019. *FUJITSU White paper: Cyber Threat Laboratory*. Technical Report. Fujitsu.

[16] Ariana Mirian, Zane Ma, David Adrian, Matthew Tischer, Thasphon Chuenchujit, Tim Yardley, Robin Berthier, Joshua Mason, Zakir Durumeric, J. Alex Halderman, and Michael Bailey. 2016. An Internet-wide view of ICS devices. *2016 14th Annual Conference on Privacy, Security and Trust, PST 2016* (2016), 96–103. https://doi.org/10.1109/PST.2016.7906943

[17] Ellen Nakashima. 2017. Why the U.S. government is moving to ban this Russian software company. https://www.washingtonpost.com/news/checkpoint/wp/2017/09/14/why-the-u-s-government-is-moving-to-ban-this-russian-software-company/

[18] Marcin Nawrocki, Matthias Wählisch, Thomas C Schmidt, Christian Keil, and Jochen Schönfelder. 2016. A survey on honeypot software and data analysis. *arXiv preprint arXiv:1608.06249* (2016).

[19] Alexandru Vlad Serbanescu, Sebastian Obermeier, and Der-Yeuan Yu. 2015. ICS threat analysis using a large-scale honeynet. In *3rd International Symposium for ICS & SCADA Cyber Security Research 2015 (ICS-CSR 2015) 3*. 20–30.

[20] Miroslaw Skrzewski. 2012. Network Malware Activity – A Viewfrom Honeypot Systems. In *Computer Networks*, A. Kwiecień, P. Gaj, and P Stera (Eds.). Springer, Berlin, 198–206.

[21] The Honeynet Project. 2001. *Know Your Enemy: Revealing the Security Tools, Tactics, and Motives of the Blackhat Community*. Addison-Wesley.

[22] Alexander Vetterl and Richard Clayton. 2018. Bitter Harvest: Systematically Fingerprinting Low-and Medium-interaction Honeypots at Internet Scale. *12th USENIX Workshop on Offensive Technologies (WOOT 18)* (2018). https://github.com/desaster/kippo

[23] Yong Wang, Jinpeng Wei, and Karthik Vangury. 2014. Bring your own device security issues and challenges. *2014 IEEE 11th Consumer Communications and Networking Conference, CCNC 2014* (2014), 80–85. https://doi.org/10.1109/CCNC.2014.6866552

[24] Ting Fang Yen and Michael K. Reiter. 2008. Traffic aggregation for malware detection. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 5137 LNCS (2008), 207–227. https://doi.org/10.1007/978-3-540-70542-0_11