



City Research Online

City St George's, University of London

Citation: Pasquale, L., Ramkumar, K., Cai, W., McCarthy, J., Doherty, G. & Nuseibeh, B. (2024). The Rocky Road to Sustainable Security. *IEEE Security & Privacy*, 22(5), pp. 82-86. doi: 10.1109/msec.2024.3429888

This is the published version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/37807/>

Link to published version: <https://doi.org/10.1109/msec.2024.3429888>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

The Rocky Road to Sustainable Security

Liliana Pasquale^{ID} and Kushal Ramkumar | Lero@University College Dublin

Wanling Cai | Lero@Trinity College Dublin

John McCarthy | Lero@University College Cork

Gavin Doherty | Lero@Trinity College Dublin

Bashar Nuseibeh | The Open University

In this column, we illustrate real-world scenarios in which modern systems cannot preserve security during operation. We examine the notion of sustainable security and discuss the challenges to engineering sustainably secure systems.

With software systems permeating our lives, we are entitled to expect that such systems are secure by design and can preserve security throughout their use and subsequent evolution. However, the complexity of cyber-physical systems can extend the attack surface in ways that cannot be foreseen at design time. Thus, it is difficult to ensure that systems are secure by design, and alternative ways to preserve security during system operation must be considered.

We argue that modern, cyberphysical, systems should preserve security during their operation. They should be capable of discovering changes that may bring unanticipated security threats (e.g., anomalies) and manage the evolution of security requirements and controls to mitigate such threats. Human intervention can

also be beneficial and sometimes essential to preserve security. For example, humans can monitor security-relevant data¹ or support the diagnosis of anomalies.²



©SHUTTERSTOCK.COM/DEEPADESIGNS

However, the success of such interventions depends on human attention and engagement, which are difficult to replenish once depleted. Thus, modern systems should be

sustainably secure: they should not only preserve security but also sustain human engagement and attention.

Why Security Cannot Be Preserved

Experience with real-world cyber-physical systems has demonstrated that it is difficult to preserve security due to vulnerable system configurations, unforeseen and third-party vulnerabilities, and the human factor.

Vulnerable system configurations: In a cyber-physical system, users could add potentially vulnerable devices during system operation. This can extend the attack surface in ways that were not expected at design time. For example, smart home users can plug in different devices during system operation (e.g., smart speakers or cameras). Even if the individual devices are secure by design, unauthorized access to one of these devices may harm other devices, the

Digital Object Identifier 10.1109/MSEC.2024.3429888
Date of current version: 13 September 2024

home, and its users. For example, in November 2018, a Google Nest camera was added to a household. A vulnerability in this device allowed an attacker to control the thermostat and turn up the temperature to 90 degrees.³

Unforeseen vulnerabilities: Newly discovered vulnerabilities can also arise. For example, an unforeseen vulnerability in Log4J⁴ allowed attackers to break into systems, steal passwords and logins, extract data, and infect networks with malicious software. Blind spots are another example of unforeseen vulnerabilities. They may be known to the security community and have fixes, but they occur due to developers' insufficient security knowledge and delayed updates. Another example of unforeseen vulnerabilities may arise when manufacturers stop shipping software updates for older devices.⁵ For example, most mobile phone vendors usually stop shipping patches for phones older than three years, making them highly vulnerable. This is also a problem for specialized

medical equipment, often controlled by obsolete computers.

Third-party vulnerabilities: Most systems rely on third-party software, which can evolve unexpectedly and introduce vulnerabilities outside the software provider or original equipment manufacturer's control to fix. Software supply chain attacks, like the infection of SolarWind's Orion platform,⁶ aim to inject malicious code into software components to compromise downstream users. The European Union Agency for Cybersecurity identified highly complex supply chains, particularly their software dependencies, as the number one threat in its forecast of cyber threats up to 2030.⁷

The human factor: People can use technology in unexpected ways that were not considered during design. For example, in a smart home, users can accidentally disclose sensitive information while using a smart speaker.⁸ Existing threat models rarely factor in how people can misuse technologies.⁹

Human intervention could mitigate some of these vulnerabilities. For example, a home tenant could be warned about a vulnerable device and temporarily disconnect it from other critical devices until the vulnerability is fixed. Security engineers could analyze data about anomalous behavior and network traffic of smart home devices to diagnose anomalies and identify appropriate security controls. Software engineers could be encouraged to develop more maintainable software systems, which will facilitate the implementation of security patches when such systems become outdated. User interfaces that provide notifications and nudges could motivate people to behave securely and avoid misuse.

An Adaptive View on Sustainable Security

Sustainable security is the cost-effective development and maintenance of "code that will go on working dependably for decades in environments that change and evolve."⁵ It

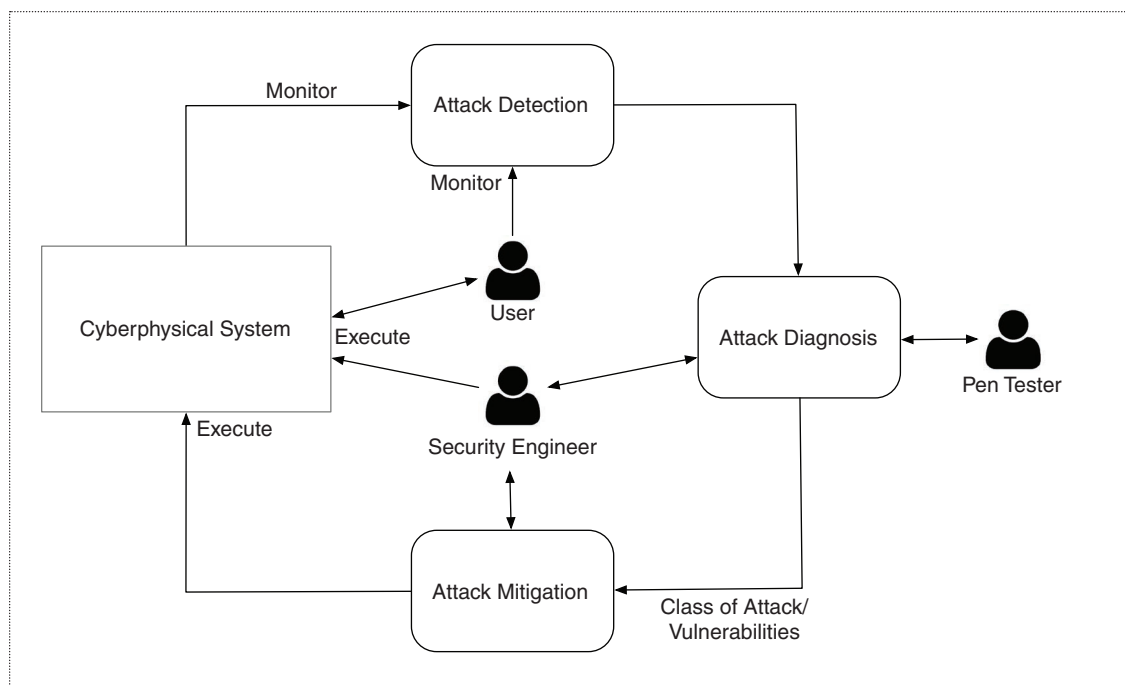


Figure 1. Our view on sustainable security.

has also been considered as the property of a network intrusion detection system to ensure continuous, reliable operation¹⁰ or the capability of organizational processes to be more resilient to insider threats.¹¹ Unlike previous work, we suggest considering sustainable security as the capability of a system to preserve security requirements throughout its use and evolution while sustaining human engagement and attention.

To be sustainably secure, systems must continuously detect, diagnose, and mitigate attacks by adapting their security controls (see Figure 1). *Attack detection* should

routers that has been exploited for a long time to deploy botnet attacks.

Attack detection, diagnosis, and mitigation cannot always be fully automated and should be supported by stakeholders, such as users, security engineers, and pen testers. For example, in a smart home, a user can complement the activities of a network anomaly detector to flag situations when a device does not exhibit instructed behavior (e.g., a smart speaker unexpectedly playing audible sounds), indicating the presence of an ultrasonic voice command attack. Users and security engineers can also execute security controls that cannot be

attack and identify the vulnerabilities exploited.

Human attention and engagement are critical to ensure the success of stakeholder interventions. However, maintaining human attention and engagement can be challenging, and once depleted, these resources are difficult to replenish. Also, stakeholders' security expertise, personality traits, experienced cognitive load, or motivation can influence their attention and engagement.¹² Thus, sustainable security requires preserving security and sustaining a positive experience for stakeholders interacting with the system. Although interactive security tools (e.g., Krüger et al.¹³) hold promise to sustain human attention, they cannot adapt their interactions and notifications depending on their stakeholders' personality, expertise, and cognitive capabilities.

Bridging the gap among attack detection, diagnosis, and mitigation is essential to address previously unknown attacks.

identify potentially unanticipated attacks using data monitored during system operation. For example, in a smart home, an anomaly detector can identify an abnormal packet rate to a device, indicating the suspected presence of an attack. *Attack diagnosis* should explain the attack by, for example, identifying the type of attack and the vulnerabilities it could exploit. For example, an abnormal packet rate can be linked to a DDoS or a botnet attack. *Attack mitigation* should use information about the type of attack to select short-term actions that stop harm (e.g., block traffic to or from malicious endpoints or reduce the traffic rate to or from a device). It should also use the information about the vulnerabilities exploited during the suspected attack to select long-term actions that fix the vulnerabilities and prevent the recurrence of the same attack. For example, CVE-2023-1389 is a vulnerability in the firmware of some Wi-Fi

automated. For example, users can update specific vulnerable devices. Security engineers can modify the network configuration for groups of households to prevent specific attacks. Moreover, users have agency in the physical world. They may also need to override the system due to requirements stronger than security, which cannot be incorporated into the system (e.g., personal safety and exceptional/emergency situations). Security engineers could support the diagnosis by identifying plausible attacks that could occur, considering the data monitored during system execution.

The discovery of the vulnerabilities exploited by an attack cannot always be performed automatically, so regular pen testing activities should be conducted for this purpose. For example, information about the type of attack obtained during the diagnosis can inform pen testers who could launch this

The Road Ahead

Several challenges should be addressed when engineering sustainably secure systems.

Coupling attack detection, diagnosis, and mitigation: Bridging the gap among attack detection, diagnosis, and mitigation is essential to address previously unknown attacks. Although anomaly detection techniques are effective for detecting previously unknown attacks, they cannot typically diagnose the type of attack and identify appropriate security controls automatically. Also, attack diagnosis techniques are not usually linked to automated approaches that can select security controls based on the attack diagnosis. Moreover, software vendors cannot always perform pen testing during system use, especially for smart homes. When performed, pen testing is not informed and configured based on the anomalies discovered during system operation and the diagnosed suspected attacks.

Managing evolving security requirements: Anomalies discovered from runtime data may not always indicate that an attack is underway, but they can suggest a change in the system configuration (added or removed components) without requiring a diagnosis of an attack. In such cases, it can be necessary to identify and manage evolving security requirements during system operation. For example, adding a vulnerable (outdated) system component to a smart home may require changing the system security requirements by, for example, changing the network topology to avoid communication among this device and other critical devices and appliances in the smart home. Logic-based learning, particularly inductive learning, identifying general rules from examples, can be useful for this purpose. For example, inductive learning has been successfully used to learn security policies from anomaly detection traces.¹⁴ Similarly, it can be used to reason about the security requirements that should be satisfied when system components change and trigger their evolution.

Sustaining human engagement: A human-machine collaborative approach would instill a culture of continuous improvement and resilience against evolving threats and attacks. An increased engagement of the stakeholders supporting attack detection, diagnosis, and mitigation can ultimately contribute to enacting effective security controls and increasing system security. To support the engagement of users, security engineers, and pen testers, there is a need to increase their situational awareness and design human-machine interactions supporting different levels of agency depending on the stakeholders' roles and expertise. This promising research direction requires multidisciplinary collaboration among

cybersecurity, AI, HCI, and psychology disciplines.

Measuring sustainable security: Finally, it is still unclear how to measure sustainable security. For example, taking inspiration from software maintenance metrics, we could evaluate sustainable security as a system's robustness to attacks (e.g., measuring the mean time to mitigate an attack or the mean time between attacks). From an ecological perspective, we could measure the reduced manual effort induced by attack detection, diagnosis, and mitigation activities. Since stakeholders' participation in the attack detection, diagnosis, and mitigation cycle is critical to preserve security, measuring their engagement and attention is also critical to assess sustainable security.¹⁵ Finally, as systems become increasingly socio-technical, representing complex human interactions with technological and social systems, we argue that sustainable security should also aim at improving the security experience of users (e.g., personalized security to individual needs and sustained engagement that improves user awareness of security). Thus, experiential trust, security, and usability should be explored to evaluate sustainable security. ■

Acknowledgment

This article was supported by Science Foundation Ireland Grant 13/RC/2094_2.

References

1. K. Yang, J. Ren, Y. Zhu, and W. Zhang, "Active learning for wireless IoT intrusion detection," *IEEE Wireless Commun.*, vol. 25, no. 6, pp. 19–25, Dec. 2018, doi: [10.1109/MWC.2017.1800079](https://doi.org/10.1109/MWC.2017.1800079).
2. A. K. Sikder, L. Babun, H. Aksu, and A. S. Uluagac, "AEGIS+: A context-aware platform-independent security framework for smart home systems," *Digital Threats Res. Practice*, vol. 2, no. 1, pp. 1–33, 2021, doi: [10.1145/3359789.3359840](https://doi.org/10.1145/3359789.3359840).

3. J. Maher. "Hacker takes over couple's smart home, plays vulgar music and raises temperature to 90 degrees." *Newsweek*. 2019. [Online]. Available: <https://www.newsweek.com/google-nest-hack-milwaukee-146>
4. D. Everson, C. Long, and Z. Zhang, "Log4shell: Redefining the web attack surface," in *Proc. Workshop Meas. Attacks Defenses Web (MAD-Web)*, 2022, pp. 1–8, doi: [10.14722/madweb.2022.23010](https://doi.org/10.14722/madweb.2022.23010).
5. R. Anderson, "Making security sustainable," *Commun. ACM*, vol. 61, no. 3, pp. 24–26, 2018, doi: [10.1145/3180485](https://doi.org/10.1145/3180485).
6. S. Peisert et al., "Perspectives on the SolarWinds incident," *IEEE Secur. Privacy*, vol. 19, no. 2, pp. 7–13, Apr. 2021, doi: [10.1109/MSEC.2021.3051235](https://doi.org/10.1109/MSEC.2021.3051235).
7. "Cybersecurity threats fast-forward 2030: Fasten your security-belt before the ride!", ENISA. 2022. [Online]. Available: <https://www.enisa.europa.eu/news/cybersecurity-threats-fast-forward-2030>
8. P. Moh, P. Datta, N. Warford, A. Bates, N. Malkin, and M. L. Mazurek, "Characterizing everyday misuse of smart home devices," in *Proc. IEEE Symp. Secur. Privacy (SP)*, San Francisco, CA, USA, 2023, pp. 2835–2849, doi: [10.1109/SP46215.2023.10179476](https://doi.org/10.1109/SP46215.2023.10179476).
9. D. Geneiatakis, I. Kounelis, R. Neisse, I. Nai-Fovino, G. Steri, and G. Baldini, "Security and privacy issues for an IoT based smart home," in *Proc. 40th Int. Conv. Inf. Commun. Technol., Electron. Microelectron. (MIPRO)*, Opatija, Croatia, 2017, pp. 1292–1297, doi: [10.23919/MIPRO.2017.7973622](https://doi.org/10.23919/MIPRO.2017.7973622).
10. G. Koutsandria, R. Gentz, M. Jamei, A. Scaglione, S. Peisert, and C. McParland, "A real-time testbed environment for cyber-physical security on the power grid," in *Proc. 1st ACM Workshop Cyber Phys. Syst. Secur. Privacy*, 2015, pp. 67–78, doi: [10.1145/2808705.2808707](https://doi.org/10.1145/2808705.2808707).

11. M. Hanse, “‘Insider threats’ and ‘supporting sustainable security’—Adding the dimension of privacy and data protection, organizational processes for supporting sustainable security (Dagstuhl Seminar 12501),” *Dagstuhl Rep.*, vol. 2, pp. 2192–5283, 2013, doi: [10.4230/DagRep.2.12.37](https://doi.org/10.4230/DagRep.2.12.37).
12. I. Rauf et al., “The case for adaptive security interventions,” *ACM Trans. Softw. Eng. Methodol.*, vol. 31, no. 1, 2022, Art. no. 9, doi: [10.1145/3471930](https://doi.org/10.1145/3471930).
13. S. Krüger et al., “CogniCrypt: Supporting developers in using cryptography,” in *Proc. 32nd IEEE/ACM Int. Conf. Automated Softw. Eng. (ASE)*, Urbana, IL, USA, 2017, pp. 931–936, doi: [10.1109/ASE.2017.8115707](https://doi.org/10.1109/ASE.2017.8115707).
14. A. Drozdov, M. Law, J. Lobo, A. Russo, and M. W. Don, “Online symbolic learning of policies for explainable security,” in *Proc. 3rd IEEE Int. Conf. Trust, Privacy Security Intell. Syst. Appl. (TPS-ISA)*, 2021, pp. 269–278, doi: [10.1109/TPSISA52974.2021.00030](https://doi.org/10.1109/TPSISA52974.2021.00030).
15. K. Doherty and G. Doherty, “Engagement in HCI: Conception, theory and measurement,” *ACM Comput. Surv.*, vol. 51, no. 5, pp. 1–39, 2018, doi: [10.1145/3234149](https://doi.org/10.1145/3234149).

Liliana Pasquale is an associate professor at University College Dublin and a funded investigator at Lero—The SFI Research Centre for Software, Dublin D04 V1W8, Ireland. Her research interests include requirements engineering and adaptive systems, focusing on security, privacy, and digital forensics. Pasquale received a Ph.D. in information and software technologies from Politecnico di Milano. She is an associate editor of *IEEE Transactions on Software Engineering*, a department editor of *IEEE Security & Privacy Magazine*, and a member of the review board of *ACM Transactions on Software*

Engineering and Methodology. Contact her at liliana.pasquale@ucd.ie.

Kushal Ramkumar is a Ph.D. student at University College Dublin and Lero—The Irish Research Centre for Software, Dublin D04V1W8, Ireland. His research interests include developing techniques to provide enduring security to cyberphysical systems such as smart homes and using logic-based learning techniques for threat diagnosis and mitigation. Ramkumar received a master of science in computer engineering from Illinois Institute of Technology. He was inducted into the IEEE Eta Kappa Nu Honors Society for exemplary academic performance during his master’s. Contact him at kushal.ramkumar@ucdconnect.ie.

Wanling Cai is a postdoctoral researcher at Trinity College Dublin and a researcher at Lero—The SFI Research Centre for Software, Dublin D02YY50, Ireland. Her research interests include human–computer interaction and human-centered AI, focusing on health technologies, recommender systems, and human-centered security and privacy. Cai received a Ph.D. in computer science from Hong Kong Baptist University. She has served on the program and organization committees of ACM conferences, such as UII, RecSys, UMAP, and MUM. Contact her at wanling.cai@tcd.ie.

John McCarthy is a professor of applied psychology at University College Cork and a principal investigator at Lero—the SFI Research Centre for Software, Cork T12CY82, Ireland. His research interests include understanding the influence of emerging social, personal, and work technologies on people’s lived experiences and using that understanding to inform design

of usable and enriching technologies. His recent work includes a project on Responsible Software Engineering and others at the core of which is Design for Care in Online Platforms. McCarthy received a Ph.D. in applied psychology from University College Cork. He has served on program committees for a number of conferences including CHI, DIS, ECCE, and COOP. Contact him at john.mccarthy@ucc.ie.

Gavin Doherty is a professor and leader of the Health Technology Design Group at the School of Computer Science and Statistics at Trinity College Dublin, Dublin D02YY50, Ireland. His research interests include supporting and extending the reach of mental health professionals, designing innovative and engaging systems that can be implemented in real-world clinical environments. Doherty received a doctorate in computer science from the University of York. He is a Distinguished Member of the ACM and is chair of the ACM Distinguished Speaker Committee. Contact him at gavin.doherty@tcd.ie.

Bashar Nuseibeh is a professor of computing and head of the Software Engineering and Design research group at The Open University, Milton Keynes MK7 6AA, U.K. His research interests include software requirements and design, engineering adaptive systems, and security and privacy. Nuseibeh received a Ph.D. in computer science from Imperial College London. He is a Fellow of The Royal Academy of Engineering and a Member of the Royal Irish Academy and Academia Europaea. Contact him at bashar.nuseibeh@open.ac.uk.