Workshop on Real-life Impacts of Security Vulnerabilities

18 April 2024

Website: https://zisc.ethz.ch/events/economics-security-workshop-2024/
Scribed Notes by Andrin Bertschi, Mark Kuhne, 17 June 2024

Event Summary

The workshop on Assessing the Real-World Impact of Security Vulnerabilities was held at ETH Zurich in April 2024, and aimed at exploring a pressing issue in cybersecurity: understanding and quantifying the true consequences of security flaws. The event was designed to bridge the gap between academic research and industry practices, fostering a multidisciplinary dialogue that includes technical, economic, and social perspectives.

In today's interconnected world, security vulnerabilities can have far-reaching implications, affecting everything from individual privacy to national security. As cyber threats continue to evolve in complexity and scale, it is imperative to develop robust models and methodologies for evaluating their impact.

The workshop aimed to initiate a conversation between different stakeholders to better assess the impact of vulnerabilities and explore tools and methodologies that could aid in this process. One key point included understanding the real-life impact of security vulnerabilities, using available tools and techniques, and learning from other fields that have faced similar challenges.

A primary outcome from both the presentations and panel discussions was the consensus on the need for an awareness campaign to mitigate the incentives for overclaiming impacts, which was identified as a main conclusion of the workshop.

Workshop Talks

Introduction by N. Asokan

Speaker: N. Asokan, University of Waterloo

Asokan opened the workshop by emphasizing the interdisciplinary nature of security vulnerabilities, involving not just technical but also economic factors. He highlighted that while many vulnerabilities reported in academic conferences are technically sound, their real-life impact is often uncertain and can be influenced by economic and practical considerations.

In his talk, Asokan first discussed the role of offensive security in identifying vulnerabilities within existing systems. Proactively searching and finding bugs is crucial for improving system defenses and developing more secure technologies. However, not all reported vulnerabilities lead to significant real-world consequences. While vulnerabilities may be technically correct, their practical impact can be limited by various factors, including economic viability and actual exploitability.

He showed examples of vulnerabilities that received significant media attention but had limited practical consequences. One such example was RFID malware, which was publicized in 2006 as a potential threat where RFID tags could serve as vectors for viruses. A malicious RFID tag, when scanned by a RFID reader, can exploit vulnerabilities in the reader's software. Despite the media hype and academic interest, this threat did not materialize into a real-world problem. Similarly, an SMS control channel attack was publicized as a severe threat to GSM networks but did not have the predicted disruptive impact in practice.

Asokan highlighted the potential opportunity costs of overpublicizing vulnerabilities that have no corresponding real-world threat: Excessive attention can lead to irrational decisions by industry, such as prematurely pulling products from the market or discouraging researchers from working on certain technologies. This overreaction can lead to significant and unnecesary economic costs and stifle innovation.

To address these issues, Asokan underlined the importance of collaboration between different stakeholders, including security researchers, industry practitioners, and economists. By working together, these groups can develop a better understanding of the real-life impacts of vulnerabilities and create frameworks for more balanced reporting and assessment. He suggested that security researchers should consider the potential economic and practical impacts of their findings and strive for balanced reporting that does not unnecessarily alarm or mislead stakeholders.

Slides: https://asokan.org/asokan/research/RealWorldVulnerabilityImpacts-Intro-April-2024.pdf

Finding, Patching, and Promoting Security Research, and what about Sustainability?

Speaker: Daniel Gruss, TU Graz

Daniel Gruss, an Associate Professor at Graz University of Technology, presented his research on the sustainability of security patches and the broader implications of security vulnerabilities in processor microarchitectures. Daniel, known for his work on side channels and transient execution attacks, highlighted the constant presence and inevitability of vulnerabilities in complex systems. He emphasized that while security patches are necessary, they often come with significant costs that can accumulate to unsustainable levels over time.

Daniel emphasized the substantial energy and performance costs associated with security patches, using the KAISER patch as a key example. This patch was developed in response to the Meltdown vulnerability discovered in 2018, a hardware flaw that allows programs to access the memory of other programs and the operating system, effectively bypassing security boundaries. Although the KAISER patch only decreased the performance of a computer by 5% per machine, the cumulative effect on a global scale could significantly raise electricity consumption. Daniel warned that by 2030, a single patch of similar gravity could potentially increase global electricity consumption by 0.4%.

Delving more into the details of processor microarchitectures, Daniel explained how side channels and transient execution attacks exploit fundamental design aspects of modern processors. These vulnerabilities arise from minimal timing differences in processing operations, which can be used to collect sensitive information. For example, side channels can leak information through various unintended signals, while transient execution attacks exploit speculative execution to access restricted data. These vulnerabilities are pervasive and challenging to defend against because they originate from inherent and necessary features of modern processor designs. This makes it difficult to eliminate them, after the processors have already been manufactured, without significantly impacting performance.

In addition to addressing technical challenges, Daniel emphasized the importance of how security research is communicated to the public. He discussed how marketing vulnerabilities can shape public perception. By creating dedicated websites, logos, and catchy names, researchers can ensure their findings are easily accessible and memorable. However, he warned that while this marketing can help to convey technical details to a broader audience, it also risks exaggerating the actual impact of the vulnerabilities.

To address the challenges of security and sustainability, Daniel advocated for fundamental changes in system design. He suggested that security measures should be integrated in a way that minimizes their impact on system performance. This involves preparing systems from the design stage to accommodate future security patches without significant performance degradation. Thus, he recommended optimizing hardware performance during manufacturing to a sweet spot where the system's efficiency and security can be balanced effectively.

Slides: https://gruss.cc/talk_ethz/slides.pdf

Modeling Vulnerabilities Based on Attack Value

Speaker: Eduardo Vela Nava, Google Zurich

Eduardo Vela Nava from Google Zurich shared insights on modeling security within software development teams. The crux of Eduardo's presentation revolved around a specialized model that he uses to assess security risks. This model involves variables like Value, Victims, Cost, Penalty, and Conviction, which collectively help in determining what Eduardo termed as the Attack Value:

(Value x Victims) - Cost - (Penalty x Conviction) = Attack Value.

Where

- **(Value x Victims)** quantifies the potential gain from an attack, considering both the value of the target and the number of victims,
- **Cost,** the expenses the attack invests to execute the attack,
- and **(Penalty x Conviction),** includes expected legal repercussions, further decreasing the net benefit of the attack.

He discussed two main types of attacks: opportunistic attacks that target large groups indiscriminately, and targeted attacks that are aimed at specific entities such as activists or governments. This distinction helps in understanding the potential impact and the necessary responses to these threats.

Further elaborating on the types of security issues, Eduardo distinguished between bugs, vulnerabilities, and exploits. A bug is an unintended behavior, a vulnerability is a bug that could have security consequences, and an exploit is the active use of a vulnerability against victims. He also explained the differences between fixing a bug (which eliminates it) and mitigation strategies (which reduce the effectiveness of potential exploits).

Eduardo also detailed the lifecycle of security management, starting from feature development by programmers, which can potentially introduce bugs, through to the detection and reporting of these bugs by testers or automated systems, and finally to the addressing and resolving of these issues. Notably, he mentioned that public awareness of a vulnerability often speeds up the resolution process, which can be crucial in preventing widespread damage.

He proposed viewing the security process as a function of time—considering how long it takes to find bugs, exploit vulnerabilities, and the time during which these exploits remain viable before they are mitigated. This temporal perspective helps in understanding and improving the efficiency of security processes.

On preventative measures, Eduardo explained Google's proactive stance in making it difficult and costly for developers to introduce vulnerabilities in the first place. This involves rigorous code reviews and testing mechanisms that catch issues as early as possible. The aim is to make vulnerabilities compile-time errors and rely on compiler enforcements rather than relying

on the programmer. This approach underlines Google's commitment to security, even at the expense of higher operational costs.

Addressing remediation efforts, he highlighted Google’s goals, both internally and publicly through Service Level Agreements (SLA), to reduce the response times to security breaches. This rapid response capability is important in managing and mitigating security issues effectively.

Eduardo also touched upon the legal and policy challenges associated with discovering vulnerabilities. He brought up initiatives like https://disclose.io, which provide legal protection for researchers, allowing them to disclose vulnerabilities without facing legal repercussions. This is vital in fostering an environment where security researchers can operate without fear of litigation.

Slides:
https://docs.google.com/presentation/d/1L_AfOWIIk6vN4S1QlEh2Oj2PFT3S3ZUk_1j47AZWWsY/edit#slide=id.g2c8bc52cd21_0_1

Quantifying Cyber Risk

Speaker: Rainer Boehme, University of Innsbruck

Rainer Boehme's presentation addressed intricacies of quantifying cyber risk, offering a systematic review of existing studies and proposing a theory for "security technology avoidance". He began by pointing the audience to over 20 years of the Workshop on the Economics of Information Security (WEIS), which laid the foundation for the review. Based on this, he addressed the naive assumption that more security inevitably leads to less harm, highlighting research that suggests otherwise.

Rainer then introduced a causal model of cyber risk that outlines how various elements interact to affect the security outcome. It considered both preventive security (e.g. antivirus software) and reactive security measures (e.g. deploying a patch to fix a vulnerability), surface exposure (e.g. many open ports in firewall), assets (e.g. database with sensitive data), the magnitude of the threat and the resulting harm.

A key part of his presentation focused on classifying cyber risk studies, distinguishing between abuse studies versus those examining market reactions. Abuse studies provide immediate, granular insights into specific security incidents, while market reaction studies offer broader economic perspectives, assessing the impact of security breaches on stock prices over time. Rainer noted the variability in results from these studies, emphasizing the importance of comprehensive data collection and analysis.

In the second part of his presentation, Rainer discussed the negative perceptions and opportunity costs associated with well-publicized vulnerabilities, coining the term "security technology avoidance." This phenomenon occurs when past negative experiences or extensive media coverage lead to mistrust and avoidance of certain technologies, potentially stalling innovation and deployment. He presented the avoidance of Intel SGX technology by researchers as a negative example. Although SGX is a robust and secure technology for isolating applications, researchers regularly avoid it when proposing new security applications since it was prone to vulnerabilities in the past which resulted in negative media coverage.

To highlight how the attention and resources dedicated to certain technologies do not always align with their practical efficiency or security, Rainer provided another negative example about digital watermarking technology. Although the core of this technology is fundamentally “broken”, it still sees tremendous interest in research and funding.

The presentation concluded with a call for more robust, interdisciplinary research to better understand the economic implications of cybersecurity measures.

Slides:
https://zisc.ethz.ch/wp-content/uploads/2024/04/Rainer_Boehme_Quantifying_Cyber_Risks.pdf

Information security vulnerabilities from an insurer’s perspective - risk transfer and the real-life financial impact on the economy and general public

Speaker: Lucas Engl, Zurich Insurance

Lukas Engl provided an analysis of information security from the perspective of insurance, focusing on the nuances of cyber insurance and the real-life impacts of cyber threats on business and the economy. In his presentation, Lukas elaborated on the scope, challenges and the necessity of cyber insurance in today's digital landscape.

Coverage and Limitations:
Lukas defines cyber insurance as a shield against unpredictable yet quantifiable threats, covering events like privacy breaches, business interruptions, and regulatory liabilities. The evaluation of insurance applications is based on three pillars:

1. **Maturity of IT Security:**
   Evaluates the company's cybersecurity measures not purely through technical audits but also from a management perspective. This allows insurers to determine the robustness of a company's defenses and capability to mitigate a cyber threat. Examples in this pillar include employee training.

2. **Business Profile:**
   This assessment considers the company's size, the type of data it handles and its industry sector to determine its vulnerability to cyber threats. To give an example, Lukas mentioned the healthcare sector. It is a high stakes environment for data security because of its valuable data (medical records) and sometimes less-than-optimal cybersecurity measures.

3. **Regulatory Compliance:**
   Examines how well a company adheres to relevant cybersecurity regulations and laws, which can impact the penalties in the event of data breaches or other security incidents. Compliance with regulatory standards is critical as it not only affects the company’s legal standing but also influences the trust customers place in it.

Lukas further discussed the reasons for denying insurance coverage, which often relate to inadequate incident response plans and other shortcomings such as missing off-site backups, or insufficient network segregation.

Economic Impact

Lukas then shifted to the broader economic impact of cyber incidents. He referenced data showing substantial financial losses and compromised data across various industries, illustrating the significant disparity in the impact between large corporations and small businesses. Smaller firms, he noted, are disproportionately affected due to similar costs but

smaller revenue bases, making recovery tougher.

The talk highlighted ransomware as a critical concern due to its high volume and impact. Lukas emphasized that smaller companies often lack basic security measures, making them vulnerable targets. He suggested that simple steps like implementing multi-factor authentication and robust backup strategies could markedly reduce risks.

Slides:
https://zisc.ethz.ch/wp-content/uploads/2024/04/vulnerabilities_insurance_ETH_presentation-final.pdf

Panel Discussion

Panel host: Shweta Shinde (ETH Zurich)
Panelists: Hans Gersbach (ETH Zurich), Kaveh Razavi (ETH Zurich), Mark Brand (Google), Anders Fogh (Intel)

The second half of the workshop covered a panel discussion. It addressed a variety of issues related to the interplay between academic research in cybersecurity and its real-world impact. The key insights and discussion points included:

1. **Research Impact vs. Real-world Application**: The panel emphasized the need for academic researchers to balance the novelty of their research with its real-world implications. CVEs (Common Vulnerabilities and Exposures) and bug bounties are commonly used as markers of impact, but they should not overshadow the academic value of discovering new classes of attacks or theoretical vulnerabilities. The example mentioned is the discovery of new attack classes, which might not yet have CVEs associated with them but are significant in advancing the field.

2. **Impact Overstatement**: There's a concern that the academic system might incentivize researchers to exaggerate the impact of their findings to secure publications, funding, or attention. Panelists discussed the need for changing reviewer guidelines to reduce the emphasis on real-world impact measured through CVEs and bug bounties, thus promoting a more honest and scientifically driven reporting culture.

3. **Influence of Bug Bounties**: Bug bounties were discussed as a flawed proxy for impact. While they do reflect the severity of vulnerabilities to some extent, they are also influenced by factors like the quality of the vulnerability report or the effort involved, which do not necessarily correlate with the potential damage or exploitability of the vulnerability in the wild or the novelty of a proposed technique. This discrepancy was highlighted as problematic because it might skew the perceived importance of different security issues. An example discussed involved a vulnerability in a software used by a small user base that posed minimal risk to the company or the public. If the report is exceptionally written, the company might issue a big bounty as a reward for the effort, despite the low risk associated with the issue. This can skew the perception of what constitutes an important security issue, for both the security community and the public eye.

4. **Challenges Attack vs. Defensive Research**: The panel noted that publishing defensive or mitigation research is challenging because it requires comprehensive proof that the defense covers all potential attacks, unlike offensive research that only needs to demonstrate a single exploit. This creates a higher burden of proof on defensive research. To give a simple example, if a researcher develops a new firewall technology, they not only must demonstrate that it can block a new type of attack, they also must

demonstrate that it performs better than all state-of-the-art firewalls in different metrics.

5. **Responsible Disclosures and Industry Responses:** The panel highlighted the complexities involved in vulnerability disclosures and the industry responses. The discussions pointed to a gap in communication and understanding between researchers and industry practitioners. Examples included how a company may respond to a disclosure: Some companies may prioritize patches and updates based on the severity of the vulnerability, while others may not recognize the importance until it is actually exploited. This can lead to frustrations among researchers who feel their findings are undervalued or ignored. When such findings are submitted to a peer-reviewed conference, reviewers may question the practical impact of the research, especially if the industry has not acknowledged the vulnerability. Ideally, the acceptance and value of a research submission should be based on its novelty and the advancement it brings to the field. However, the industry's initial response to a vulnerability disclosure can disproportionately sway the peer-review process, affecting whether a paper is accepted, rejected, or requires major revisions.