On the Mismatch of Disclosure Practices in Security and Privacy Research

Simon Koch*, Jannik Hartung†, Rainer Böhme*, and David Klein‡
* *University of Innsbruck, Email: {simon.koch,rainer.boehme}@uibk.ac.at*
† *TU Braunschweig, Email: jannik.hartung@tu-braunschweig.de*
‡ *Max Planck Institute for Security and Privacy, Email: david.klein@mpi-sp.org*

***Abstract*—Responsible disclosure is a long-standing practice and, effectively, a requirement for security papers that detect vulnerabilities. However, the privacy community seems to be different. Based on a snapshot of 19 privacy and 23 security papers, there appears to be a difference in practice regarding the notification of the responsible parties when privacy violations are found. Our data shows that only 4 privacy papers report disclosure in contrast to 16 security papers.**

This raises the question of whether privacy disclosure should also become an established practice for privacy research. While on the surface, this intention seems reasonable, it might not be. We discuss how well the argument for security disclosure translates to privacy research and touch on multiple perspectives, including the researcher's role, the distinction between privacy and security research, and the ethical obligations arising from research results. Even though we cannot offer a definitive right or wrong answer to the question, we intend to start a discussion that allows us, as a community, to critically reflect on our role in society and the impact of our results.

1. Introduction

There is no responsible disclosure requirement for privacy research, and it is necessary to discuss whether this should change. This lack of an obligation is in stark contrast to security research, where the responsible disclosure of discovered vulnerabilities is an established practice and a requirement for the peer-review process across all major security venues. This is justified by ethical considerations aimed at minimizing any harm resulting from the exploitation of reported security vulnerabilities. The discrepancy between the two related fields raises the question of whether the difference should stand or if a community-wide disclosure policy for privacy research is required. Similar arguments could be used to justify this requirement for privacy research, or different considerations could apply.

To address this question, we first take stock of the current state of affairs. To do so, we analyze the disclosure policies stated in the Call for Papers (CfP) of the main 10 security and privacy conferences since 2018. We want to understand the official policies regarding the disclosure of both security vulnerabilities and privacy issues. Subsequently, we move towards an analysis of the implementation of the official policy by sampling 23 security and 19 privacy papers within the same time frame and analyzing them for any reported disclosure. We finish by analyzing how well the responsible disclosure process from security translates to privacy.

Our results indicate that the majority of CfPs do not address specific disclosure requirements, and even when they do, privacy is never explicitly mentioned. At best, no specification is given of which types of discoveries must be disclosed, thereby indirectly including privacy. The analysis of security and privacy papers mirrors these observations. It provides strong indications that the security community has a strict requirement for disclosure. In contrast, the privacy community reports fewer disclosures, with a wider variety of processes employed. Our subsequent analysis of how well the reasoning and process of responsible disclosure for security translates to privacy provides possible arguments for why this mismatch exists, as at multiple steps, the translation breaks down or comes with significant caveats. Thus, we not only provide first evidence for the mismatch between security and privacy, but also argue that this mismatch might be reasonable and that a separate discussion within the privacy community is required.

In summary, we have two main contributions:

- An analysis of the formal disclosure policy based on CfPs and its implementation based on reported disclosures of 23 security and 19 privacy papers.
- An analysis of how well the responsible disclosure process of security translates to privacy, showing multiple mismatches and caveats.

The remainder of the paper is structured as follows: We first introduce the differences and intersections between security and privacy research, followed by an overview of the responsible disclosure process (Section 2). Then we describe our analysis of CfPs and the results of our paper sampling, followed by a summary of the observed implementation of disclosure (Section 3). Based on the documented mismatch, we present an analysis of how the disclosure process for security does not translate well to privacy (Section 4) and conclude with a short summary of our main results and arguments (Section 5).

2. Background

To establish the foundation for our question on whether there is a mismatch between the disclosure practices of

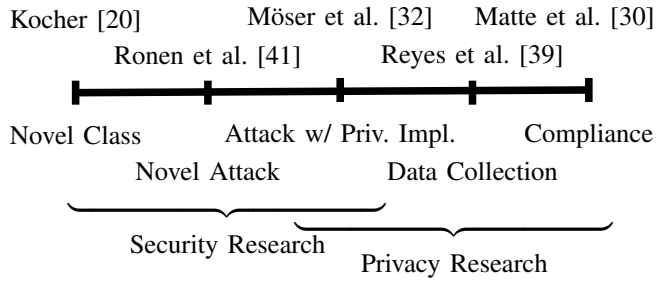

Figure 1. A spectrum of computer science research showing how research from pure security (left) intersects and moves to pure privacy (right). Example papers are cited above the bold line.

security and privacy research, we first need to establish commonality between the two topics (2.1) and provide a short explanation of what disclosure actually is (2.2).

2.1. Security and Privacy Research on a Spectrum

The security and privacy research communities overlap. This is evident from the leading venues for either. The renowned security venues ACM CCS, NDSS, or USENIX Security also feature papers that cover privacy. IEEE S&P even has both security and privacy in its name, suggesting a shared community. We note that the term *privacy* has many definitions. We use it to describe the protection of individuals from potential harm caused by the misuse of their personal data. To visualize the shared community and the intersection of research, we grouped security and privacy research into five areas and positioned them on a spectrum visualized in Figure 1.

The leftmost research area is the discovery of *novel classes* of security vulnerabilities, such as side-channel attacks on cryptographic algorithms, e.g., Kocher [20], thereby opening up a new area of security research. Next comes the application of known classes of security vulnerabilities to discover *novel attacks*, such as Ronen et al. [41], who leveraged a side-channel attack to take over IoT light bulbs. Both of these research areas fall within the security domain. Privacy concerns are, at best, a downstream effect. This, however, changes with *attacks with privacy implications*, where security vulnerabilities are used or discovered that have clear implications for the affected users' privacy. An example of such research is Möser et al. [32], who leveraged combinatoric attacks to deanonymize transactions in the cryptocurrency Monero. Subsequently, we move into the realm of pure privacy, focusing on *data collection* practices. While security research tools are deployed, such as network traffic interception, the goal is to document and analyze observed data-collection practices that affect user privacy. Reyes et al. [39] for example, analyze mobile applications targeted at children and discover excessive data collection that may violate privacy laws. At the rightmost end of the spectrum, we have privacy *compliance* research, focusing on whether design decisions and data collection practices comply with current legislation. The difference to data collection is that compliance research is always positioned within legislation, whereas research on data collection has implications on its own. An example for compliance research is Matte et al. [30], who analyze websites on whether their cookie consent banner implementation is compliant with the consent requirements of the GDPR.

Given the overlap on the spectrum of topics and of the research communities, one could expect a shared approach towards responsible disclosure by both research areas, whether squarely within security or within privacy. This is a premise that needs to be verified.

2.2. The Disclosure Process

There are three main methods of disclosure [38]. Each carries its own advantages and disadvantages. *Private disclosure* means disclosing the vulnerability privately to the responsible parties, without subsequent publication. While this ensures that responsible parties are informed and that no malicious third party is made aware of the discovered vulnerability, there is also no incentive or enforcement that the vulnerability will be resolved. A responsible party could ignore the report, leaving the vulnerability unaddressed for other people to find and potentially abuse. *Full disclosure* is the polar opposite of private disclosure and means that the details of the discovered vulnerability are made public immediately. This forces the responsible party to act on the report, but also gives any malicious actor a window during which the vulnerability is known but not resolved. Finally, *responsible disclosure* presents a halfway point between the two extremes. The first step is to privately report a discovered vulnerability, with a deadline after which the vulnerability is made public. Thus, the responsible party has a timeframe to preempt any abuse by fixing the vulnerability, but cannot simply ignore the report. This mode of disclosure is the accepted standard practice in the security community for any discovered vulnerability during research projects.

Security disclosures largely follow the same set of steps, though the specific actions taken during each step might vary depending on the chosen disclosure type:

(1) Identification. The first step is identifying a security issue. This covers conceptualizing a possible attack vector and testing it, whether in a dummy application or in the real world.

(2) Reproduction. The second step is to ensure that a vulnerability can be exploited reproducibly. This may include crafting a reliable payload or thoroughly documenting the steps and setup required for a successful exploitation.

(3) Responsible party. After identifying an issue and ensuring reproducibility, the responsible parties who need to be notified have to be identified. These fall into four high-level categories: the *supplier* develops external parts of the application that lead to the vulnerability (e.g., SDK developer), the *developer* either uses vulnerable external components or introduces the vulnerability themselves (e.g., the App developer), the *platform* distributes the vulnerable application (e.g., the App Store), and the *authorities* are

government agencies responsible for the markets in which the vulnerable application is distributed (e.g., the BSI in Germany[1] or CISA in the U.S.[2]).

(4) Disclosure. Finally, the identified responsible party is contacted with a report on the issue, including a description of how to reproduce the vulnerability. This initial contact may then extend into a follow-up or dialog if the responsible party does not respond, or to answer additional questions as they arise while addressing the vulnerability.

3. Delta in Security and Privacy Disclosures

To assess the gap between established security and privacy disclosure practices, we take two snapshots spanning the established security and privacy venues. We selected these venues based on the core ranking,[3], requiring them to be ranked A or A* and to be popular in the security or privacy community. This led to the selection of IEEE S&P, ACM CCS, NDSS, USENIX Security, PETS, IEEE Euro S&P, RAID, Asia CCS, ACSAC, and ESORICS.

First, we analyzed the Call for Papers (CfPs) of our selected venues from 2018 to 2026, if available, scrutinizing the corresponding documents for disclosure requirements. Second, we filtered papers from 2018 to the present (Q1 2026) that reported security or privacy issues and analyzed their texts for any mentions of disclosure.

3.1. Differences in What Venues Require

To understand presumed established disclosure practices, we analyzed the CfPs for our selected venues. The goal was to identify any disclosure requirements and determine whether they distinguish between security vulnerabilities and privacy issues. To achieve this goal, we visited the venue pages from 2018 to 2026. If a CfP was no longer publicly available, we used the Wayback Machine. If references were made to a separate ethics or disclosure page, we also included those in our analysis. We scanned each in-scope document for mentions of disclosure requirements and their accompanying wording to determine whether they cover only security, both security and privacy, or do not specify their focus. If only vulnerability disclosure is mentioned, we consider this to pertain only to security vulnerabilities, as privacy issues are not necessarily vulnerabilities and may arise from a range of problems not captured and, consequently, are not intended by this security-focused term. Table 1 provides an overview of the venues and our classification of their CfPs across the years.

IEEE S&P: During the in-scope time span, the IEEE S&P went through two different CfPs phrasings touching on disclosure. The first, used between 2018 and 2020, discussed both *Human Subjects and Ethical Considerations* in a single section. It explicitly referenced vulnerability disclosures and was ambiguous towards privacy disclosure as it stated *[t]he same applies if the submission deals with personal identifiable information (PII) or other kinds of sensitive data* in reference to the previous sentence touching on vulnerability disclosure. The subsequent CfP starting in 2021 and being used until now, introduced a split between *Ethical Considerations for Vulnerabilitiy Disclosure* and *Ethical Considerations for Human Subject Research* and reused the PII phrasing and now explicitly requires to detail *the steps the authors have taken to mitigate harms to the persons identified* but stops implying that disclosure is required. Overall, we found only an indication of disclosure requirements for discovered privacy issues in the CfPs from 2018 to 2020.

ACM CCS: Until last year the ACM CCS did not have a dedicated ethics section in their CfP and referenced ethical considerations with the single phrase *[f]or papers that might raise ethical concerns, authors are expected to convince reviewers that proper procedures (such as IRB approval) have been followed and due diligence has been made to minimize potential harm*. Starting this year (2026), the CCS requires an explicit ethics section for *[p]apers that raise ethical concerns, such as those involving human subjects, user data, or real-world vulnerability analysis*. This section *should discuss the balance of risks vs. benefits and the steps taken to minimize potential harm (e.g., responsible disclosure, data anonymization)*, which does not necessarily require disclosure but does suggest it as an option without specifying whether it only applies to security or also to privacy issues.

USENIX Security: The USENIX Security shared the same ethics considerations as the IEEE S&P between 2018 and 2022, containing the same ambiguity. In 2023, they expanded the ethics section and stated that they *expect authors to carefully consider and address the potential harms associated with carrying out their research, as well as the potential negative consequences that could stem from publishing their work*. They do single out security vulnerabilities and reference disclosure as a valid mitigation for harm, but drop any reference to PII. In 2025, another change to the CfP is made, and now the CfP explicitly states that USENIX Security *expect[s] authors to carefully and proactively consider and address potential negative outcomes*, mentioning financial loss as an example. The CfP now neither explicitly mentions security or privacy issues, but remains general in tone. Finally, for 2026, another change is made, requiring authors to conduct a *stakeholder-based ethics analysis or should justify the use of an alternative approach to considering the ethics of their work*. The mentioned stakeholders include companies, and the harms mentioned still include financial loss, but now also include additional considerations such as *Beneficience* and *Respect for Law and Public Interest*. The text does not specifically focus on security or privacy.

NDSS: The NDSS only changed their ethics and disclosure-related part of the CfP in the past two years (i.e., for 2025 and 2026). Prior, it simply stated that *If the paper reports a potentially high-impact vulnerability, the authors should discuss their plan for responsible disclosure*. Starting in 2025, the conference added a reference to an ERB review

1. https://www.bsi.bund.de/EN/Das-BSI/Auftrag/auftrag_node.html
2. https://www.cisa.gov/about)
3. https://portal.core.edu.au/conf-ranks/

of ethically flagged papers and to the Menlo Report, but did not change the phrasing. For their 2026 CfP, the NDSS further expanded its ethical section but still kept disclosure specifications limited to security vulnerabilities.

PETS: The PETS used the same ethics section for 2018 and 2019. While it states *follow the basic principles of ethical research*, it does not mention any form of disclosure, whether for privacy or security. Starting in 2020 until today (2026), the PETS added explicit mentions for *cryptographic weaknesses, software exploits, and privacy attacks* but still did not mention any form of disclosure requirements.

ACSAC: Between 2018 and 2022 the ACSAC did not have any mention for ethical considerations either on *Call for Submissions* page nor the specific call for *Technical Papers*. Starting in 2023, the Call for Papers included an *Ethical Considerations* section stating *papers should discuss the steps taken to avoid negatively affecting any third parties, whether an institutional ethics committee reviewed the research, or how the authors plan to responsibly disclose the vulnerabilities to the appropriate software/system responsible party or owners before publication*.

IEEE Euro S&P: Between 2018 and 2021, the Euro S&P only had a disclaimer mentioning required ethics approval for work with human subjects. Starting in 2022 until 2025 the CfP added a section called *Proactive Prevention of Harm* requiring the authors to *carefully consider and address the potential harms associated with carrying out their research, as well as the potential negative consequences that could stem form publishing their work* and *following responsible disclosure practices* without specifying whether it only applies to security or privacy research. For 2026, the CfP for the *Proactive Prevention of Harm* was significantly shortened and only states that it *expect[s] authors to carefully consider and address the potential harms associated with carrying out their research, as well as the potential negative consequences that could stem from publishing their work*.

ACM Asia CCS: Between 2018 and 2023, the ACM Asia CCS had an *Ethics Considerations* section only stating that *The authors should take care of clarifying any potential ethical and legal concerns to their results, highly critical vulnerabilities or exploits, etc. The authors should provide evidence that they have thoroughly considered such issues*. Starting in 2024, minor adjustments now state that *[t]he authors should clarify any potential ethical and legal concerns to their results, highly critical vulnerabilities or exploits [...] elaborating how they have thoroughly considered such issues*.

RAID: The CfP for 2018 did not contain any ethics requirements. Starting in 2019 until 2025, the CfP included a section on *Human Subjects and Ethical Considerations* requiring a discussion of possible disclosure steps for security vulnerabilities and stating that *the same applies if the submission deals with personally identifiable information or other kinds of sensitive data.* The CfP for 2026 is not yet publicly available as of 2026-01-19.

ESORICS: The first CfP we were able to access was from 2019, which did not contain any reference to ethical or disclosure obligations and has not changed until today (2026).

	Disclosure Requirements as Specified by CfP								
Venue	'18	'19	'20	'21	'22	'23	'24	'25	'26
IEEE S&P	●	●	●	◐	◐	◐	◐	◐	◐
ACM CCS	○	○	○	○	○	○	○	○	⊗
USENIX Security	●	●	●	●	●	◐	◐	⊗	⊗
NDSS	**N/A**	◐	◐	◐	◐	◐	◐	◐	◐
PETS	○	○	○	○	○	○	○	○	○
ACSAC	○	○	○	○	○	◐	◐	◐	**N/A**
IEEE Euro S&P	○	○	○	○	⊗	⊗	⊗	⊗	○
ACM Asia CCS	○	○	○	○	○	**N/A**	○	○	○
RAID	○	⊗	⊗	⊗	⊗	⊗	⊗	⊗	**N/A**
ESORICS	**N/A**	○	○	○	○	○	○	○	○

TABLE 1. OVERVIEW ON THE CFPS FOR OUR SELECTED VENUES. IT VISUALIZES WHETHER THEY REQUIRE DISCLOSURE FOR SECURITY VULNERABILITIES (◐), BOTH (●), MENTION DISCLOSURE BUT DO NOT SPECIFY WHETHER FOR SECURITY OR PRIVACY (⊗), OR DO NOT MENTION DISCLOSURE AT ALL (○). **N/A** INDICATES THAT THE CFP IS EITHER NOT YET PUBLIC OR WAS INACCESSIBLE BOTH VIA THE REGULAR HOMEPAGE AND THE WAYBACK MACHINE.

3.2. Differences in What Authors Report

To analyze not only the official approaches outlined in the venues' Call for Papers but also the community's practices, we examined a set of recent privacy papers and a corresponding set of security papers. Our goal was to review each paper for mentions of disclosure or developer notification to quantify any observed differences in frequency. We filtered papers from our initial selection of venues, going back to 2018, by searching for keywords in both the title and the abstract.

We required papers to actively search for security vulnerabilities or privacy issues to be in scope, i.e., papers proposing mitigations or doing theoretical considerations are out of scope, as no disclosure would be required based on their results. To filter privacy papers, a paper had to use the terms *GDPR* or *CCPA* and any variation of the root word *violat*. For security, we required the term *vulnerability* and any variation of the root word *detect*. We leveraged the tool top4grep [23] that allows to build a database of papers of selected security conferences and search their title and abstracts for a list of keywords. It fetches the metadata of the papers from the publisher websites and uses regular expressions. Using the tool, we applied our search criteria and retrieved the list of matching papers. Finally, we reviewed the in-scope papers and read their abstracts to exclude those that are out of scope.

As we had significantly more matches for security papers than for privacy papers, we sampled the available security papers. We grouped each year into batches based on the conferences' disclosure policies and randomly selected a single paper from each batch. For example, in 2022, there were two venues only requiring disclosure for security vulnerabilities (IEEE S&P, NDSS), three having a disclosure

ID	Title	Venue	Year	Discl.
P01	Analyzing the AI Nudification Application Ecosystem [13]	USENIX Security	2025	○
P02	Navigating Cookie Consent Violations Across the Globe [46]	USENIX Security	2025	○
P03	You Can't Trust Your Tag Neither: Privacy Leaks and Potential Legal Violations within the Google Tag Manager [31]	IEEE Euro S&P	2025	●
P04	Johnny Can't Revoke Consent Either: Measuring Compliance of Consent Revocation on the Web [16]	PETS	2025	◐
P05	Automated Large-Scale Analysis of Cookie Notice Compliance [6]	USENIX Security	2024	⊗
P06	Wear's my Data? Understanding the Cross-Device Runtime Permission Model in Wearables [55]	IEEE S&P	2024	◐
P07	PolicyChecker: Analyzing the GDPR Completeness of Mobile Apps' Privacy Policies [53]	ACM CCS	2023	○
P08	CHKPLUG: Checking GDPR Compliance of WordPress Plugins via Cross-language Code Property Graph [43]	NDSS	2023	○
P09	Do Opt-Outs Really Opt Me Out? [7]	ACM CCS	2022	●
P10	Freely Given Consent?: Studying Consent Notice of Third-Party Tracking and Its Violations of in Android Apps [35]	ACM CCS	2022	●
P11	Automating Cookie Consent and GDPR Violation Detection [5]	USENIX Security	2022	○
P12	Are iPhones Really Better for Privacy? A Comparative Study of iOS and Android Apps [21]	PETS	2022	○
P13	Checking Websites' GDPR Consent Compliance for Marketing Emails [22]	PETS	2022	○
P14	On dark patterns and manipulation of website publishers by CMPs [48]	PETS	2022	○
P15	Keeping Privacy Labels Honest [19]	PETS	2022	○
P16	Viopolicy-Detector: An Automated Approach to Detecting GDPR Suspected Compliance Violations in Websites [36]	RAID	2022	○
P17	Share First, Ask Later (or Never?) Studying Violations of GDPR's Explicit Consent in Android Apps [34]	USENIX Security	2021	●
P18	Do Cookie Banners Respect my Choice? : Measuring Legal Compliance of Banners from IAB Europe's Transparency and Consent Framework [30]	IEEE S&P	2020	○
P19	Panoptispy: Characterizing Audio and Video Exfiltration from Android Applications [37]	PETS	2018	◐
S01	Bamboozling Certificate Authorities with BGP [2]	USENIX Security	2018	○
S02	VulDeePecker: A Deep Learning-Based System for Vulnerability Detection [26]	NDSS	2018	⊗
S03	Sereum: Protecting Existing Smart Contracts Against Re-Entrancy Attacks [40]	NDSS	2019	⊗
S04	RVFuzzer: Finding Input Validation Bugs in Robotic Vehicles through Control-Guided Testing [18]	USENIX Security	2019	●
S05	Static Detection of Uninitialized Stack Variables in Binary Code [12]	ESORICS	2019	○
S06	Automated Discovery of Cross-Plane Event-Based Vulnerabilities in Software-Defined Networking [50]	NDSS	2020	●
S07	A Stealthier Partitioning Attack against Bitcoin Peer-to-Peer Network [49]	IEEE S&P	2020	●
S08	Speculative Probing: Hacking Blind in the Spectre Era [14]	ACM CCS	2020	○
S09	Black Widow: Blackbox Data-driven Web Scanning [10]	IEEE S&P	2021	●
S10	ReDoSHunter: A Combined Static and Dynamic Approach for Regular Expression DoS Detection [25]	USENIX Security	2021	●
S11	Time to Rethink the Design of Qi Standard? Security and Privacy Vulnerability Analysis of Qi Wireless Charging [52]	ACSAC	2021	●
S12	FuzzUSB: Hybrid Stateful Fuzzing of USB Gadget Stacks [17]	IEEE S&P	2022	●
S13	Mining Node.js Vulnerabilities via Object Dependence Graph and Query [24]	USENIX Security	2022	◐
S14	Cross Miniapp Request Forgery: Root Causes, Attacks, and Vulnerability Detection [54]	ACM CCS	2022	●
S15	LOKI: State-Aware Fuzzing Framework for the Implementation of Blockchain Consensus Protocols [28]	NDSS	2023	●
S16	An Unbiased Transformer Source Code Learning with Semantic Vulnerability Graph [15]	IEEE Euro S&P	2023	○
S17	CoCo: Efficient Browser Extension Vulnerability Detection via Coverage-guided, Concurrent Abstract Interpretation [56]	ACM CCS	2023	●
S18	Undefined-oriented Programming: Detecting and Chaining Prototype Pollution Gadgets in Node.js Template Engines for Malicious Consequences [27]	IEEE S&P	2024	●
S19	CSRFing the SSO Waves: Security Testing of SSO-Based Account Linking Process [3]	IEEE Euro S&P	2024	●
S20	What All the PHUZZ Is About: A Coverage-guided Fuzzer for Finding Vulnerabilities in PHP Web Applications [33]	ACM Asia CCS	2024	●
S21	YuraScanner: Leveraging LLMs for Task-driven Web App Scanning [44]	NDSS	2025	●
S22	FIXX: FInding eXploits from eXamples [47]	USENIX Security	2025	●
S23	Parcel Mismatch Demystified: Addressing a Decade-Old Security Challenge in Android [8]	ACM CCS	2025	●

TABLE 2. A LIST OF SAMPLED PRIVACY AND SECURITY PAPERS SELECTED BASED ON CFP DISCLOSURE POLICY AND PUBLICATION DATE. THE *Discl.* COLUMN SHOWS IF ALL (●), NO (○), OR ONLY PARTS (◐) OF THE RESULTS WERE DISCLOSED, OR IF DISCLOSURE WAS ONLY DISCUSSED (⊗).

policy not specifying if it only applies to security (USENIX Security, IEEE Euro S&P, RAID), and five venues without any disclosure policy (ACM CCS, PETS, ACM Asia CCS, and ESORICS). This resulted in three batches for 2022, leading to three randomly sampled in-scope papers.

After we had a fixed selection of papers, we searched them for mentions of discovered vulnerabilities or issues and whether a disclosure or contact with the developers or other responsible parties occurred. We noted any such mention as well as the context thereof. If the notification was a disclosure to the developer or other responsible party, we counted it as a disclosure; if not all responsible parties were notified, we counted it as a partial disclosure; and if there was no mention of a disclosure, we counted it accordingly. Table 2 provides an overview of the selected papers, their publication venue and date, and finally their disclosure approach.

We deliberately took a positive stance on disclosure requirements, especially for the privacy side. This might seem like we are preempting the question of whether disclosure by privacy papers is required, but it must not be understood that way. We only apply two objective criteria: **(a)** does the paper state to find privacy issues within their research subject, and **(b)** do the researchers report any form of disclosure or contact with the developers? If (a) applies, the paper is in scope, whereas the answer to (b) decides the classification in Table 2. We provide a thorough discussion of the included papers in the next section. If we report that a paper has made no disclosure, this must not be understood as a judgment but simply as an objective fact.

Privacy papers: We identified a total of 19 privacy papers matching our selection criteria between 2018 and 2025. The most prominent venue was PETS (6), followed by USENIX Security (5). Out of the 19 privacy papers, 4 reported a disclosure to all affected responsible parties, 3 reported partial disclosure, and 12 did not state any form of

disclosure. We classified 3 privacy papers as partial disclosure because the authors do not have a blanket disclosure statement. Pan et al. [37] discuss three case studies, but only state for one that they notified the Google privacy team, who are not the developers of the corresponding app, but are in a position to effect change. Yeke et al. [55] only report a single observed issue to the Google Privacy Team. Kancherla et al. [16] only state that they notified the 23 explicitly named companies, implying that there are unnamed companies that have not been notified.

Security papers: We selected a total of 23 security papers based on our methodology between 2018 and 2025. 2018 is the only year where we had only two venue groups to sample from, namely, no required disclosure by the CfP and an inclusive statement for disclosure. The remaining years, we had three covering all observed CfP disclosure statements. Out of the 23 security papers, 16 report a disclosure to all affected responsible parties, 1 reports partial disclosure, and 6 did not state any form of disclosure. We classified a single security paper as partial disclosure because Li et al. [24] validated only a subset of the detected vulnerabilities, and, based on the text, the disclosure depended on that validation.

Limitations: While our selection of papers provides a first snapshot of the effective disclosure differences between security and privacy, it does come with two main limitations:

(1) We are selecting papers only based on title and abstract, and also using specific keywords. Especially on the privacy side, requiring the mentioning of privacy legislation (GDPR or CCPA) likely shifts the focus of the selected titles more towards compliance. This may exclude work on non-compliance privacy questions, thereby biasing the overall picture. Nonetheless, all selected papers are part of either the privacy or security research community, and while they may form a subgroup within these communities, they do provide a window into the larger community.

(2) We are relying on the documentation given within the paper to determine whether disclosure has happened. It is possible that a paper does not report any disclosure procedure, but did disclose and is thus misclassified by us. We consider this unlikely, as we encountered privacy papers actively reporting disclosures, indicating no inherent bias within the privacy community to make silent disclosures.

3.3. Observed Perspectives on Disclosure

Overall, the security community appears to agree that disclosing discovered vulnerabilities is necessary and should be reported in research papers. The supermajority of papers (16/23) report responsible disclosure of discovered vulnerabilities. Most papers report that disclosure was directed to the developers of the affected products. Only Yang et al. [54] report the discovered vulnerabilities to the platform (Tencent) rather than to individual developers, most likely because the vulnerabilities were found across a large number of apps. No paper reports any actions intended to ensure that a vulnerability is being patched, besides providing a report to the responsible parties.

In addition to the already high number of papers that provide responsible disclosure, two of those that do not report a disclosure still discuss disclosure. Li et al. [26] state that all discovered vulnerabilities were independently patched in the newest software version, and thus no disclosure was necessary. Rodler et al. [40] explain that, because of the blockchain's anonymity, disclosure was not possible. Finally, the paper that reports only partial disclosure also provides a reason for their actions, an infeasible workload, showing that they have considered the issue.

The privacy community, on the other hand, exhibits a more fractured picture. Only a minority of papers (4/19) report disclosure, with the majority of papers not discussing disclosure at all. Among the leading conferences, all selected PETS papers have no or only partial disclosure, whereas at USENIX Security, one out of five papers reports disclosure. The same pattern holds for the other conferences: some papers report disclosure, while others do not. This distribution shows that, even at established security conferences with clear patterns of disclosure for security vulnerabilities, the privacy sub-community is not in agreement on whether disclosure is required, leaving it to the authors' discretion.

Even more interesting are the different perspectives offered in favor of disclosure. Nguyen et al. [34] explicitly state that *since disclosing the findings to authorities (e.g., regulators, Google) might cause financial harm to developers, we consciously decided not to involve authorities but rather notify developers directly to remedy compliance issues*, indicating that they consider reporting infractions to the authorities potentially harmful to the developers. This contrasts with Pan et al. [37], who reported one of their discovered issues to the Google privacy team. While not reporting any disclosure, Bouhoula et al. [6] discusses the option and state, *we emphasize that the violations observed by our automated procedure cannot be directly taken by a court or DPA to enforce fines*, implying that if such a disclosure were possible, they would have done so despite the potential fines. Finally, we have Kancherla et al. [16], who state that *we notified 23 companies that own the domains that we explicitly mention in the main text. In sharing our findings, we took inspiration from Maass et al. [29] who showed that mentioning GDPR and its fines and sending notifications from a legal academic significantly increases the remediation rate of website owners*, thus employing psychological manipulation to effect change.

In summary, we observe that, unlike the security community, there is no agreement in the privacy community on whether disclosure is necessary and, if so, how it should be conducted. While some authors actively try to avoid harm to developers, others risk financial impact, imply a willingness to take enforcement actions through the authorities, or even employ soft enforcement measures themselves. This shows that we need to discuss whether and, if so, how disclosure of discovered privacy issues should occur, to ensure a community standard we all agree on.

4. A Discussion on Change

We have documented a discrepancy between the disclosure practices in security and privacy research. Based on our results and accounting for the limitations of our method, it is fair to conclude that privacy research lacks a disclosure standard that researchers adhere to, resulting in disclosures, if they occur, varying widely across chosen approaches. This state of affairs raises the question of whether we need to push for change, be it toward developing a standard for the responsible disclosure of privacy issues or against conducting such disclosure campaigns. We address this question by going through the security disclosure process as described in Section 2.2 and analyzing how well the requirements and underlying thought process map to privacy at each step.

4.1. Identification

The first step in the disclosure chain—identifying something noteworthy—applies to both security and privacy research. Necessarily, a privacy paper contains some novel insights into the behavior of systems that affect a user's individual privacy or rights. However, in security, an identified vulnerability carries an a priori assumption of harm, meaning that any unresolved vulnerability can and will be abused by malicious third-party actors. Thus, the question of whether disclosure is required to prevent harm is always answered positively, a necessary prerequisite for responsible disclosure. While this does also hold for the center of the spectrum, i.e., security vulnerabilities with privacy implications (ref. Figure 1), it does not necessarily apply to for pure privacy research on the right end of the spectrum.

> 👁 **Observation 1:** Security issues with privacy implications carry a reasonable harm assumption as they involve a security vulnerability.

At the rightmost point of Figure 1—compliance research (ref. Section 2.1)—there cannot be a universal assumption of harm. A compliance violation may simply mean that some law has not been obeyed. Violating an existing law does not, in itself, imply harm, as laws vary across jurisdictions, sometimes in opposite ways. This issue is exacerbated by the fact that the interpretation of law can be political and is subjective unless there are rulings by the highest authority with jurisdiction over the subject. But even then, the law's applicability remains local. Thus, establishing a universal disclosure requirement for all privacy research would force computer scientists to make judgments about laws with which they are unfamiliar. Such judgments would not necessarily stand up in court, might be motivated by non-existent harm, and might not even be applicable. Consequently, the justification for mandatory responsible disclosure breaks down for privacy compliance research.

> 👁 **Observation 2:** There cannot be a universal harm assumption for privacy compliance research.

This assessment, however, changes when we look at the remaining area of privacy research. In the case of data collection on personal attributes, e.g., political affiliation or geolocation data, there is a potential for harm to affected individuals. Without disclosure, but with publication, malicious third parties could become aware of the wealth of information and attempt to gain access, subsequently abusing the knowledge gained against affected individuals. Whereas in the case of a responsible disclosure to the responsible parties, the problematic behavior could be rectified, future abuse mitigated, and thus further harm prevented. However, this assessment requires that the data collected does invite harm if leaked to the wrong actors, an assumption that might not hold universally for all personal data collected, and that needs to be assessed on an individual basis.

> 👁 **Observation 3:** There is a reasonable harm assumption for data collection.

This observation lets us further explore the responsible disclosure chain for privacy issues related to data collection.

4.2. Reproduction

After identifying a potentially harmful issue, the second step is to prepare evidence to accompany the responsible disclosure notification, usually in the form of a reproduction chain. This step is rather easy with security vulnerabilities, as any exploit that has been successfully tested against a test system had a sequence of steps or used a payload. However, it can become very difficult for observed data collection.

Even though a researcher observes data collection, they do so passively by analyzing network traffic or other exfiltration methods. While this shows that data is being processed by the device and sent somewhere, it does not necessarily imply storage or processing on the receiver's end. However, storage or processing on the receiving end is required to notify a responsible party, else the previously established harm assumption breaks down. Alternatively, a very strong attacker model could justify a universal responsible disclosure requirement, such as the assumption that a machine-in-the-middle can intercept and read the transmitted information. This leaves the researcher with either weak evidence or the need to obtain stronger evidence of the receiver's end storage and processing of data prior to responsible disclosure. In the absence of strong evidence, a universal responsible disclosure requirement rests on a weak foundation, casting doubt on it. However, if the researchers are required to obtain stronger evidence, we run into technical and possibly ethical problems. Inferring data storage or processing requires probing the receiving black box, for example, through an ablative study of the transmitted data. This comes with technical and corresponding ethical risks, as it is uncertain how receiving servers might react to manipulated data values. Thus, a universal responsible disclosure requirement would force researchers to conduct invasive and potentially harmful actions, or conduct disclosure on weak evidence, which would rather necessitate a harm-benefit consideration.

Observation 4: Observed data collection is only weak evidence; obtaining strong evidence is technically and ethically challenging.

4.3. Responsible Party

In security research, responsible disclosure is targeted at the responsible party, meaning either the developer or the platform distributing the vulnerable software. While this statement seems trivial at first glance, even for security research, identifying and notifying all responsible parties is non-trivial. In the context of security, Stock et al. [45] and Böhme et al. [4] have reported that identifying and successfully reaching the responsible parties is non-trivial for web and blockchain applications, respectively. Utz et al. [51] report that these problems are exacerbated for privacy disclosure.

Observation 5: Identifying responsible parties for privacy research can be more challenging than for security research.

Finally, the choice of party to disclose privacy issues to is non-trivial. While the common choice of target for security can be replicated for privacy research, it comes with significant caveats, as privacy violations, unlike security violations, may also be in violation of the law and subject to fines. Thus, the choice of party to disclose to is political: notifying the developer might circumvent the law and adopt a rather activist approach, whereas notifying the authorities would be an enforcement approach. Either approach is problematic as it puts the researcher in the position to decide between harming society (from missed opportunities to collect fines) or developers (by inflicting fines). Both actions create externalities in the sense that the harm would not have occurred in the same way without the researchers' involvement. Thus, an obligation to disclose, whether to developers or authorities, could push researchers into a role with ethical implications that is opposed to their own beliefs.

Observation 6: Mandatory responsible disclosure could push privacy researchers into an activist or enforcement role.

4.4. Disclosure

Past work on the effects of security vulnerability disclosure has shown mixed, long-term effects. Arora et al. [1], for example, report that patched vulnerabilities attracted more exploitation than secret or only published vulnerabilities. Results by Finifter et al. [11] showed that vulnerability reward programs are cheaper than hiring a full-time employee, suggesting that externalizing security can be economical. Also, zero-day markets have been the focus of past discussion among researchers (ref. Egelman et al. [9]), discussing the question of whether vulnerability reward programs or zero-day markets are beneficial, resulting in a diverse set of arguments both in favor and against. Arguments include that companies might use such programs to outsource product testing and the risk of a cobra effect.[4] While the results for security vulnerabilities likely do not directly translate to privacy, it is important to note that there are possible downsides, too. If responsible parties cultivate the impression that if something is wrong, a researcher will notify them about the problem, they may engineer less privacy by design. Such a shift to reactive privacy means that, across the board, individuals enjoy less privacy than in a regime where responsible parties have incentives to apply precaution.

Observation 7: Mandatory responsible disclosure could lead to reactive privacy, reducing the overall level of privacy.

Finally, past research has also shown that identifiers associated with discovered vulnerabilities are misused or misunderstood by reviewers as proxies for impactful research [42]. Thus, more and broader obligations to disclose and report on disclosure results could shift the focus of reviews away from novelty and soundness towards successful reportability. This could discourage the type of innovative research that society needs most.

Observation 8: Broadening the scope of mandatory responsible disclosure could impede scientific progress.

5. Conclusion

We have started with the hypothesis that, although security and privacy research intersect both as research topics and in communities, the two differ in their culture and approaches to disclosure. Our analysis of the CfPs and papers at the top 10 security and privacy venues since 2018 supported this hypothesis. While the majority of CfPs either made no specific mention of disclosure or left the target of the disclosure unspecified, the corresponding papers exhibited a clear pattern: *security papers reported responsible disclosure, whereas privacy papers did not*. Even in the case of reported disclosures in privacy papers, the described approaches differed widely, further cementing the fact that not only is there no universally agreed-upon requirement, but also no standard for disclosure practices.

We followed up on our inventory of the current state of affairs by reviewing the reasoning and process for responsible security disclosure, trying to map them to privacy. During this process, we observed that the parallels between security and privacy are weak, and *a naive transfer of the disclosure requirement for security research onto the whole of privacy research could have unintended consequences*.

We conclude that the question of mandatory responsible disclosure for privacy research requires a broader community discussion. It cannot be reduced to copying established practices from security. This work should therefore be regarded as an informed starting point for a discussion that is overdue. Clearly, our initial set of arguments needs to be expanded at the workshop and in future work.

4. http://en.wikipedia.org/wiki/Cobra_effect

References

[1] Ashish Arora, Anand Nandkumar, and Rahul Telang. Does information security attack frequency increase with vulnerability disclosure? An empirical analysis. *Information Systems Frontiers*, 2007.

[2] Henry Birge-Lee, Yixin Sun, Anne Edmundson, Jennifer Rexford, and Prateek Mittal. Bamboozling Certificate Authorities with BGP. In *USENIX Security Symposium (USENIX Sec)*, 2018.

[3] Andrea Bisegna, Matteo Bitussi, Roberto Carbone, Luca Compagna, Silvio Ranise, and Avinash Sudhodanan. CSRFing the SSO Waves: Security Testing of SSO-Based Account Linking Process. In *IEEE European Symposium on Security and Privacy (EuroS&P)*, 2024.

[4] Rainer Böhme, Lisa Eckey, Tyler Moore, Neha Narula, Tim Ruffing, and Aviv Zohar. Responsible vulnerability disclosure in cryptocurrencies. *Communications of the ACM*, 63(10):62–71, 2020.

[5] Dino Bollinger, Karel Kubicek, Carlos Cotrini, and David Basin. Automating Cookie Consent and GDPR Violation Detection. In *USENIX Security Symposium (USENIX Sec)*, 2022.

[6] Ahmed Bouhoula, Karel Kubicek, Amit Zac, Carlos Cotrini, and David Basin. Automated large-scale analysis of cookie notice compliance. In *USENIX Security Symposium (USENIX Sec)*, 2024.

[7] Duc Bui, Brian Tang, and Kang G. Shin. Do Opt-Outs Really Opt Me Out? In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2022.

[8] Sheng Cao, Hao Zhou, Songzhou Shi, Yanjie Zhao, and Haoyu Wang. Parcel Mismatch Demystified: Addressing a Decade-Old Security Challenge in Android. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2025.

[9] Serge Egelman, Cormac Herley, and Paul C. Van Oorschot. Markets for zero-day exploits: Ethics and implications. In *NSPW '13: New Security Paradigms Workshop*, 2013.

[10] Benjamin Eriksson, Giancarlo Pellegrino, and Andrei Sabelfeld. Black Widow: Blackbox Data-driven Web Scanning. In *IEEE Symposium on Security and Privacy (SP)*, 2021.

[11] Matthew Finifter, Devdatta Akhawe, and David Wagner. An Empirical Study of Vulnerability Rewards Programs. In *22nd USENIX Security Symposium (USENIX Security 13)*, 2013.

[12] Behrad Garmany, Martin Stoffel, Robert Gawlik, and Thorsten Holz. Static Detection of Uninitialized Stack Variables in Binary Code. In *European Symposium on Research in Computer Security (ESORICS)*. 2019.

[13] Cassidy Gibson, Daniel Olszewski, Natalie Grace Brigham, Anna Crowder, Kevin R. B. Butler, Patrick Traynor, Elissa M. Redmiles, and Tadayoshi Kohno. Analyzing the AI nudification application ecosystem. In *USENIX Security Symposium (USENIX Sec)*, 2025.

[14] Enes Göktas, Kaveh Razavi, Georgios Portokalidis, Herbert Bos, and Cristiano Giuffrida. Speculative Probing: Hacking Blind in the Spectre Era. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2020.

[15] Nafis Tanveer Islam, Gonzalo De La Torre Parra, Dylan Manuel, Elias Bou-Harb, and Peyman Najafirad. An Unbiased Transformer Source Code Learning with Semantic Vulnerability Graph. In *2023 IEEE 8th European Symposium on Security and Privacy (EuroS&P)*, 2023.

[16] Gayatri Priyadarsini Kancherla, Nataliia Bielova, Cristiana Santos, and Abhishek Bichhawat. Johnny Can't Revoke Consent Either: Measuring Compliance of Consent Revocation on the Web. In *Privacy Enhancing Technologies Symposium (PETS)*, 2025.

[17] Kyungtae Kim, Taegyu Kim, Ertza Warraich, Byoungyoung Lee, Kevin R. B. Butler, Antonio Bianchi, and Dave Jing Tian. FuzzUSB: Hybrid Stateful Fuzzing of USB Gadget Stacks. In *IEEE Symposium on Security and Privacy (SP)*, 2022.

[18] Taegyu Kim, Chung Hwan Kim, Junghwan Rhee, Fan Fei, Zhan Tu, Gregory Walkup, Xiangyu Zhang, Xinyan Deng, and Dongyan Xu. RVFUZZER: Finding input validation bugs in robotic vehicles through control-guided testing. In *USENIX Security Symposium (USENIX Sec)*, 2019.

[19] Simon Koch, Malte Wessels, Benjamin Altpeter, Madita Olvermann, and Martin Johns. Keeping Privacy Labels Honest. *Privacy Enhancing Technologies Symposium (PETS)*, 2022.

[20] Paul C. Kocher. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. *Advances in Cryptology — CRYPTO '96*, 1996.

[21] Konrad Kollnig, Anastasia Shuba, Reuben Binns, Max Van Kleek, and Nigel Shadbolt. Are iPhones Really Better for Privacy? A Comparative Study of iOS and Android Apps. *Privacy Enhancing Technologies Symposium (PETS)*, 2022.

[22] Karel Kubíček, Jakob Merane, Carlos Cotrini, Alexander Stremitzer, Stefan Bechtold, and David Basin. Checking Websites' GDPR Consent Compliance for Marketing Emails. *Privacy Enhancing Technologies Symposium (PETS)*, 2022.

[23] Kyle-Kyle and Peace-Maker. top4grep. URL https://github.com/peace-maker/top4grep/tree/abstracts. visited: 10-02-2026.

[24] Song Li, Mingqing Kang, Jianwei Hou, and Yinzhi Cao. Mining Node.js Vulnerabilities via Object Dependence Graph and Query. In *USENIX Security Symposium (USENIX Sec)*, 2022.

[25] Yeting Li, Zixuan Chen, Jialun Cao, Zhiwu Xu, Qiancheng Peng, Haiming Chen, Liyuan Chen, and Shing-Chi Cheung. ReDoSHunter: A Combined Static and Dynamic Approach for Regular Expression DoS Detection. In *USENIX Security Symposium (USENIX Sec)*, 2021.

[26] Zhen Li, Deqing Zou, Shouhuai Xu, Xinyu Ou, Hai Jin, Sujuan Wang, Zhijun Deng, and Yuyi Zhong.

VulDeePecker: A Deep Learning-Based System for Vulnerability Detection. In *Network and Distributed System Security Symposium (NDSS)*, 2018.

[27] Zhengyu Liu, Kecheng An, and Yinzhi Cao. Undefined-oriented Programming: Detecting and Chaining Prototype Pollution Gadgets in Node.js Template Engines for Malicious Consequences. In *IEEE Symposium on Security and Privacy (SP)*, 2024.

[28] Fuchen Ma, Yuanliang Chen, Meng Ren, Yuanhang Zhou, Yu Jiang, Ting Chen, Huizhong Li, and Jiaguang Sun. LOKI: State-Aware Fuzzing Framework for the Implementation of Blockchain Consensus Protocols. In *Proceedings 2023 Network and Distributed System Security Symposium*, 2023.

[29] Max Maass, Alina Stöver, Henning Pridöhl, Sebastian Bretthauer, Dominik Herrmann, Matthias Hollick, and Indra Spiecker. Effective Notification Campaigns on the Web: A Matter of Trust, Framing, and Support. In *30th USENIX Security Symposium (USENIX Security 21)*, 2021.

[30] Célestin Matte, Nataliia Bielova, and Cristiana Santos. Do Cookie Banners Respect my Choice? : Measuring Legal Compliance of Banners from IAB Europe's Transparency and Consent Framework. In *IEEE Symposium on Security and Privacy (SP)*, 2020.

[31] Gilles Mertens, Nataliia Bielova, and Vincent Roca. You Can't Trust Your Tag Neither: Privacy Leaks and Potential Legal Violations within the Google Tag Manager. In *IEEE European Symposium on Security and Privacy (EuroS&P)*, 2025.

[32] Malte Möser, Kyle Soska, Ethan Heilman, Kevin Lee, Henry Heffan, Shashvat Srivastava, Kyle Hogan, Jason Hennessey, Andrew Miller, Arvind Narayanan, and Nicolas Christin. An Empirical Analysis of Traceability in the Monero Blockchain. *Proceedings on Privacy Enhancing Technologies*, 2018.

[33] Sebastian Neef, Lorenz Kleissner, and Jean-Pierre Seifert. What All the PHUZZ Is About: A Coverage-guided Fuzzer for Finding Vulnerabilities in PHP Web Applications. In *Proceedings of the 19th ACM Asia Conference on Computer and Communications Security*, 2024.

[34] Trung Tin Nguyen, Michael Backes, Ninja Marnau, and Ben Stock. Share First, Ask Later (or Never?) Studying Violations of GDPR's Explicit Consent in Android Apps. In *USENIX Security Symposium (USENIX Sec)*, 2021.

[35] Trung Tin Nguyen, Michael Backes, and Ben Stock. Freely Given Consent? Studying Consent Notice of Third-Party Tracking and Its Violations of GDPR in Android Apps. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2022.

[36] Haoran Ou, Yong Fang, Yongyan Guo, Wenbo Guo, and Cheng Huang. Viopolicy-Detector: An Automated Approach to Detecting GDPR Suspected Compliance Violations in Websites. In *International Symposium on Research in Attacks, Intrusions and Defenses (RAID)*, 2022.

[37] Elleen Pan, Jingjing Ren, Martina Lindorfer, Christo Wilson, and David Choffnes. Panoptispy: Characterizing Audio and Video Exfiltration from Android Applications. *Privacy Enhancing Technologies Symposium (PETS)*, 2018.

[38] Open Worldwide Application Security Project. Vulnerability disclosure cheat sheet. URL https://cheatsheetseries.owasp.org/cheatsheets/Vulnerability_Disclosure_Cheat_Sheet.html. visited: 10-02-2026.

[39] Irwin Reyes, Primal Wijesekera, Joel Reardon, Amit Elazari Bar On, Abbas Razaghpanah, Narseo Vallina-Rodriguez, and Serge Egelman. "Won't Somebody Think of the Children?" Examining COPPA Compliance at Scale. *Proceedings on Privacy Enhancing Technologies*, (3):63–83, 2018. doi: 10.1515/popets-2018-0021.

[40] Michael Rodler, Wenting Li, Ghassan O. Karame, and Lucas Davi. Sereum: Protecting Existing Smart Contracts Against Re-Entrancy Attacks. In *Proceedings 2019 Network and Distributed System Security Symposium*, 2019.

[41] Eyal Ronen, Adi Shamir, Achi-Or Weingarten, and Colin OFlynn. IoT Goes Nuclear: Creating a ZigBee Chain Reaction. In *2017 IEEE Symposium on Security and Privacy (SP)*, 2017.

[42] Moritz Schloegel, Daniel Klischies, Simon Koch, David Klein, Lukas Gerlach, Malte Wessels, Leon Trampert, Martin Johns, Mathy Vanhoef, Michael Schwarz, Thorsten Holz, and Jo Van Bulck. Confusing Value with Enumeration: Studying the Use of CVEs in Academia. In *34th USENIX Security Symposium (USENIX Security 25)*, 2025.

[43] Faysal Hossain Shezan, Zihao Su, Mingqing Kang, Nicholas Phair, Patrick William Thomas, Michelangelo Van Dam, Yinzhi Cao, and Yuan Tian. CHKPLUG: Checking GDPR Compliance of WordPress Plugins via Cross-language Code Property Graph. In *Network and Distributed System Security Symposium (NDSS)*, 2023.

[44] Aleksei Stafeev, Tim Recktenwald, Gianluca De Stefano, Soheil Khodayari, and Giancarlo Pellegrino. YuraScanner: A Task-driven Web Application Scanner. In *Network and Distributed System Security Symposium (NDSS)*, 2025.

[45] Ben Stock, Giancarlo Pellegrino, Christian Rossow, Martin Johns, and Michael Backes. Hey, You Have a Problem: On the Feasibility of Large-Scale Web Vulnerability Notification. In *25th USENIX Security Symposium (USENIX Security 16)*, 2016.

[46] Brian Tang, Duc Bui, and Kang G. Shin. Navigating cookie consent violations across the globe. In *USENIX Security Symposium (USENIX Sec)*, 2025.

[47] Neil P. Thimmaiah, Yashashvi J. Dave, Rigel Gjomemo, and V. N. Venkatakrishnan. FIXX: FInding eXploits from eXamples. In *USENIX Security Symposium (USENIX Sec)*, 2025.

[48] Michael Toth, Nataliia Bielova, and Vincent Roca. On

dark patterns and manipulation of website publishers by CMPs. *Privacy Enhancing Technologies Symposium (PETS)*, 2022.

[49] Muoi Tran, Inho Choi, Gi Jun Moon, Anh V. Vu, and Min Suk Kang. A Stealthier Partitioning Attack against Bitcoin Peer-to-Peer Network. In *2020 IEEE Symposium on Security and Privacy (SP)*, 2020.

[50] Benjamin E. Ujcich, Samuel Jero, Richard Skowyra, Steven R. Gomez, Adam Bates, William H. Sanders, and Hamed Okhravi. Automated Discovery of Cross-Plane Event-Based Vulnerabilities in Software-Defined Networking. In *Network and Distributed System Security Symposium (NDSS)*, 2020.

[51] Christine Utz, Matthias Michels, Martin Degeling, Ninja Marnau, and Ben Stock. Comparing Large-Scale Privacy and Security Notifications. *Proceedings on Privacy Enhancing Technologies*, 2023.

[52] Yi Wu, Zhuohang Li, Nicholas Van Nostrand, and Jian Liu. Time to Rethink the Design of Qi Standard? Security and Privacy Vulnerability Analysis of Qi Wireless Charging. In *Annual Computer Security Applications Conference (ACSAC)*, 2021.

[53] Anhao Xiang, Weiping Pei, and Chuan Yue. PolicyChecker: Analyzing the GDPR Completeness of Mobile Apps' Privacy Policies. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2023.

[54] Yuqing Yang, Yue Zhang, and Zhiqiang Lin. Cross Miniapp Request Forgery: Root Causes, Attacks, and Vulnerability Detection. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2022.

[55] Doguhan Yeke, Muhammad Ibrahim, Güliz Seray Tuncay, Habiba Farrukh, Abdullah Imran, Antonio Bianchi, and Z. Berkay Celik. Wear's my Data? Understanding the Cross-Device Runtime Permission Model in Wearables. In *IEEE Symposium on Security and Privacy (SP)*, 2024.

[56] Jianjia Yu, Song Li, Junmin Zhu, and Yinzhi Cao. CoCo: Efficient Browser Extension Vulnerability Detection via Coverage-guided, Concurrent Abstract Interpretation. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2023.