

Cybersecurity has become inseparable from everyday life; from hospitals and universities to the infrastructures that power science and government. Yet despite massive investments in automation, analytics, and defense technologies, major security failures continue to occur. These failures rarely result from missing tools alone; they emerge when people, processes, and technologies fall out of alignment or when the people responsible for security are constrained by the very systems meant to support them.

My research highlights these gaps through a human-centered lens. I ask two complementary questions: (1) *How do people shape security from the bottom up, through their expertise, collaboration, and everyday interactions?* and (2) *How does security shape people from the top down, through the infrastructures and policies that govern their work?*

I study cybersecurity as a sociotechnical system that integrates human reasoning, organizational processes, and computational infrastructure. Methodologically, my work blends social science and security practice by using qualitative and mixed-methods techniques including interviews, surveys, content analysis, and process tracing.

Together, these questions structure my research across two levels, *practice and structure*, linking human reasoning with institutional security norms and structure. My research collectively underscores *security as alignment work*, the continuous process of reconciling human reasoning, organizational practices, and institutional systems. This human-centered perspective informs practical improvements within critical systems and infrastructures, strengthening organizational resilience, supporting the people who rely on them, and advancing human well-being.

### **Main Topic 1: How do People Shape Security in Practice?**

Humans form a critical layer within organizational cybersecurity. They may be analysts triaging and escalating incidents, managers prioritizing security investments and overseeing operations, white-hat hackers experimenting with new exploits, researchers collaborating across institutions, or leaders translating policy into practice. Each plays a distinct yet interconnected role to ensure overall security posture. My work in this area examines how resilience is constructed and sometimes compromised through human and organizational practices at different layers.

At the operational layer, our CCS 2019 paper [1] examined *how Security Operations Centers (SOCs) function from the perspective of the people who run them: analysts and managers*. Despite SOCs being the core of enterprise defense, serious incidents continue to occur even in well-resourced organizations. Beyond improving technology, we wanted to understand what SOC professionals perceived as issues resulting into inefficiencies. We turned our investigative lens to the everyday human factors of SOC operations. We found that managers and analysts valued certain key aspects of SOC operations differently: analysts emphasized visibility and log quality, whereas managers prioritized automation, response speed, and performance metrics. These mismatched priorities were sometimes sharply opposed. The study reframes SOC performance as a problem of differing mental models and shows how analysts and managers value and make sense of the same systems in divergent ways.

At the individual level, our IEEE S&P 2024 paper [2] examined *how security expertise unfolds in practice; specifically, how hackers make and recover from mistakes while solving binary exploitation challenges*. We hypothesized that much of this effort stems from the mistakes hackers make while exploiting the targets; understanding those matter because they reveal how experts reason through uncertainty and adapt when things go wrong. To observe authentic real-world hacking behavior, we analyzed 30 screencasts from 11 hackers uploaded to YouTube. Mistakes consumed nearly 40% of total work time; most stemming from cognitive (memory lapses and missteps) rather than technical causes. Yet these errors also revealed adaptive expertise: hackers debugged collaboratively and adjusted strategies. The findings highlight that learning to recognize and recover from mistakes is central to developing expertise, suggesting practical opportunities for hacker education and debugging support.

At the institutional level, our IEEE S&P 2025 work [3] examined how research institutions manage security and collaboration within Research Computing Infrastructures (RCIs)—the high-performance systems that now underpin data-driven science and enable large-scale, collaborative research across disciplines and institutions. Through interviews with 24 researchers and system administrators, we found that everyday research often relies on informal, trust-based sharing practices, and inherited permissions that keep projects agile but create lasting security exposures. Administrators described difficulty balancing openness and compliance, while researchers viewed rigid controls as barriers to productivity. Our findings highlight that securing RCIs requires designing systems that make security a shared and sustainable practice across technical and human boundaries, revealing that *trust, while convenient, is flawed and a vulnerability in itself*.

At the strategic level, in my dissertation research (under review), I explore how Zero Trust principles are being interpreted and implemented across academic, industry, and government organizations. Across the broader cybersecurity community, the move away from implicit trust has crystallized in the adoption of Zero Trust architectures; an approach that redefines security by assuming no implicit trust and ensures continuous verification. Through interviews with 27 cybersecurity professionals, we found that while Zero Trust has become a policy mandate for government organizations, its meaning and application vary widely across sectors. For some, it represents a genuine rethinking of security culture; for others, it is a branding exercise driven by vendors. Many grapple with how to reconcile its ideals with existing organizational norms, legacy systems, and human workflows. Rather than treating Zero Trust as a purely

technical or policy directive, my work examines how its implementation interacts with organizational culture, existing infrastructures, and human constraints and what it takes to make “continuous verification” viable in practice. This work extends the notion of resilience beyond defense and recovery, toward the capacity to continuously evolve security culture in response to shifting technologies, mandates, and threats.

More recently, I extended this line of inquiry into human–AI teaming by examining how large language models (LLMs) enter and reshape security workflows. I collaborated on a controlled human study of novices and experts reverse engineers working with LLM-assistance [4]. We found that LLM support significantly improved novice comprehension while offering limited benefits for experts and occasionally introducing misleading suggestions. The results show that AI functions not as a replacement for expertise, but as a conditional collaborator whose impact depends on task structure and patterns of reliance.

### **Main Topic 2: How does Security Shape People and Institutions?**

As education has become increasingly dependent on interconnected educational technologies (EdTech), governance has struggled to keep pace with expanding data flows and rising breach frequency.

To investigate how such decisions are made in practice, our CCS 2024 work [5] examined the *EdTech acquisition process* through interviews with people in leadership positions within learning enterprises. We found that procurement decisions were driven by pedagogical, accessibility, and cost considerations, not the security and safety of vendor tools. Once adopted, these tools became deeply embedded in institutional networks with little visibility into data exchanges among vendors and their sub-vendors. While contracts can restrict vendors’ use of institutional data, the lack of visibility leaves HEIs unaware of potential misuse until after the fact, creating a governance structure that exists but struggles to act. Off-boarding processes are similarly fraught with uncertainty: institutions rarely have the power or technical means to verify data deletion. This raised a troubling question: if institutions already struggle to govern the tools they formally acquire, *what happens when technologies enter classrooms through informal channels?*

A follow-up study [6] surveyed 375 educators across U.S. K-12 and higher education institutions to investigate the landscape of tools adopted by instructors outside formal procurement. Educators chose tools largely for instructional utility and student engagement rather than their data protection and security features, and many described knowingly bypassing institutional policy to better support their classrooms. These choices, while well-intentioned, create invisible data flows between unvetted vendors, expanding exposure to breaches and non-compliance with FERPA and HIPAA. Together, these studies reveal that governance gaps in education are not isolated policy failures but systemic conditions embedded in everyday institutional practice.

Building on this work, I am currently mentoring a high school student through an 8-month internship period on a project examining online discourse surrounding a major educational technology breach. By analyzing discussions on platforms such as Reddit, we investigate how affected communities seek advice, assign responsibility, and articulate frustration in the aftermath. This project extends my work on governance by examining how breach consequences are experienced and interpreted by end users.

### **Main Topic 3: Meta Research and Science of Security**

Another dimension of my research focuses *inward* on how the security community itself produces, evaluates, and sustains its own knowledge. My IEEE S&P 2022 paper [7] on the cybersecurity peer-review process was the first to investigate how the security research community decides what counts as good science and how those decisions shape what becomes accepted knowledge. Through interviews with 21 program committee members and chairs across top-tier conferences, we examined how reviewers interpret core evaluation criteria, assume reviewing responsibilities, and define review quality. We found that novelty was universally valued; however, its meaning varied widely, revealing the lack of shared standards for evaluating research and fueling anecdotal perceptions of “randomness” in outcomes. Efforts by authors to game the system added further complexity into this notion which resulted into extra reviewing load while authors try to *get lucky*. These dynamics were further complicated by systemic pressures such as record submission volumes over rolling deadlines and overlapping service-related commitments. By empirical examination, this study helped scientifically and systematically voice a range of long-standing community concerns, transforming anecdotal frustrations into documented evidence that the field could analyze and build upon. Moreover, many advisors include this paper in graduate courses and lab reading groups as an accessible introduction to how scientific cybersecurity research and its peer-review process operate; a system that is often a black box to students entering the field.

### **Other Research**

Alongside my primary work on organizational security, governance, and meta-research, my work on the adoption of OnlyFans [8] examined why creators without prior experience in sex work joined the platform and how its affordances such as boundary setting, privacy, and self-governance enabled agency within stigmatized digital labor. My USENIX Security 2024 paper [9] on digital safety and privacy investigated how OnlyFans creators experience and mitigate risks such as harassment, censorship, and deplatforming, revealing how informal security practices evolve in response to

stigma, prominence, and uncertainties in platform policies. Together, these studies extend my broader research on security governance to the context of digital platforms where individuals, rather than institutions, bear the responsibility of managing risk and trust.

### Future Work

My future research will deepen and expand the human-centered understanding of cybersecurity that anchors my doctoral work. I will expand my meta-research agenda to examine how AI systems influence research practices within computer security. I am particularly interested in how researchers perceive and integrate LLMs into human factors research, how such use shapes study considerations and analysis practices, and how peer review processes adapt to AI-assisted scholarship.

During the postdoctoral tenure, I will investigate how strategic paradigms such as Zero Trust are operationalized across organizational levels—from executives to frontline practitioners. Importantly, the rapid rise of AI/LLM is reshaping the cybersecurity landscape, fueling a digital arms race in which both attackers and defenders harness the same technologies. Modern cybersecurity paradigms such as Zero Trust already aim to minimize implicit trust by continuously validating identities and access requests. A central paradox emerges: *to strengthen security, we must trust AI systems to act intelligently, yet also mistrust them enough to retain control, accountability, and human judgment.* These tensions define the space my research seeks to address: “how trust in AI can be systematically designed, calibrated, and governed in modern cybersecurity systems?” My investigation aims to answer following questions: How can systems quantify and recalibrate trust as AI behavior evolves? What mechanisms allow agents to earn, lose, or regain privilege based on performance and context? How do executives, policymakers, employees, and end-users align expectations of AI accountability and acceptable autonomy? What governance structures foster transparency and shared responsibility without constraining innovation? How do cognitive load, fatigue, and stress shape security decision-making in AI-assisted operations? What role does human well-being play in cyber defense posture?

As organizations adopt autonomous AI agents that learn, reason, and act across boundaries, the assumptions underlying these architectures begin to strain. Existing models of verification and access control presume static identities, predictable behaviors, and clearly defined system boundaries; all of which dissolve when AI systems can reinterpret goals, delegate authority, and make context-dependent decisions. At the same time, humans across different skills, expertise, and organizational roles must learn to collaborate with systems whose reasoning is partially opaque and whose speed outpaces deliberate reflection. This combination of *technical autonomy* and *human cognitive vulnerability* introduces new risks: misplaced trust, automation bias, accountability gaps, emotional strain, and fatigue. As AI systems become embedded in defensive environments, I aim to examine how responsibilities are redistributed in increasingly “autonomous” security settings.

By examining multi-stakeholder aspects in emerging ecosystems, I aim to make the human layer more resilient, and by extension, strengthen the resilience of organizations. This reflects my broader research ethos: to redefine *the human not as the weakest link, but as an indispensable part of the solution.*

## References

- [1] Faris Bugra Kokulu, Ananta Soneji, Tiffany Bao, Yan Shoshitaishvili, Ziming Zhao, Adam Doupé, and Gail-Joon Ahn. Matched and mismatched socs: A qualitative study on security operations center issues. In *Proceedings of the 2019 ACM SIGSAC conference on computer and communications security*, pages 1955–1970, 2019.
- [2] Irina Ford, Ananta Soneji, Faris Bugra Kokulu, Jayakrishna Vadayath, Zion Leonahenahe Basque, Gaurav Vipat, Adam Doupé, Ruoyu Wang, Gail-Joon Ahn, Tiffany Bao, et al. “watching over the shoulder of a professional”: Why hackers make mistakes and how they fix them. In *2024 IEEE Symposium on Security and Privacy (SP)*, pages 350–368. IEEE, 2024.
- [3] Souradip Nath, Ananta Soneji, Jaejong Baek, Tiffany Bao, Adam Doupé, Carlos Rubio-Medrano, and Gail-Joon Ahn. “It’s almost like Frankenstein”’: Investigating the Complexities of Scientific Collaboration and Privilege Management within Research Computing Infrastructures. In *2025 IEEE Symposium on Security and Privacy (SP)*, pages 2995–3013. IEEE Computer Society, 2025.
- [4] Zion Leonahenahe Basque, Samuele Doria, Ananta Soneji, Wil Gibbs, Adam Doupé, Yan Shoshitaishvili, Eleonora Losiouk, Ruoyu Wang, and Simone Aonzo. Decompiling the synergy: An empirical study of human–llm teaming in software reverse engineering. 2026.
- [5] Easton Kelso, Ananta Soneji, Sazzadur Rahaman, Yan Shoshitaishvili, and Rakibul Hasan. Trust, Because You Can’t Verify: Privacy and Security Hurdles in Education Technology Acquisition Practices. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, pages 1656–1670, 2024.

- [6] Easton Kelso, Ananta Soneji, Syed Zami-Ul-Haque Navid, Yan Shoshitaishvili, Sazzadur Rahaman, and Rakibul Hasan. Investigating the Security & Privacy Risks from Unsanctioned Technology Use by Educators. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, pages 1–6, 2025.
- [7] Ananta Soneji, Faris Bugra Kokulu, Carlos Rubio-Medrano, Tiffany Bao, Ruoyu Wang, Yan Shoshitaishvili, and Adam Doupé. “Flawed, but like democracy we don’t have a better system”’: The Experts’ Insights on the Peer Review Process of Evaluating Security Papers. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 1845–1862. IEEE, 2022.
- [8] Vaughn Hamilton, Ananta Soneji, Allison McDonald, and Elissa M Redmiles. “Nudes? Shouldn’t I charge for these?”: Motivations of New Sexual Content Creators on OnlyFans. In *Proceedings of the 2023 CHI conference on human factors in computing systems*, pages 1–14, 2023.
- [9] Ananta Soneji, Vaughn Hamilton, Adam Doupé, Allison McDonald, and Elissa M Redmiles. “I feel physically safe but not politically safe”’: Understanding the Digital Threats and Safety Practices of {OnlyFans} Creators. In *33rd USENIX Security Symposium (USENIX Security 24)*, pages 1–18, 2024.