USENIX

THE ADVANCED COMPUTING SYSTEMS ASSOCIATION

SoK: Beyond Burnout – A Gap Analysis of Psychological Harm in Cybersecurity Work

Jessica Phelan and Ivan Flechais, *University of Oxford*

https://www.usenix.org/conference/usenixsecurity26/presentation/phelan

This paper is included in the Proceedings of the 35th USENIX Security Symposium.

August 12–14, 2026 • Baltimore, MD, USA

ISBN 978-1-939133-58-8

Open access to the Proceedings of the 35th USENIX Security Symposium is sponsored by

جامعة الملك عبدالله للعلوم والتقنية

King Abdullah University of Science and Technology

SoK: Beyond Burnout – A Gap Analysis of Psychological Harm in Cybersecurity Work

Jessica Phelan
University of Oxford

Ivan Flechais
University of Oxford

Abstract

Burnout among cybersecurity professionals drives workforce attrition, but its consequences extend beyond staffing shortages. Degraded judgment and human error, particularly during incident response, can be catastrophic. Yet the sources and forms of psychological harm across cybersecurity roles remain poorly understood. This Systematization of Knowledge (SoK) paper examines the literature on wellbeing and psychological harm in cybersecurity work. Using a scoping review across five databases, we review 101 papers and synthesize 57 covering wellbeing research on cybersecurity practitioners and four adjacent professions selected for structural similarity. We find that research on offensive cyber roles emphasizes harm driven by operational conditions, while the work itself is often experienced as meaningful. In contrast, research on defensive roles highlights stressors intrinsic to the work, including sustained vigilance, exposure to human malice, and invisible success metrics. Adjacent professions facing similar demands consistently report psychological harm extending beyond burnout. How such harm manifests in cybersecurity practitioners remains unclear, with existing studies measuring burnout, cognitive workload, fatigue, clinical distress, or wellbeing, but no studies to date assessing whether psychological harms beyond these constructs are present in this population. We identify cybersecurity-specific barriers to processing occupational stress, including confidentiality constraints, limited institutional support, cultures discouraging help-seeking, and temporal dynamics that prevent psychological closure. Our findings provide a structured foundation for research to explore cybersecurity interventions calibrated to the psychological harms the work actually produces.

1 Introduction

Given today's rapidly evolving threat landscape and dependence on cybersecurity professionals to defend against increasingly adaptive adversaries, it seems logical to invest in the workforce responsible for such a crucial job. Despite this importance, industry surveys consistently report high burnout rates and persistent talent shortages [19, 35, 52]. Taxonomies exist for the harms cyber-attacks inflict on organizations and individuals [2], but the psychological harms experienced by practitioners who defend against these attacks have received less systematic attention. More recently, researchers have begun examining individual-level factors including burnout, fatigue, cognitive overload, and exhaustion. This is a welcome development, but it raises a more fundamental question: burnout in which roles, and under what conditions?

Cybersecurity encompasses a wide range of roles that vary significantly in the nature of the work, from offensive operations that initiate action against targets to defensive work that responds to adversary-initiated events. These distinctions matter because different modes of work may produce distinct stressors, distinct psychological harms, or require different intervention approaches, yet no research has systematically examined whether and how they diverge. Without this differentiation, "cybersecurity burnout" remains a monolithic label applied across contexts that may share little beyond a common professional domain.

The limited research that does exist on cybersecurity practitioner wellbeing underscores this problem. Studies tend to focus on offensive roles in military contexts (NSA and USCYBERCOM tactical operators) or on crisis-oriented defensive roles (SOC analysts, incident responders) in civilian settings. Sustained operational roles such as protection and compliance work are largely unexamined, and no study has compared populations directly across role types. Despite this fragmentation, the reported prevalence of burnout across cybersecurity professions and the documented stressors available in the literature suggest that the harms practitioners experience bear structural similarities to those in adjacent professions such as emergency services, healthcare, content moderation, and air traffic control [17, 30, 32].

The occupational health literature on these adjacent professions is far more mature and consistently documents psychological harm beyond burnout alone, including trauma-related conditions that imply different processing pathways and re-

quire different interventions [3, 9, 21, 23]. Whether cybersecurity practitioners experience similar harm has not been assessed. This gap matters. If cybersecurity work produces harm beyond burnout, interventions designed for burnout alone may prove insufficient. Moreover, if processing pathways differ across constructs, the barriers practitioners face may obstruct recovery in ways that standard workplace wellness programs do not address. This review examines what we know, and what remains unknown, about the psychological harms cybersecurity practitioners face and how they might be mitigated.

RQ1: What are the documented stressors and sources of psychological harm in the existing cybersecurity literature? What gaps remain in understanding practitioners' lived experience?

RQ2: To what extent can research on burnout and psychological harm in adjacent professions be applied to cybersecurity?

RQ3: What mitigations for occupational stress and psychological harms exist in the literature for cybersecurity and adjacent professions? What remains empirically uninvestigated?

2 Background

This section introduces psychological theories that contextualize our findings. These frameworks were not part of our formal search criteria but provide essential grounding for interpreting documented stressors and harms.

2.1 Job Demands-Resources Model

The Job Demands-Resources (JD-R) model [6, 8] provides the dominant explanatory framework in occupational health psychology and has been referenced in recent cybersecurity practitioner research [14, 41]. The model posits two parallel psychological processes. The *health impairment process* shows how high job demands (workload, time pressure, emotional demands) deplete mental and physical resources, producing chronic exhaustion that progresses to burnout and negative outcomes. The *motivational process* operates in reverse: job resources (autonomy, feedback, social support) foster intrinsic motivation and work engagement. Critically, job resources buffer the impact of demands on strain. This interaction effect means resources become especially protective when demands are high.

2.2 Conservation of Resources Theory

Hobfoll's Conservation of Resources (COR) theory begins from the basic tenet that individuals strive to obtain, retain, and protect the things they centrally value [15, 16]. In a work setting, these valued resources span not only conditions such as job security and adequate tooling, but also aspects such as time, expertise, and a sense of professional competence. These resources are continually expended on the job and must be replenished for practitioners to function. COR posits that stress occurs when valued resources are threatened, lost, or when effort fails to produce expected gains. COR's most powerful contribution is the *loss spiral* mechanism: initial resource loss creates vulnerability to further loss, triggering an accelerating downward cascade. This explains how chronic workplace demands, without adequate resource replenishment, compound over time into exhaustion and breakdown.

2.3 Burnout

Maslach's burnout model [28, 29] defines the syndrome through three components: emotional exhaustion (feeling depleted), depersonalization or cynicism (detached response to work), and reduced personal accomplishment (declining self-efficacy). Research consistently shows a sequential progression: emotional exhaustion develops first, precipitating depersonalization as a coping mechanism, with reduced personal accomplishment emerging as cynicism persists.

Burnout is the most extensively studied construct in cybersecurity practitioner research, measured with validated instruments including the MBI-GS, BAT, and CBI [7, 32, 40, 53]. Cognitive workload and fatigue have been measured using the NASA-TLX and Samn-Perelli Fatigue Scale in tactical operator and experimental samples [10, 36, 44]. Clinical distress has been measured in a single study of U.S. Air Force cyber operators using the OQ-45.2 [7], and wellbeing and flow have been measured in a single pilot study of SOC practitioners using the Secure Flourish Index and Short Flow Scale [53]. Section 4.4 returns to what these instruments do and do not cover.

3 Methods

This SoK paper employs scoping review methodology [4, 31] to map the extent and nature of research on psychological harms in cybersecurity work. Scoping reviews are appropriate when a field is too fragmented for systematic review or meta-analysis, and when the aim is to identify gaps and research priorities rather than synthesize effect sizes. We synthesized findings using thematic analysis organized by our research questions [20]. Given the underdeveloped state of this literature, our goals are to: (a) map what is currently known about psychological harms in cybersecurity practitioners and adjacent professions, (b) identify which constructs have been measured and which remain uninvestigated, and (c) surface gaps warranting primary research.

3.1 Search Strategy

We searched three databases (ACM Digital Library, IEEE Xplore, and PubMed) for publications between 2005 and 2025. Preliminary searches of Google Scholar and SOLO/Bodleian Libraries confirmed substantial overlap with primary databases and were used to validate coverage rather than for formal screening. The 2005 lower bound reflects the substantial evolution in cybersecurity roles, threat landscapes, and working conditions. Older literature is unlikely to reflect contemporary considerations. Psychological theory used for background (Section 2) is not bound by this date range.

Search terms included:

- **Primary cybersecurity terms:** "cybersecurity" AND ("burnout" OR "stress" OR "wellbeing" OR "mental health" OR "psychological" OR "trauma" OR "PTSD")
- **Role-specific terms:** ("SOC analyst" OR "incident response" OR "incident responder" OR "CISO" OR "security operations") AND ("stress" OR "burnout" OR "wellbeing" OR "trauma")
- **Adjacent profession terms:** ("content moderation" OR "air traffic control" OR "emergency dispatch" OR "healthcare worker" OR "mental health professional" OR "therapist" OR "counselor" OR "psychiatrist") AND ("burnout" OR "trauma" OR "PTSD" OR "vicarious trauma" OR "secondary traumatic stress") AND ("occupational" OR "workplace")

Search strategies were adapted to database-specific syntax (e.g., PubMed's [tiab] field tags). Additional sources were identified through backward citation searching from key papers, conducted as a second search round after the database search. The first author conducted the formal database searches and first-pass title and abstract screening. The second author was actively involved throughout, reviewing inclusion and exclusion decisions, the thematic analysis of abstracts, and emergent themes. Formal inter-rater reliability and filtration efficacy were not quantitatively measured, as scoping review methodology does not require this [31].

3.2 Inclusion and Exclusion Criteria

For cybersecurity sources, we included peer-reviewed empirical studies, doctoral dissertations, and government technical reports that collected primary data from cybersecurity practitioners and measured at least one psychological outcome (burnout, stress, fatigue, distress, wellbeing) or cognitive workload. We also screened for any cybersecurity studies that had assessed clinical psychological constructs beyond those listed above using validated clinical instruments, applying the same search terms used for our adjacent profession criteria. This screen was conducted to determine whether such studies existed in the cybersecurity literature, not to presuppose that any particular construct applies to cybersecurity practitioners. Section 4.4 reports what this search surfaced. We also included experimental studies examining cybersecurity tasks with non-practitioner samples for their theoretical contribution to understanding task demands. Conceptual papers and industry surveys were retained for context but classified separately from empirical evidence.

Sources were excluded if they: (1) examined end-user security fatigue rather than practitioner wellbeing; (2) were vendor white papers or industry reports that lacked transparent reporting of sampling, measures, and analytic approach; or (3) fell outside the 2005–2025 date range.

A key challenge is that search terms like "cybersecurity stress" surface papers spanning distinct populations (security practitioners, general IT staff, and users subject to security policies) that face fundamentally different occupational demands. Sources that addressed these broader populations were retained for contextual insight where relevant.

For adjacent profession sources, we included peer-reviewed empirical studies and systematic reviews examining occupational stress, burnout, PTSD, secondary traumatic stress, or vicarious trauma in four comparison professions: emergency services, healthcare, content moderation, and air traffic control. These professions were selected because they share structural features hypothesized to drive psychological harm in cybersecurity work: sustained vigilance under uncertainty (air traffic control), repeated exposure to human suffering or disturbing content (healthcare, content moderation), acute incident response under time pressure with high-stakes consequences (emergency services), and chronic workforce shortages that intensify individual workload (all four). By examining professions where occupational psychology research is more mature, we aimed to identify harm mechanisms and protective factors that may transfer to cybersecurity contexts.

3.3 Selection Results

Figure 1 summarizes the selection process. Database searching across ACM Digital Library, IEEE Xplore, and PubMed produced the initial pool of records, followed by a second search round of backward citation searching from key papers identified during the initial review. Records from both rounds were combined into a single screening pool of 2,335 records. After removing 254 duplicates, 2,081 records were screened against the inclusion criteria. Of these, 1,980 were excluded as not directly relevant to our research questions, leaving 101 sources for review: 43 cybersecurity-specific sources, 47 adjacent profession sources, and 11 sources addressing burnout and psychological theory across occupations.

In a *breadth pass*, we coded the 101 abstracts to identify emerging themes and to surface which sources addressed each research question. In a *depth pass*, sources whose themes aligned with the research questions were fully coded, producing the 57 sources carried into synthesis.

Of the 57 sources informing the synthesis, 24 are cybersecurity sources that directly addressed one of the research

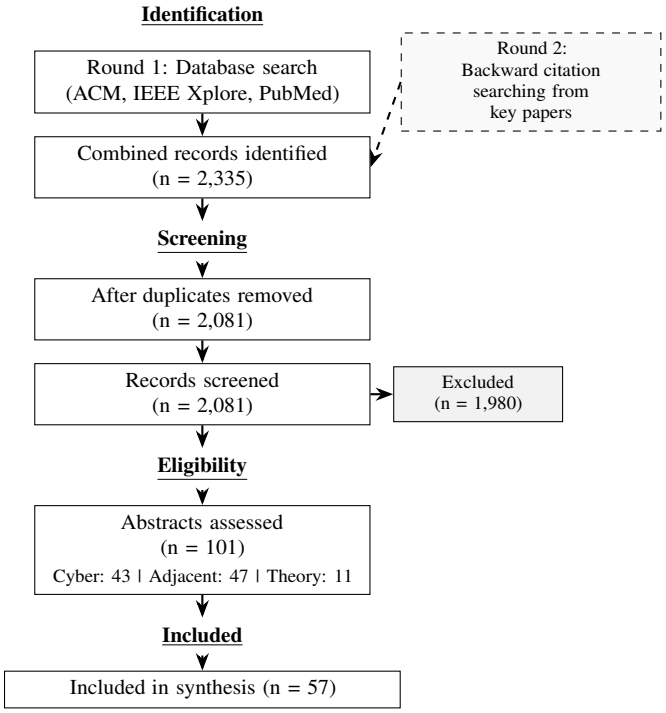

Figure 1: PRISMA flow diagram. Backward citation searching from key papers identified during the initial review served as a second search round; records from both rounds were combined into the screening pool of 2,335.

questions. Table 1 presents these by study type.

3.4 Analytical Approach

The two-stage coding process produced themes organized by each research question. For RQ1 (cybersecurity stressors), we produced three organizing themes: the offensive/defensive distinction in how harm is studied, the seven stressor dimensions documented within defensive work (which functioned as sub-themes under the second), and the questions that remain uninvestigated. For RQ2 (adjacent profession comparison), we produced mappings from each defensive stressor dimension to the harms documented in four adjacent professions, plus a parallel analysis of structural differences that limit transfer. For RQ3 (recovery barriers), we produced four barriers to processing occupational stress and a mitigation-gap analysis comparing cybersecurity to adjacent professions.

4 RQ1: Psychological Harm in Cybersecurity Work

This section examines what the cybersecurity literature does and does not reveal about psychological harm among practitioners. We organize findings around three themes: the distinction between offensive and defensive roles that shapes which questions have been asked; the stressors and harms that have been documented within defensive work; and the questions that remain uninvestigated. We reserve the comparison with adjacent professions for RQ2, and the discussion of recovery barriers and mitigations for RQ3.

4.1 The Structure of Cybersecurity Work: Offensive vs. Defensive Roles

Cybersecurity encompasses roles that differ fundamentally in their relationship to adversaries. Offensive roles (military cyber operators, penetration testers, red teamers) involve proactive engagement where operators initiate action against targets. Defensive roles (SOC analysts, incident responders, threat intelligence analysts) involve reactive postures where practitioners respond to adversary-initiated events. This distinction appears to shape both the nature of stressors experienced and the harms that result.

Research on offensive roles derives primarily from military populations (U.S. Air Force, NSA, USCYBERCOM), with civilian offensive practitioners (penetration testers, red teamers) largely unstudied. Research on defensive roles spans military and civilian contexts but focuses heavily on crisis-oriented work (incident response, SOC operations) rather than sustained defensive functions (security architecture, vulnerability management, threat intelligence).

Spears (2023) interviewed 14 cybersecurity professionals spanning security analysts, GRC, penetration testing, and security leadership [48]. Participants who had switched from incident response to less reactive roles, whether internal GRC positions or external consulting, reported stress levels dropping from 4 to 2 on a 5-point scale. Although Spears does not frame this as an offensive-defensive distinction, one participant articulated it directly: “My job is to pen test...My job is not to defend. [Conversely, if] my job is to make sure all of our security systems and things are working to defend us against attack. That’s stressful. That’s hard." [48].

Table 2 summarizes key differences between role types as documented in the literature. Readers should note that this comparison synthesizes findings across separate research populations using different methods. No study has directly compared offensive and defensive practitioners.

4.2 What Has Been Documented in Offensive Roles

Research on offensive cyber operators consistently identifies operational conditions, rather than the adversarial nature of the work, as primary stressors. In the first empirical study of burnout among cyber warfare operators, Chappelle et al. (2013) found that 26% of active duty U.S. Air Force operators (n=376) and 15% of civilian cyber contractors (n=156) reported high emotional exhaustion, with qualitative analysis identifying “shift work, shift changes, and hours worked"

rather than cyber warfare stressors as the primary sources of occupational stress (p. 1) [7]. However, "nature of work" appeared among top stressors for cyber operators but not for a non-cyber military control group (n=795), suggesting something about cyber work itself may contribute beyond typical operational demands [7]. Paul and Dykstra (2017) similarly pointed to operational conditions, finding that fatigue among NSA tactical operators increased 16% and frustration 12% over the course of operations, driven primarily by operation length [36].

Rangarajan et al. (2025) interviewed 10 former NSA tactical cyber operators with an average of 8.4 years of active service [39]. Despite experiencing burnout severe enough to leave their roles, eight of ten participants "spoke enthusiastically about having meaningful work with a sense of purpose and satisfaction," with three describing the mission as "addicting" or "riding the dragon" (p. 8) [39]. Notably, none cited pay as a reason for leaving. The desire for better work-life balance drove departures. The authors suggest that "overwhelming love of the mission may have masked negative environmental factors and personal impacts" (p. 9) [39].

The pattern of operational conditions driving harm and work content providing meaning has implications for intervention design. As the work is seen as rewarding, interventions targeting workload, scheduling, and work-life balance may be more effective than those addressing work content.

4.3 What Has Been Documented in Defensive Roles

Research on defensive practitioners suggests a different pattern, in which stressors appear tied not only to operational conditions but to the nature of the work itself.

Emotional Labor and Exhaustion. Three studies examining defensive practitioners document a consistent pattern: the work demands sustained emotional regulation under extreme conditions, and the resulting exhaustion drives practitioners not just to burnout, but out of the field entirely.

Hollis (2023), examining 22 incident responders who had worked major incidents using interpretative phenomenological analysis, identified emotional labor as a previously unexamined dimension in cybersecurity work, noting that "the concept of emotional labor was not found in any of the cybersecurity literature reviewed" (p. 87) [17]. Practitioners described managing their own fear and exhaustion while projecting confidence to stakeholders, what Hollis (2023) termed "rapid succession micro resiliencies" (p. 85) [17]. This emotional regulation occurred alongside extreme physical demands. One participant reported working "100-150 hours a week for the course of a better part of a year and a half" (p. 45), while another stated "I risk dying in the chair" (p. 38) [17]. The pressure to continue was so intense that one participant's team member suffered a heart attack during a 72-hour continuous response yet kept working [17]. Mott et al. (2024) found similar patterns among 83 ransomware response professionals, where participants described neglecting basic needs under sustained pressure: "I forgot to drink, eat ... one of [my team] was hospitalized for a few days, just through not caring for themselves" (p. 8) [30]. Nepal et al. (2024) provide quantitative evidence of the scope, finding that over half of 35 incident responders at a major technology corporation (54%, n=19) met burnout criteria, with exhaustion as the dominant symptom [32].

What distinguishes this from standard burnout reporting is the severity of the aftermath. Participants did not describe gradual disengagement, they described transformation. Hollis (2023) found that "these incidents transform everybody who's in it" (p. 40), with participants reporting acute emotional distress: "I would go cry in the bathtub every night...I really wanted to never work in cybersecurity again" (p. 33) [17]. Mott et al. (2024) documented a participant who described "PTSD every time I walked through the office door ... I was at times suicidal" (p. 8) [30]. Departure was framed not as career progression but as escape: "I did 18 [months] and I was done. I was burned out...I literally took a 50% pay cut to go do something else" (p. 34), with another observing that "if you make it two or three years, you've made it to veteran status" (p. 34) [17].

Participants who held both cybersecurity and traditional first responder roles reported that the psychological responses were "exactly the same," but that cybersecurity responders faced additional challenges: limited institutional mental health support, low professional recognition, incidents lasting months rather than hours, and confidentiality constraints preventing disclosure even to family [17]. This comparison is significant because it comes from practitioners with direct experience in both contexts rather than from researchers drawing analogies across separate literatures.

Sustained Vigilance. Cybersecurity monitoring demands continuous attention under conditions that consistently degrade human performance. Sawyer et al. (2014) demonstrated this experimentally, testing 24 participants on a 40-minute simulated cyber defense task and finding significant performance degradation over time alongside exceptionally high mental workload scores on the NASA Task Load Index [44].

This experimental finding aligns with what practitioners report. Chappelle et al. (2013) identified the "ubiquitous" nature of cyber threats as a distinct stressor, noting that operators must contend with the realization that many "adversarial cyber threats are ever present, requiring operators to continually race to expand their scope of capabilities within their assigned cyberspace domain" (p. 4) [7]. Sundaramurthy et al. (2015), embedded in a corporate SOC for six months, observed analysts performing repetitive monitoring procedures to the point where "they were starting to feel like robots performing the same tasks everyday" (p. 8) [50]. Tariq et al. (2025) found that SOC analysts face false positive rates as high as 40%, creating a task environment that simultaneously demands constant

attention and systematically fails to reward it [51].

Weight of Responsibility. Defensive cybersecurity practitioners bear an asymmetric burden, in that defenders must succeed every time while attackers only need to succeed once [17, 30]. Spears (2023) found this was a major theme among practitioners, with one participant describing "knowing the responsibility of being the responsible party for the client information, knowing that we would have to send the breach notification letters, go on the CMS wall of shame. It was a pretty heavy responsibility" [48]. This pressure is compounded by personal accountability. Mott et al. (2024) found that ransomware incidents placed existential organizational threats on the shoulders of IT and cybersecurity staff, with CISOs unable to reach senior executives for decisions and staff bearing responsibility for outcomes they could not fully control [30].

Nominet (2019) found that only 60% of CISOs believed their CEO accepted a breach as inevitable, while nearly a third (32%) believed they would lose their job or receive an official warning following a breach [35]. Taken together, these findings suggest practitioners are held individually accountable for preventing an outcome which may be unavoidable.

Unpredictable Demands. The adversary controls when incidents occur, and this unpredictability extends practitioners' stress well beyond the workplace. Nepal et al. (2024) found that random incident timing created chronic anticipatory stress, with one participant noting "the largest aspect of stress in my life is the uncertainty of the day in relation to raising a family" (p. 21) and another describing being forced to "cancel planned events with your children because of an escalated case that has come in" (p. 21) [32]. The frustration was captured by one participant who admitted: "I scream into my pillow at night in frustration so people can't hear me" (p. 23) [32]. Hollis (2023) documented the same dynamic at longer timescales. Incidents arrived on holidays, at night, and without warning, with one participant observing that "security doesn't sleep and attacks certainly know when IT teams and security teams are at their best and when they're not" (p. 124) [17]. Chappelle et al. (2013) identified short-notice tasking and manning shortages as primary operational stressors among military cyber operators [7], suggesting this dimension spans both military and civilian contexts. What distinguishes unpredictable demands from simple overwork is the inability to plan recovery. Practitioners cannot schedule rest with confidence because the next incident may arrive at any moment, a dynamic that COR theory would predict accelerates resource depletion into loss spirals (Section 2.2) [16].

Invisible Labor and Success Metrics. Successful defense is, by definition, invisible. When practitioners prevent an incident, nothing happens, and the work that produced that outcome goes unrecognized. Tariq et al. (2025) documented the structural conditions behind this invisibility: SOC analysts face high-volume alerts with false positive rates as high as 40%, with industry surveys reporting that 51-54% of SOC teams feel overwhelmed and analysts spend 25-27% of their time handling alerts that produce no actionable outcome [51]. Sundaramurthy et al. (2015) found that this dynamic extended to how analysts' work was measured, with one participant noting "we feel that we are not doing security monitoring in the SOC... we are just working to generate numbers for higher management" (p. 355) and another describing spending hours investigating an alert only to discover it was a false positive, effort that existing metrics failed to capture [50]. Nepal et al. (2024) found that 73.7% of burned-out incident responders had received no recognition in the preceding six months [32]. Across these studies the pattern is consistent. The work is demanding, but its successful execution is structurally invisible to the organization. In JD-R terms (Section 2.1), this represents a systematic depletion of a key job resource, recognition and feedback, precisely when demands are highest.

Exposure to Human Malice. Unlike most high-stress professions where harm arises from accidents, system failures, or natural events, defensive cybersecurity work involves sustained immersion in evidence of deliberate adversarial intent. Mott et al. (2024) documented this concretely, showing that ransomware responders processed extortion demands, negotiated with attackers, reviewed stolen data, and managed communications from threat actors designed to intimidate and coerce [30]. Hollis (2023) similarly found that responders routinely processed breach data and evidence of malicious activity, with the adversarial nature of the work shaping how practitioners experienced incidents [17]. This dimension has received the least direct investigation in the cybersecurity literature. Whether exposure to adversarial intent produces psychological harm distinct from those produced by equivalent workload without an adversarial component is an open question. Section 5 examines what adjacent professions facing analogous exposure patterns have documented.

Latent and Adaptive Threats. Defenders operate under a form of uncertainty that most professions do not face: threats cannot be confirmed as neutralized, and adversaries actively evolve in response to defensive measures. This requires what industry practice terms an "assume compromise" posture, operating as if systems are already breached [7, 17, 30]. Spears (2023) found that practitioners described "a persistent anxiety that they may miss some important technical detail that enables an attack to occur, or that they may be busy with other tasks and be too slow to notice a cyber attack before significant damage is done" (p. 5) [48]. The consequence was that "even when an IR worker is not physically on the job, they are mentally on alert" (p. 5) [48]. Chappelle et al. (2013) documented a parallel dynamic among military operators, for whom the need to sustain "continual expertise in the rapidly changing technologies and functions of adversary networks" meant that the threat landscape was never static, and competence required constant adaptation (p. 4) [7]. This inability to achieve psychological closure distinguishes cybersecurity from professions where threats have defined endpoints. A fire-

fighter can confirm that a fire is extinguished. An emergency physician can discharge a patient. A cybersecurity defender can never fully confirm that a network is clean. The implication, explored further in Section 6, is that this dimension functions not only as a stressor but as a barrier to recovery.

Table 3 identifies these seven stressor dimensions that emerge from the defensive cybersecurity literature as potentially distinctive stressor sources. Evidence strength varies considerably. Several dimensions with theoretical plausibility lack direct empirical investigation.

These dimensions do not operate independently. Sustained vigilance is compounded by unpredictable demands that prevent planned recovery. Invisible success metrics deplete the recognition resources that might buffer those demands. Emotional labor intensifies under the weight of responsibility. Latent and adaptive threats prevent the psychological closure that might allow recovery between incidents. And exposure to human malice, though the least investigated dimension, distinguishes the environment in which all other stressors are experienced from that of most high-stress professions. The interactions between dimensions may matter as much as any individual dimension, but no study has examined them as a system. Section 7 examines these interactions in detail.

4.4 What Remains Unmeasured in Cybersecurity

The most consequential finding is that there is not yet a body of knowledge robust enough to systematize the wellbeing and workplace longevity of people doing cybersecurity work. The literature contains documented stressors and qualitative accounts of severe distress. What it does not yet contain is the kind of cumulative empirical material that would let the field draw confident conclusions about what is happening to its workforce, or what to do about it. The gap itself is the finding.

Lack of medical or clinical studies. None of the cybersecurity studies we reviewed were designed as medical or clinical investigations. The most clinically oriented work in the literature, Chappelle et al.'s (2013) screening of U.S. Air Force cyber warfare operators using the OQ-45.2 [7], was explicitly framed by its authors as a screening assessment rather than a diagnostic study. The authors themselves note that further work is needed to determine whether elevated screening scores correspond to functional impairment. The cybersecurity literature has measured burnout extensively, with cognitive workload, fatigue, clinical distress, and wellbeing measured in a smaller number of studies using validated instruments (see Section 2.3). What these instruments capture is real and worth taking seriously. What they may not capture is less clear. Qualitative accounts repeatedly surface descriptions that the constructs currently in use were not designed to characterise [17, 30], and practitioners themselves reach for vocabulary drawn from outside the burnout literature to describe their experience. Whether what lies beyond is a more severe form of burnout, a constellation of constructs imported from adjacent occupational health literatures, or something specific to cybersecurity work that does not yet have a name, is an open empirical question. Different harm profiles would imply different intervention pathways, and without measurement the field cannot calibrate interventions to the harms actually in play.

No comparative work across role types. The cybersecurity workforce includes offensive, defensive, and mixed roles, and within defensive work it includes crisis-oriented responders and sustained protection and compliance functions (Section 4.1). These are substantially different jobs. Protection and compliance work is steady-state, governance-oriented, and structured around continuous risk reduction rather than incident response. It is likely to produce a different cognitive profile than crisis-oriented defensive work, with different attentional demands, different time horizons, and different relationships to acute stress. Yet existing studies sample within these subgroups but not across them, leaving the antecedents, presentation, and trajectory of harm in each context essentially uncompared. Whether the seven stressor dimensions documented in defensive work generalize to offensive practitioners, whether harm in penetration testing looks like harm in incident response, whether sustained protection roles produce a different kind of strain than crisis response, these are open empirical questions. The reviewed literature does not let the field answer them.

Limited evidence on why practitioners do not speak about it. Qualitative accounts consistently show that practitioners experience confidentiality constraints, professional culture stigma, and institutional silence around their distress [5, 17, 30, 32]. But the mechanisms by which silence shapes harm have not been empirically tested in cybersecurity populations. Whether it slows recovery, whether it concentrates harm in particular subgroups, whether it deters help-seeking even where help is available, these questions remain open. The silence is documented, but its effects on outcomes are not. We return to these barriers in greater detail in RQ3 (Section 6), where we examine how they may also obstruct the natural processing of occupational stress.

5 RQ2: Lessons from Adjacent Professions

RQ1 established what the cybersecurity literature does and does not reveal about practitioner harm. This section examines what adjacent profession research can, and cannot, contribute. We selected four comparison professions (emergency services, healthcare, content moderation, and air traffic control) because they share specific structural features hypothesized to drive harm in cybersecurity: sustained vigilance under uncertainty, exposure to human suffering or disturbing content, high-stakes judgment under time pressure, and chronic workforce shortages. By examining professions with more mature occupational health research, we aim to identify harm mecha-

nisms that may transfer to cybersecurity and constructs that warrant investigation in this population.

These professions are not perfect comparisons. We deliberately widened the scope to four professions sharing specific structural features, selecting them where occupational-health research is mature and where mitigations for workplace harm have been developed. Loss of life is not commensurate with loss of profit, but common stressors, harms, and mitigation infrastructure remain relevant points of comparison. The stakes, work structures, and exposure pathways differ from cybersecurity in ways we make explicit in Section 5.3.

Our analytical approach maps stressors to harms: when workers in adjacent professions face the same stressor dimensions identified in RQ1 (Table 3), what psychological harms result? This mapping allows us to generate hypotheses about harms cybersecurity practitioners may experience but which remain unmeasured.

5.1 Similar Structural Demands Consistently Produce Harm

The premise that certain occupations carry inherent psychological risk is well-established. Lee et al. (2020) reviewed 31 studies examining occupational PTSD and found prevalence rates of 8.4% to 41.1% depending on the definition of PTSD, type of traumatic event, length of exposure, and occupational group [23]. The review identifies a mechanism relevant to cybersecurity: chronic repeated exposure to traumatic events of lower magnitude can produce mental health impacts equivalent to single catastrophic events [23].

Spence et al. (2024) documented a dose-response relationship among content moderators, with 53% scoring above the clinical threshold for psychological distress and 47.6% at levels associated with clinical depression [47]. Content moderation shares key features with some cybersecurity work: exposure to evidence of human malice, high-volume processing demands, and invisible success.

However, the exposure-harm relationship is not straightforward. Devilly et al. (2009) found that among mental health professionals, work-related factors (support, role clarity, workload) predicted harm better than trauma exposure itself [9]. The authors conclude that "the impact of working with traumatized patients is overestimated" (p. 384) [9].

While it remains debated within the trauma literature whether exposure alone drives harm, evidence suggests that accumulated chronic micro-experiences can lead to significant distress and opportunities to process stress are essential. If trauma exposure does not reliably produce harm even among therapists treating trauma survivors, cybersecurity's claim to distinctiveness cannot rest solely on exposure to adversarial content. Instead, the relevant questions concern frequency and severity of exposure, and whether the combination of multiple stressor dimensions produces harms or barriers to processing those stressors impede recovery.

5.2 Mapping Stressor Dimensions to Documented Harms

Sustained Vigilance and Cognitive Degradation. Air traffic control is the most extensively studied vigilance-intensive profession. Controllers must maintain continuous situational awareness across multiple simultaneous information streams while any lapse in attention carries immediate safety consequences. This demand structure closely parallels the monitoring responsibilities of SOC analysts, who track network traffic, triage alerts, and maintain awareness of system states across extended shifts.

Air traffic controllers show higher burnout than firefighters despite rating their life situation as less stressful, suggesting habitual vigilance normalizes the demands that produce harm [26, 27]. Alaydi and Ng (2024) confirmed a direct negative relationship between mental workload and job performance across 324 controllers ($\beta = -0.264$, $p < 0.001$), with mindfulness and social support moderating but not eliminating the effect [1].

Exposure to Disturbing Content. Content moderation provides direct evidence for workers who are routinely exposed to disturbing material produced by other humans. Spence et al. (2023) documented intrusive thoughts, avoidance of children, cynicism, and emotional detachment among moderators exposed to child sexual abuse material, a symptom profile the authors mapped onto PTSD symptom domains [46]. Critically, this trauma is screen-mediated. Moderators encounter events through digital artifacts, not physical proximity.

Invisible Labor and Emotional Suppression. Smith et al. (2019) found that emergency dispatchers described a perceived "invisible nature" of their role and a divide between themselves and field personnel [45]. They were required to maintain "dispassionate authority" while processing calls involving child abuse, suicide, and verbally aggressive callers (p. 622) [45]. Those who received abusive calls showed higher emotional exhaustion, depersonalization, and intention to quit. Despite this, dispatchers were often excluded from organizational support. After September 11, 2001, New York City call-takers could not initially access the trauma services designated for firefighters [45]. This pattern suggests that when the emotional labor required by a role is neither recognized nor supported, harm may become systemic rather than individual.

Unpredictable Demands and Recovery Failure. Lewis-Schroeder et al. (2018) reported substantially higher PTSD prevalence among first responders than the general population, with irregular schedules and absent recovery periods among the risk factors [24]. Hruska and Barduhn (2021) showed through a daily diary study of EMS workers that recovery activities and meaning-making protected against depression, suggesting the capacity to recover is mechanistically important rather than incidental [18].

Harm Degrades Professional Judgment. Research in adjacent professions suggests psychological harm may de-

grade judgment in ways particularly relevant to security work. Regehr and LeBlanc (2017) examined PTSD symptoms among emergency workers and found effects emerged in professional judgment and risk assessment rather than routine task performance [42]. Child protection workers with elevated PTSD symptoms were paradoxically less likely to determine a child was at risk, suggesting desensitization to threat indicators [42]. Li et al. (2025) found burnout was the primary predictor of safety behavior among air traffic controllers, with relative importance nearly five times greater than job satisfaction [25]. These findings suggest that cybersecurity practitioner wellbeing is not merely a human resources concern but a determinant of organizational security posture. Analyst burnout or trauma may manifest through missed threats, reduced proactive investigation, or degraded risk assessment.

5.3 What Does Not Translate: Structural Differences

Despite several structural similarities, it is critical to address where the comparison breaks down. Findings from adjacent professions require adaptation before they can responsibly inform cybersecurity intervention design. The structural differences below shape what can and cannot be transferred.

The stakes differ in kind, not just degree. Loss of life is not commensurate with loss of profit, proprietary data, or operational disruption. Emergency services, healthcare, and air traffic control workers operate under stakes where human survival is at issue, and that fact shapes both the social legitimacy of their distress and the institutional infrastructure that has grown up around it. Cybersecurity practitioners work under stakes that are nevertheless serious (financial ruin, regulatory consequence, exposure of sensitive personal data, and in some sectors public-safety implications) but socially read as less acute. This asymmetry may affect how practitioner distress is recognized, how interventions are funded, and how stigma operates. Adjacent-profession evidence about the harms that follow from exposure may translate. Adjacent-profession evidence about the legitimacy granted to that harm, and the resources made available to treat it, likely does not.

First-responder exposure differs from screen-mediated exposure. Adjacent professions vary in the immediacy of their exposure to human suffering. Emergency services and healthcare workers experience direct physical exposure. Content moderators and most cybersecurity practitioners experience screen-mediated exposure. The mechanisms of harm may overlap but are not identical, and evidence calibrated to one form of exposure may need adjustment before being applied to the other. Content moderation is the closest analogue, but content moderation research itself is comparatively young.

Teamwork norms and physical co-presence differ. Emergency responders work in physically co-located teams with strong norms of mutual support and post-incident debriefing. Healthcare teams operate within established peer-support and supervision structures. Air traffic controllers work under regulated staffing and rest regimes. Cybersecurity work, particularly in remote or distributed teams and during after-hours incidents, frequently lacks these structural supports. The mitigation infrastructure that adjacent professions have developed assumes contexts of co-presence and shared exposure that cybersecurity work may not provide.

The relationship to the source of harm differs. Adjacent professions face harm from accidents, illness, natural events, or, in content moderation and some emergency-service contexts, human malice. Cybersecurity is unusual in that the source of harm is also the source of professional identity and livelihood. The very threats that distress practitioners are also what makes their roles necessary and valued. This ambivalence has no clean analogue in the comparison professions and limits the transferability of intervention strategies developed elsewhere.

Common stressors, harms, and mitigation infrastructure remain relevant points of comparison even when the stakes differ. The value of the comparison is that it surfaces questions worth addressing in cybersecurity populations. The limit of the comparison is that adjacent-profession evidence cannot substitute for cybersecurity-specific data.

5.4 RQ2 Summary

Adjacent profession research establishes that structural conditions similar to those documented in defensive cybersecurity consistently produce psychological harm beyond burnout. However, direct transfer is complicated by uncertainty about which harm constructs apply to cybersecurity practitioners, by uncertainty about the exposure thresholds that drive harm, and by the structural differences laid out in Section 5.3. Burnout has been measured in cybersecurity practitioners and is clearly present. Whether other harms are also present, and if so what shape they take, remains unanswered with direct implications for intervention design.

6 RQ3: Mitigations and Barriers to Processing Occupational Stress

The preceding sections established two things. RQ1 documented stressors and harms in cybersecurity practitioner research while identifying gaps in measurement, role coverage, and the silence around practitioner distress. RQ2 established that adjacent professions facing similar stressors consistently develop psychological harm and have built mitigation infrastructure in response, though structural differences limit direct transfer. This section examines what mitigations the cybersecurity literature currently proposes, and whether cybersecurity presents distinctive barriers to processing occupational stress that conventional mitigations may not address.

6.1 Mitigations

The cybersecurity literature proposes mitigations that cluster around organizational and leadership interventions: workload and scheduling optimization [7, 32], leadership engagement that addresses root causes rather than symptoms [5, 7, 50, 53], recognition and professional development [32, 50], resilience training and supportive team cultures [32], dedicated mental health infrastructure [7, 17, 30, 53], and incident-specific welfare considerations such as staff rotation and peer debriefing [30]. Adjacent professions have developed more mature mitigation infrastructure, addressed in Section 6.2.

In many high-stress occupations, interventions can target exposure directly. Defensive cybersecurity may lack this option. The adversary dictates incident timing, duration, and intensity. Organizations cannot defer a ransomware attack or schedule breaches for convenient hours. If exposure cannot be meaningfully reduced, then processing capacity becomes the primary lever for preventing cumulative harm. This makes the barriers documented below particularly consequential.

6.2 Barriers

Confidentiality Constraints and Isolation. Research on processing difficult occupational experiences suggests verbal disclosure facilitates cognitive integration [22]. The most relevant comparison group with regard to barriers may be content moderators, who share key features with cybersecurity work including exposure to evidence of human malice, high-volume processing demands, and invisible success [49]. Steiger et al. (2021) found that content moderators face structural barriers to processing difficult experiences, describing their work environments as creating a "Big Brother environment" that erodes their "sense of humanity" (p. 2) [49].

Cybersecurity practitioners face analogous barriers through security clearances, classification requirements, incident confidentiality norms, and professional cultures that discourage discussing operational stress. Hollis (2023) documented this explicitly. Participants described strict enforcement: "I had been warned very strictly I was not to speak to anybody... I can't talk to anybody" (p. 127) and "you sign the oath of secrecy. You have to be very, very careful about that" (p. 127) [17]. The result was that discussing their work was simply "not something you can talk about" (p. 51) [17].

The consequences of this constraint were significant. Hollis (2023) found that "these thoughts and feelings were trapped within the responder because no verbal outlet exists for any specific emotion or stress due to confidentiality requirements" (p. 85) [17]. "Participants lamented not being able to talk to anyone, which created a sense of isolation and lack of support" (p. 85), and "described the requirement to keep cybersecurity incidents confidential and relayed details they have held onto for years" (p. 85) [17]. If verbal disclosure facilitates cognitive integration, then structural barriers to disclosure may transform routine occupational stress into compounded harm. Those who held both cybersecurity and first responder roles noted that while traditional responders could debrief with colleagues and therapists, cybersecurity responders often could not discuss incidents even at an abstract level [17]. Nor does the constraint end with the incident. Mott et al. (2024) found regulatory and legal processes extended psychological strain for years beyond incident resolution, with practitioners bound by confidentiality even as they continued experiencing distress [30].

Temporal Dynamics Preventing Closure. Recovery from occupational stress requires psychological disengagement, periods where work can be mentally set aside. Cybersecurity practitioners may lack access to such periods.

Incidents resist closure. Hollis (2023) noted that first responder incidents typically last hours while cybersecurity incidents extend for weeks or months [17]. Even after technical containment, practitioners cannot confirm threat eradication. The "assume breach" mentality required for competent defense means vigilance never fully relaxes [17]. Mott et al. (2024) found regulatory aftermath extended strain for years beyond incident resolution [30].

Between incidents, on-call expectations persist. The adversary's control over timing means practitioners cannot predict when demands will arrive or plan recovery periods with confidence. This dimension functions as both a stressor and barrier. The temporal structure of the work creates strain while simultaneously preventing recovery from that strain.

Professional Culture. Mott et al. (2024) documented environments demanding staff "suck it up and sort it out" (p. 10) [30]. Nepal et al. (2024) found 73.7% of burned-out participants received no recognition in six months [32]. Arora and Hastings (2024) found that cybersecurity experts are reluctant to report their struggles to management, perpetuating a cycle of silence and neglect [5].

Professional status compounds this barrier. Hollis (2023) found that participants who held both cybersecurity and first responder roles experienced different social responses to their work. One described: "One thing about being in cybersecurity is you're a dork. So, you know, one gets respect; the other one gets mocked... that difference can be pretty substantial when you're trying to deal with grief and trauma" (p. 60) [17]. The participant noted that this lack of respect has "a much bigger potential impact on individuals" processing similar experiences (p. 61) [17].

Similar cultural barriers exist in adjacent professions. Fogarty et al. (2021), for instance, documented beliefs about self-reliance among firefighters [12]. Whether cybersecurity culture is more or less conducive to help-seeking than comparable professions remains unknown.

Limited Institutional Infrastructure. Emergency services have incident debriefing protocols, peer support programs, and clinical pathways developed over decades [12]. Healthcare has wellness programs and duty hour regulations [43]. Air

traffic control has mandatory rest periods and fatigue risk management systems [25]. Content moderation, though newer, has seen companies implement technological interventions like image blurring and mandatory wellness time [49]. Cybersecurity appears to lack equivalent infrastructure.

Both Hollis (2023) and Mott et al. (2024) noted that some incident response firms had developed in-house confidential trauma counseling, but this appears exceptional [17, 30]. Generic employee assistance programs may not suit classified or confidential work contexts, and most therapists are unfamiliar with cybersecurity-specific constraints. Even practitioners willing to seek help may find appropriate support unavailable.

6.3 RQ3 Summary

Cybersecurity practitioners face four barriers to processing occupational stress. Confidentiality constraints prevent them from talking about their experiences. Limited institutional infrastructure means support may be unavailable even when sought. Professional culture discourages help-seeking in the first place. And temporal dynamics (extended incidents, regulatory aftermath, on-call expectations) prevent the psychological closure that enables recovery. Table 4 summarizes evidence for each barrier.

7 Discussion

The cybersecurity literature documents stressors, qualitative accounts of severe distress, and indications that practitioners themselves reach for clinical vocabulary to describe their experience. Our analysis identifies a clear gap in knowledge, highlighting that there is not yet enough empirical material to let the field draw confident conclusions about what is happening to its workforce, or what to do about it. Here, we discuss what that gap implies: how does harm present across cybersecurity roles, why practitioner harm should be treated as a security concern rather than a wellbeing concern, what workers and employers can reasonably expect of each other, and what a feasible research agenda would look like.

7.1 The Offensive-Defensive Distinction

This review surfaced a potential distinction between offensive and defensive cybersecurity roles. Offensive work (studied primarily in military contexts) appears to involve a different psychological profile. Operational conditions drive harm, but the work itself is experienced as deeply meaningful, even "addicting" [39]. Practitioners leave for work-life balance, not because of the work's content.

Defensive work appears different. Stressors seem tied not only to operational conditions but to occupational demands: chronic vigilance that never confirms safety; invisible success that provides no positive feedback; exposure to evidence of human malice; or asymmetric stakes where defenders must succeed every time. If these features are intrinsic to defensive work, then workload reduction alone may be insufficient.

Within defensive work, a further distinction matters: protection and compliance roles operate at a steady state and are governance-oriented, structured around continuous risk reduction whereas crisis-oriented roles such as incident response operate under acute time pressure with periodic extreme demands. The cognitive profile, antecedents, and presentation of harm are likely to differ across these contexts. No study has directly compared offensive, crisis-oriented defensive, and sustained protection roles, and this requires further research. If the distinctions hold, "cybersecurity wellbeing" is not a single problem requiring a single solution, and interventions calibrated to one mode of work may not transfer to another.

7.2 How Harm in Cybersecurity May Differ

The individual stressor dimensions identified in this paper appear across adjacent professions (Table 3). What appears unusual in cybersecurity, particularly in crisis-oriented defensive cybersecurity roles, is their combination together with the structural barriers that may prevent recovery. Sustained vigilance is compounded by unpredictable demands that prevent planned recovery. Structurally invisible success metrics deplete the recognition resources that might buffer those demands. Latent and adaptive threats prevent the psychological closure that might allow recovery between incidents. Confidentiality constraints block verbal processing of those non-closing experiences. Temporal dynamics extend incidents for months and regulatory aftermath for years.

The implication is that the antecedents and presentation of harm in cybersecurity may not map cleanly onto burnout as conventionally measured. Burnout in this population may be the surface presentation of something more complex underneath. We do not claim that the severe psychological harms documented in adjacent professions are present in cybersecurity practitioners on the basis of qualitative self-descriptions alone. We do note that the practitioners themselves reach for clinical vocabulary, in a literature otherwise organised around burnout, and that this is a signal worth investigating. The most distinctive risk in defensive cybersecurity may not be any individual stressor, but the combination of stressors and the structural barriers that prevent recovery from them. Adjacent professions have developed debriefing protocols, peer support, and clinical pathways. Cybersecurity has not.

7.3 Why This Is a Security Problem in Addition to a Wellbeing Problem

Research in adjacent professions demonstrates that psychological harm degrades professional judgment in ways particularly relevant to security work. Regehr and LeBlanc (2017) found that emergency workers with elevated PTSD symptoms showed impaired risk assessment [42]. Li et al. (2025) found

burnout was the primary predictor of safety behavior among air traffic controllers, with relative importance nearly five times greater than job satisfaction [25]. If similar dynamics apply to cybersecurity, the externalities of leaving practitioner harm unaddressed extend well beyond practitioner welfare.

The security externality is the most obvious. Symptoms of psychological distress may manifest as missed indicators of compromise, reduced proactive investigation, degraded risk assessment, or errors during incident response, precisely when sound judgment matters most. Sustained vigilance failures and high false-positive rates that drive alert fatigue may compound when practitioners are also experiencing exhaustion or cognitive avoidance. The profession that defends organizations against adaptive adversaries cannot afford impaired defenders, and yet the literature provides limited insight into estimating how impaired the defending workforce is.

The boundary between work and personal security is also unusually thin in this profession. An air traffic controller leaves the radar behind at the end of a shift, but a cybersecurity practitioner goes home to manage their own passwords, networks, and devices, and remains a target while doing so. The off-hours that would normally allow recovery never fully arrive, so strain accumulated on the job may follow the practitioner home and compound rather than easing. A depleted defender whose personal security habits have slipped becomes a softer target, and a compromise at home can spill over into the organization.

Aside from these direct risks, industry surveys consistently document a global cybersecurity workforce shortage [19], and the qualitative literature gives one explanation: practitioners leave the field as a form of escape rather than career progression [17, 30, 39]. Each departure removes hard-to-replicate tacit knowledge that documentation only partially captures. If practitioner harm contributes to attrition, then unaddressed harm functions as a hidden, recurring tax on organizational security posture, paid in institutional knowledge.

The longest-horizon cost is to the profession itself. Cybersecurity defenders are rarely recognized for the work they do. A prevented incident leaves traces in logs and alerts, but little that registers beyond the technical team, and the work that produced the non-event remains largely invisible to the wider organization. The profession also sits in an ambivalent relationship to the threats that define it, with attacks serving simultaneously as the source of distress, professional recognition, and livelihood [17]. The cumulative effect is a profession in which longevity, growth, and self-respect are harder to sustain than they need to be. Taken together, these externalities reframe practitioner wellbeing from a human-resources concern to a structural concern for the security ecosystem, the labor market, and the profession itself.

7.4 What Workers and Employers Can Reasonably Expect

Workers can reasonably expect resources commensurate with documented demands: recognition that accounts for structurally invisible labor, confidentiality-compatible mental health support, peer debriefing access, realistic vigilance windows with mandated recovery, and protection from open-ended on-call expectations. Employers should reasonably expect competent sustained performance, not infinite resilience. Wellbeing infrastructure should be treated as a security control rather than a discretionary human-resources benefit.

Different stakeholders carry different parts of the response. Practitioners may benefit from validation of lived experience and from evidence that supports conversations with employers about workload, scheduling, and access to support. Security leadership can draw on the recovery barriers documented in Section 6.2 to inform arguments for confidentiality-compatible mental health support, peer debriefing, and mandated recovery periods similar to those required in other high-stress adjacent professions.

For researchers, the next step is primary research that tests which psychological constructs apply to cybersecurity, explores role-specific differences, and identifies supportive interventions. Policymakers can treat practitioner wellbeing as a security concern in the legal, regulatory, and procurement frameworks they shape, including allowing confidentiality regimes to accommodate clinical disclosure to credentialed providers, incentivizing organizational investment in cybersecurity workforce health, and funding research that intersects occupational health and cybersecurity.

7.5 Research Agenda

Four research priorities follow from this paper. The first priority is to establish which harm constructs are present in cybersecurity populations, and at what prevalence. Until we can ascertain the extent to which constructs are present, interventions cannot be properly calibrated to the harm actually in play.

The second priority is to test the mechanisms hypothesized in this paper. Do confidentiality constraints predict worse outcomes controlling for exposure? Does access to confidentiality-compatible support moderate harm? Does organizational recognition buffer the impact of structurally invisible success? Mechanism research is what will distinguish whether barriers slow recovery or fundamentally change the nature of harm.

The third priority is to quantify the negative externalities that are likely present. Research linking practitioner wellbeing measures to security-relevant outcomes (incident detection rates, false-negative rates, judgment under simulated incident pressure, attrition timelines) would establish whether the externality argument holds in cybersecurity populations.

The fourth priority is to evaluate interventions under the constraints cybersecurity work imposes. Trial designs will need to accommodate confidentiality, secure data channels, and adaptation of approaches from adjacent professions to the specific barriers documented in Section 6. Adapting peer debriefing for confidentiality-constrained contexts, evaluating cybersecurity-literate clinical pathways, and testing organizational recognition designs are defensible starting points.

All four priorities share a methodological challenge that amounts to a Catch-22. Harm cannot be addressed until it is measured, and it cannot be measured until practitioners disclose it. But disclosure is blocked by the very practices the research would scrutinize. Confidentiality requirements that block verbal disclosure also limit research participation, and professional cultures that stigmatize help-seeking also stigmatize participation in studies framed around mental health. The resulting absence of evidence is then invoked as a reason to leave those practices untouched. Left unchallenged, this loop guarantees that the harm remains structurally unmeasurable, and therefore deniable, indefinitely.

We believe the field should accept short-term pain for long-term gain. Breaking the loop will require researchers, practitioners, organizations, and professional bodies working in coordination. Opening credentialed, secure channels for clinical and research disclosure carries real near-term costs, including legal and contractual exposure, the discomfort of measurement, and the ethical weight of asking practitioners to re-engage with distressing experiences, even in controlled research environments where some degree of harm must be present to be studied at all. But these are bounded, governable costs, and the alternative is to keep paying an unbounded, unmeasured one. Cybersecurity is, after all, a discipline that institutionalized controlled harm. Penetration tests and red-team exercises deliberately attack systems precisely because controlled exposure is cheaper than uncontrolled failure. A profession whose core purpose is to break its own systems in order to build more secure ones should be willing to examine its own workforce in order to sustain it.

7.6 Limitations

This review has substantial limitations. Although two authors were involved throughout screening and synthesis, formal inter-rater reliability was not computed. Rigor was supported instead through pre-specified inclusion criteria, structured decision logging, and backward citation searching.

The cybersecurity-specific literature remains sparse, and much of the available evidence derives from small qualitative samples that may overrepresent severe cases. The offensive-defensive distinction rests on non-comparable populations studied with different methods. The defensive evidence base draws predominantly from crisis-oriented roles (incident response, SOC analysis) and may not generalize to sustained protection and compliance work, which remains substantially under-researched. Adjacent profession comparisons are constrained by the structural differences laid out in Section 5.3.

The central argument linking documented stressors to hypothesized harm beyond burnout remains inferential, with no direct measurement in cybersecurity populations. The cybersecurity and adjacent profession literature search covered publications between 2005 and 2025, with the 2005 lower bound excluding earlier literature unlikely to reflect contemporary occupational demands. Finally, this paper focused specifically on cybersecurity practitioners rather than attempting a comprehensive synthesis of burnout or psychological theory.

These limitations reflect the underdeveloped state of research on a population facing a documented occupational crisis. The scarcity of rigorous evidence for a profession of this scale is itself a finding, and it is the central reason this paper is presented as a gap analysis rather than a conventional systematization.

8 Conclusion

Defensive cybersecurity is essential, expanding, and permanent. Organizations require protection, and practitioners will continue to provide it. That is precisely why understanding the psychological demands of this work matters. If the profession is necessary and growing, then so too is the imperative to understand what it costs practitioners and how to support them.

The limited evidence suggests practitioners experience burnout, fatigue, and cognitive depletion consistent with adjacent professions. Whether psychological harms beyond burnout are also present remains unmeasured. What the evidence does suggest is that barriers to processing warrant prioritized research attention. Confidentiality constraints foreclose disclosure. Limited institutional infrastructure means support may be unavailable even when sought. Professional culture discourages help-seeking. And temporal dynamics (extended incidents, regulatory aftermath, on-call expectations) prevent the psychological closure that enables recovery. Whether these barriers simply slow recovery or fundamentally change the nature of harm in defensive cybersecurity remains unknown.

Security systems and incident response practices implicitly assume infinite human resilience. When that assumption fails, the consequences fall on individual practitioners. The practitioners who defend organizations against adaptive adversaries deserve research attention commensurate with their importance. Ignoring occupational harm in an essential profession does not make it more sustainable. It makes attrition, burnout, and degraded judgment inevitable. This paper maps the terrain and provides a foundation for the empirical work that must follow.

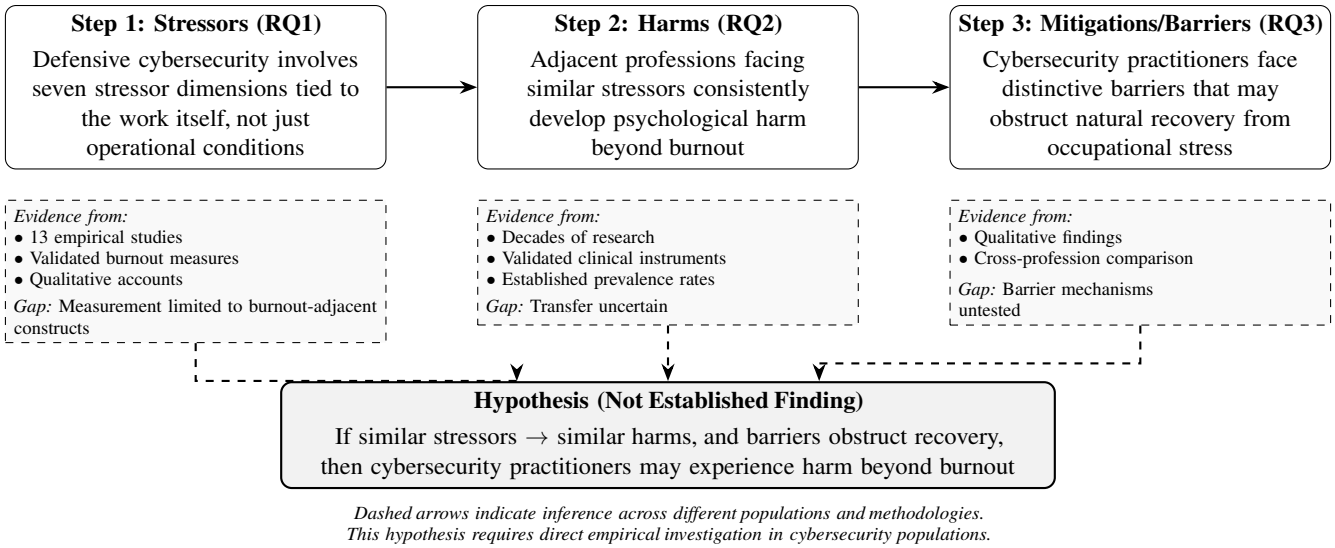

Figure 2: Documented stressors to hypothesized harm. Steps 1–3 summarize findings from each research question. Dashed arrows indicate that the hypothesis follows inferentially from combining evidence across domains, not from direct measurement in cybersecurity populations. The gaps noted in each step represent priorities for future research.

Ethical Considerations

This paper synthesizes existing published research and does not involve primary data collection from human subjects. All sources cited are publicly available peer-reviewed publications, dissertations, or publicly released reports. The decision to submit for publication is to review the existing literature and highlight identified gaps, and surface structural issues in the profession that have not been systematically examined. All source material is already in the public domain, and we believe the potential for this work to inform improvements to practitioners' wellbeing outweighs potential negative implications.

We followed our institution's ethical review process to determine whether the project met the criteria for formal ethical review. Research based on secondary synthesis of publicly available peer-reviewed literature, where no human subjects are contacted and no primary data is collected, is classified as not requiring ethical review. Reading was paced over several weeks allowing for enough time to identify and remediate any concerns that might arise. Regular meetings between the two authors provided a space to discuss the material and the themes emerging from it.

Open Science

This paper is a Systematization of Knowledge (SoK) synthesizing exisiting literature. All sources supporting this scoping review are publicly available to enable audit, reproduction, and extension of our search and selection process. The release comprises the database search strategy, the inclusion and exclusion criteria, the 101 records that advanced to abstract assessment and the 57 records carried into synthesis in BibTeX format, and the literature matrix recording extracted information for the 57 synthesized sources across 24 cybersecurity and 33 adjacent-profession papers.

The full artifact is available for review at https://doi.org/10.5281/zenodo.20401452.

Acknowledgements

We thank our reviewers and the shepherd for their helpful feedback on the framing of this SoK. This research is supported in part by the Engineering and Physical Sciences Research Council grant SCULI: Securing Convergent Ultra-large Scale Infrastructures [EP/Z531315/1].

References

[1] Alaydi, B., & Ng, S.-I. (2024). Mitigating the negative effect of air traffic controller mental workload on job performance: The role of mindfulness and social work support. *Safety*, *10*(1), 20. https://doi.org/10.3390/safety10010020

[2] Agrafiotis, I., Nurse, J., Goldsmith, M., Creese, S., & Upton, D. (2018). A taxonomy of cyber-harms: Defining the impacts of cyber-attacks and understanding how they propagate. *Journal of Cybersecurity*, *4*(1). https://doi.org/10.1093/cybsec/tyy006

[3] American Psychiatric Association. (2022). *Diagnostic and statistical manual of mental disorders* (5th ed., text rev.).

[4] Arksey, H., & O'Malley, L. (2005). Scoping studies: Towards a methodological framework. *International Journal of Social Research Methodology*, *8*(1), 19–32. https://doi.org/10.1080/1364557032000119616

[5] Arora, S., and Hastings, J. D. (2024). A Survey-Based Quantitative Analysis of Stress Factors and Their Impacts Among Cybersecurity Professionals. Sep. 18, 2024, arXiv: arXiv:2409.12047. https://doi.org/10.48550/arXiv.2409.12047

[6] Bakker, A. B., & Demerouti, E. (2017). Job demands–resources theory: Taking stock and looking forward. *Journal of Occupational Health Psychology*, *22*(3), 273–285. https://doi.org/10.1037/ocp0000056

[7] Chappelle, W., McDonald, K., Christensen, J., Prince, L., Goodman, T., Thompson, W., & Hayes, W. (2013). *Sources of occupational stress and prevalence of burnout and clinical distress among U.S. Air Force cyber warfare operators*. Defense Technical Information Center. https://apps.dtic.mil/sti/tr/pdf/ADA584653.pdf

[8] Demerouti, E., Bakker, A. B., Nachreiner, F., & Schaufeli, W. B. (2001). The job demands-resources model of burnout. *Journal of Applied Psychology*, *86*(3), 499–512. https://doi.org/10.1037/0021-9010.86.3.499

[9] Devilly, G. J., Wright, R., & Varker, T. (2009). Vicarious trauma, secondary traumatic stress or simply burnout? Effect of trauma therapy on mental health professionals. *Australian & New Zealand Journal of Psychiatry*, *43*(4), 373–385. https://doi.org/10.1080/00048670902721079

[10] Dykstra, J., & Paul, C. L. (2018). Cyber Operations Stress Survey (COSS): Studying fatigue, frustration, and cognitive workload in cybersecurity operations. *Proceedings of the USENIX Workshop on Cyber Security Experimentation and Test (CSET)*. https://www.usenix.org/conference/cset18/presentation/dykstra

[11] Ehlers, A., & Clark, D. M. (2008). Posttraumatic stress disorder: The development of effective psychological treatments. *Nordic Journal of Psychiatry*, *62*(Suppl 47), 11–18. https://doi.org/10.1080/08039480802315608

[12] Fogarty, A., Steel, Z., Ward, P. B., Boydell, K. M., McKeon, G., & Rosenbaum, S. (2021). Trauma and mental health awareness in emergency service workers: A qualitative evaluation of the Behind the Seen education workshops. *International Journal of Environmental Research and Public Health*, *18*(9), 4418. https://doi.org/10.3390/ijerph18094418

[13] Greenlee, E., Funke, G., Warm, J., Sawyer, B., Finomore, V., Mancuso, V., Funke, M., & Matthews, G. (2016). Stress and workload profiles of network analysis: Not all tasks are created equal. In D. Nicholson (Ed.), *Advances in Human Factors in Cybersecurity* (pp. 153–166). Springer. https://doi.org/10.1007/978-3-319-41932-9_13

[14] Gupta, V., Rangarajan, A., & Nobles, C. (2024). Burnout in the cybersecurity profession: A scoping review. *Proceedings of the Midwest Association for Information Systems Conference*.

[15] Hobfoll, S. E. (1989). Conservation of resources: A new attempt at conceptualizing stress. *American Psychologist*, *44*(3), 513–524. https://doi.org/10.1037/0003-066X.44.3.513

[16] Hobfoll, S. E., Halbesleben, J., Neveu, J.-P., & Westman, M. (2018). Conservation of resources in the organizational context: The reality of resources and their consequences. *Annual Review of Organizational Psychology and Organizational Behavior*, *5*, 103–128. https://doi.org/10.1146/annurev-orgpsych-032117-104640

[17] Hollis, T. R. (2023). *All quiet on the digital front: The unseen psychological impacts on cybersecurity first responders* [Doctoral dissertation, University of South Florida].

[18] Hruska, B., & Barduhn, M. S. (2021). Dynamic psychosocial risk and protective factors associated with mental health in Emergency Medical Service (EMS) personnel. *Journal of Affective Disorders*, *282*, 9–17. https://doi.org/10.1016/j.jad.2020.12.130

[19] ISC2. (2024). *2024 ISC2 Cybersecurity Workforce Study*. https://www.isc2.org/Insights/2024/10/ISC2-2024-Cybersecurity-Workforce-Study

[20] M. E. Kiger and L. Varpio, "Thematic analysis of qualitative data: AMEE Guide No. 131," *Medical Teacher*, vol. 42, no. 8, pp. 846–854, Aug. 2020. https://doi.org/10.1080/0142159X.2020.1755030

[21] Kim, J., Chesworth, B., Franchino-Olsen, H., & Macy, R. J. (2022). A scoping review of vicarious trauma interventions for service providers working with people who have experienced traumatic events. *Trauma, Violence & Abuse*, *23*(5), 1437–1460. https://doi.org/10.1177/1524838021991310

[22] Köhler, M., Goebel, S., & Pedersen, A. (2019). PTSD severity among emergency personnel: An investigation based on the Ehlers and Clark cognitive model. *Psychological Trauma: Theory, Research, Practice, and Policy*, *11*(6), 677–684. https://doi.org/10.1037/tra0000466

[23] Lee, W., Lee, Y.-R., Yoon, J.-H., Lee, H.-J., & Kang, M.-Y. (2020). Occupational post-traumatic stress disorder: An updated systematic review. *BMC Public Health*, *20*(1), 768. https://doi.org/10.1186/s12889-020-08903-2

[24] Lewis-Schroeder, N. F., Kieran, K., Murphy, B. L., Wolff, J. D., Robinson, M. A., & Kaufman, M. L. (2018). Conceptualization, assessment, and treatment of traumatic stress in first responders: A review of critical issues. *Harvard Review of Psychiatry*, *26*(4), 216–227. https://doi.org/10.1097/HRP.0000000000000176

[25] Li, X., Huang, Z., Peng, X., Zhang, J., Liu, Y., Luo, Y., & Guo, L. (2025). Job satisfaction, burnout, and safety behavior in air traffic controllers: A mediation analysis and decision tree insights. *Frontiers in Public Health*, *13*, 1626328. https://doi.org/10.3389/fpubh.2025.1626328

[26] Makara-Studzińska, M., Załuski, M., Jagielski, P., Wójcik-Małek, D., & Szelepajło, M. (2021). An exploration of perceived stress, burnout syndrome, and self-efficacy in a group of Polish air traffic controllers and maritime navigators: Similarities and differences. *International Journal of Environmental Research and Public Health*, *18*(1), 53. https://doi.org/10.3390/ijerph18010053

[27] Makara-Studzińska, M., Załuski, M., Biegańska-Banaś, J., Tyburski, E., Jagielski, P., & Adamczyk, K. (2021). Perceived stress and burnout syndrome: A moderated mediation model of self-efficacy and psychological comfort among Polish air traffic controllers. *Journal of Air Transport Management*, *96*, 102105. https://doi.org/10.1016/j.jairtraman.2021.102105

[28] Maslach, C., Schaufeli, W. B., & Leiter, M. P. (2001). Job burnout. *Annual Review of Psychology*, *52*(1), 397–422. https://doi.org/10.1146/annurev.psych.52.1.397

[29] Maslach, C., & Leiter, M. P. (2016). Understanding the burnout experience: Recent research and its implications for psychiatry. *World Psychiatry*, *15*(2), 103–111. https://doi.org/10.1002/wps.20311

[30] Mott, G., Turner, S., Nurse, J. R. C., Pattnaik, N., MacColl, J., Huesch, P., & Sullivan, J. (2024). 'There was a bit of PTSD every time I walked through the office door': Ransomware harms and the factors that influence the victim organization's experience. *Journal of Cybersecurity*, *10*(1), tyae013. https://doi.org/10.1093/cybsec/tyae013

[31] Munn, Z., Peters, M. D. J., Stern, C., Tufanaru, C., McArthur, A., & Aromataris, E. (2018). Systematic review or scoping review? Guidance for authors when choosing between a systematic or scoping review approach. *BMC Medical Research Methodology*, *18*(1), 143. https://doi.org/10.1186/s12874-018-0611-x

[32] Nepal, S., Hernandez, J., Lewis, R., Chaudhry, A., Houck, B., Knudsen, E., Rojas, R., Tankus, B., Prafullchandra, H., & Czerwinski, M. (2024). Burnout in cybersecurity incident responders: Exploring the factors that light the fire. *Proceedings of the ACM on Human-Computer Interaction*, *8*(CSCW1), 27:1–27:35. https://doi.org/10.1145/3637304

[33] Nobles, C. (2022). Stress, burnout, and security fatigue in cybersecurity: A human factors problem. *HOLISTICA – Journal of Business and Public Administration*, *13*(1), 49–72. https://doi.org/10.2478/hjbpa-2022-0003

[34] Nobles, C. (2019). Establishing Human Factors Programs to Mitigate Blind Spots in Cybersecurity. https://aisel.aisnet.org/mwais2019/22/

[35] Nominet. (2019). *Life inside the perimeter: Understanding the modern CISO*. https://media.nominet.uk/wp-content/uploads/2019/02/12130924/Nominet-Cyber_CISO-report_FINAL-130219.pdf

[36] Paul, C., & Dykstra, J. (2017). Understanding operator fatigue, frustration, and cognitive workload in tactical cybersecurity operations. *Journal of Information Warfare*, *16*(2), 1–11. https://www.jstor.org/stable/26502752

[37] Pearlman, L. A., & Saakvitne, K. W. (1995). *Trauma and the therapist: Countertransference and vicarious traumatization in psychotherapy with incest survivors*. W.W. Norton.

[38] Pearlman, L. A. (2012). Vicarious trauma. In *Encyclopedia of Trauma: An Interdisciplinary Guide* (pp. 783–786). SAGE Publications.

[39] Rangarajan, A., Nobles, C., Dykstra, J., Cunningham, M., Robinson, N., Hollis, T., Paul, C. L., & Gulotta, C. (2025). A roadmap to address burnout in the cybersecurity profession: Outcomes from a multifaceted workshop. arXiv:2502.10293. https://doi.org/10.48550/arXiv.2502.10293

[40] Reeves, A., Pattinson, M., & Butavicius, M. (2024). The sleepless sentinel: Factors that predict burnout and sleep quality in cybersecurity professionals. *Information and Computer Security*, *32*(4), 477–491. https://doi.org/10.1108/ICS-11-2023-0222

[41] Reeves, A. , Calic, D., & Delfabbro, P. (2024). How to De-CyFa the actor-observer bias in cybersecurity fatigue: Building the CyFa measure of attribution styles and mitigation strategies. *Computers & Security*, *vol. 150*, p. 104179. https://doi.org/10.1016/j.cose.2024.104179

[42] Regehr, C., & LeBlanc, V. R. (2017). PTSD, acute stress, performance and decision-making in emergency service workers. *Journal of the American Academy of Psychiatry and the Law*, *45*(2), 184–192.

[43] Reith, T. P. (2018). Burnout in United States healthcare professionals: A narrative review. *Cureus*, *10*(12), e3681. https://doi.org/10.7759/cureus.3681

[44] Sawyer, B. D., Finomore, V. S., Funke, G. J., Mancuso, V. F., Funke, M. E., Matthews, G., & Warm, J. S. (2014). Cyber vigilance: Effects of signal probability and event rate. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *58*(1), 1771–1775. https://doi.org/10.1177/1541931214581369

[45] Smith, E. C., Holmes, L., & Burkle, F. M., Jr. (2019). Exploring the physical and mental health challenges associated with emergency service call-taking and dispatching: A review of the literature. *Prehospital and Disaster Medicine*, *34*(6), 619–624. https://doi.org/10.1017/S1049023X19004990

[46] Spence, R., Bifulco, A., Bradbury, P., Martellozzo, E., & DeMarco, J. (2023). The psychological impacts of content moderation on content moderators: A qualitative study. *Cyberpsychology: Journal of Psychosocial Research on Cyberspace*, *17*(4). https://doi.org/10.5817/CP2023-4-8

[47] Spence, R., Bifulco, A., Bradbury, P., Martellozzo, E., & DeMarco, J. (2024). Content moderator mental health, secondary trauma, and well-being: A cross-sectional study. *Cyberpsychology, Behavior, and Social Networking*, *27*(2), 149–155. https://doi.org/10.1089/cyber.2023.0298

[48] Spears, J. L. (2023) Job stress in the cybersecurity incidence response work role. https://devusec.de/2023-conference-wsiw/downloads/wsiw2023-DePaul-final4.pdf

[49] Steiger, M., Bharucha, T. J., Venkatagiri, S., Riedl, M. J., & Lease, M. (2021). The psychological well-being of content moderators: The emotional labor of commercial moderation and avenues for improving support. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–14. https://doi.org/10.1145/3411764.3445092

[50] Sundaramurthy, S., Bardas, A. G., & Case, J. (2015). A human capital model for mitigating security analyst burnout. *Proceedings of the Symposium on Usable Privacy and Security (SOUPS)*.

[51] Tariq, S., Baruwal Chhetri, M., Nepal, S., & Paris, C. (2025). Alert fatigue in Security Operations Centres: Research challenges and opportunities. *ACM Computing Surveys*, *57*(9), 1–38. https://doi.org/10.1145/3723158

[52] Tines. (2023). *Voice of the SOC 2023*. https://www.tines.com/reports/voice-of-the-soc-2023/

[53] Thimmaraju,K. , Rispens, S. I., and Ahn, G.-J. (2025). Human Performance in Security Operations: A Survey on Burnout, Well-Being and Flow State Among Practitioners *Proceedings 2025 Workshop on Security Operation Center Operations and Construction, San Diego, CA, USA: Internet Society, 2025* doi:10.14722/wosoc.2025.23002

[54] Vielberth, M., Böhm, F., Fichtinger, I., & Pernul, G. (2020). Security Operations Center: A systematic study and open challenges. *IEEE Access*, *8*, 227756–227779. https://doi.org/10.1109/ACCESS.2020.3045514

Table 1: Overview of Cybersecurity Literature by Study Type

Study	N	Population	Type	Key Contribution
Empirical Studies of Practitioners				
Chappelle et al. (2013) [7]	532	USAF cyber operators	Quant	First large-scale study; operational stressors
Sundaramurthy et al. (2015) [50]	26	SOC analysts	Qual	Human capital model; vicious cycles
Paul & Dykstra (2017) [36]	126	NSA/USCYBERCOM operators	Quant	Fatigue/frustration increases over operations
Dykstra & Paul (2018) [10]	126	NSA/USCYBERCOM operators	Quant	COSS instrument development
Hollis (2023) [17]	22	Incident responders	Qual	Emotional labor; confidentiality barriers
Spears (2023) [48]	14	Cybersecurity professionals	Qual	Weight of Responsibility; inability to unplug; exhaustion
Arora and Hastings (2024) [5]	50	Cybersecurity professionals	Quant	Reluctance to report struggles; cycle of silence
Reeves et al. (2024) [40]	119	Australian cyber professionals	Quant	Gender/role predict burnout; sleep quality
Mott et al. (2024) [30]	83	Ransomware IR professionals	Qual	PTSD symptoms; regulatory aftermath
Nepal et al. (2024) [32]	35	Corporate IR	Mixed	54% burnout; telemetry validation
Reeves et al. (2024b) [41]	506	Cybersecurity professionals/managers	Quant	Mitigations
Rangarajan et al. (2025) [39]	10	Former NSA operators	Qual	Mission as addicting; Work-life balance
Thimmaraju et al. (2025) [53]	19	SOC practitioners	Quant	Burnout/wellbeing/flow among SOC practitioners
Experimental Studies (Non-Practitioner Samples)				
Sawyer et al. (2014) [44]	24	Recruited participants	Exp	Vigilance decrement in cyber monitoring
Greenlee et al. (2016) [13]	71	Novice analysts	Exp	Triage vs. escalation stress profiles
Conceptual and Review Papers				
Agrafiotis et al. (2018) [2]	–	–	Taxon	Cyber-harms including psychological impacts
Nobles (2019) [34]	–	–	Persp	Formal human factors programs
Vielberth et al. (2020) [54]	–	–	SysRev	SOC building blocks; PPTGC framework
Nobles (2022) [33]	–	–	Framew	Security fatigue taxonomy; HRCO framework
Gupta et al. (2024) [14]	–	–	Scoping	Burnout definitions/antecedents; 15 papers
Tariq et al. (2025) [51]	–	–	SysRev	Alert fatigue; A2C framework
Industry Surveys (Non-Peer-Reviewed)				
Nominet (2019) [35]	408	CISOs (USA/UK)	Survey	91% moderate-high stress
Tines (2023) [52]	900	SOC analysts	Survey	63% burnout prevalence
ISC2 (2024) [19]	15,852	Global workforce	Survey	Workforce gap; skills shortage

Type: Quant = Quantitative; Qual = Qualitative; Exp = Experimental; Taxon = Taxonomy; Framew = Framework; SysRev = Systematic Review; Persp = Perspective. Paul & Dykstra (2017) and Dykstra & Paul (2018) report on the same dataset but are listed separately as they make distinct contributions.

Table 2: Nature of Work Comparison: Offensive vs. Defensive Cybersecurity Roles

Characteristic	**Offensive/Tactical Roles**	**Defensive/Reactive Roles**	**Key Sources**
Primary stressors	Operational conditions: shift work, operation length, hours worked	Work content *and* operational conditions; includes emotional labor, confidentiality barriers, invisible success (see Table 3 for full framework)	[7, 10, 36] vs. [17, 30, 32, 51]
Relationship to mission	High meaning and purpose; work described as "addicting"; mission masks negative factors	Mixed; invisible success undermines sense of accomplishment	[7, 36, 39] vs. [17, 30, 32, 51]
Temporal pattern	Discrete operations with defined endpoints; fatigue accumulates within operations	Extended incidents (weeks to months); continuous monitoring between incidents; regulatory aftermath extends strain for years	[10, 36] vs. [17, 30]
Adversary dynamics	Proactive engagement; operators initiate action against targets	Reactive posture; adversary adapts to defensive responses; success generates future sophisticated threats	[7, 36] vs. [17, 30]
Disclosure capacity	Limited by classification, but incidents are discrete and concludable; disclosure barriers not examined in offensive literature	Limited by confidentiality; incidents may remain legally constrained for years; "not something you can talk about"	[7, 36, 39] vs. [17, 30]
Burnout rates	26% emotional exhaustion (active duty); 15% (contractors)	54% burnout (incident responders, SIB and BAT); 63% (SOC analysts per industry surveys)	[7] vs. [32, 51, 52]
Primary departure reason	Work-life balance; not pay or mission dissatisfaction	Escape from harm; pay cuts accepted to leave; physical health crises	[7, 36, 39] vs. [17, 30]

Note: Offensive/tactical evidence derives primarily from military cyber operators (NSA, U.S. Air Force); defensive evidence derives from incident responders and SOC analysts across military and civilian contexts. Direct comparative studies are absent. This table synthesizes findings across separate research populations. The offensive literature is notably smaller and focused on operational factors; the defensive literature increasingly examines work content itself as a stressor source.

Table 3: Stressor Dimensions in Defensive Cybersecurity Work (What the Work Demands)

Dimension	**Definition**	**Evidence**	**Sources**
Sustained Vigilance	Continuous monitoring depletes cognitive resources over time; performance degrades during extended watch periods	Strong	[7, 17, 30, 32, 44, 50, 51]
Weight of Responsibility	Accountability for protecting organizational assets and personal data; defenders must succeed every time while attackers need only succeed once	Strong	[7, 17, 30, 32, 48]
Unpredictable Demands	Random incident timing prevents planning and creates chronic anticipatory stress; personal commitments disrupted without warning	Strong	[7, 17, 30, 32, 50]
Invisible Labor & Success Metrics	Successful defense is invisible to the organization; high false positive rates mean substantial work produces no visible outcome; recognition absent until failure occurs	Strong	[17, 32, 50, 51]
Emotional Labor & Exhaustion	Managing own fear and exhaustion while projecting confidence to stakeholders; "rapid succession micro resiliencies" required during active incidents	Moderate	[17, 30, 50]
Exposure to Human Malice	Immersion in evidence of adversarial intent: ransomware demands, breach data, extortion, attacker communications	Moderate	[17, 30, 51]
Latent & Adaptive Threats	Defenders cannot confirm when threats are neutralized; adversaries study defensive responses and evolve tactics; competent defense requires assuming compromise	Moderate	[7, 17, 30, 44, 51]

Note: Evidence strength reflects availability in current cybersecurity literature. 'Strong' indicates multiple studies with direct measurement or consistent qualitative findings across multiple studies; 'Moderate' indicates qualitative findings across two or more studies

Table 4: Recovery Barriers in Defensive Cybersecurity Work (What Prevents Processing)

Barrier	**Definition**	**Evidence**	**Sources**
Confidentiality Constraints & Isolation	Classification, NDAs, and legal proceedings prevent disclosure even to family and therapists; verbal processing structurally blocked	Strong	[7, 17, 30, 48]
Temporal Dynamics Preventing Closure	Extended incidents (weeks to months), regulatory aftermath (years), and on-call expectations prevent psychological disengagement; recovery windows unavailable	Strong	[17, 30, 32, 48]
Professional Culture	Norms demand practitioners absorb harm ("suck it up"); help-seeking stigmatized; low professional recognition	Moderate	[17, 30, 32]
Limited Institutional Infrastructure	No peer support programs, debriefing protocols, or clinical pathways tailored to cybersecurity; inadequate or limited workplace counseling; therapists unfamiliar with work context	Moderate	[17, 30, 32]

Note: Evidence strength reflects availability in current cybersecurity literature. 'Strong' indicates multiple studies with direct evidence; 'Moderate' indicates consistent qualitative findings across studies. These barriers may explain why cybersecurity practitioners report harms similar to adjacent professions while facing worse outcomes: the harms may be familiar, but what prevents recovery is not.