Non-Disruptive Disruption: An Empirical Experience of Introducing LLMs in the SOC

Francis Hahn*
University of South Florida
fhahn@usf.edu

Mohd Mamoon*
University of Kansas
mohdmamoon@ku.edu

Alexandru G. Bardas
University of Kansas
alexbardas@ku.edu

Michael Collins
University of Southern California – ISI
mcollins@isi.edu

Jaclyn Lauren Dudek
University of Kansas
jdudek@ku.edu

Daniel Lende
University of South Florida
dlende@usf.edu

Xinming Ou
University of South Florida
xou@usf.edu

S. Raj Rajagopalan
Resideo Technologies
siva.rajagopalan@resideo.com

***Abstract*—Security Operations Centers (SOCs) are high-stress, time-critical environments in which analysts manage multiple concurrent tasks and depend heavily on both technical expertise and effective communication. This paper examines the integration of Large Language Model (LLM) technologies into an operational SOC using an anthropological, fieldwork-based approach. Over a six-month period, two computer science graduate researchers were embedded within a corporate SOC, guided by an internal advocate, to observe workflows and assess organizational responses to emerging technologies. We began with an initial demonstration of an LLM-based incident response tool, followed by sustained participant observation and fieldwork within the incident response and vulnerability management teams. Drawing on these insights, we co-developed and deployed an LLM-based SOC companion platform supporting root cause analysis, query construction, and asset discovery. Continued in-situ observation was used to evaluate its impact on analyst practices. Our findings show that anthropological and sociotechnical approaches, coupled with practitioner co-creation, can enable the nondisruptive introduction of LLM companion tools by closely aligning development with existing SOC workflows.**

I. Introduction

A Security Operations Center (SOC) is an operational unit in which security analysts continuously engage with both technology and people, and where the effectiveness of processes and communication skills directly shapes workflow efficacy [1], [2]. To better understand the dynamics within a SOC in the context of emerging technologies, we embedded two computer science graduate students in an industry SOC. A designated member of the organization guided them, serving both as a mentor in navigating the organizational environment and as a source of legitimacy until individual trust was established. Through participant observation and fieldwork, we examined the tacit and explicit knowledge flows within the organization and co-created tools with practitioners that synthesized shared expertise to meet the needs of the operational environment. Using a sociotechnical approach, we show how embedded, trust-centered collaboration can support the effective and sustainable introduction of new technologies, such as Large Language Models (LLMs), in SOC settings.

Despite long-standing efforts to improve SOC effectiveness and reduce burnout [3] through improved processes, tools, and training, research consistently identifies a persistent gap between academic work and operational SOC practice (e.g., [4], [5]). Ethnographic studies show that tools designed without accounting for analysts' experiential and iterative reasoning often fail in practice [6], leading to limited adoption or abandonment [5], [7]. This reflects a recurring mismatch between practitioners' needs and researchers' assumptions.

Bridging the research-practitioner gap is vital, especially for emerging Artificial Intelligence (AI) and LLM-based SOC tools [8]–[10]. Although these systems promise efficiency gains, their success depends on analyst trust, workflow alignment, and organizational context. Co-developing solutions with SOC practitioners improves integration, adoption, and perceived usefulness, while giving researchers access to real incident workflows and decision-making that are otherwise difficult to study [11]. Prior work has shown that self-reported security behaviors (such as surveys) are reliable for certain measures but insufficient for others, particularly when tasks are complex or embedded within broader workflows [12].

Achieving meaningful collaboration, however, is non-trivial. SOC analysts operate in high-pressure, alert-driven environments with limited time to engage external researchers [2], [11], [13]. Much operational knowledge remains tacit and is shared informally through day-to-day interaction rather than formal documentation [6], [7]. While not all SOCs exhibit the same degree of insularity, gaining access to authentic operational practices still requires trust, time, and sustained presence. Embedded collaborations that capture real SOC work are therefore rare and difficult to execute, but when successful, they reveal rich sociotechnical insights into how people, tools, and processes jointly shape security operations [6], [11], [14].

*Francis Hahn and Mohd Mamoon contributed equally to this work and share first authorship.

Workshop on Security Operation Center Operations and Construction (WOSOC) 2026
23 - 27 February 2026 , San Diego, CA, USA
ISBN 978-1-970672-03-9
https://dx.doi.org/10.14722/wosoc.2026.23015
www.ndss-symposium.org

Over a six-month engagement, our graduate students (field-workers) participated in on-boarding and training, attended meetings, and collaborated closely with analysts and managers. This sustained presence enabled us to observe workflows, tool usage, communication patterns, and breakdowns from an insider perspective while minimizing disruption to ongoing operations. Sustained day-to-day interaction (along the lines of [15]) fostered trust, allowing access to candid discussions and operational practices that would otherwise remain out of view from research teams.

This engagement enabled the iterative design and deployment of an LLM-based SOC companion platform that integrated multiple tools alongside industry-standard SOC technologies and supported a co-created, multi-stage effort for embedded research in operational environments. Tools were introduced incrementally, allowing analysts to shape their functionality and integration, resulting in strong engagement and a deliberately non-disruptive adoption that augmented rather than replaced existing workflows.

II. Related work

Related work is organized into three areas: embedded sociotechnical studies of SOC practice, AI and LLM-based tools for security operations, and analyst-centered research on trust, explainability, and workflow integration.

A. Embedded Sociotechnical Studies in SOCs

Prior research at the intersection of cybersecurity, sociotechnical systems, and researcher-practitioner collaboration highlights persistent gaps between academic research and operational security work. Ethnographic studies of SOCs demonstrate that security operations are complex, practice-driven, and shaped by social dynamics as much as technical tools. Sundaramurthy et al. [6], through a 3.5-year multi-site anthropological study, model SOC work as an activity-theoretic system in which analysts continually navigate contradictions between tools, procedures, and organizational expectations; they argue that meaningful innovation requires embedded researchers who understand the culture and tacit practices of analysts. Subsequent fieldwork reinforces this view. Jones et al. [14] showed how COVID-19-driven shifts toward endpoint-centric monitoring increased operational load, introduced new privacy tensions, and exacerbated analyst burnout, underscoring the sociotechnical fragility of modern SOCs.

Beyond SOCs, Tuladhar's insider ethnography shows that security practices become organizational norms only when learned in context and reinforced through hands-on collaboration with security advocates [16]. Vielberth et al.'s [2] systematic review similarly finds that SOC research is fragmented, relying on interviews or isolated case studies and lacking holistic sociotechnical frameworks. Greig et al. [17] further identify persistent gaps between formal security policies and everyday practice. The common theme is sociotechnical: security is enacted through people's stories, workflows and peer norms, not just through technology.

B. AI and LLM-based Tools for Security Operations

Recent work has begun exploring how LLMs and other AI methods can support SOC tasks. Srinivas et al.'s [18] survey and taxonomy of AI in SOCs catalogs applications ranging from log summarization and triage to report generation and vulnerability management, but notes persistent issues (e.g. interpretability, robustness, legacy integration, hallucinations) that hamper operational fit. Empirical studies reinforce these points: Mink et al. [8] found that practitioners view machine learning (ML) tools as augmenting, not replacing, existing rule-based workflows and demand contextual explanations beyond raw feature highlights. Likewise, Binbeshr et al.'s [10] review reports that ML techniques currently dominate SOC research and highlights challenges in data quality, model interpretability and integration, even as AI-driven approaches promise improved detection accuracy and scalability.

In the LLM context, Singh et al.'s [19] 10-month longitudinal study of 45 analysts showed that GPT-4 was chiefly used as an on-demand sense-making aid (for interpreting commands and contextualizing telemetry) rather than for final decisions. Relatedly, Microsoft's Copilot for Security guided response (CGR) system [20] has demonstrated strong triage and remediation performance at scale using aggregate metrics, but its impact on individual analysts' workflows unexamined.

C. Analyst-centered and Sociotechnical Factors

Rastogi et al [21] showed that generic explainability techniques often miss analysts' needs in high-pressure SOC environments. Nyre-Yu et al. [22] report that a deployed explainable-AI tool was scarcely used and did not improve analyst accuracy; they stressed on aligning tool outputs with users' roles, workflows and trust thresholds. Similarly, Mink et al. [8] found security practitioners value low false positives and desire explanations that help verify alerts in context. These insights echo broader human-AI collaboration research which emphasizes maintaining user authority and embedding AI outputs into existing workflows. In particular, trust-centered design and strategies for sustained adoption remain open problems. Bridging this gap will require iterative, embedded research in live SOCs to ensure AI companions truly fit analysts' needs and constraints.

III. Methodology

We conducted an ethnographic study using participant observation, fieldwork, and co-creative practices to develop a grounded understanding of organizational dynamics and social structures, see Figure 1. Through close engagement, we examined how SOC analysts carry out their work, how existing tools are integrated into daily workflows, and how analysts interact with one another and their environment. By working on peripheral tasks [23], we encountered situations [24] where we engaged directly with the SOC analysts. These interactions complemented with interviews and informal conversations and allowed us to experience firsthand the practical challenges analysts face in their everyday operations.

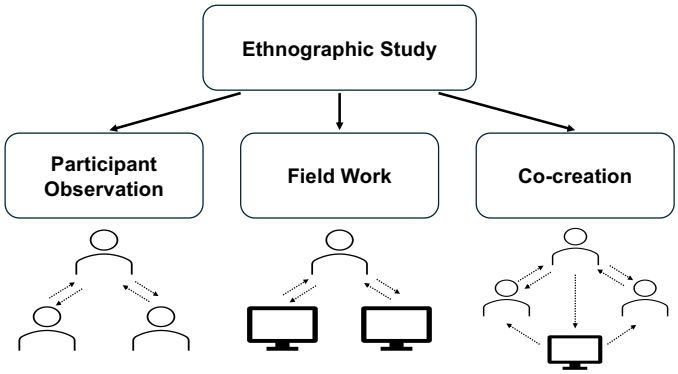

Fig. 1: Overview of the Adopted Ethnographic Approach – How participant observation, fieldwork, and co-creation were used to examine analyst workflows, interactions, and tool use within an operational SOC.

A. Internal Advocate

This study began with two graduate students embedded within an organizational setting, where they were guided by an advocate who was a member of the organization, or an internal advocate (IA). The IA started our engagement by setting up meetings with various members of the organization exposing the fieldworkers to a wide breadth of positions within the organization, broadening their social circle and giving perspective to the landscape [25]. This guided entry into organization had two critical benefits **(1)** it gave the fieldworkers the opportunity to understand how to engage with the organization's members and **(2)** it informed the members of the organization of legitimacy of these unknown outsiders [26], thus establishing the initial seed of trust. The IA also met frequently with the fieldworkers in both one-on-one sessions and group settings allowing them to inquire about observed practices and to learn about organizational intricacies from a trusted-insider who was external to the study.

B. Participant Observation

The fieldworkers used participant observation to develop an external perspective on these teams while gradually establishing trust through ongoing social interaction. The work environment was completely remote, which posed challenges for participant observation due to the absence of informal "water-cooler" interactions, spontaneous conversations, or opportunities to casually observe colleagues at work. Instead, the fieldworkers conducted participant observation through Microsoft Teams meetings, shared screens, and text-based communication. To support this process, fieldworkers relied on artifacts such as organizational hierarchy maps to understand the company's social and operational structure, identify relevant stakeholders, and coordinate meetings. These tools enabled the fieldworkers to systematically engage with a broad range of critical personnel across teams that were geographically and continentally distributed.

C. Fieldwork

After approximately one month of observing and engaging with members of the security team, the fieldworkers developed an initial understanding of how work was done and identified several recurring pain points. However, to deeply understand these challenges and explore potential solutions, it was necessary for the fieldworkers to engage directly in the analysts' work. The fieldworkers requested access to relevant databases, servers, API systems, and submitted a formal proposal seeking permission to execute code and deploy LLM models on the organization's infrastructure. This role-based fieldwork enabled to gain hands-on experience with a broad range of tools, tasks, and procedures encountered in daily security operations. Beyond informing subsequent discussions with the security team, this experience provided first-hand insight into how operational challenges arise, how they are shaped by constraints, and how they propagate across other workflows and responsibilities within the organization.

D. Co-creation

Our approach to co-creation followed an iterative design that heavily relied upon continued engagement with the security team members to leverage their tacit knowledge which was made explicit through conversation and engagement to inform and direct the creation of these solutions [27]. This approach utilized numerous instances of employee testing and demo sessions to present and gather information on how useful they considered the tools, what issues existed, and how to develop it further to be a tool that could be integrated into their day-to-day workflows. We deliberately incorporated co-creation into the study to ensure the development of the tools embedded the synthesized experience of the fieldworkers and the analysts. Co-creation in the context of our study was viewed through two complementary lenses: **(1)** a process in which researchers and stakeholders collaboratively engage in the ideation, planning, implementation, and evaluation of new services and systems to enhance the practical impact of research outcomes [28]; and **(2)** the joint generation of knowledge through close collaboration between academics and stakeholders from other sectors [29].

IV. Observing the Issues within the SOC

To better understand the SOC, the fieldworkers developed an initial LLM tool (ContextDemo) for demonstration purposes. This was the first step towards showing the security team what the LLMs could do for them, but presented in a way that would allow us to understand the needs and how the analysts wanted to use the technologies. It was developed to be the first window into the analysts workflows, methodologies, and building trust by positioning themselves as capable and knowledgeable for developing actionable tools specific to their needs. The way ContextDemo was designed was through a single agent process connected to numerous vector stores and two SOC tools. ContextDemo was developed prior to beginning the embedded research from information provided by the organization based on the expectations of our positions.

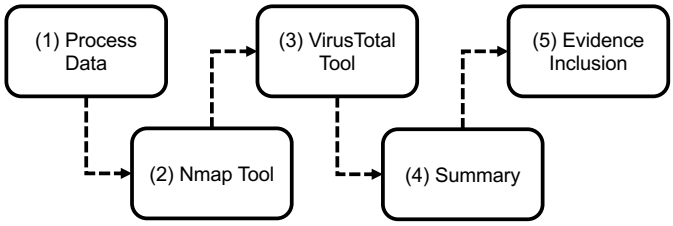

Fig. 2: ContextDemo Workflow Visualization – Portrays the actions the tool takes before producing an output.

A. Demo Work

*1) **Process Data**:* ContextDemo was given access to a "ticket" which described the type of investigation, which included minor relevant details, a collection of documents representing data from a mock organization, and three vector stores containing data from CWE, CVE, and MITRE STIX. These vector stores were created to show the security team the ability to attach various repositories of information to the LLM for retrieval. We pre-processed the data to collect names and relevant IP addresses by correlating the names with a mock organizational Excel spreadsheet which mapped users to workstations with various IP addresses. It was designed to show that the LLM could be designed to pull in information from various Excel sheets.

*2) **Nmap Tool**:* The Nmap tool [30] performed a fast scan on Docker containers modeling the data pairs identified during pre-processing, demonstrating that the LLM and tool design can be orchestrated for task completion.

*3) **VirusTotal Tool**:* The VirusTotal tool would use any URL collected from the pre-processing to scan for phishing links or malware deployment. The VirusTotal tool was designed to show that the LLM could utilize API libraries to help complete tasks.

*4) **Identification and Summarization of Indicators of Compromise (IoCs)**:* The first pass was purely the LLM generated summarization of the SOC ticket and the accompanying documents with a system message directing the purpose to identify IoCs. The LLM would use the information stored in the vector stores to support the findings.

*5) **Analyze Summaries and Provide Evidence**:* In the second pass, the system substantiated the initial summary with tool-generated evidence, demonstrating ContextDemo's ability to combine static information with dynamic runtime data to support trustworthy outputs. Subsequent discussions involved SOC analysts in incident response and vulnerability management, alongside AI governance and legal teams and managers at multiple levels.

B. Critical Conversations

*1) **Chief Information Security Officer (CISO)**:* The first critical discussion was with the organization's CISO as this gave the research a top-level priority on how to properly address the issues the organization faced. To understand directly the CISO was asked what their definition of a productive workflow was, their response directed us to three areas of focus: SIEM-based threat detection, incidence response, and vulnerability management. This helped us understand the critical nature of these areas by describing that any lapse in these areas due to communication or the technologies presents an opportunity for risk. There was continued discussion on how having a tooling platform that would complement the existing technologies would help improve the workflows by enabling rapid response with minimal stop points.

*2) **SOC Director**:* The engagement with the SOC director focused initially on learning the organization's security vendor platform and on identifying opportunities to design automation workflows using the capabilities already available to the SOC. As discussions progressed, there was group discussion on how AI-based tools could be leveraged to further enhance these automation workflows. In parallel, the fieldworkers examined the classification of security incidents and alerts within the security vendor's platform to inform alert tuning efforts, relying on input from the SOC to identify specific operational gaps that could be meaningfully improved. Through these iterative conversations and brainstorming sessions, the team converged on solution areas related to asset discovery and end-of-life (EOL) remediation. The SOC director emphasized that while both tasks are critical to the organization's overall security assurance, they are often among the most time-consuming due to the need to process large volumes of documentation, database records, and frequently changing application dependencies.

*3) **Legal Team**:* Speaking with the organization's legal team on actionable requirements following a security event that exceeds a threshold of severity which involved: data loss, data alteration, or data injection, lead to a discussion on Root-Cause-Analysis or RCA. RCA is an event where multiple members of various teams come together, under a strict window of time to analyze and understand what were the "root causes" of the incident and who was involved. It was noted that this can often be a very exhausting process both mentally and emotionally. There was an identification of usefulness for an LLM or automation tool that could conduct this work to aid in speeding up the RCA process by reducing the amount of human-interaction involved and thus minimizing the cognitive load on it's participants.

*4) **SOC Analysts**:* The analysts wanted to know more about how ContextDemo (the demonstration tool) worked and how it could be integrated into their workflow. Leveraging what we had learned from the legal team and the SOC director, we began asking the analysts to walk us through their process and day-to-day tasks. In the discussions the fieldworkers focused on the tasks of asset discovery and RCA. Asset discovery is a critical task for ensuring the vendor's security platform had visibility on the organization's various assets. This walk-through of their day-to-day task of managing databases of information, using various Python scripts, and correlating data across various spreadsheets to create a summary report for an asset owner to read was identified as something that could be automated using an LLM and would save them

countless working hours. RCA was a little more challenging for them to walk us through due to two factors: **(1)** RCA was not commonly done due to the incidents being few and far between, and **(2)** the sensitivity of the data regarding the events made it difficult for the analysts to reproduce the events. However, the analysts were able to simulate the process by walking the fieldworkers through how they use the security vendor's platform and various database systems for searching for relevant log and event data to corroborate claims made during an RCA event. What we found throughout the walk-through for the RCA event was similar to asset discovery, that there is an overwhelming amount of data which requires consideration and no consolidated platform to consider it.

V. Fieldwork

After the first month of conversations the fieldworkers wanted to get a hands-on perspective of what the conversations with the analysts had brought to their attention. To do this, the embedded researchers worked directly as SOC analysts performing the work to learn both the formal and informal processes which lead to the pain-points for conducting incidence response and the challenges faced when developing solutions as members of the vulnerability management team.

A. Acting within the SOC

As SOC analysts the researchers were on-boarded as members of the organization's SOC. The first step was being assigned tickets for events labeled as "alerts". Alerts were organization specific investigations, the difference between an "alert" and an "investigation" was that the organization's security platform vendor provided SOC services to handle the investigations which were generated by the security platforms detection services and sent to the system information event management (SIEM) platform. Investigations were defined as events handled by L1/L2 analysts, while cases requiring higher levels of expertise or access to information unavailable to the security vendor's SOC were escalated to alerts for internal handling by the organization.

The fieldworkers handled alerts for one month before adding additional tasks and duties to their workflow. During the work as SOC analysts the fieldworkers learned three critical aspects of the work: **(1)** the amount of social interaction required to perform these investigations, **(2)** the challenges with generating queries for the vendor's security platform using a proprietary query language, and **(3)** how the information to confidently close a ticket required correlating information across various sources of data. There were numerous interactions with the SOC team to confirm when alerts were appropriate for closure, validate methodologies for handling different alert causes, and understand analyst query construction practices.

B. Managing Vulnerabilities

Another role the fieldworkers acted in was Vulnerability Management. This role covered a range of tasks, but what the analysts focused on was aiding with maintaining the organization's security vendor's agent system which provided visibility to asset event logs and system information for producing investigation and alert tickets. This was described as a critical task by both the SOC director and the SOC analysts who worked in this area. The agents used to report to the security platform could be in one of the three following states:

- **Active:** The state of actively reporting to the security platform within the last four hours.
- **Stale:** The state of not having reported to the security platform between 4 and 18 hours.
- **Not Reporting:** The state of having no record of an asset, indicating the agent had not been properly installed there.

To identify the state of the assets requires the use of the security vendor's API system, which required permitted access to the organization's jump server and permitted access to the security vendor's platform by the organization's platform manager. Once given the appropriate access and permissions the next step was to write scripts to retrieve the necessary information from the security vendor's platform using their APIs. After collecting the necessary information we found the logging data did not output the documented state and that the documentation made suggestions for labeling the states, this was confirmed in a meeting with the security vendor's support team through the security vendor's ticketing service.

The next step was pre-processing on the data to identify a way to label each agent as a particular state. Once we had the created programs for collecting data from the API and pre-processing to match the labeling suggestions we would meet with the vulnerability management team members to discuss how to take this information and report on the assets. Because many of the assets are hosted in various cloud networks and the data regarding ownership is hosted in separate databases this required a large amount of effort to find the database which housed the owner's contact info, which required additional permission requests and guidance on how to securely handle data on both a development and production server. The fieldworkers observed that vulnerability management work requires a combination of cybersecurity expertise, software development proficiency, and effective communication.

While many workflow issues were observed during our fieldwork, we found that the issues of social interactions, query building, and distributed data were among the highest ranking when considering issues with the workflow.

*1) **Issues with Social Interactions**:* An observed issue with social interactions found during the fieldwork as SOC analysts was that when conducting an alert investigation there are often times where confirmation of certain events observed the alert logs is required from the effected users. Due to a variety of reasons there can be delays in communications which result in challenges with the closing of the alert or a need to lock an account during the presence of a false positive. Several challenges arose, including remote communication driven by the organization's multinational footprint, workloads that delayed responses, and the time required to troubleshoot technical issues with users who have varying technical backgrounds.

*2) **Issues with Query Building**:* While working as SOC analysts the main feature for investigating log data was through

the security vendor's SIEM platform which used a proprietary query language. While the security vendor provided documentation for use, the types of queries often required were more advanced than the provided documentation could assist with, which introduced noise and additional time to prepare queries. The security vendor provided an AI query building tool, however due to the alerts involving data that the security vendor was not permitted access to this limited the capabilities of the tool in aiding with building queries for the organization's analysts. A key observation that arose with query building is that if sufficient queries could not be crafted within a reasonable amount of time, the fallback to this was to directly communicate with users.

3) ***Issues with Distributed Data****:* An issue that was shared by both the SOC and vulnerability management teams was the issue of distributed data. For the SOC analysts this distribution of resources introduced additional time spent on alerts by needing to search various databases spread across multiple applications or service providers. As a SOC analyst a piece of information stored in one of many distributed data locations could be critical to accurately or quickly labeling an alert as a true positive or false positive. For the vulnerability management team the issues with distributed data sources was emphasized greater as a source of repetitive and mundane actions. This was particularly evident when identifying whether the security vendor's platform agents were processing data correctly and, when failures occurred, determining the responsible asset owner.

VI. Co-creating New Tools

The process of tool-building went alongside participant observation and the fieldwork which quickly lead to usable insights. During the initial design phase there were three approaches to our considerations on what tools to build **(1)** the embedded researchers and SOC analysts decided on tool designs based on discussions over repeated demonstration of ContextDemo and during participant observations and **(2)** other tool designs were identified through the fieldwork where points in the workflow would lead to discussion on how to make improvements or automate the process. The third observation **(3)** was that as the embedded researchers began developing and deploying tools, members of the SOC would actively make requests for additional capabilities to be integrated into the SOC companion platform. In response to these requests and observed needs, the fieldworkers prioritized the development of LLM-based agents for root cause analysis (RCA), asset discovery, and query building.

A. Root Cause Analysis

The fieldworkers and SOC analysts jointly developed the RCA tool through an iterative process in which they collaborated to define requirements, refine functionality, and evaluate the tool's outputs. The fieldworkers took lead in implementation (programming the prototypes), while the analysts provided continuous input and domain guidance. The RCA tool was prototyped as an analyst-facing interface integrating a large language model that ingests an incident's alerts and relevant logs to give a structured incident summary, impacted entities, root cause identification, contributing factors, remediation actions, and preventive measures for the future. During development, early builds were shared with a small group of analysts to get rapid feedback on accuracy and usefulness. This tight feedback loop ensured the first prototypes aligned with analyst expectations before wider testing. As one senior analyst noted during a demo, ***"the initial version is on the right track"***, and they offered concrete suggestions for improvement (such as refining the user interface and ensuring the LLM's recommendations were interpretable).

Development then entered an iterative refine-and-test cycle. The prototypes were deployed within the SOC, and analysts used them on real incidents over multiple iterations. After each trial, feedback was gathered through observations, think-aloud sessions, and debrief interviews, then rapidly incorporated changes. Each iteration brought the tools closer to the SOC's needs. By the second demo, the analyst's feedback had grown more positive (they reported that LLM suggestions were becoming more relevant and saw clear time-saving benefits), coupled with remaining requests like improving the UI/UX and adding some guardrails for the LLM outputs. This participatory refinement ensured the solutions evolved in-sync with user expectations. A pain point was the limited access to historical RCA reports. Formal RCA documents are rarely produced, typically only for severe or escalated incidents. As a result, our limited access constrained our ability to ground early prototypes in extensive real-world reference material or to train and validate the tool against a large corpus of completed analyses. This scarcity reflects a broader reality of SOC practice, where RCA documentation is often inconsistent, informal, or de-prioritized under operational pressure.

The fieldworkers remained on-site to support initial use of the tool and to troubleshoot any issues. By embedding the new RCA tool into routine operations, real-world adoption was observed. The close collaboration up to this point paid off: analysts were receptive and even eager to incorporate the new aids, since they had co-designed the tools and could see the time-saving potential in their day-to-day tasks.

B. Asset Discovery

The asset discovery tool was developed to aid the vulnerability management team by autonomously detecting and reporting on agents which were labeled as **stale** or **not reporting**. A team member described the necessity of the tool based on the anticipated amount of work for the following year caused by reported workflow speeds due to AI assisted development. The asset discovery agent was discussed by the analysts and fieldworkers with the desired capability to leverage the LLMs ability to parse and efficiently make connections between large amounts of data that is often widely distributed across platforms and tedious for humans to recognize patterns in logs that is often challenging for humans to interpret [31].

To achieve this functionality the agent was designed to autonomously make the necessary API queries based on the

code developed during the fieldwork performed as part of the vulnerability management team. The goal was to look for updates to the asset agent logs, generate a notification of the issue, produce remediation steps based on the operating system, draft an email based on the findings and then store for the team member to review before the agent would send to the owner of the asset. The design of the tool was through the analysts and fieldworkers discussing various approaches by sharing YouTube videos and documentation snippets describing techniques and functionality. This resulted in a tool in which the only required manual effort from the user is the review and acceptance of the generated content. Meetings were held for one month to demonstrate progress on the integration of the API code, after which text-based communication became the primary mode of contact.

The initial demonstration was presented to the vulnerability management team, who found the tool beneficial and noted that the summarization performed as expected. An issue that arose with the development of the asset discovery tool was that the SOC analysts who were apart of the vulnerability management team were often busy being part of two teams and as async meetings regarding the design and development increased the focused was lost. The de-syncing of communications caused a lag in development due to the disruption required to actualize the tools full capabilities.

C. Query Builder Agent

The query builder agent was not explicitly requested; rather, it was discovered during fieldwork and subsequently described as a highly valued capability. The fieldworkers discovered this tool during their time working as analysts. Shortly after the engagement with the organization the security vendor launched a query builder tool readily available for use alongside their SIEM service. However, based on their hands-on experience and their understanding of the data accessible to the vendor, the fieldworkers concluded that an internally developed query builder could offer a superior user experience. The internally developed tool differed from the vendor-provided solution by iteratively engaging analysts in discussions about relevant internal documentation, thereby enabling controlled access to contextual information not shared with the security vendor. This training process relied exclusively on supplying data to the LLM's short-term context through retrieval-augmented generation (RAG) and prompt engineering.

The RAG system was populated with documents that incorporated queries developed by analysts and fieldworkers through ongoing discussions of shared and role-specific use cases. These documents described investigation steps and associated queries and were jointly authored by the analysts and fieldworkers. A challenge emerged during the data creation and ingestion phase: the process introduced a disruption to analyst workflows by requiring the explicit documentation of procedures that were largely tacit in practice. Although analysts were generally willing to contribute during discussions, document creation did not become a sustained priority until the SOC director was made aware of the lag in the creation process and elevated the task as an explicit requirement.

Once analysts observed the force-multiplying effect of user-supplied information, they directly requested a feature that would allow them to provide contextual information from alert logs to help populate query fields. Enabling analysts to supply additional context to the LLM when crafting queries reduced the need to manually parse and encode context-specific search parameters. This shift marked a transition from disruptive co-creation and integration to a largely non-disruptive process. One SOC analyst noted that this style of tool development was unfamiliar, stating that it ***"is not the type of tool we are used to have developed for us"***.

VII. Discussion

We view this six-month embedded collaboration as an approach to breaking the barriers between industry professionals and academic researchers by becoming encultured [32], [33] through an embedded research effort where the graduate researchers learn the norms, practices and intricacies of working within an organization and its SOC. We recognized how imperative it is to understand the nuanced details of the work to truly understand the needs for developing impactful solutions using the emerging technologies of today. To achieve the knowledge required to build effective tools and conduct this research, we adopted a Zone of Proximal Development (ZPD) perspective [34], positioning ourselves as learners embedded in close collaboration with knowledgeable analysts.

A. Bootstrapping Trust

The role of the internal advocate (IA) proved critical to establishing trust. The necessity of this role became evident early in the engagement, when the IA was not initially included. Shortly thereafter, referencing the IA by name or including them on email threads was sufficient to obtain approvals, as the researchers were able to leverage the trust already established between the IA and the organization. At several points during the engagement, team members reported that they met with the IA to review and discuss plans proposed by the researchers, with approval granted through these discussions. As the engagement progressed, the IA remained involved, but we observed a shift: team members increasingly contacted the researchers directly, and access no longer required explicit approval from the IA. This transition signaled that the trust vested in the IA had been effectively imparted to the researchers. As a result, researchers were able to engage more seamlessly in their work without disrupting the IA's ongoing responsibilities.

B. Disruption and Non-Disruption

The activities of participant observation, fieldwork, and co-creation can be seen as three separate yet parallel journeys [35] to understand the cognitive aspects of being a SOC analyst in a large-scale organizational setting. These activities enabled us to identify networks of interaction [36] among humans, technologies, and processes in which knowledge is implicitly

embedded, providing access to valuable and largely undocumented tacit knowledge [37]. This understanding supported the non-disruptive development and deployment of the LLM technologies introduced into the SOC. This approach is characterized by a trust-centered engagement with continuous feedback loops. Our approach distinguishes itself from the perceived typical vendor-driven or top-down technology deployments, by emphasizing co-creation and aligning to workflows over one-size-fits-all solutions.

A key outcome of the field work in Section V was the discovery of the SOC processes as they are actually implemented, including informal practices and undocumented workarounds that diverged from prescribed procedures. Surfacing these lived workflows was essential to understand where existing tools constrained the analyst's work and to avoid premature automation. This process-level understanding directly informed the subsequent co-development of LLM tools that aligned with real operational practice.

To understand progress, adoption, and disruption caused by the tools provided to the SOC, fieldworkers relied on the feedback of analysts. As one analyst put it, ***"any kind of automation that reduces the workload of analysts is welcome"***, indicating a strong desire for well-targeted tooling. Through brainstorming sessions and workflow sketching, we combined the analysts' domain knowledge with the fieldworkers' technical expertise to formulate solutions that were both useful and technically feasible. Our decision to develop the RCA tool, query builder, and asset discovery tools were well received by many of the analysts despite the initial set of issues identified. The issues of the criticality of time during investigations introduced by the CISO were echoed by an analyst during the platform's post-deployment discussions, where ***"time sinks... communicating with users and waiting for responses and confirmation on the suspicious activity ... reducing that time is critical"***.

Further discussions with the analyst were noted on how the developed tools provided a meaningful reduction to their time spent on an event, which we viewed as a characteristic of non-disruptive integration. Analysts had reported to the fieldworkers that ***"...we can get more rich data to make decisions. Reducing the need to reach out to users, by helping us generate queries that provide more information"***. However, an initial drawback or occurrence of disruption to the analysts was a wide-spread lack of understanding how to properly use the tool and transparency on the data the tools had available to them. This prompted the fieldworkers to return to developing various UI/UX elements to elevate usage, as well as develop a process allowing the analysts to provide their own data to store for later use and provide the team as a whole access to a collective knowledge source. Despite the drawbacks of the tool it was found to be a positive force-multiplier to the analysts' workflows.

Another characteristic of non-disruptive integration was the recognition that our tools were filling gaps created by the wide-range of tools and resources which require interaction for any given task. We observed that during this time the amount of tools and resources required to work had reached double-digits in quantity. One of the analysts mentioned that an issue with their work was having to dig through and correlate data across multiple sources. The motivation for developing a tool that brought together multiple sources of information into a single location aimed to reduce the cognitive load by having to manually access the resources. In another discussion with an analyst, following two demonstration iterations and in the context of planning further tool development, the analyst remarked with respect to the proposed features that ***"reducing the number of features will make for a truly useful and performant application."***

C. A Transferred Trust

Importantly, the analysts framed the collaboration itself as impactful. Enthusiasm for the project's continuity was evident as the team members repeatedly inquired about next steps and notifying the embedded researchers on the disruptive nature of early deployments, this signaled a notable degree of early trust by the organizational collaborators and a desire to keep the momentum going. We believe this to be a large factor in our trust building due to the common occurrence of security tools being dropped into the SOC with limited ways of modifying their usage to fit organization specific needs. By introducing disruptions into the workflow and receiving feedback on them to make them non-disruptive we established a sense of agency among the analysts which allowed them developmental control over these companion tools. An analyst remarked that they had ***"never seen a project where [outsiders] come in and build tools with us according to our needs"***, highlighting how rare and welcomed this embedded co-development was in their experience.

Taken together, these findings indicate that success cannot be assessed solely through productivity metrics. Qualitative factors such as sustained engagement with researchers, increased analyst confidence, and expanded internal capabilities proved equally important. SOC leadership identified complementary criteria, noting that, in the words of the director and CISO, ***"successful initiatives must (1) be replicable and capable of being institutionalized, and (2) produce outputs that yield corporate value through tools that are genuinely usable in the operational environment."***

We had concerns that once the embedding work had concluded, tool adoption would conclude with it. To analyze continued usage a follow-up was conducted a month after the fieldworkers' departure from the organization, it was mentioned by 2 of 3 analysts that the tools were still in use. One analyst mentioned during discussions for additional models to be made available. Another analyst mentioned ***"We use these things, when we need to use it, it helps us learn more"***, this was a strong signal that the tools were augmenting the workflow and not disrupting it. The analyst stating that the tools are there when they need them informed us that the tools are not something they relied on to get the job done, but something they used to help get the job done, signaling a companion aid style of usage instead of a dependence style

of usage [38]–[40]. A desire for further development and as needed usage were consistent with observed usage during initial tool deployment, signaling that tools were still in-use and being considered.

By building solutions with the analysts rather than for them, the co-building phase established a foundation of mutual trust and relevance. In contrast to the perceived vendor model where a one-size-fits-all product is introduced into the SOC, our embedded co-development led to tools that were "owned" by the team, both conceptually and physically by uploading the tools' source code to the organization's GitHub repository. This collaborative groundwork set the stage for future adoption, as participants felt personal investment in the technology and confidence that it would genuinely fit their workflows.

VIII. Conclusion

This work demonstrates that the process by which LLM tools are introduced can be as important as the tools themselves. A co-creative, embedded approach fostered analyst trust, ownership, and alignment with existing workflows by augmenting established practices rather than displacing them, resulting in non-disruptive adoption. Trust building rested on two complementary mechanisms: an internal advocate who initially bootstrapped trust and provided direction, and embedded researchers who established credibility by participating as analysts during fieldwork and responding directly to workflow disruptions caused by early tool iterations.

Through this ethnographic engagement, fieldworkers trained in both computing and anthropology enabled the SOC to gain applications tailored to its operational needs, while the fieldworkers themselves became encultured in the SOC domain, enhancing their capacity to engage with future innovations. Although this work reflects a single organization over a limited period, its core contribution is to show how embedded, trust-based co-development can shift AI tool adoption from a disruptive intervention to a sustainable, contextually grounded improvement in SOC practice.

IX. Future Work

There is a clear need for longitudinal work to understand the long-term effects of using LLM-based tools in security operations. This includes evaluating the security risks introduced by deploying LLM- and AI-based systems into organizational pipelines and assessing how such tools may themselves become targets for exploitation. A systematic examination of how reliance on LLMs to support or close investigation tickets affects analysts' long-term skills, judgment, and decision-making is also necessary to better understand the psychological and cognitive implications of LLM use in cybersecurity. In parallel, continued embedded, fieldwork-based cybersecurity research is critical for understanding how trust in LLM-based tools is established, sustained, and recalibrate when failures or unexpected behaviors occur.

X. Ethical Considerations

This study involved security analysts and managerial staff and was conducted in accordance with established ethical research standards. Research activities occurred between July and November 2025 and were approved by the Institutional Review Boards (IRBs) of the participating universities. The IRBs granted a waiver of written consent; all participants instead provided informed oral consent and were given IRB-approved protocol documentation describing the study's purpose, procedures, and participant rights, including the ability to withdraw at any time without penalty.

XI. LLM Usage Considerations

LLMs were only used for editorial purposes in this manuscript, and all outputs were inspected by the authors to ensure accuracy and originality.

XII. Acknowledgments

This work is supported by the National Science Foundation under awards no. 2143393, no. 2235102, and the Office of Naval Research under award no. N00014-23-1-2538. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of these agencies.

We are grateful to our organizational collaborators for welcoming us into their environments and for their partnership in enabling this work through shared learning and co-creation.

References

[1] C. Zimmerman, *Ten Strategies of a World-Class Cybersecurity Operations Center*. MITRE Corporation, 2014. [Online]. Available: https://books.google.co.uk/books?id=yieHoAEACAAJ

[2] M. Vielberth, F. Böhm, I. Fichtinger, and G. Pernul, "Security operations center: A systematic study and open challenges," *IEEE Access*, vol. 8, pp. 227 756–227 779, 2020. [Online]. Available: https://api.semanticscholar.org/CorpusID:230513062

[3] K. Thimmaraju, S. I. Rispens, and G.-J. Ahn, "Human performance in security operations: a survey on burnout, well-being and flow state among practitioners," in *Proc. 2025 Workshop on Security Operations Center Operations and Construction (WOSOC 2025)*, 2025, pp. 2–4.

[4] F. B. Kokulu, A. Soneji, T. Bao, Y. Shoshitaishvili, Z. Zhao, A. Doupé, and G.-J. Ahn, "Matched and mismatched socs: A qualitative study on security operations center issues," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 1955–1970. [Online]. Available: https://doi.org/10.1145/3319535.3354239

[5] J. M. Haney, C. C. IV, and S. M. Furman, "Towards bridging the Research-Practice gap: Understanding Researcher-Practitioner interactions and challenges in Human-Centered cybersecurity," in *Twentieth Symposium on Usable Privacy and Security (SOUPS 2024)*. Philadelphia, PA: USENIX Association, Aug. 2024, pp. 567–586. [Online]. Available: https://www.usenix.org/conference/soups2024/presentation/haney

[6] S. C. Sundaramurthy, J. McHugh, X. Ou, M. Wesch, A. G. Bardas, and S. R. Rajagopalan, "Turning contradictions into innovations or: How we learned to stop whining and improve security operations," in *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. Denver, CO: USENIX Association, Jun. 2016, pp. 237–251. [Online]. Available: https://www.usenix.org/conference/soups2016/technical-sessions/presentation/sundaramurthy

[7] R. Werlinger, K. Hawkey, D. Botta, and K. Beznosov, "Security practitioners in context: Their activities and interactions with other stakeholders within organizations," *International Journal of Human-Computer Studies*, vol. 67, no. 7, pp. 584–606, 2009. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1071581909000354

[8] J. Mink, H. Benkraouda, L. Yang, A. Ciptadi, A. Ahmadzadeh, D. Votipka, and G. Wang, "Everybody's Got ML, Tell Me What Else You Have: Practitioners' Perception of ML-Based Security Tools and Explanations," in *2023 IEEE Symposium on Security and Privacy (SP)*, 2023, pp. 2068–2085.

[9] S. Tariq, M. B. Chhetri, S. Nepal, and C. Paris, "A2c: A modular multi-stage collaborative decision framework for human-ai teams," 2024. [Online]. Available: https://arxiv.org/abs/2401.14432

[10] F. Binbeshr, M. Imam, M. Ghaleb, M. Hamdan, M. A. Rahim, and M. Hammoudeh, "The rise of cognitive socs: A systematic literature review on ai approaches," *IEEE Open Journal of the Computer Society*, vol. 6, pp. 360–379, 2025.

[11] S. C. Sundaramurthy, A. G. Bardas, J. Case, X. Ou, M. Wesch, J. McHugh, and S. R. Rajagopalan, "A human capital model for mitigating security analyst burnout," in *Proceedings of the Eleventh USENIX Conference on Usable Privacy and Security*, ser. SOUPS '15. USA: USENIX Association, 2015, p. 347–359.

[12] R. Wash, E. Rader, and C. Fennell, "Can people self-report security accurately?: Agreement between self-report and behavioral measures," 05 2017, pp. 2228–2232.

[13] A. Reeves and D. Ashenden, "Understanding decision making in security operations centres: building the case for cyber deception technology," *Frontiers in Psychology*, vol. 14, p. 1165705, 2023. [Online]. Available: https://doi.org/10.3389/fpsyg.2023.1165705

[14] K. R. Jones, D. A. Brucker-Hahn, B. Fidler, and A. G. Bardas, "Work-From-Home and COVID-19: Trajectories of endpoint security management in a security operations center," in *32nd USENIX Security Symposium (USENIX Security 23)*. Anaheim, CA: USENIX Association, Aug. 2023, pp. 2293–2310. [Online]. Available: https://www.usenix.org/conference/usenixsecurity23/presentation/jones

[15] J. Lave and E. Wenger, *Situated Learning: Legitimate Peripheral Participation*, ser. Learning in Doing: Social, Cognitive and Computational Perspectives. Cambridge University Press, 1991.

[16] A. Tuladhar, D. Lende, J. Ligatti, and X. Ou, "An analysis of the role of situated learning in starting a security culture in a software company," in *Seventeenth Symposium on Usable Privacy and Security (SOUPS 2021)*. USENIX Association, Aug. 2021, pp. 617–632. [Online]. Available: https://www.usenix.org/conference/soups2021/presentation/tuladhar

[17] A. Greig, K. Renaud, and S. Flowerday, "An ethnographic study to assess the enactment of information security culture in a retail store," in *2015 World Congress on Internet Security (WorldCIS)*, 2015, pp. 61–66.

[18] S. Srinivas, B. Kirk, J. Zendejas, M. Espino, M. Boskovich, A. Bari, K. Dajani, and N. Alzahrani, "Ai-augmented soc: A survey of llms and agents for security automation," *Journal of Cybersecurity and Privacy*, vol. 5, no. 4, 2025. [Online]. Available: https://www.mdpi.com/2624-800X/5/4/95

[19] R. Singh, S. Tariq, F. Jalalvand, M. B. Chhetri, S. Nepal, C. Paris, and M. Lochner, "Llms in the soc: An empirical study of human-ai collaboration in security operations centres," 2025. [Online]. Available: https://arxiv.org/abs/2508.18947

[20] S. Freitas, J. Kalajdjieski, A. Gharib, and R. McCann, "Ai-driven guided response for security operation centers with microsoft copilot for security," 2024. [Online]. Available: https://arxiv.org/abs/2407.09017

[21] N. Rastogi, S. Pant, D. Dhanuka, A. Saxena, and P. Mairal, "Too much to trust? measuring the security and cognitive impacts of explainability in ai-driven socs," 2025. [Online]. Available: https://arxiv.org/abs/2503.02065

[22] M. Nyre-Yu, E. Morris, M. R. Smith, B. Moss, and C. Smutz, "Explainable ai in cybersecurity operations: Lessons learned from xai tool deployment," *Proceedings 2022 Symposium on Usable Security*, 2022. [Online]. Available: https://api.semanticscholar.org/CorpusID:253156531

[23] J. Lave and E. Wenger, *Situated learning: Legitimate peripheral participation*. Cambridge university press, 1991.

[24] L. A. Suchman, *Plans and situated actions: An inquiry into the idea of human-machine communication*. University of California, Berkeley, 1984.

[25] Y. Engeström and A. Sannino, "Studies of expansive learning: Foundations, findings and future challenges," *Introduction to Vygotsky*, pp. 100–146, 2017.

[26] J. Van Maanen, *Tales of the field: On writing ethnography*. University of Chicago Press, 2011.

[27] P. Ehn, "Work-oriented design of computer artifacts," Ph.D. dissertation, Arbetslivscentrum, 1988.

[28] K. Anderson, M. Foster, C. Freeman, and I. Scott, "A multifaceted intervention to reduce inappropriate polypharmacy in primary care: research co-creation opportunities in a pilot study," *The Medical journal of Australia*, vol. 204, pp. S41–S, 04 2016.

[29] T. Greenhalgh, C. JACKSON, S. Shaw, and T. Janamian, "Achieving research impact through co-creation in community-based health services: Literature review and case study: Achieving research impact through co-creation," *The Milbank Quarterly*, vol. 94, pp. 392–429, 06 2016.

[30] Nmap Project. (2025) Nmap: Network mapper. Accessed: 2025-12-27. [Online]. Available: https://nmap.org

[31] S. Hastings, C. Bolger, P. Shumway, and T. Moore, "Transforming raw authentication logs into interpretable events," in *Workshop on SOC Operations and Construction (WOSOC 2024). https://dx. doi. org/10.14722/wosoc*, 2024.

[32] D. Lende and G. Downey, Eds., *The encultured brain: an introduction to neuroanthropology*. MIT Press, Jan. 2012.

[33] S. D. Blum, "Why don't anthropologists care about learning (or education or school)? an immodest proposal for an integrative anthropology of learning whose time has finally come," *American anthropologist*, vol. 121, no. 3, pp. 641–654, 2019.

[34] M. Cole and S. SCRIBNER, "Vygotsky, lev s.(1978): Mind in society. the development of higher psychological processes," 1978.

[35] E. Hutchins, *Cognition in the Wild*. MIT press, 1995.

[36] L. Argote and P. Ingram, "Knowledge transfer: A basis for competitive advantage in firms," *Organizational behavior and human decision processes*, vol. 82, no. 1, pp. 150–169, 2000.

[37] M. Polanyi, "The tacit dimension," in *Knowledge in organisations*. Routledge, 2009, pp. 135–146.

[38] D. A. Norman, "The 'problem'with automation: inappropriate feedback and interaction, not 'over-automation'," *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, vol. 327, no. 1241, pp. 585–593, 1990.

[39] B. Shneiderman, "Human-centered artificial intelligence: Reliable, safe & trustworthy," *International Journal of Human–Computer Interaction*, vol. 36, no. 6, pp. 495–504, 2020.

[40] R. Parasuraman, T. B. Sheridan, and C. D. Wickens, "A model for types and levels of human interaction with automation," *IEEE Transactions on systems, man, and cybernetics-Part A: Systems and Humans*, vol. 30, no. 3, pp. 286–297, 2000.