TALLINN UNIVERSITY OF TECHNOLOGY
School of Information Technologies

Elguj Yusifbayli

Use of AI-based security tool assistance for improving the technical analysis of incidents in Security Operations Centre

Master's thesis

Supervisor:
Professor
Risto Vaarandi

Tallinn 2025

TALLINNA TEHNIKAÜLIKOOL
Infotehnoloogia teaduskond

Elguj Yusifbayli

TEHISINTELLEKTIL PÕHINEVATE TURVATÖÖRIISTADE KASUTAMINE INTSIDENTIDE TEHNILISE ANALÜÜSI TÄIUSTAMISEKS TURBEOPERATSIOONIDE KESKUSES

Magistritöö

Juhendaja:
Professor
Risto Vaarandi

Tallinn 2025

Author’s declaration of originality

I hereby certify that I am the sole author of this thesis. All the used materials, references to the literature and the work of others have been referred to. This thesis has not been presented for examination anywhere else.

Author: Elguj Yusifbayli

18.05.2025

Abstract

This study investigates the effectiveness of Microsoft Security Copilot, an AI-assisted security tool, in supporting SOC analysts during technical incident investigations. It compares traditional manual analysis with MSC-assisted workflows in a simulated SOC environment, using realistic alerts and telemetry across enterprise scenarios that include scripted threats, malware, and known attack patterns.

The analysts participated in two rounds of investigation using a custom 5W1H-based incident analysis template: one round was conducted manually, while the other utilised prompts within MSC, with analysts making decisions based on multiple different prompts and MSC-supported console results. Key performance metrics such as investigation time and analyst confidence were evaluated, alongside qualitative feedback and observations.

The findings indicate that MSC significantly reduces investigation time, allowing analysts to focus more on decision-making rather than manual data parsing. The 5W1H approach improved clarity and completeness, particularly for less experienced analysts. Feedback from the analysts highlights MSC's potential to streamline investigations and enhance efficiency; however, human oversight remains essential due to current limitations in contextual awareness.

The quality of the prompts and AI-human collaboration were identified as crucial factors influencing the effectiveness of MSC. While some participants encountered challenges with prompt design and response interpretation, most agreed that AI-assisted security tools like MSC, when used thoughtfully, can provide meaningful value in modern SOC environments.

Keywords: Microsoft Security Copilot, Security Operations Centre, AI in SOC, Cybersecurity, AI-assisted security tools, incident investigation

Annotatsioon

Käesolev uurimistöö käsitleb Microsoft Security Copilot tõhusust, mis on tehisintellektil põhinev turvavahend, toetades SOC analüütikuid tehniliste intsidentide uurimisel. Töö võrdleb traditsioonilist käsitsi tehtavat analüüsi MSC abil toetatud töövoogudega simuleeritud SOC-keskkonnas, kasutades realistlikke hoiatusi ja telemeetriat ettevõtte stsenaariumites, mis hõlmavad skriptitud ohte, pahavara ja teadaolevaid ründeviise.

Analüütikud osalesid kahes uurimisvoorus, kasutades kohandatud 5W1H-põhist intsidentide analüüsimise mall: üks voor viidi läbi käsitsi ning teine MSC abil, kus otsuseid tehti erinevate küsimusmallide ja MSC konsoolitulemuste põhjal. Hindamiskriteeriumideks olid uurimisele kulunud aeg ja analüütikute kindlustunne, mida täiendati kvalitatiivse tagasiside ja tähelepanekutega.

Tulemused näitavad, et MSC vähendab märkimisväärselt uurimisaega, võimaldades analüütikutel keskenduda otsuste tegemisele, mitte käsitsi andmete läbitöötamisele. 5W1H-lähtekoht parandas selgust ja täielikkust, eriti vähem kogenud analüütikute puhul. Osalejate tagasiside toob esile MSC potentsiaali tõhustada uurimisprotsesse ja parandada töö efektiivsust; samas jääb inimlik järelevalve endiselt oluliseks piiratud kontekstitundlikkuse tõttu.

Küsimusmallide kvaliteet ja inimese ning tehisintellekti koostöö osutusid võtmeteguriteks MSC tõhususe määramisel. Kuigi mõned osalejad kogesid raskusi küsimuste koostamise ja vastuste tõlgendamisega, nõustus enamik, et sellised AI-toega turvavahendid nagu MSC võivad läbimõeldult kasutatuna pakkuda väärtuslikku tuge kaasaegsetes SOC-keskkondades.

Abbreviations and terms:

SOC - Security Operations Centre

AI - Artificial Intelligence

IEEE - Institute of Electrical and Electronics Engineers

CGR - Microsoft Copilot for Security Guided Response

XDR - Extended detection and response

EDR - Endpoint detection and response

MS- Microsoft

MSC-Microsoft Copilot for Security and Microsoft Security Copilot are interchangeable.

SIEM - Security Information and Event Management

LLM - Large Language Model

NL -Natural Language

NLP -Natural Language Processing

KQL - Kusto Query Language

NL to KQL - Natural language to KQL for advanced hunting

MITRE ATT&CK - MITRE Adversarial Tactics, Techniques, and Common Knowledge

SCU - Security Compute Unit

Command and Control – C2

Table of Contents

List of Figures

1 Introduction

1.1 Introduction

With the rapid development and expansion of digital technologies, crucial in numerous domains, security threats and cybersecurity vulnerabilities are continually evolving. In response to these advancements, techniques for mitigating potential threats and applying appropriate intervention methods are also being updated in parallel with technological progress as cybersecurity concerns intensify. Alongside these advancements, artificial intelligence (AI) is becoming increasingly prevalent in many fields [1-3].

Such digital development challenges organisations, making it crucial for companies to tackle and defend against digital threats for the security and sustainability of their services. Over the last decade, organisations have increasingly established their own Security Operations Centres (SOCs) or outsourced SOC services to counter malicious actors and digital threats, thereby strengthening their timely security measures [4-5].

Establishing a SOC also brings additional challenges, one of which is the crucial task of detecting malicious actions and responding to incidents promptly and appropriately. However, despite the prevalence of technical solutions in the SOC space, several challenges still need to be addressed. Security automation issues and a shortage of specialised experts and technical skills often lead to delayed or unnoticed incident responses [6-7]. This is primarily due to the time-consuming nature of analysing cyber incidents, identifying security threats, and producing accurate reports.

Moreover, many SOC teams within organisations are small and often lack experience and technical resources, resulting in ongoing gaps in the SOC environment. These challenges can be categorised into three main areas, which are further elaborated in the literature review:

- Issues with security automation
- A shortage of specialised experts and technical skills
- Analysis gaps due to fatigue from 24/7 shift work and burnout

In this context, this thesis aims to assess whether leveraging generative AI-powered solutions can enhance the agility of SOC teams, reduce operational risks, and improve the technical analysis of security incidents. To achieve this, Microsoft Security Copilot (MSC) is evaluated to determine its effectiveness and impact on the SOC team's performance during incident analysis and response.

1.2 Motivation

SOCs are critical in safeguarding organisations against cyber threats. Despite the availability of advanced tools and services, many SOCs struggle to utilise their capabilities fully due to several barriers. These include a lack of technical expertise, limited resources, and operational complexities. These challenges often result in delayed responses and inconsistent outcomes in defending against evolving cyber threats.

The ongoing challenges have resulted in an increasing academic emphasis on developing more efficient, scalable, and human-enhancing solutions. This thesis aims to experimentally assess AI-based tools, specifically MSC, to support SOC analysts and enhance investigation outcomes. By leveraging AI assistance, the research aims to determine whether it can reduce the time to insight, increase accuracy, and enhance decision-making processes during incident investigations. Addressing these challenges is crucial for improving the effectiveness and efficiency of SOCs, ultimately contributing to better cybersecurity practices.

1.3 Research Goal

The primary goal of this thesis is to empirically assess the impact of MSC on SOC analyst performance during incident investigations. Specifically, the research aims to:

- Assess the functionality and effectiveness of specific features of MSC in supporting incident analysis.
- Evaluate differences in investigation time, classification accuracy, and confidence levels between manual and MSC-assisted investigations.
- Identify areas where Copilot enhances investigation quality and where it may introduce limitations.

- Gather qualitative feedback from SOC analysts regarding the usability, trustworthiness, and practical value of MSC through hands-on experiments.

By achieving these objectives, this research aims to provide valuable insights into the role of AI in modern SOC workflows. It seeks to conduct an evidence-based assessment of the effectiveness of MSC as an AI-assisted security tool.

Furthermore, the study will analyse how MSC supports human decision-making and enhances the quality and efficiency of the incident analysis process.

1.4 Research Scope

The study replicates realistic SOC conditions using a researcher's production tenant within a controlled experimental environment. Participants will analyse a variety of security incidents both manually and with the assistance of MSC. These incidents will cover common attack types such as web exploitation, malware, and lateral movement. The investigation process will be structured using the 5W1H framework and impact and classification assessments to ensure consistency and comparability.

This research aims to evaluate the effectiveness of MSC during the incident investigation phase within SOC environments. Instead of covering full incident response or remediation, this focuses on the technical analysis of security incidents, which includes identifying root causes, establishing timelines, determining affected assets, and classification.

This research exclusively examines the technical aspects of incident investigation without addressing AI's policy, legal, or business implications in cybersecurity. It assumes that analysts operate in a controlled, simulated SOC environment with full access to all necessary tools.

The primary purpose is to determine how AI-based security tools, such as MSC, assist in the technical analysis of security incidents and support cyber threat investigation within SOC environments.

1.5 Novelty

This thesis presents an empirical study on the impact of MSC within the SOC. While Microsoft has introduced Security Copilot through technical industry-based materials, there is limited academic research evaluating its practical effectiveness in SOC environments. This study aims to contribute to this emerging field through real analyst feedback and structured incident analysis.

Core Contribution Areas:

- The research addresses significant challenges faced by SOCs, including alert fatigue and skill shortages, and evaluates how Security Copilot helps mitigate these issues.
- This research contributes to the academic community by providing additional empirical data on the impact of MSC on SOC performance.
- It compares investigations conducted manually with those assisted by Security Copilot, measuring improvements in investigation time and accuracy.
- The study explores how Security Copilot enhances human decision-making and emphasises the collaboration between AI and human analysts.
- Additionally, it identifies areas where Copilot improves the quality of investigations and where it may have limitations.
- Furthermore, the study includes feedback from SOC analysts, providing insights into the tool's usability and practical value in SOC workflows.

2 Research Problem

2.1 Research Problem

As cybersecurity threats grow more complex and persistent, organisations must adopt proactive and adaptive strategies to detect, analyse, and respond to malicious activities

effectively. While preventive measures are essential, attackers often operate over extended periods, making early detection and comprehensive incident analysis critical components of a robust defence strategy.

In response, many organisations have established SOCs with centralised teams and systems that monitor, investigate, and mitigate security threats. However, SOC operations present ongoing challenges, including skill shortages, analyst fatigue, limited automation, and difficulties maintaining investigation speed and accuracy. These issues are widely recognised in academic literature and industry reports [4,8-9].

The rise of artificial intelligence has prompted the cybersecurity field to explore AI-based solutions for enhancing visibility, optimising investigation workflows, and supporting overwhelmed SOC teams. One such security tool is MSC, a generative AI assistant designed to assist with alert summarisation, contextual interpretation, and natural language-based querying. It supports security professionals in various use cases, including incident response, threat hunting, intelligence gathering, and posture management [10-11].

Although MSC shows considerable promise in augmenting SOC workflows, there is a noticeable lack of academic research evaluating its real-world impact on technical incident investigations [10].

This research aims to address that gap by conducting an empirical study on how MSC affects analyst performance during the incident investigation phase.

2.2 Research Questions

The research emphasised the importance of formulating questions to define and narrow down the specific area of focus:

Key Question: How does Microsoft Security Copilot enhance the performance of SOC analysts during the technical investigation of security incidents?

Sub Questions:

- What are the most common challenges SOC teams face during incident investigations?

- Which features of Microsoft Security Copilot are designed to address these challenges, and how do they impact investigation quality and speed?
- To what extent does Microsoft Security Copilot help resolve the limitations found in traditional SOC workflows?
- What are the perceived benefits and limitations of Microsoft Security Copilot, based on hands-on experience from SOC analysts?

3 Literature Review

3.1 SOC

Security attacks are becoming more complex and sophisticated. To address these evolving threats, investing in preventive measures and developing intelligent and integrated monitoring capabilities is essential as part of an incident response program [4,8-9].

However, a breach does not necessarily result in immediate negative consequences for the business, as attackers often require time to achieve their objectives beyond gaining unauthorised access. Detecting and preventing such activities is essential. Therefore, it is essential for organisations to establish a SOC [4,8-9].

Organisations must establish a SOC and team to respond rapidly to cybersecurity incidents [4,8-9]. However, establishing a SOC introduces additional challenges, as mentioned in the introduction. The literature review will help clarify these remaining issues and provide further insights into addressing them.

3.1.1 SOC Models

According to the findings, multiple SOC operational models are presented. Establishing a SOC should be tailored to the specific needs, conditions, team capabilities, regulations and different factors of the organisation [4,12].

However, it is crucial for companies to precisely evaluate these factors to determine the most suitable SOC operational model for their unique circumstances. This assessment

ensures that the chosen model aligns with the organisation's security objectives and operational requirements, whether the SOC will be in-house, outsourced, co-managed, hybrid or different models [4,12].

3.1.2 Components of SOC

While SOC-related topics encompass a wide range of issues, it is essential to focus on the critical components of a SOC. According to data from ManageEngine and other referenced sources, the fundamental elements of a SOC can be categorised into three core areas: people, processes, and technology [4,13-16].

- People - This component involves SOC team members responsible for detecting, investigating, and responding to security events. Their roles include containing threats, managing crises, and ensuring effective communication during incidents, which are vital for the SOC's success [4-5,13-16].

- Process – This area includes the documentation, procedures, playbooks, and workflows that guide the SOC's operations. Well-defined processes ensure consistency and efficiency in incident handling, enabling smooth and effective operations [4-5,13-16].

- Technology – This component encompasses the tools used by the SOC to monitor events and gather threat intelligence, such as Security Information and Event Management systems, Intrusion Detection and Prevention Systems, Endpoint Detection and Response tools, and threat intelligence platforms. These technologies are essential for detecting and analysing security threats, allowing for timely responses [4-5,13-16].

These components represent the essential elements of a SOC, though they do not encompass every aspect. While people, processes, and technology are vital components, other factors, such as governance and compliance, also play a significant role [4,15].

3.1.3 SOC and Challenges

As noted in section 3.1.2, three key factors are essential for establishing a SOC. All our findings, which stated their research, surveys, and various analyses, indicate that SOCs are primarily influenced by these three factors. All studies reviewed have addressed the

same issues, and these challenges continue to persist. These investigations highlight the ongoing need for research and the application of new ideas in this field. Some of the factors outlined in the documents indicate the following challenges [5-7,13-20].

- Staffing Shortages: Finding and retaining qualified cybersecurity professionals is a significant challenge.
- Alert Fatigue: SOC analysts often deal with many alerts, many of which are false positives.
- Slow Response Times: Incident response time is essential to minimise damage.
- Massive Data Management: Analysing large volumes of security data is complex and resource intensive.
- Limited Visibility: Extensive network monitoring is required to avoid blind spots.
- Multiple Security Tools: Security tools that are complex and numerous need to be simplified and optimised.
- Automation: Efficient use of automation tools reduces manual workloads and improves response times.

Numerous factors can be added to this list of challenges, including budget constraints and the need to adapt to and integrate new security technologies [5-7,13-20]. These challenges emphasise the necessity for SOCs to continually improve their defences against cyber threats.

3.2 SOC and Artificial Intelligence

When looking into the challenges of SOCs, it becomes clear that many aspects depend on human factors. Organisations face new threats and challenges continuously. These situations increasingly demand the analysis, correlation, and automation of huge amounts of rules, among other tasks that rely heavily on human skills and expertise [6,17,21].

The identified challenges have led to further research exploring different approaches and solutions to enhance the detection and response capabilities of SOCs. An important and

growing trend in this field is the use of artificial intelligence power, which aims to effectively close gaps against modern cybersecurity threats [18-19].

Based on the latest reports from Vectra and the SANS Institute, security experts are positive about AI-powered tools and are increasing their investments in this area. The 2024 surveys indicate that these tools are already replacing legacy incident detection and response capabilities.

Experts believe AI-powered tools will continue positively impacting the security industry and help SOC teams improve visibility and detect and respond capability [1,7,22].

3.2.1 Microsoft Copilot for Security

The evolution of AI has led to the development of AI-powered tools, such as Microsoft Copilot for Security [10-11]. This tool is currently assisting SOC teams in several ways:

- Incident Summarisation: By leveraging generative AI, Copilot for Security can quickly transform complex security alerts into concise, actionable summaries. This improves organisational communication and enables faster response times, leading to more streamlined decision-making [10-11].

- Impact Analysis: The tool utilises AI-driven analytics to evaluate the potential impact of security incidents. It provides insights into affected systems and data, helping teams prioritise their response efforts effectively. This information is crucial for preventing large-scale attacks, such as ransomware campaigns [10-11].

- Reverse Engineering of Scripts: Copilot allows analysts to understand attacker actions by analysing complex command-line scripts and translating them into clear, natural language explanations. This process efficiently extracts and connects indicators found in the script to their respective entities within the environment [10-11].

- Guided Response: The tool offers actionable, step-by-step guidance for incident response, covering triage, investigation, containment, and remediation. It includes relevant deep links to recommended actions, facilitating quicker responses [10-11].

3.2.2 Trend Companion

Trend Micro's AI Companion is another AI-powered security tool to enhance cybersecurity operations to assist engineers.

Trend Companion helps security teams by providing real-time risk assessments, facilitating threat detection and response, and proactive risk mitigation [23].

3.2.3 Charlotte AI

CrowdStrike's Charlotte AI is an AI-powered assistant designed to enhance cybersecurity operations within the CrowdStrike Falcon platform.

Charlotte AI assists security analysts by automating the detection and triage process, speeding up threat investigations, and simplifying complex security queries with a user-friendly natural language interface [24].

3.2.4 Gemini AI

Google's Gemini AI is designed for various purposes and comes in multiple versions tailored for different tasks, similar to Microsoft Copilot. Gemini AI is integrated into Google Security Operations as a generative assistant to support cybersecurity analysts. It facilitates natural language searches, rule generation, case summarisation, and playbook assistance.

While it offers AI-enhanced capabilities such as automated detection logic and support for threat intelligence, Gemini also provides a broad range of functionalities to strengthen cybersecurity processes [25].

3.3 SOC and Prompt Engineering

As AI security tools become more advanced and integrated into security workflows, the importance of analysts interacting with them effectively continues to grow. This underscores the relevance of prompt engineering as a key focus in this study.

3.3.1 The Importance of Prompt Engineering Skills in SOC

The rise of generative AI has introduced a new skill set known as prompt engineering, which requires analysts to effectively communicate with tools based on large language

models (LLMs). For cybersecurity professionals, mastering prompt engineering represents a significant shift in how threats are identified, analysed, and defended against in modern security environments [26].

Creating structured and purposeful prompts is essential for obtaining relevant and actionable insights. As AI-powered tools become increasingly integrated into SOC workflows, cybersecurity engineers must develop proficiency in prompt engineering to maximise the effectiveness of these solutions and ensure accurate, high-quality investigative outcomes.

Prompt engineering involves crafting precise and context-aware inputs that guide AI systems to generate accurate and useful responses. Research shows that effective prompts significantly enhance the relevance of AI outputs, while poorly designed prompts can lead to vague or misleading results [27].

3.4 Similar Research and Papers

Due to the emerging nature of the field, there is currently a limited body of academic research specifically focused on MSC and similar AI-assisted security tools.

Most available literature originates from industry blogs, technical whitepapers, or product documentation rather than peer-reviewed academic sources.

One notable academic contribution is the paper titled “AI-Driven Guided Response for Security Operation Centres with Microsoft Copilot for Security”, which presents the design and deployment of Copilot Guided Response (CGR), a machine learning-driven architecture integrated into Microsoft Defender XDR and deployed across enterprise environments. CGR assists SOC analysts in three key areas: investigation, triaging, and remediation [10].

In addition, the paper “A User-Centred Security Evaluation of Copilot” (2023) investigates human interaction with GitHub Copilot, focusing on usability, trust, and analyst reliance. While not specific to MSC, the findings highlight those users found AI assistance beneficial but remained cautious about fully trusting its outputs, emphasising the continued importance of human oversight in security tasks [28].

3.5 Security Tool Selection for the Experiment

The integration of AI into cybersecurity is evolving rapidly and has become a highly research-driven area. While an increasing number of AI-assisted security tools are emerging, only a limited selection currently offers meaningful natural language interaction and AI-powered capabilities tailored to SOC environments.

During the research phase, multiple AI-based security tools were evaluated. Some were tested hands-on, while others were assessed through documentation reviews or limited trial demonstrations.

Certain tools offered narrow assistive functionalities, such as alert triage or summarisation, while others aimed to provide broader, end-to-end support for SOC environments.

The selection process followed a defined set of criteria to identify tools that could effectively support SOC workflows and were accessible for in-depth assessment. These criteria included:

- Support for natural language interaction
- Ease of integration with existing SOC platforms
- Usability and analyst experience
- Functional coverage across the incident lifecycle
- Accessibility for long-term research
- Examination or recognition within academic or industry contexts

As detailed in Section 3.2, following an initial evaluation of a broader set of tools, four AI-assisted security solutions were selected for detailed comparison based on these criteria (**Table 1**).

In addition to the evaluation criteria, the methodology and scope of data collection played a critical role in the comparative assessment. When applied to incident analysis and correlation tasks, LLMs depend heavily on access to structured, high-quality data.

Many AI-assisted security tools are tightly coupled with their native ecosystems, meaning their effectiveness is often limited to the data generated within those platforms. From this perspective, unified SOC solutions that support a wide range of data sources, such as Microsoft's integrated Defender XDR and Sentinel suite, offer a distinct advantage.

Broader data visibility not only enhances the accuracy of AI-assisted analysis but also improves the overall operational effectiveness of SOC teams (**Table 1**).

Evaluation Criterion	Microsoft Security Copilot	Gemini in Google SecOps	Trend Micro Companion	CrowdStrike Charlotte AI
AI Assistance Capabilities	Provides advanced NLP for queries and interaction, automated event summarisation, guided incident response, and playbook suggestions	Offers robust NLP functionalities, automated summarisation, rule generation, and AI-assisted playbooks	Features foundational NLP capabilities tested during the basic release phase. Advanced triage and IOC analysis were not available during the evaluation	Incorporates NLP focusing on alert triage; tightly integrated with the CrowdStrike Falcon ecosystem
Raw Log Data Ingestion (Windows/Linux)	Supports comprehensive raw log ingestion via Microsoft Sentinel and Defender XDR	Supports broad raw log ingestion using Chronicle forwarders and connectors	It does not natively ingest raw log data from diverse third-party sources; it focuses on trend-managed telemetry	It relies on Falcon agent telemetry but does not support native raw OS-level log collection
SIEM and Log Correlation Capabilities	Provides full SIEM functionality and advanced correlation through Microsoft Sentinel	Delivers complete SIEM and correlation capabilities via the Chronicle platform	Limited to internal XDR data correlation; not designed as a general-purpose SIEM	Partial log correlation is limited to internal Falcon data; it is not a standalone SIEM
Endpoint Telemetry Collection and Analysis	Native, in-depth telemetry via Defender for Endpoint and Defender XDR	It requires integration with third-party EDR solutions, but lacks native endpoint telemetry collection	Collects endpoint telemetry via Trend agents with visibility into threats	Strong endpoint telemetry via Falcon EDR is a core capability of the platform
Accessibility for Research Purposes	Thoroughly tested using a private licensed tenant; no enterprise barriers	Enterprise-level access is required; it is not publicly available for academic research	Tested during limited release. The basic version is available, but advanced features are inaccessible during testing	Not accessible without an enterprise contract or sales agreement

Evaluation Criterion	Microsoft Security Copilot	Gemini in Google SecOps	Trend Micro Companion	CrowdStrike Charlotte AI
Representation in Academic Literature	Referenced in both peer-reviewed research and industry publications	Gemini is cited in AI research; no academic studies on use in Google SecOps	No academic presence found at the time of review	No academic presence found at the time of review
Scope of Product Evaluation in the Present Study	Fully evaluated; no significant access restrictions during testing	It has not been tested; access is unavailable within the research timeframe	The basic version was evaluated; advanced features were not accessible during the study period	It has not been tested; access is unavailable within the research timeframe

Table 1. Comparison of AI-Assisted Security Tools for SOC Operations [23-25,29].

Following the comparative evaluation presented in Table 1, it is important to contextualise the assessment of each tool.

The capabilities of MSC and Trend Micro Companion were evaluated through hands-on testing. In contrast, information about Google SecOps's Gemini and CrowdStrike's Charlotte AI was analysed using publicly available vendor documentation, as these platforms required enterprise contracts or sales engagements for access.

While each of the reviewed tools offers valuable features for SOC environments, many are tightly bound to proprietary commercial ecosystems, limiting their accessibility for independent academic research. Several require enterprise-level licensing and often lack the integration flexibility or experimental usability provided by MSC.

MSC was selected as the primary focus of this study due to its unique combination of accessibility, advanced AI-driven features, and seamless integration with Microsoft's security ecosystem. Its consumption-based pricing model enabled academic experimentation without the constraints of long-term licensing.

Furthermore, its native interoperability with Microsoft Defender XDR and Microsoft Sentinel allowed for smooth deployment and practical application in a simulated SOC environment.

These characteristics made MSC the most suitable candidate for evaluating the real-world impact of AI-assisted analysis in incident investigation workflows.

3.6 5W1H Approach

The 5W1H framework (What, Who, When, Where, Why, and How) is a widely adopted method for structured information gathering and problem-solving. It is applied across multiple domains, including cyber threat analysis, forensic investigation, and digital evidence processing.

This framework enables a systematic approach to questioning, which enhances situational awareness, supports time-sensitive decision-making, and improves traceability in complex investigations.

5W1H has been used in cybersecurity and forensics to classify incident elements, prioritise actions, and reduce false positives in time-critical scenarios. Its use in digital forensics has also contributed to developing models that reduce analyst workload and improve the accuracy of evidence processing [30-32].

In this research, the 5W1H framework serves as the foundation for a structured incident analysis template, standardising how incidents are investigated manually and through the assistance of MSC.

4 Methodology

This study uses a hybrid methodology that combines qualitative and quantitative approaches to evaluate the effectiveness of MSC during incident investigations in a SOC environment.

The methodology is based on an empirical experiment that compares traditional manual investigations with those enhanced by Security Copilot, all conducted within a controlled laboratory environment.

The research evaluates how MSC influences investigation outcomes, including investigation time, classification accuracy, analyst confidence, and perceived usability, within the technical analysis phase of security incidents.

4.1 Experimental Environment

The experimental environment was established using a Microsoft Azure-based cloud tenant and a small-scale on-premises virtual infrastructure. A detailed overview of the lab design, tooling, and infrastructure components is provided in Chapter 5.

Simulated attack scenarios were executed through red team activity, which included techniques such as malware execution, lateral movement, and script-based exploitation. These activities were detected by EDR sensors (e.g., Microsoft Defender for Endpoint), while system-level telemetry was collected using Sysmon and Azure Log Analytics agents.

All logs and alerts were ingested into Microsoft Defender XDR and Azure Sentinel, which served as the central SOC-level monitoring and investigation platforms throughout the study.

4.2 Study Participants

The study involved 7 to 15 cybersecurity specialists from various sectors, including finance, government, IT, and education. While some participants contributed by generating red team activity and assisting with the SOC environment setup, a total of nine analysts took part in the incident investigation phase.

All participants received preparation materials and guidance about the investigation structure and toolset.

Each analyst completed two phases:

- Manual investigations conducted with standard SOC tools
- AI-assisted investigations using MSC

All participants provided informed consent. Investigations were anonymised using unique codes, and participants could choose to be named in the thesis contributions section. No actual organisational or sensitive production data was utilised.

4.3 Investigation Workflow

A consistent investigative workflow was followed using a custom Excel-based incident analysis template, structured around the 5W1H framework (Who, What, When, Where, Why, How).

This methodology ensured uniformity across manual and Copilot-assisted analyses.

All investigations included the following components:

- Incident classification (e.g., True, False Positive)
- Impact level assessment (e.g., High, Medium)
- Confidence scoring on a scale from 1 to 5
- Summary of findings

The incident analysis template is explained in detail in Chapter 5.

4.4 Data Collection

4.4.1 Initial Assessment

Before the experiment, all analysts completed an initial assessment to:

- Confirm informed consent
- Record participant roles (e.g., investigation or red teaming)
- Outline cybersecurity experience and industry background
- Identify familiarity with security tools (e.g., SIEM, EDR)

Determine experience with AI-assisted security tools, such as MSC.

This information enhances participant profiling and facilitates analysis of feedback and outcomes. A summary of the survey structure is provided in Chapter 5.

4.4.2 Metrics Collection

Metrics collected during both phases included:

- Investigation Time
- Classification Outcome (e.g., True/False Positive)
- Impact Level
- Confidence Score (on a scale of 1 to 5)

4.4.3 Qualitative Feedback

After the Copilot-assisted investigation phase, analysts provided qualitative feedback on the following aspects:

- User-friendliness of the interface and interaction process
- Confidence in the suggestions and guidance provided by MSC
- Perceived improvements in efficiency during the investigation workflow
- Insights regarding the clarity or complexity of the investigation process

4.5 Assessment Strategy

To ensure balanced exposure, incident scenarios were rotated among analysts. Each incident was examined both manually and with Copilot by different participants, facilitating a fair comparison.

The analysis concentrated on:

- Time taken to analyse each investigation
- Accuracy of incident classification (compared to predefined ground truth)
- Analyst-reported confidence
- Report completeness and any use of MITRE ATT&CK mappings (if applicable)

4.6 Validation

All incidents in the experiment were predefined with known ground truth, allowing for the objective validation of classification outcomes.

The lab environment and investigation structure were standardised across all participants, with consistent templates and tools, and equal briefing to minimise skill variation.

Real-time feedback was collected after each phase to avoid recall bias. The multi-dimensional validation process combined structured data with analyst experiences, focusing on accuracy and usability.

4.7 Ethical Considerations and Data Handling

All experiments were conducted within a controlled research environment using the researcher's private Microsoft Cloud tenant and self-contained virtual SOC infrastructure. No real organisational or production-sensitive data was utilised during the study.

Participation in the study was voluntary, with all participants providing informed consent. They were briefed on the experimental process, data collection methods, and their right to anonymity. Each analyst’s investigation data was anonymised using unique participant codes to ensure confidentiality.

Participants could be named in the thesis contribution section or remain anonymous. No results were linked to personal identifiers, and no sensitive operational data was collected, stored, or shared.

These measures ensured adherence to ethical standards concerning privacy, consent, and responsible data handling, in line with accepted academic and industry practices.

5 Experimental Environment and Process

This chapter provides a comprehensive overview of the experimental environment used for conducting the MSC evaluation. It includes the setup of the infrastructure, the

implementation of the 5W1H methodology for structured analysis, and the flow of data and incident assignment logic among participants.

5.1 Infrastructure Setup

The experimental environment was built using a hybrid infrastructure combining Microsoft Azure (cloud platform) and an on-premises virtual machine setup (Fig. 1).

- The cloud platform hosted various simulated services and endpoints with telemetry agents installed, enabling alert generation and logging
- The on-premises virtual lab simulated internal enterprise services, user behaviour, and attacker activity
- A red team across both infrastructures executes simulated attacks

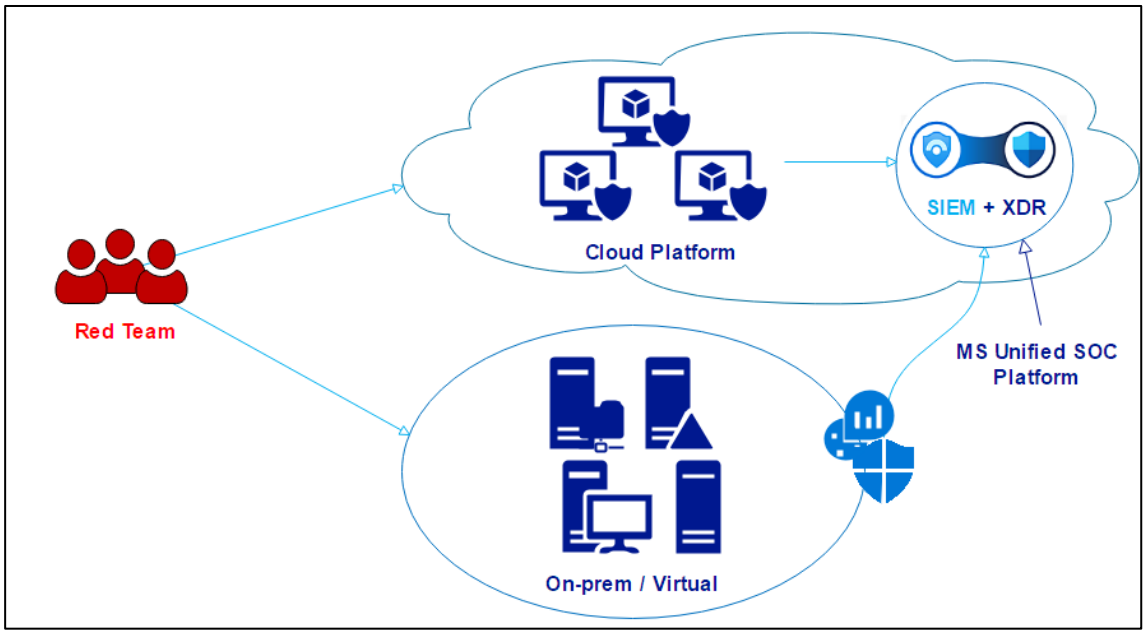

Fig. 1 Attack Simulation and SOC Design

Note: The "MS Unified SOC Platform" refers to Microsoft's converged security architecture that unifies SIEM and XDR capabilities through a single platform. It integrates Microsoft Sentinel (SIEM), Microsoft Defender XDR, and Copilot for Security to deliver end-to-end detection, investigation, and response capabilities across cloud and on-premises infrastructure [29].

The Attack Simulation and SOC Design diagram outlines the structure of the simulated environment, which involves over 15 endpoints across various roles and platforms (Fig. 1).

The simulated attack campaigns triggered alerts corresponding to tactics such as credential dumping and lateral movement, which were observed and documented by the SOC analysts. These scenarios are further mapped to MITRE ATT&CK techniques [21] in Appendix 2.

Logs and alerts were collected using:

- Microsoft Defender for Endpoint for EDR-based threat detection
- Sysmon for deep system telemetry
- Azure Monitor Agent (AMA) to forward telemetry to Log Analytics
- Microsoft Defender XDR and Azure Sentinel served as the central SIEM/SOAR platform for investigation and correlation of alerts

Data was collected using Microsoft Defender sensors, the Azure Monitor Agent, and Sysmon. During the experiment, over 300 alerts were logged, and more than 200 incidents were generated, with varying levels of severity.

Visual statistics from Microsoft Sentinel are provided below to illustrate active alert coverage and incident distribution by severity (Figs. 2–3):

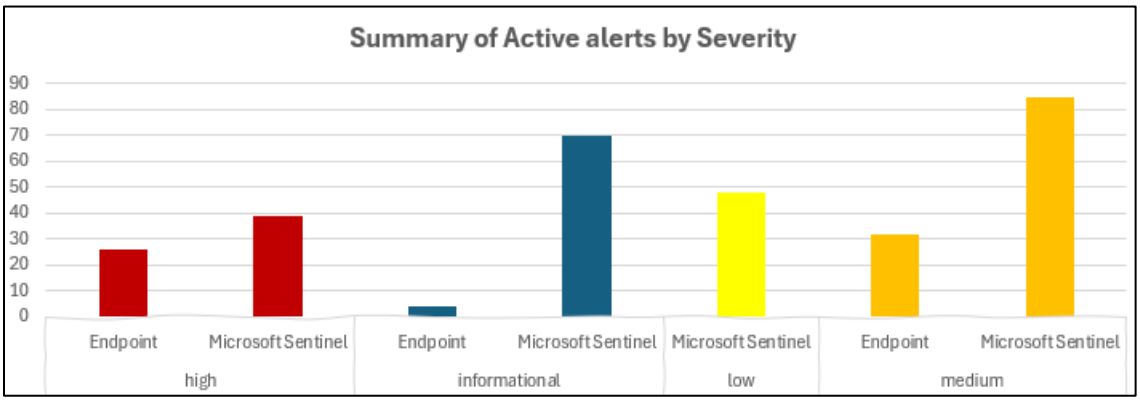

Fig. 2 Summary of Active Alerts by Severity

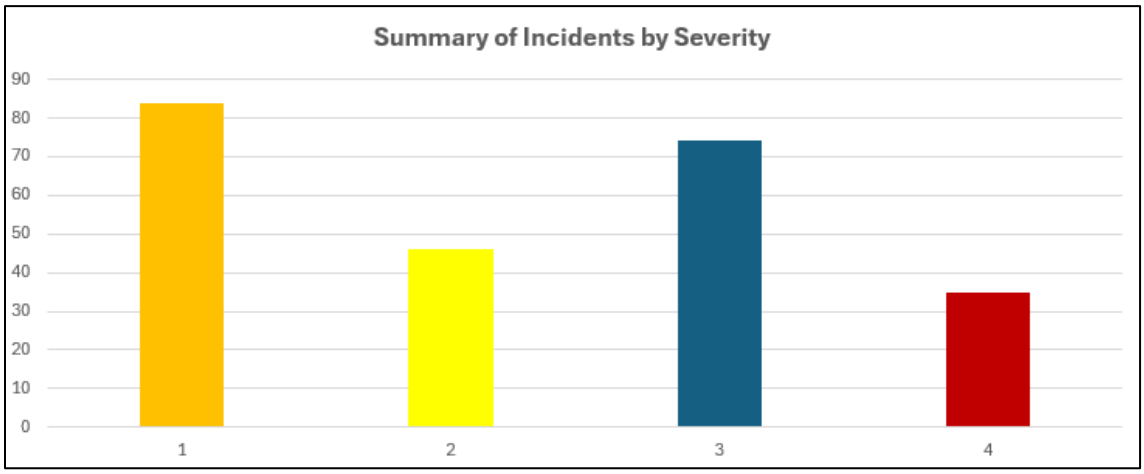

Fig. 3 Summary of Incidents by Severity

To ensure a comprehensive analysis, over 110 built-in and custom SOC alert rules, based on industry best practices, were implemented. These rules were crucial in generating meaningful alerts and achieving accurate detection coverage during the experiment (Fig. 4).

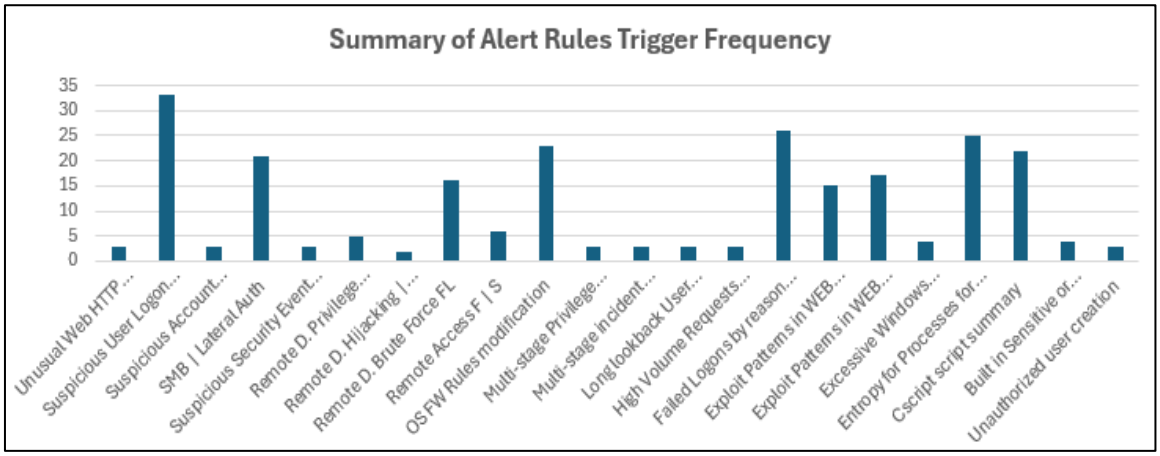

Fig. 4 Summary of Alert Rules Trigger Frequency

5.2 5W1H-Based Incident Analysis Template

Incident analysis is inherently a dynamic and non-linear process. To maintain consistency and ensure thorough evaluation, a custom-designed incident analysis template was developed based on the 5W1H framework. This structured approach enabled standardised documentation and comparison of both manual and MSC-assisted investigations, facilitating clear and coherent analysis across all participants (Fig. 5).

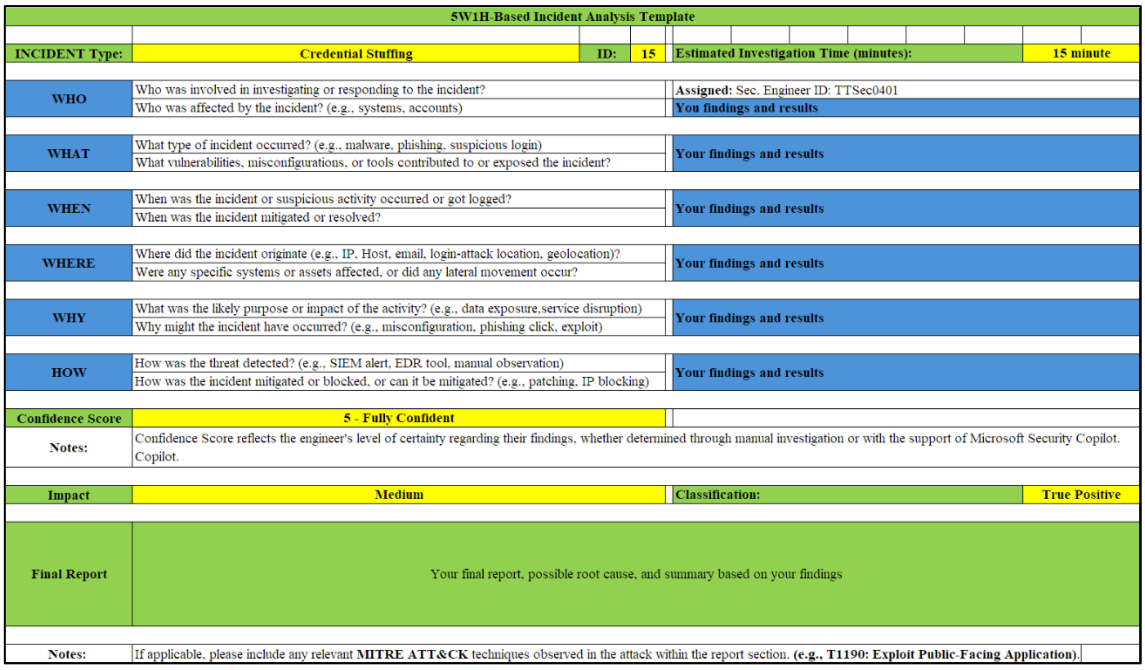

5W1H-Based Incident Analysis Template					
INCIDENT Type:	**Credential Stuffing**	**ID:**	**15**	**Estimated Investigation Time (minutes):**	**15 minute**
WHO	Who was involved in investigating or responding to the incident? Who was affected by the incident? (e.g., systems, accounts)			**Assigned:** Sec. Engineer ID: TTSec0401 **You findings and results**	
WHAT	What type of incident occurred? (e.g., malware, phishing, suspicious login) What vulnerabilities, misconfigurations, or tools contributed to or exposed the incident?			**Your findings and results**	
WHEN	When was the incident or suspicious activity occurred or got logged? When was the incident mitigated or resolved?			**Your findings and results**	
WHERE	Where did the incident originate (e.g., IP, Host, email, login-attack location, geolocation)? Were any specific systems or assets affected, or did any lateral movement occur?			**Your findings and results**	
WHY	What was the likely purpose or impact of the activity? (e.g., data exposure,service disruption) Why might the incident have occurred? (e.g., misconfiguration, phishing click, exploit)			**Your findings and results**	
HOW	How was the threat detected? (e.g., SIEM alert, EDR tool, manual observation) How was the incident mitigated or blocked, or can it be mitigated? (e.g., patching, IP blocking)			**Your findings and results**	
Confidence Score	**5 - Fully Confident**				
Notes:	Confidence Score reflects the engineer's level of certainty regarding their findings, whether determined through manual investigation or with the support of Microsoft Security Copilot. Copilot.				
Impact	**Medium**			**Classification:**	**True Positive**
Final Report	Your final report, possible root cause, and summary based on your findings				
Notes:	If applicable, please include any relevant **MITRE ATT&CK** techniques observed in the attack within the report section. **(e.g., T1190: Exploit Public-Facing Application).**				

Fig. 5 5W1H-Based Incident Analysis Template

The incident analysis template minimised cognitive and procedural overhead by focusing on core investigation components:

- Who was involved in responding to the incident, and who was affected by it
- What type of activity or threat occurred
- When the event occurred and was logged
- Where it originated or spread
- Why the event occurred (intent/cause)
- How it was detected and mitigated

The incident analysis template also included sections for additional questions, such as:

- Estimated Investigation Time
- Incident Type, Classification and Impact Level
- Final Report Summary and MITRE ATT&CK Techniques (if applicable)
- Analyst's ID and Confidence Score

This structured method ensured comparability across participants and allowed for clearer analysis of the results from traditional and AI-supported investigations.

5.3 Data Collection and Sharing with Participants

This section outlines the methods used for data collection and sharing among participants. The primary goal was maintaining clearness and consistency throughout the investigative process while gathering important metrics for later analysis.

5.3.1 Participant Survey and Profiling

Before the experiment began, all analysts were required to complete a baseline survey that collected the following information:

- Informed consent
- Preferred anonymity status
- Current role (e.g., SOC Analyst, Red Teamer, Security Engineer)
- Level of cybersecurity experience (e.g., Less than 1 year, 1–2 years, etc.)
- Industry/sector (e.g., Education, Healthcare, Technology)
- Familiarity with SOC tools (e.g., SIEM, EDR, SOAR)
- Previous use of AI-based security tools (e.g., MSC and others)

This data helped and facilitated the mapping of participants, allowing the researcher to correlate backgrounds with feedback and investigation outcomes (Figs. 6-13)

The initial survey results show that most participants regularly use EDR and SIEM tools, as well as scripting or query languages, which are essential technologies for effective SOC operations (Figs. 6–7).

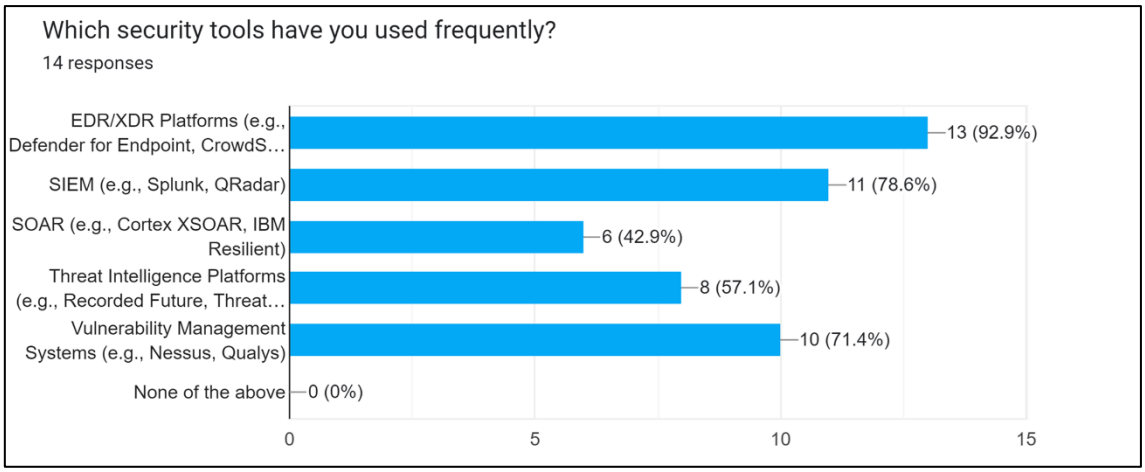

Fig. 6 Familiarity with Security Tools

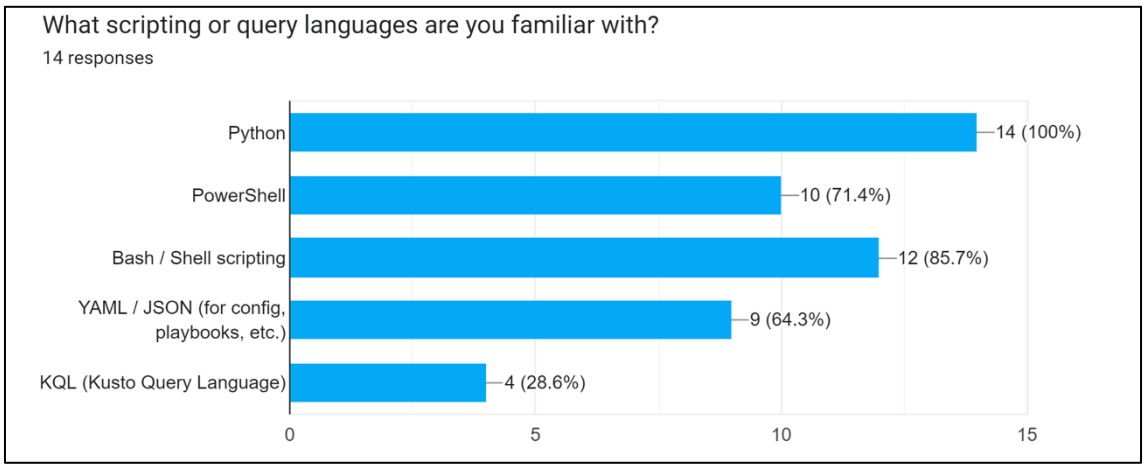

Fig. 7 Familiarity with Scripting or Query Languages

However, approximately 70% of the participants reported limited or no prior experience with AI-assisted security tools, including MSC (Figs. 8–9).

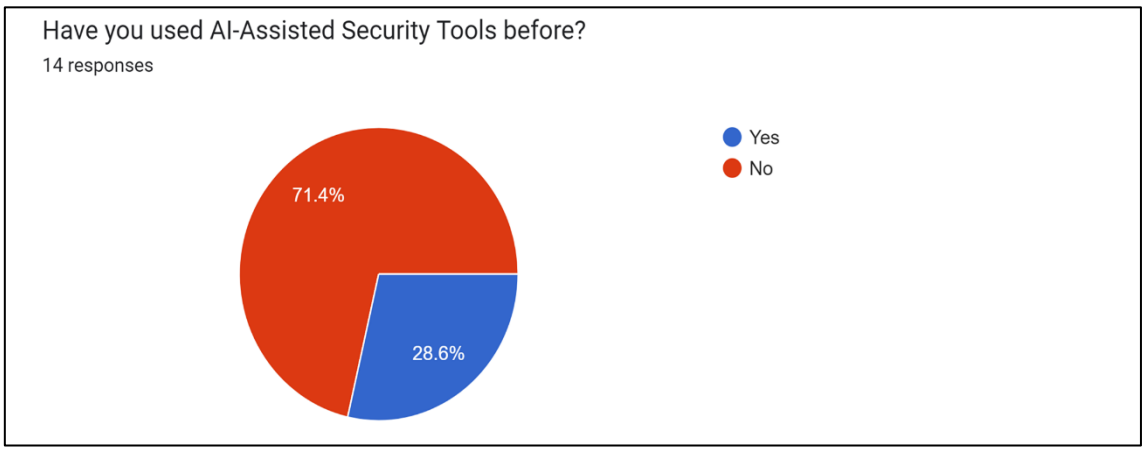

Fig. 8 Familiarity with AI-Assisted Security Tools

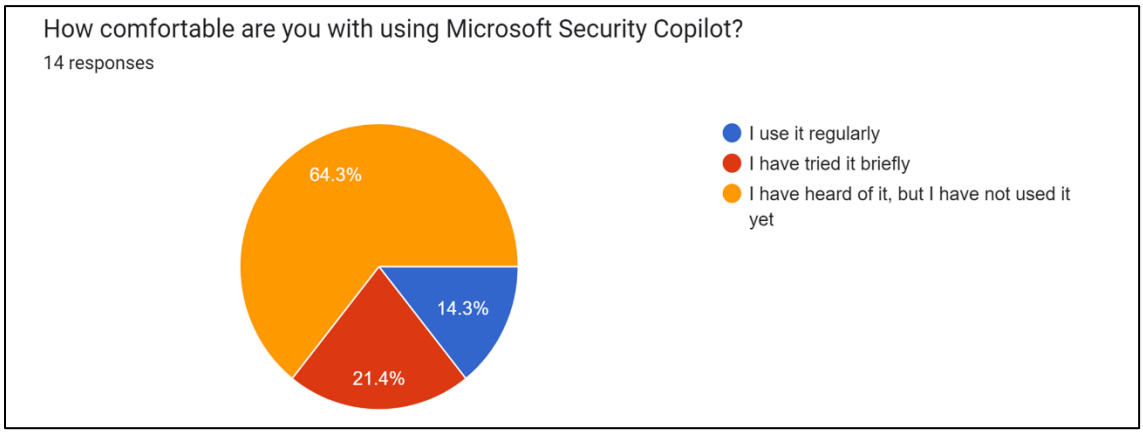

Fig. 9 Familiarity with MSC

Participants represented a variety of industries, including banking, finance, information technology, and the public sector (Fig. 10).

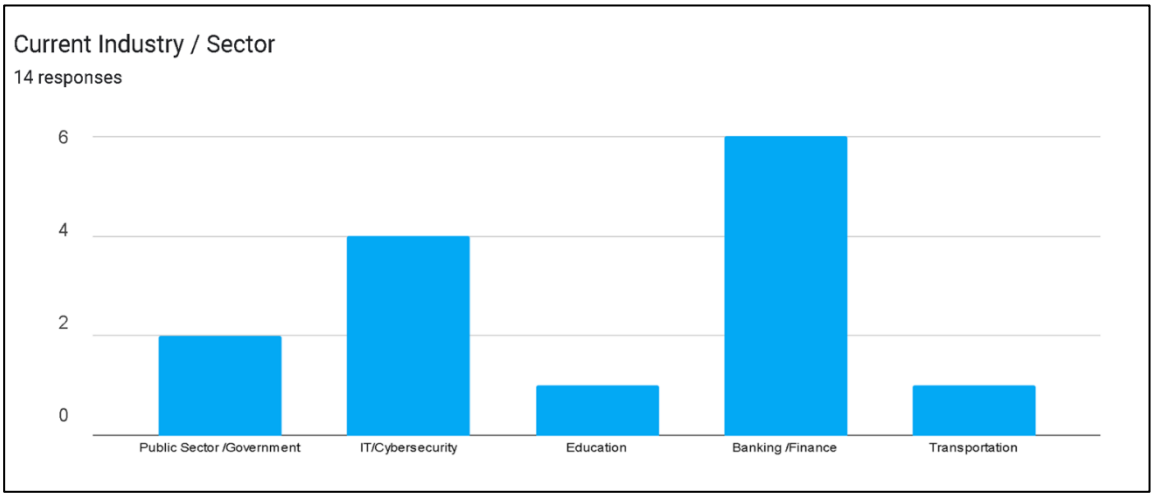

Fig. 10 Participant Distribution by Current Industry/Sector

In terms of experience, approximately 50% of respondents had less than five years of cybersecurity experience (Fig. 11).

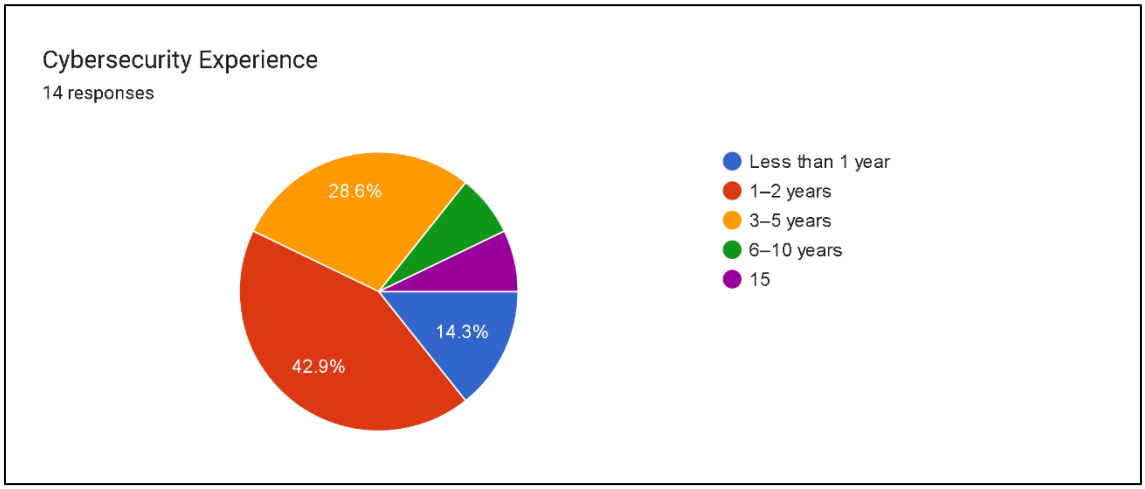

Fig. 11 Years of Cybersecurity Experience

Additionally, 78% of participants consented to having their names listed in the thesis contribution section (Fig. 12-13).

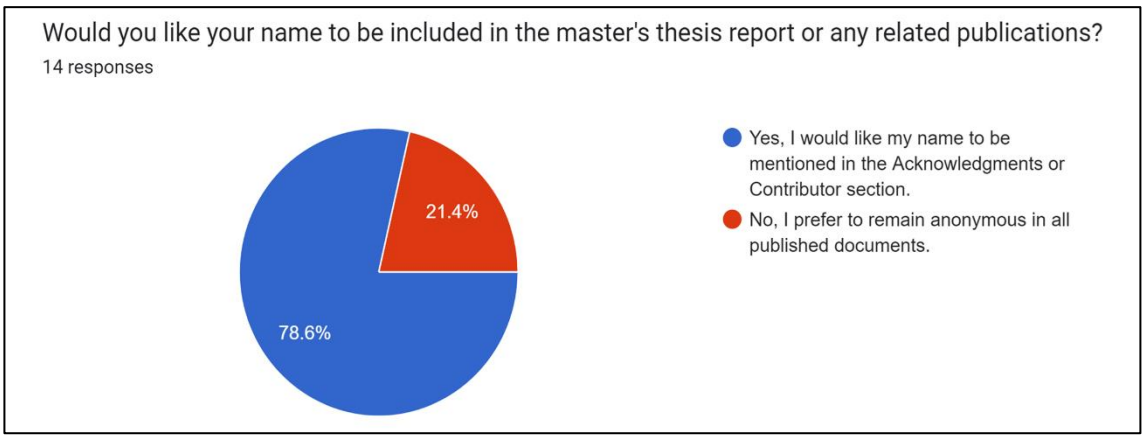

Fig. 12 Contributors' Consent: Name Disclosure Preferences

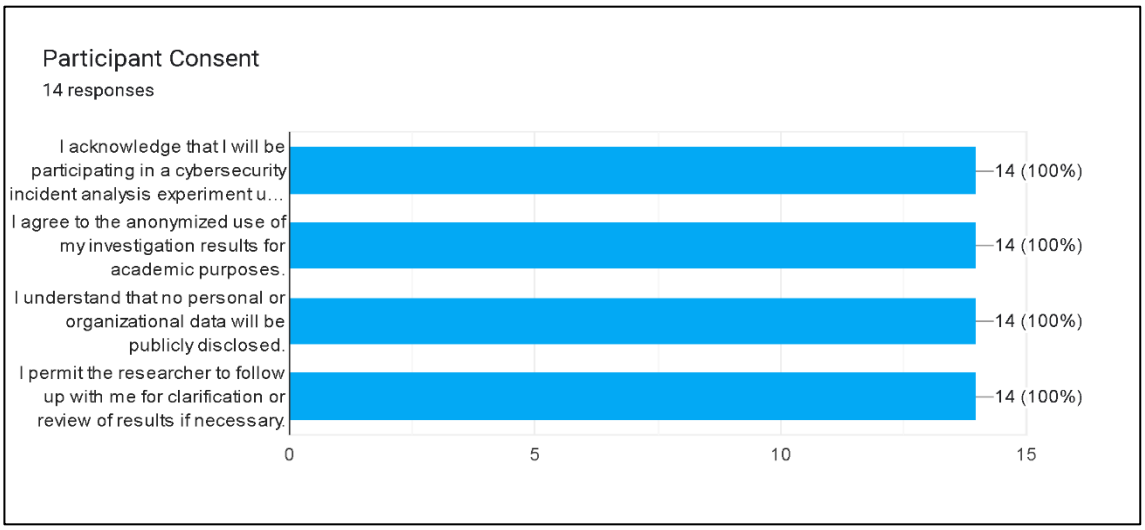

Fig. 13 Contributors' Consent: Participation and Data Use Agreement

5.3.2 Investigation Strategy and Reporting

Incidents were assigned using live system incident IDs, as generated by the SOC platform during the experiment. Participants examined the same and similar characteristics of incident IDs in both rounds to ensure balance and quality in the investigations. (Fig. 14).

Participants were assigned two rounds of incident investigation:

- Manual investigation using standard SOC tools
- AI-assisted investigation using MSC

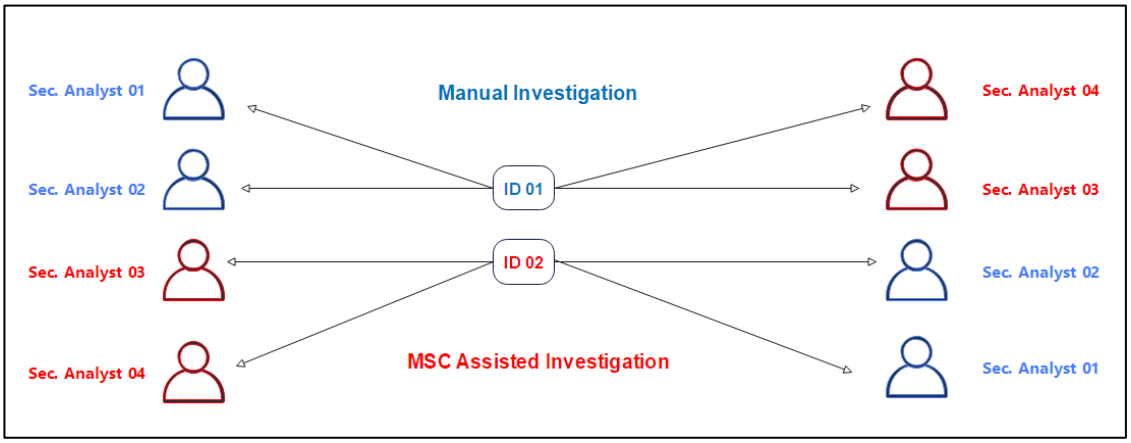

Fig. 14 Incident Assignment Approach

Each participant's investigation was tracked using a structured incident analysis template (capture 5.2), and final reports were compiled based on their shared results.

To prevent participant fatigue and ensure focused evaluations, each participant was assigned ten incidents, five of which they investigated manually and five with the assistance of MSC. This balanced approach helped maintain a consistent workload while allowing for a comparative analysis between traditional investigation methods and those assisted by MSC.

Analysts' identities were anonymised using unique ID (e.g., TTSec0401, TTSec0402) to ensure participant confidentiality. They received their assigned incidents directly within the system and documented their findings using a structured incident analysis template and reporting format.

5.4 Technical Constraints and Boundaries

The experimental phase concentrated solely on the incident analysis stage, utilising three core MSC plugins: Microsoft Defender XDR and Microsoft Sentinel, NL to KQL. These tools are part of an actively evolving ecosystem. For instance, in March 2025, Microsoft announced the addition of new agents and expanded plugin integrations, along with several enhancements to Microsoft Threat Intelligence features [33-35].

Due to this rapid development, evaluating every new component released during the experimental timeline was not feasible. Future research is encouraged to explore the capabilities of newly introduced integrations and feature enhancements, particularly as MSC expands its plugin support across external and hybrid platforms [33-35].

Additional technical limitations of the experiment included:

Resource constraints: The experimental lab was deployed on a small-scale on-premises virtual infrastructure and Microsoft Cloud environment, which was affected by subscription-based licensing and cloud cost restrictions. This limited the scalability and volume of incidents that could be tested.

Learning Curve for Security Tools: MSC was relatively new to most participants. Limited prior exposure may have influenced confidence and efficiency during MSC-assisted investigations.

6 Findings and Discussion

In this section, we present and analyse the key findings from the experiment, including user feedback, observations, and the comparative results of manual investigations versus those assisted by MSC.

6.1 Functionality

During the experimental phase, it became evident that MSC offers a wide range of functionalities extending beyond incident analysis. Its capabilities include extended incident response, threat hunting, intelligence gathering, and posture management.

MSC integrates numerous plugins, supporting both Microsoft-native services, such as Azure AI Search, Azure Firewall, Microsoft Defender Threat Intelligence, Microsoft Sentinel, Microsoft Intune, and Microsoft Purview, and third-party platforms such as Abuseipdb, Censys, CheckPhish, CyberArk Privilege Cloud, and Cybersixgill.

These integrations enable MSC to correlate telemetry data, enrich investigations with threat intelligence, and assist in identifying phishing threats and other security risks. The security tool's modular design and extensibility through plugins underscore its potential as a centralised support system for a wide range of SOC and IT-related tasks.

Although MSC supports a broad spectrum of use cases, only three core plugins were actively used in this study, selected for their relevance to incident analysis and data interpretation in a small-scale experimental environment.

It is essential to highlight that MSC's extensive feature set positions it as a powerful and versatile security tool in modern SOC environments.

6.2 Usability and Effectiveness

This section presents the key strengths of the AI-assisted security tool, MSC, observed during the experimental period. It includes usability factors, technical advantages, and the overall effectiveness of the tool in supporting SOC analysts.

6.2.1 Setup and Role-Specific Interface

From both a SOC perspective and in the context of LLM-based tools, our literature review and findings confirm that skills shortages and technical complexity remain common challenges. Effective use of LLMs typically requires robust infrastructure, specific data ingestion, and a baseline of technical expertise.

However, MSC is well-integrated into Microsoft’s unified SOC platform and does not require advanced skills for initial setup or regular operational use.

The security tool is easy to set up and can be activated with only a few simple configurations, making it accessible for use in operational environments. Deployment is initiated through Microsoft Azure or from the main console by assigning a resource group and specifying a capacity name and region.

Analysts can provision security compute units (SCUs) with a flexible usage model, including on-demand overage options, making the tool scalable based on actual workload demand.

MSC offers a user-friendly interface with prompts and role-specific options, including CISO, SOC Analyst, Threat Intelligence Analyst, and IT Administrator. This allows analysts to tailor their specific areas and interactions based on their functional responsibilities, thereby improving focus and relevance during investigations (Fig. 15).

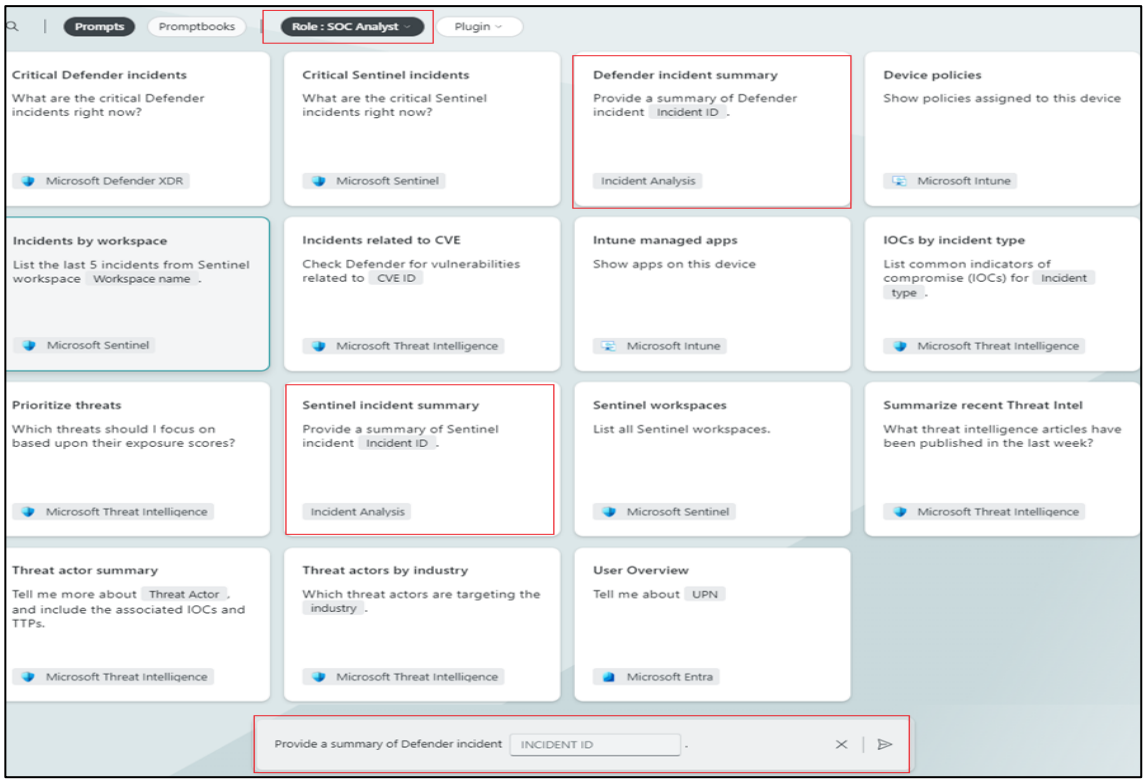

Fig. 15 Role-Specific Interface

6.2.2 Integration, Summarisation and Response

MSC improves SOC operations by combining natural language prompts with backend functionalities from platforms like Microsoft Defender XDR. During investigations, analysts benefit from not only results generated by prompts but also from dynamically created incident summaries (Fig. 16).

These summaries provide a clear, timeline-based overview of the attack chain, affected entities, and critical activities, all displayed directly within the integrated SOC console. They are accompanied by AI-generated guided response suggestions, such as device isolation or account disabling, which help analysts take informed action more efficiently (Fig. 16).

This unified approach helps analysts minimise oversight, reduce manual effort, and improve their understanding of the investigation context. AI-generated guidance offers actionable insights while allowing analysts to maintain control over the decision-making process.

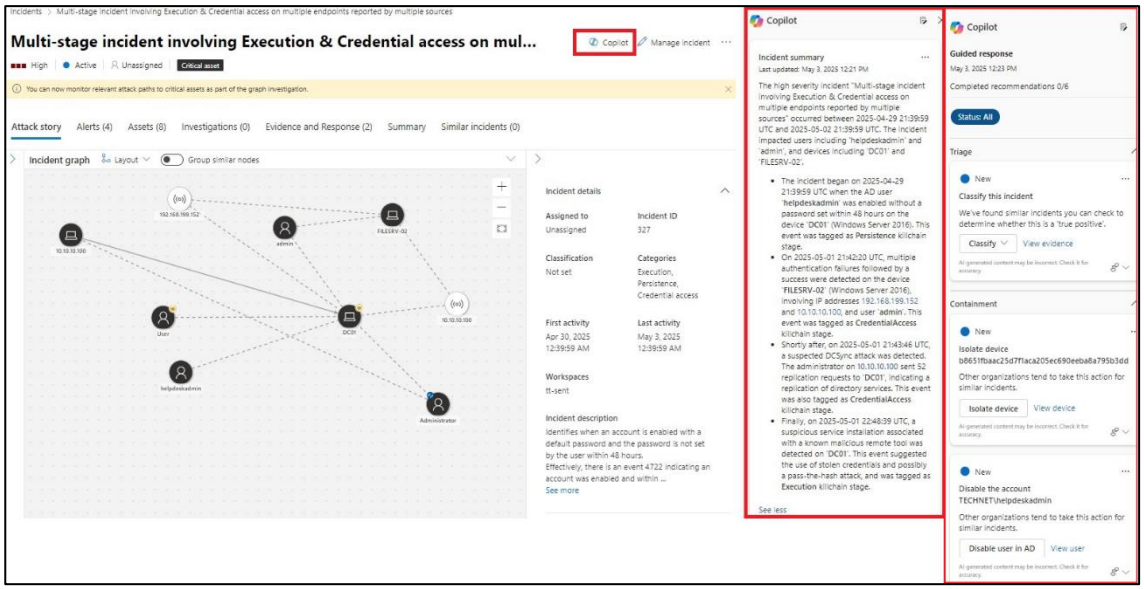

Fig. 16 Integration, Summarisation and Response

6.2.3 Prompt Sharing and Collaboration

MSC allows for both the use of pre-defined prompt templates and the creation of custom prompts. These can be saved and shared across the SOC team, supporting collaboration and knowledge transfer. This functionality enables the standardisation of workflows, alignment with use cases, and consistency across investigations, which is beneficial especially in larger teams or rotational shifts.

In addition, during the experiment, this functionality was actively used to distribute the 5W1H-based investigation prompt. The shared prompt was made accessible to all participants directly through MSC's prompt sharing interface, which helped ensure a consistent starting point for investigations and reduced the time spent on individual prompt formulation. This demonstrated the practical value of the prompt sharing collaboration feature in improving efficiency and reducing variation across analysts.

6.2.4 Script Interpretation and SOC Analyst Support

A key advantage observed during the experiment was that MSC have a better ability to interpret and explain suspicious scripts commonly used in cyberattacks, such as those related to command-and-control, lateral movement, or data exfiltration. These scripts, which are often obfuscated or encoded, were automatically translated into clear natural language explanations (Fig. 17).

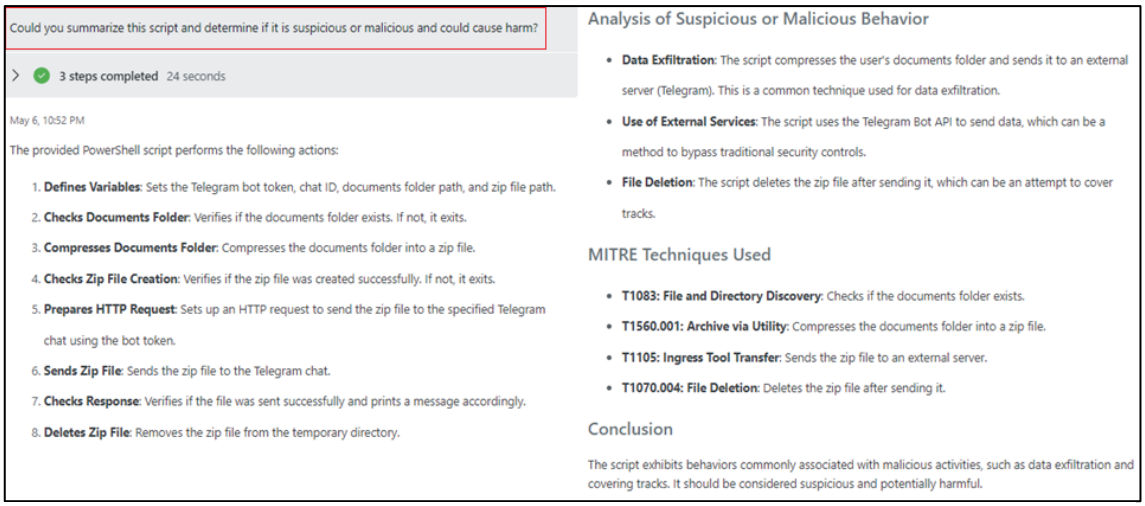

Fig. 17 Script Interpretation and SOC Analyst Support

This capability significantly assisted analysts, especially those without advanced scripting or reverse engineering skills, by bridging gaps in technical proficiency. As a result, SOC teams were able to accelerate impact assessments, deliver more accurate responses, and rely less on highly specialised skills.

6.2.5 Enhancing Time Efficiency

MSC demonstrably reduced investigation time during the experiment. By offering prompt-driven summaries, targeted recommendations, and streamlined visibility, analysts were able to focus more on decisions rather than repetitive technical querying.

The structured 5W1H prompt template used in the experiment provided clearer initial context compared to Microsoft's default summarisation prompt. This often led to better scoping of follow-up questions and reduced the number of iterative queries needed to complete an investigation.

The time savings were particularly noticeable in routine incident triage and were supported by quantitative comparisons between manual and MSC-assisted investigations.

6.2.6 Augmenting Analyst Decisions with AI

A key finding from the experiment was the collaborative relationship between the analyst and the AI. While the MSC provided clear summaries and intelligent suggestions, the final decisions, such as isolating a host, remained with the human analyst. This approach preserved the integrity of SOC workflows and ensured compliance with organisational

security standards. Instead of replacing human judgment, the MSC enhanced it, allowing for quicker and more confident decision-making. This highlighted the importance of maintaining critical thinking and human oversight in AI-supported environments.

6.2.7 Data Accuracy and Privacy

Unlike many custom or in-house LLM implementations, MSC is built on a foundation of secure data handling and trusted infrastructure. All prompt interactions occur within Microsoft's protected environment, with robust safeguards designed to prevent the inadvertent exposure of sensitive organisational information.

MSC's contextual understanding is strengthened by Microsoft's extensive threat intelligence and integrated security frameworks. This approach not only enhances response accuracy and reduces hallucinations but also fosters greater confidence in AI-generated insights. Together, these practices ensure that data integrity and confidentiality are maintained throughout every interaction with MSC.

6.3 Limitations and Drawbacks

This section outlines the challenges and limitations encountered while using MSC, including practical concerns, user observations, and technical gaps identified during the experiment.

6.3.1 Learning Curve and SOC Analyst Adaptation

One of the main limitations identified during the experiment was the steep learning curve associated with effectively using prompts in MSC. Unlike general-purpose LLMs, MSC relies heavily on well-structured and contextually relevant natural language input to generate accurate and actionable responses. This reliance posed significant challenges for users who were unfamiliar with AI tools or prompt-driven workflows.

A key issue noted during the research was that MSC does not offer dynamic guidance or support for the investigation period to offer suggested questions or idea creation, effective prompts, unlike platforms such as general-purpose LLMs like MS Copilot or ChatGPT.

For less experienced analysts, particularly those new to SOC operations or prompt-based systems, the lack of features like prompt suggestions, auto-complete options, or example queries made it more difficult to formulate inputs that would yield useful results.

This challenge was also evident during the feedback collection phase. Several contributors mentioned that they had to repeatedly rephrase their prompts or use external LLM tools to achieve the desired outcomes. While most users were able to adapt over time and became more proficient in interacting with MSC, the initial cognitive load was considerable.

6.3.2 Contextual Challenges and Fragmented Data Handling

A key limitation observed during the experiment was MSC's difficulty in maintaining and handling fragmented alert data during the prompt period. While MSC is capable of summarising incidents based on telemetry from sources like Microsoft Defender XDR and Sentinel, its ability to deliver meaningful insights is heavily dependent on the quality and completeness of backend data.

In more complex attack scenarios, such as multi-stage lateral movements or command-and-control activities, the incident data was sometimes fragmented or lacked the necessary detail in the originating alerts. For instance, MSC returned generic results without critical visibility when the underlying alert, though valid, was not well-populated with contextual metadata (e.g., process, path). This limitation is evident in cases where the security tool failed to fully interpret fragmented alerts (Fig. 18).

Fig. 18 Contextual Challenges and Fragmented Data Handling

6.3.3 Cost Constraints

One of the challenges identified with AI-assisted security tools is that, despite their innovative capabilities, they may not be financially accessible to smaller organisations due to the infrastructure and operational effort needed for implementation. Based on this research, several key aspects of MSC's billing model were identified.

MSC operates on a consumption-based billing system, using Security Compute Units (SCUs) as the basis for measurement. Organisations provision SCUs in line with their expected workloads, with billing applied hourly for the allocated capacity. To manage unexpected demand, overage capacity can also be configured, which is billed according to actual usage.

While this flexible model enables dynamic scaling, it also introduces potential challenges in cost management. For example, when analysts initiate multiple sessions or use resource-intensive plugins, SCU consumption can increase rapidly, potentially resulting in higher costs if the limit is not properly configured.

Additionally, billing occurs in full-hour blocks, regardless of actual utilisation, which may lead to cost inefficiencies. For smaller organisations, these cost dynamics can create obstacles to adoption. The need to provision and manage SCUs, along with the risk of unexpected overage charges, highlights the importance of careful capacity planning and monitoring to ensure cost-effective use of MSC.

6.4 Contributors' Results

In this section, we evaluate the findings from the experiment using both quantitative and qualitative methods. The analysis focuses on the results of contributors' investigations and includes a comparison between manual and AI-assisted (MSC) investigation rounds.

Key performance indicators, such as investigation time and analyst confidence levels, are presented alongside insights gathered from structured surveys and open-ended feedback.

The results also assess the perceived usability of MSC, the practical effectiveness of the 5W1H investigation approach, and the influence of prompt design, specifically comparing standard prompts with the structured 5W1H template. Contributors’ feedback further highlights the strengths and limitations of MSC from a practitioner's point of view.

To ensure consistency across both manual and MSC-assisted investigation phases, all SOC analysts were instructed to follow the 5W1H structure while documenting their work.

Investigation time and analyst confidence were measured quantitatively, while additional factors, including classification support, usability, and prompt effectiveness, were evaluated qualitatively through participant feedback and the final survey.

These insights are used to address the research questions introduced in Chapter 2 and provide a basis for the discussion in Section 6.6.

6.4.1 Manual Investigation Results

During the manual investigation phase, the first measured metric was the time spent per incident, recorded in minutes. This aimed to estimate how long, on average, each SOC analyst required to complete an investigation using standard tools, without support from MSC (Fig. 19).

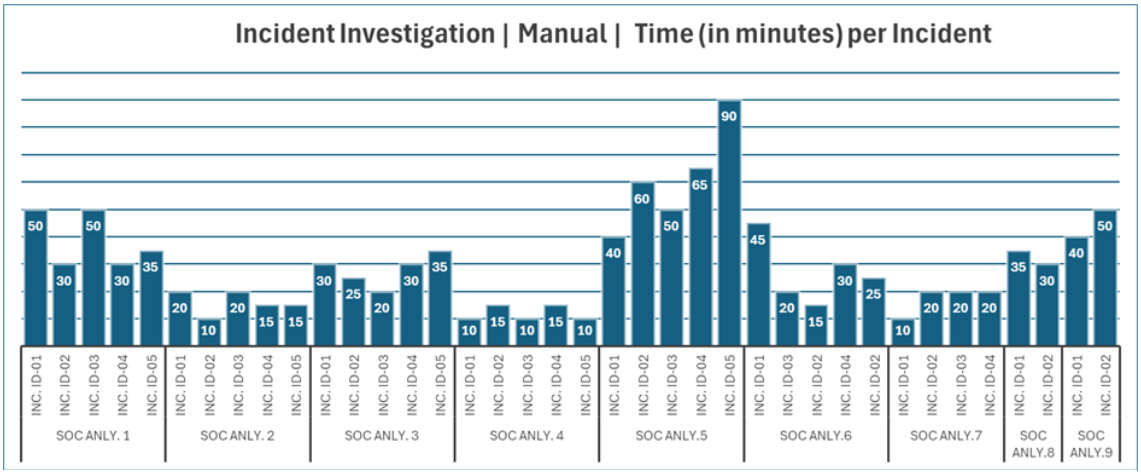

Fig. 19 Incident Investigation | Manual | Time (in minutes) per Incident

While all contributors participated voluntarily, two SOC analysts chose to investigate only two incidents per round instead of the assigned five. One analyst investigated four incidents, with a fifth excluded from the dataset due to an unusually long investigation time exceeding 200 minutes.

This exclusion was made to keep the overall time comparison consistent between the manual and MSC-assisted phases. The remaining analysts completed five incidents each as instructed.

In total, 42 valid manual investigation time entries were collected across nine SOC analysts. This dataset forms the basis for comparing investigation time between manual and MSC-assisted phases.

Following the individual time measurements, the total time spent by each SOC analyst during the manual investigation phase was calculated. This step aimed to assess the overall time investment required for completing all assigned incidents without AI assistance. The cumulative time across all valid manual investigation entries amounted to **1140 minutes** (Fig. 20).

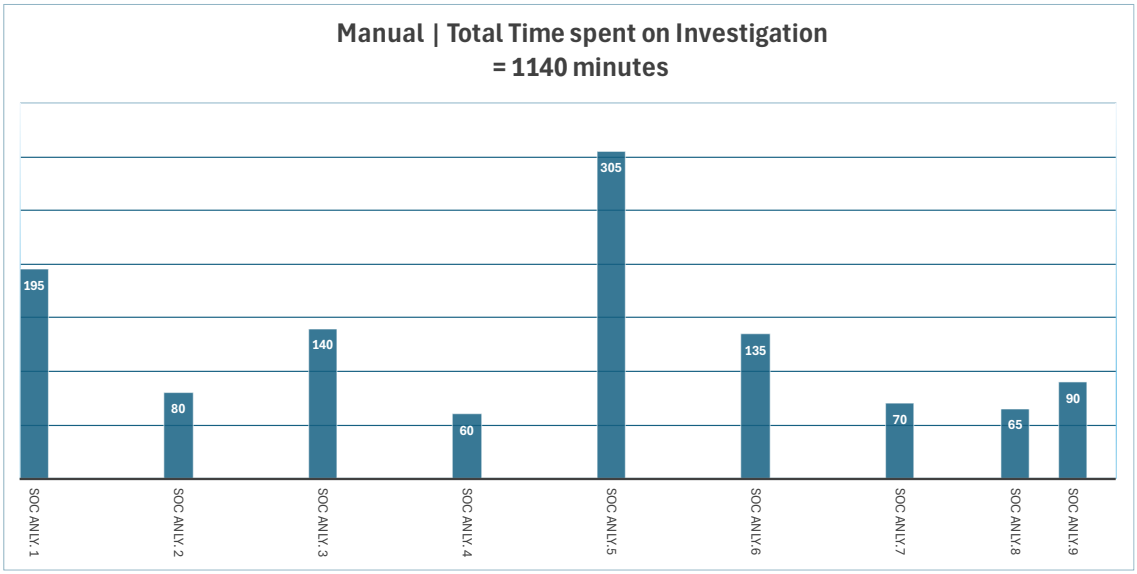

Fig. 20 Manual | Total Time spent on investigation

This total was based on 42 valid entries across 9 SOC analysts. While investigation time varied by analyst and incident complexity, the data provides a clear overview of the workload and time commitment involved in traditional investigation workflows.

6.4.2 MSC Assisted Investigation

In the MSC-assisted investigation phase, the same measurement approach was followed as in the manual phase. The investigation time per incident was recorded in minutes to evaluate how the use of MSC impacted analyst performance.

Each SOC analyst was asked to investigate a predefined set of incidents using MSC, following the same 5W1H structure to maintain consistency. A total of 42 valid investigation time entries were collected, matching the manual phase for balanced comparison (Fig. 21).

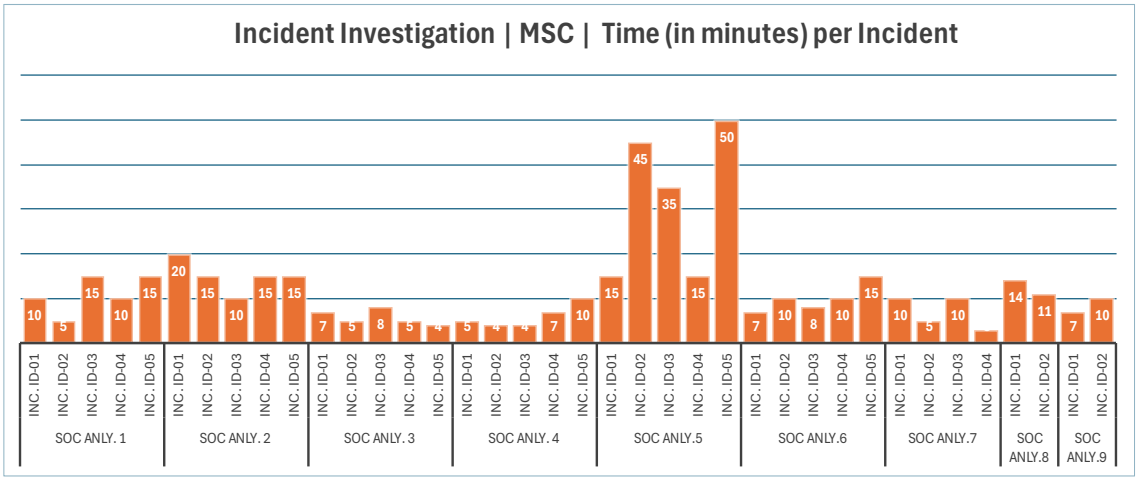

Fig. 21 Incident Investigation | MSC | Time (in minutes) per Incident

The cumulative time recorded across all MSC-assisted investigations amounted to **469 minutes** based on 42 valid entries (Fig. 22)

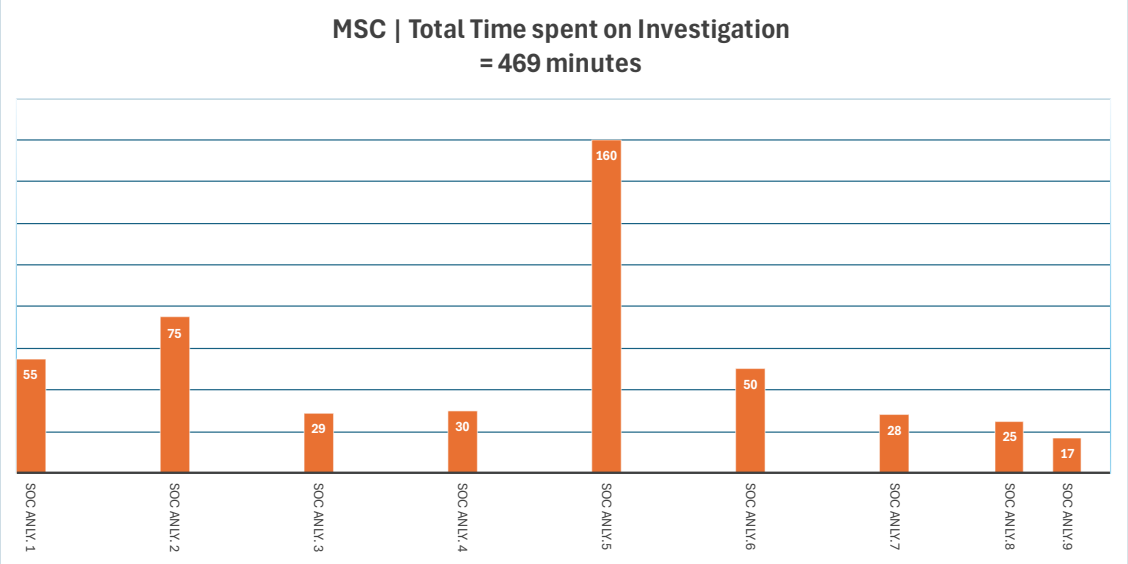

Fig. 22 MSC | Total Time spent on investigation

6.4.3 Manual vs MSC Comparison

In this section, a direct comparison is made between manual and MSC-assisted investigations to assess the impact of AI support on investigation efficiency.

The comparison focuses on the time spent per SOC analyst, across all incident entries, and highlights the overall difference in total time investment (Fig. 23).

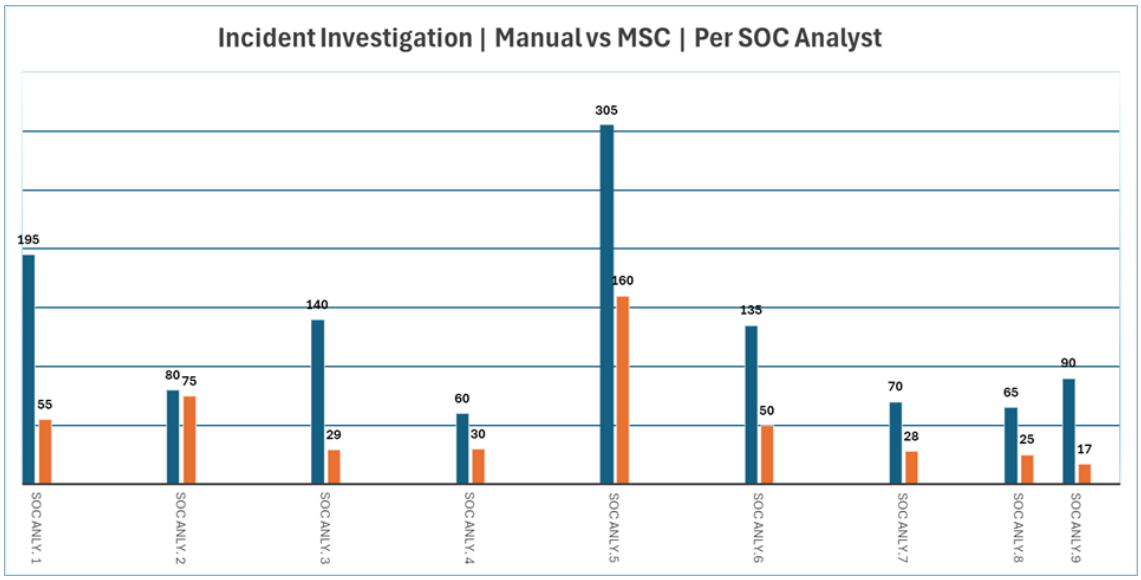

Fig. 23 Incident Investigation | Manual vs MSC | Per SOC Analyst

The cumulative time spent on manual investigations was **1140 minutes**, whereas the total for MSC-assisted investigations was **469 minutes**, covering the same number of incidents (42 entries). This reflects a total time **reduction of over 58%** when using MSC (Fig. 24).

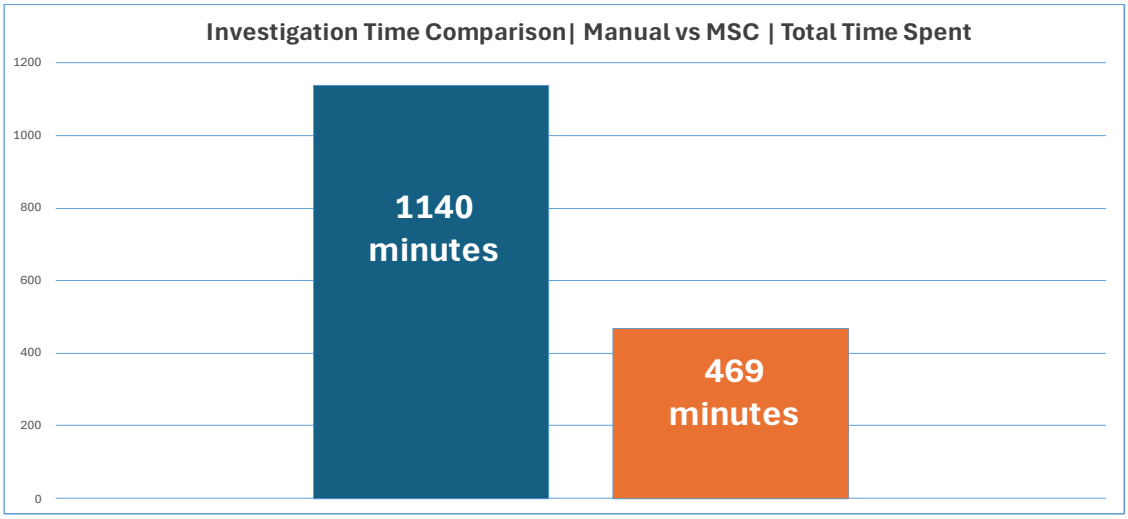

Fig. 24 Investigation Time Comparison| Manual vs MSC | Total Time Spent

Most SOC analysts experienced time savings during MSC-assisted investigations. However, individual results varied depending on factors such as working style, familiarity with security tools, and proficiency in formulating effective prompts.

While some analysts achieved only modest improvements, others were able to reduce their investigation time by more than 50%. Notably, several junior SOC analysts demonstrated significant gains in efficiency.

These findings indicate that MSC can significantly enhance operational performance, but the degree of its benefits relies on how effectively each analyst engages with and adapts to AI-assisted investigation workflows.

6.4.4 Confidence Score Comparison

In addition to time-based performance, this study also evaluated analyst confidence in incident investigation outcomes. SOC analysts rated their confidence on a 5-point scale after each investigation round.

A comparison of confidence scores from the manual phase and the MSC-assisted phase showed that most analysts reported similar or slightly increased confidence when using the MSC (Fig. 25).

The results are presented in a table that directly compares manual and MSC-assisted investigations (Fig. 25).

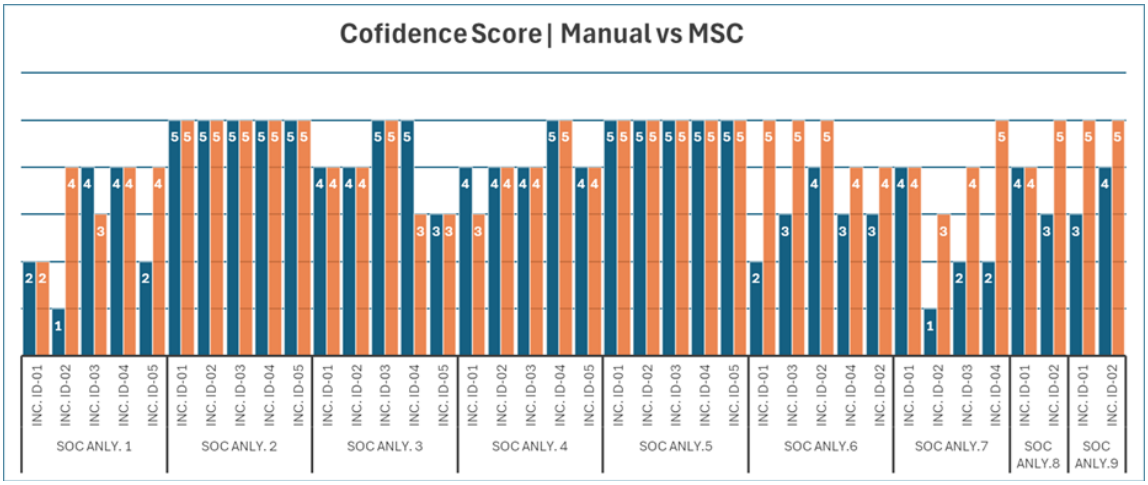

Fig. 25 MSC | Confidence Score | Manual vs MSC

Across the dataset, 12 incidents showed a clear increase in confidence scores when analysts used MSC. The increase was particularly noticeable among less experienced analysts, as also supported by qualitative feedback.

Overall, while the difference in scores was not large across all participants, the results indicate that MSC has a positive impact on perceived confidence, especially for analysts with less experience or lower familiarity with standard investigation tools.

6.5 Contributors' Feedback

In the final phase of this research, a structured feedback survey was conducted to evaluate SOC analysts' experiences with MSC and the 5W1H prompt approach.

The survey consisted of 15 structured questions and 2 open-ended questions, designed to assess both the manual investigation phase and the use of AI-assisted security tools with prompt templates. The responses to the structured questions were analysed and summarised into two tables, while the full open-ended responses have been included in Appendix 3 for reference.

The structured questions addressed various aspects of MSC, with a focus on classification accuracy, contextual understanding, natural language interaction, and prompt effectiveness. Table 2 presents responses related to usability and investigation quality, while Table 3 covers broader topics such as perceived workload impact, the necessity of prompt engineering skills, and the materials provided during the experimental period.

Question	Data collection type	Average (1–5)	Strongly Agree / 5	Short Summary
1. Did MSC help reduce investigation time?	Likert (1–5)	4.0	5 out of 9	The score was 4 to 5, indicating a positive trend in time-saving efforts observed.
2. Did MSC enhance confidence in incident analysis completeness?	Likert (1–5)	3.4	2 out of 9	Overall, there have been mostly moderate improvements in confidence.
3. Did MSC enhance classification accuracy (e.g. true/false positives)?	Likert (1–5)	4.2	4 out of 9	The support for improving classification is effective.
4. Was the MSC user-friendly for investigations?	Likert (1–5)	4.2	5 out of 9	Positive feedback received.
5. MSC's natural language understanding and result usefulness?	Likert (1–5)	3.7	2 out of 9	Moderate satisfaction was reported, although some weaknesses were noted.
6. Did you use any external LLM-based tools (e.g., ChatGPT) to formulate prompts?	Yes/No	–	4 Yes, 5 No	Almost half of the participants used external tools like ChatGPT to help formulate their prompts.
7. Did you utilise any contextual methods during the prompt period (e.g., alert details)?	Yes/No	–	7 Yes, 2 No	Many analysts indicated that they utilised contextual methods to improve their prompt inputs.
8. Was the 5W1H template helpful in manual investigation structure?	Likert (1–5)	4.3	5 out of 9	There is a strong preference for a structured approach.
9. Did the 5W1H-style prompts improve MSC response structure and accuracy?	Likert (1–5)	4.6	6 out of 9	Many SOC Analysts found it to be very effective.
10. Did MSC assist in identifying things overlooked manually?	Yes/No/Unsure	–	6 Yes, 2 Not Sure	Many SOC Analysts found that the MSC improved visibility.

Table 2. Usability & Effectiveness | Contributors' Results

When analysing the detailed results from **Table 2: Usability & Effectiveness, Q1** indicates that contributors generally agreed MSC helped reduce investigation time, with 5 out of 9 analysts rating it 5 out of 5, and an average score of 4.0. This finding aligns with the quantitative analysis in Section 6.4.3, where measurable time reductions were observed during AI-assisted investigations.

Both qualitative and quantitative data suggest that AI-based security tools like MSC positively impact investigation efficiency and enable faster incident response.

In **Q2**, a moderate improvement in analyst confidence was reported (average score: 3.4). This aligns with the confidence trends discussed in Section 6.4.4, confirming that while MSC supports investigation workflows, human validation remains essential, especially in complex or ambiguous scenarios.

Q3 further supports these findings, with analysts recognising improvements in classification accuracy (average score: 4.2), demonstrating MSC's effectiveness in helping distinguish between true and false positives. A similar positive outcome is reflected in **Q4**, where usability received a high average score of 4.2, highlighting MSC's accessible interface and ease of use during investigations.

Moving to **Q5**, **Q6**, and **Q7**, the results underscore the critical importance of prompt formulation. While MSC's natural language understanding received moderate satisfaction (**Q5**, avg. 3.7), many analysts relied on external tools such as ChatGPT (**Q6**) to help construct effective prompts.

Q7 confirms that using contextual methods (e.g., logs or alert summaries) improves the quality of prompt input and the relevance of MSC responses. These findings highlight the importance of integrating prompt engineering with cybersecurity expertise for the effective use of AI tools in SOC operations.

Q8 and **Q9** focus on structured investigation methods. The 5W1H template was found highly effective in supporting manual investigation structure (**Q8**, avg. 4.3), while structured prompts were shown to improve MSC's output clarity and accuracy (**Q9**, avg. 4.6). These results validate the experimental decision to implement structured questioning and reflect participants' preference for methodical investigative approaches.

Finally, **Q10** reveals that MSC contributed to identifying aspects that may have been overlooked during manual investigations, with 6 out of 9 analysts reporting improved visibility.

This benefit was particularly valuable for junior SOC analysts or those less familiar with traditional investigation tools and backend platforms such as Microsoft Sentinel or Defender XDR. In such cases, the AI-assisted guidance provided by MSC helped bridge experience gaps, offering clearer summaries and structured insights that simplified the investigation process.

Responses to **Q11** through **Q15** in **Table 3: AI Collaboration and Support** broaden the discussion beyond tool effectiveness to include considerations such as skill requirements, human-AI collaboration, and the experimental setup. These findings underscore the organisational and cognitive factors that influence the successful adoption of MSC.

In **Q11**, all analysts agreed that AI-assisted security tools like MSC can reduce the workload in the SOC and enhance response efficiency, confirming the perceived operational value of these security tools.

Question	Data collection type	Short Summary
11. Are AI-assisted security tools like MSC beneficial for reducing SOC workload and improving response?	Likert (1–5)	7 Agree, 2 Strongly Agree
12. Does effective use of MSC require both security knowledge and prompt formulation skills?	Likert (1–5)	5 Strongly Agree, 3 Agree, 1 Neutral
13. Did prompt clarity affect MSC response quality and usefulness?	Likert (1–5)	3 Strongly Agree, 5 Agree, 1 Neutral
14. Is a successful incident investigation with MSC dependent on human–AI collaboration?	Likert (1–5)	5 Strongly Agree, 3 Agree, 1 Neutral
15. Did you receive sufficient materials, support, and environment to complete your investigation properly?	Yes/No	All 9 respondents answered Yes

Table 3. AI Collaboration and Support | Contributors' Results

Q12 reinforces a key theme from earlier feedback: the effective use of MSC relies not only on general security knowledge but also on prompt engineering skills. This is also discussed in **Q13**, where the majority of analysts concurred that the clarity of the prompt significantly impacts the quality and usefulness of MSC's responses.

Q14 emphasises the necessity of a human–AI collaboration model. Analysts highlighted that the tool is most effective when guided by human interpretation and oversight. This

reinforces the view that maintaining human-AI collaboration is essential for ensuring reliable investigation outcomes.

Finally, **Q15** confirms that the experimental conditions were adequate; all nine contributors agreed they had the necessary materials, guidance, prepared video materials and technical setup to conduct their investigations confidently and independently.

In addition to the structured survey results, participants also provided open-ended feedback that offered deeper insight into their practical experiences with MSC. These qualitative responses emphasised both the benefits and limitations of using AI-assisted security tools in SOC operations.

Below are selected quotes that represent a range of perspectives from the contributors:

"It would truly improve the day-to-day life of SOC analysts. Especially for junior analysts, it could help filter false positives and reduce uncertainty during triage."

"MSC is helpful but still maturing. It gathers data quickly but doesn't always understand context. Analysts still need to validate results manually."

"My experience with MSC was largely positive, great for streamlining investigations and reducing repetitive tasks. It allows SOC teams to focus on decisions instead of raw data."

"I wouldn't consider AI tools as great assistance. Experienced analysts still perform better. MSC lacks full incident comprehension."

"It could be beneficial for newer analysts, but there's a risk of over-relying on AI. It's fast, but it's not always right yet."

"I believe this tool will become every SOC analyst's future friend, like ChatGPT. It reduces workload and helps explain complex malicious scripts quickly."

Beyond the structured survey responses, participants also provided open-ended feedback that highlighted the practical impact of the 5W1H investigation template.

As supported by responses to **Q8** and **Q9**, several analysts found that the structured approach helped optimise their workflow, improve the clarity of investigations, and reduce oversight in both manual and MSC-assisted phases. The 5W1H format enabled systematic thinking and clearer reporting, especially when time was limited or when less experienced analysts needed additional support.

Below are selected quotes that represent a range of perspectives from the contributors:

"It served as a very effective tool for structuring investigations. It ensures that all angles are considered and helps gather and analyse information in a logical way."

"The template was highly effective in guiding my work through a structured and focused process. It significantly improved the quality and coherence of my investigation results."

These reflections demonstrate that many contributors view AI-assisted security tools as a positive addition to SOC workflows, particularly in terms of efficiency and guidance for less experienced analysts. At the same time, several noted that MSC's current limitations, such as gaps in contextual understanding and reliance on well-structured prompts, mean that human oversight remains essential in high-quality incident response.

6.6 Summarising and Answering Key Questions

In the final section of the discussion and findings, we address the research questions outlined in Chapter 2. This is achieved by utilising insights gathered from the literature review, feedback from contributors, and experimental results.

Both quantitative metrics, such as investigation time and confidence scores, and qualitative observations from SOC analysts were taken into account to provide comprehensive answers. This approach ensures that the evaluation of the AI-assisted security tool, MSC, reflects both performance-based outcomes and practical experiences.

Key Question: How does Microsoft Security Copilot enhance the performance of SOC analysts during the technical investigation of security incidents?

MSC improves SOC analyst performance by reducing investigation time, enhancing the clarity of incident analysis, and streamlining report generation. It enables analysts, especially less experienced ones, to focus more on decision-making by offering structured summaries, guided prompts, and quick access to relevant data. The 5W1H structure proved particularly useful in improving investigation flow and consistency.

However, the full benefit of MSC depends on skilled human oversight and effective prompt formulation.

Sub-question 1. What are the most common challenges SOC teams face during incident investigations?

As highlighted in the literature review, key challenges include alert fatigue, manual data correlation, time pressure, and a shortage of skilled analysts. These factors slow down investigations and increase the risk of oversight.

Sub-question 2. Which features of Microsoft Security Copilot are designed to address these challenges, and how do they impact investigation quality and speed?

MSC addresses these issues through summarisation, natural language querying, script interpretation, easy integration with Defender/Sentinel data, and the use of structured prompts like 5W1H. These features reduce cognitive load, improve data access, and help analysts respond more quickly. In the experiment, they contributed to a measurable reduction in investigation time and increased confidence in analyst decision-making.

Sub-question 3. To what extent does Microsoft Security Copilot help resolve the limitations found in traditional SOC workflows?

MSC improved visibility and summarisation but could not fully replace human expertise. However, it does not eliminate the need for human judgment, particularly in complex or ambiguous scenarios. Prompt quality, backend data, and SOC analyst skills continue to affect the outcome.

Sub-question 4. What are the perceived benefits and limitations of Microsoft Security Copilot, based on hands-on experience from SOC analysts?

Analysts reported that MSC improved efficiency, provided clear guidance, and reduced repetitive workload. It was especially beneficial for less experienced analysts and supported better incident visibility.

However, limitations included the need for well-formulated prompts, incomplete contextual awareness, and a dependency on backend data quality.

7 Limitations and Future Work:

This chapter highlights the key limitations encountered during the study and offers recommendations for future research. The goal is to reflect on the challenges faced during the experimental process and to suggest ways to expand and enhance this work in broader contexts.

7.1 Limitations

This study faced several limitations, primarily related to infrastructure, tool scope, and participant availability.

As described in Section 5.4, the experimental environment was built using a small-scale, cloud-based and on-premises lab setup integrated with Microsoft cloud resources. This configuration was affected by subscription-based licensing and cost constraints, which limited the scale and diversity of incident scenarios that could be simulated.

A key limitation of this research was the evaluation of only a single AI-based security tool, MSC. While this enabled a focused assessment of its practical application within SOC workflows, it did not allow for comparative analysis with other AI-assisted security platforms.

Although AI-assisted security tools are still emerging and currently offer a limited number of functionally mature solutions tailored to SOC operations, the rapid pace of development in this field suggests that comparative evaluations will become increasingly relevant in future research.

The final participant pool included nine SOC analysts. Although this is a relatively small number, it is consistent with other exploratory studies in cybersecurity and human factors research, where recruitment is often constrained by the availability of qualified professionals and the sensitive nature of the subject matter.

Additionally, the voluntary nature of the study, along with the significant time commitment required for multi-stage incident investigations, further contributed to the recruitment challenges.

However, this limitation is not uncommon; comparable studies have successfully engaged similarly small groups, including those with nine [36], thirteen [37], and twenty-one [38] participants. These studies demonstrate that careful selection and deep engagement of participants can still yield valuable insights, even in small-scale evaluations.

This study contributes to that body of work by providing actionable, real-world perspectives on AI-assisted SOC workflows.

7.2 Future Work

Building on this study's findings, several directions for future research can be proposed. First, evaluations of MSC should include larger, more complex SOC environments to assess tool performance under real-world conditions.

Second, side-by-side comparisons with other AI-assisted security tools would offer a more comprehensive understanding of MSC's relative strengths and limitations. As AI integration in cybersecurity continues to evolve, comparative studies across platforms will be essential for guiding SOC decision-makers on security tool adoption and integration strategies.

Additionally, this study focused primarily on MSC's support for technical incident analysis. However, MSC supports a broader range of SOC functions, including threat hunting, posture management, and intelligence gathering, which were not explored in depth. Future research could investigate these additional capabilities, assessing how they contribute to different stages of security operations and their impact on SOC performance and team workflows.

Lastly, future studies would benefit from involving a larger and more diverse participant group, enabling statistical comparisons across experience levels, organisational types, and investigation styles. This will enhance generalisability and provide deeper insights into how various analyst profiles engage with AI-assisted security tools in different operational contexts.

8 Conclusion:

This master's thesis investigated the use of MSC, an AI-assisted security tool, in enhancing the technical analysis of incidents within SOC environments. The study combined both manual and MSC-assisted investigation phases in a realistic SOC setup, supported by structured prompts, enterprise tools, and simulated attack scenarios.

The results show that MSC can significantly reduce investigation time, improve visibility, and help SOC analysts analyse more effectively. However, the effectiveness of MSC was highly dependent on prompt quality and backend data. In complex or multi-stage incidents, human judgment remained essential.

Overall, this study concludes that the AI-assisted security tool MSC can serve as a valuable assistant in SOC workflows, particularly when integrated with structured investigation frameworks and operated by trained analysts.

Acknowledgements:

I would like to express my sincere gratitude to my supervisor, Professor Risto Vaarandi, for his invaluable guidance, support, and feedback throughout this research. His insights and encouragement played a crucial role in shaping the direction and quality of this thesis.

I would also like to extend my sincere thanks to the Cybersecurity Engineers who participated in the experimental phase of this research. Their time, dedication, and professional insights were essential in enabling the practical evaluation of MSC.

The contributors include:

Anastasija Mihnenko
Orkhan Hasanzade
Sabuhi Safarov
Rennet Tamm
Mahammad Garayev
Tural Rzayev
Farid Yusubov
Aghadadash Guliyev
Ramiz Museyibov
Turqut Valiyev
Musa Salamov

Finally, I am deeply thankful to my family and children for their unwavering support and motivation throughout my academic journey. Their patience and encouragement helped me persevere and complete this milestone.

9 References:

[1] K. Palmgren, D. Parsons, and SANS Institute, “SANS 2024 AI survey: AI and Its Growing Role in Cybersecurity: Lessons Learned and path forward,” Sans.org. [Online]. Available: https://www.sans.org/webcasts/sans-2024-ai-survey-ai-and-its-growing-role-in-cybersecurity-lessons-learned-and-path-forward/. [Accessed: 02-Nov-2024].

[2] W. Chen and J. Zhang, “Elevating security operations: The role of AI-driven automation in enhancing SOC efficiency and efficacy,” JAMM, vol. 8, no. 2, pp. 1–13, 2024.

[3] M. Ozkan-Okay et al., “A comprehensive survey: Evaluating the efficiency of artificial intelligence and machine learning techniques on cyber security solutions,” IEEE Access, vol. 12, pp. 12229–12256, 2024.

[4] M. Vielberth, F. Bohm, I. Fichtinger, and G. Pernul, “Security operations Centre: A systematic study and open challenges,” IEEE Access, vol. 8, pp. 227756–227779, 2020.

[5] “The essentials of Security Operations Centres (SOC),” Intersecinc.com. [Online]. Available: https://www.intersecinc.com/blogs/the-essentials-of-security-operations-Centres-soc. [Accessed: 01-Nov-2024].

[6] J. Tilbury and S. Flowerday, “Humans and automation: Augmenting security operation Centres,” J. Cybersecur. Priv., vol. 4, no. 3, pp. 388–409, 2024.

[7] K. Palmgren, SANS Institute, and A. Kim, “SANS 2024 SOC survey: Facing Top Challenges in security operations,” Sans.org. [Online]. Available: https://www.sans.org/webcasts/sans-2024-soc-survey-facing-top-challenges-in-security-operations/. [Accessed: 01-Nov-2024].

[8] J. Muniz, G. McIntyre, and N. AlFardan, Security Operations Centre: Building, operating, and maintaining your SOC. Cisco Press, 2015.

[9] “Building and securing the modern security operations Centre (SOC),” International Journal of Business Intelligence and Big Data Analytics, vol. 5, no. SOC, pp. 1–15, 2022.

[10] S. Freitas, J. Kalajdjieski, A. Gharib, and R. McCann, “AI-driven guided response for security operation Centres with Microsoft copilot for security,” arXiv [cs.LG], 2024.

[11] “What is Microsoft Copilot for security?,” Microsoft.com. [Online]. Available: https://learn.microsoft.com/en-us/copilot/security/microsoft-security-copilot. [Accessed: 03-Nov-2024].

[12] Arctic Wolf, “3 types of security operations Centre models,” Arctic Wolf, 06-Jul-2023. [Online]. Available: https://arcticwolf.com/resources/blog/five-types-of-security-operations-Centre-models/. [Accessed: 27-Oct-2024].

[13] ManageEngine, “ManageEngine Log360,” ManageEngine Log360. [Online]. Available: https://www.manageengine.com/log-management/siem/components-of-security-operations-Centre-soc.html. [Accessed: 27-Oct-2024].

[14] H. Patel, “Security Operations Centre explained: Components, setup, and key benefits,” WPG Consulting, 22-Nov-2023. [Online]. Available: https://wpgc.io/blog/security-operations-Centre-benefits/. [Accessed: 01-Nov-2024].

[15] Y. Baddi, M. A. Almaiah, O. Almomani, and Y. Maleh, The Art of Cyber Defense: From Risk Assessment to Threat Intelligence. Boca Raton, FL: CRC Press, 2024.

[16] M. Majid and K. Ariffi, “Success factors for cyber security operation Centre (SOC) establishment,” in Proceedings of the Proceedings of the 1st International Conference on Informatics, Engineering, Science and Technology, INCITEST 2019, 18 July 2019, Bandung, Indonesia, 2019.

[17] S. A. Chamkar, Y. Maleh, and N. Gherabi, “The human factor capabilities in security operation Centre (soc),” EDPACS, vol. 66, no. 1, pp. 1–14, 2022.

[18] K. Zidan, A. Alam, J. Allison, and A. Al-sherbaz, “Assessing the challenges faced by security operations centres (SOC),” in Lecture Notes in Networks and Systems, Cham: Springer Nature Switzerland, 2024, pp. 256–271.

[19] D. Moorthy, “2024 SME security workload impact report,” Coro Cybersecurity, 17-Apr-2024.

[20] L. N. Kaliyaperumal, Ed., The Evolution of Security Operations and Strategies for Building an Effective SOC, vol. 5, 2021. ISACA Journal.

[21] S. A. Chamkar, Y. Maleh, and N. Gherabi, “Security Operations Centres: Use case best practices, coverage, and gap analysis based on MITRE adversarial tactics, techniques, and common knowledge,” J. Cybersecur. Priv., vol. 4, no. 4, pp. 777–793, 2024.

[22] “Research Report 2024 State of Threat Detection,” Vectra.ai. [Online]. Available: https://www.vectra.ai/resources/2024-state-of-threat-detection. [Accessed: 02-Nov-2024].

[23] “Trend Companion,” Trend Micro. [Online]. Available: https://www.trendmicro.com/en_us/business/technologies/ai-companion.html. [Accessed: 12-Dec-2024].

[24] “CrowdStrike unveils Charlotte AI Detection Triage for faster SOC triage,” Crowdstrike.com. [Online]. Available: https://www.crowdstrike.com/en-us/press-releases/crowdstrike-delivers-next-breakthrough-in-ai-powered-agentic-cybersecurity-with-charlotte-ai-detection-triage/. [Accessed: 13-Apr-2025].

[25] “Gemini in Google SecOps,” Google Cloud. [Online]. Available: https://cloud.google.com/chronicle/docs/secops/gemini-chronicle. [Accessed: 14-Apr-2025].

[26] “The Cyber Security Professional’s Guide to Prompt Engineering,” Checkpoint.com. [Online]. Available: https://www.checkpoint.com/resources/items/white-paper-the-cyber-security-professionals-guide-to-prompt-engineering. [Accessed: 05-May-2025].

[27] A. Bozkurt and R. C. Sharma, “Generative AI and prompt engineering: The art of whispering to let the genie out of the algorithmic world.” Zenodo, 2023.

[28] O. Asare, M. Nagappan, and N. Asokan, “A user-Centreed security evaluation of copilot,” in Proceedings of the IEEE/ACM 46th International Conference on Software Engineering, 2024, pp. 1–11.

[29] “What is Microsoft’s unified security operations platform?” Microsoft.com. [Online]. Available: https://learn.microsoft.com/en-us/unified-secops-platform/overview-unified-security. [Accessed: 07-Apr-2025].

[30] K. Kim, S. Yoon, D. Lee, J. Jang, H. Oh, and D. Shin, “Study on prioritization of actions by classifying and quantifying cyber operational elements using 5W1H method,” IEEE Access, vol. 10, pp. 74765–74778, 2022.

[31] M. Girdhar, J. Hong, Y. You, T.-J. Song, and M. Govindarasu, “Cyber-attack event analysis for EV charging stations,” in 2023 IEEE Power & Energy Society General Meeting (PESGM), 2023, pp. 1–5.

[32] S. Grigaliunas, J. Toldinas, A. Venckauskas, N. Morkevicius, and R. Damasevicius, “Digital evidence object model for situation awareness and decision making in digital forensics investigation,” IEEE Intell. Syst., vol. 36, no. 5, pp. 39–48, 01 Sep-Oct 2021.

[33] V. Jakkal, “Microsoft unveils Microsoft Security Copilot agents and new protections for AI,” Microsoft Security Blog, 24-Mar-2025. [Online]. Available: https://www.microsoft.com/en-us/security/blog/2025/03/24/microsoft-unveils-microsoft-security-copilot-agents-and-new-protections-for-ai/. [Accessed: 13-Apr-2025].

[34] “Securely integrate On-Prem and Self-Hosted VM instances of Splunk with Microsoft Security Copilot,” Microsoft.com. [Online]. Available: https://techcommunity.microsoft.com/blog/securitycopilotblog/securely-integrate-on-prem-and-self-hosted-vm-instances-of-splunk-with-microsoft/4402551. [Accessed: 13-Apr-2025].

[35] “What’s new in Microsoft Security Copilot?,” Microsoft.com. [Online]. Available: https://learn.microsoft.com/en-us/copilot/security/whats-new-copilot-security. [Accessed: 13-Apr-2025].

[36] A. Nganga, G. Nganya, M. Lützhöft, S. Mallam, and J. Scanlan, “Bridging the gap: Enhancing maritime vessel cyber resilience through security operation centers,” Sensors (Basel), vol. 24, no. 1, 2023.

[37] T. Braun, I. Pekaric, and G. Apruzzese, “Understanding the process of data labeling in cybersecurity,” in Proceedings of the 39th ACM/SIGAPP Symposium on Applied Computing, 2024.

[38] B. A. Alahmadi, L. Axon, and I. Martinovic, “99% false positives: A qualitative study of SOC analysts’ perspectives on security alarms,” USENIX Secur Symp, pp. 2783–2800, 2022.

Appendix 1 – Non-exclusive licence for reproduction and publication of a graduation thesis[1]

I Elguj Yusifbayli

1. 1. Grant Tallinn University of Technology free licence (non-exclusive licence) for my thesis “Use of AI-based security tool assistance for improving the technical analysis of incidents in Security Operations Centre” supervised by Risto Vaarandi.

    1.1 to be reproduced for the purposes of preservation and electronic publication of the graduation thesis, incl. to be entered in the digital collection of the library of Tallinn University of Technology until expiry of the term of copyright.

    1.2 to be published via the web of Tallinn University of Technology, incl. to be entered in the digital collection of the library of Tallinn University of Technology until expiry of the term of copyright.

2. I am aware that the author also retains the rights specified in clause 1 of the non-exclusive licence.

3. I confirm that granting the non-exclusive licence does not infringe other persons' intellectual property rights, the rights arising from the Personal Data Protection Act or rights arising from other legislation.

[1] The non-exclusive licence is not valid during the validity of access restriction indicated in the student's application for restriction on access to the graduation thesis that has been signed by the school's dean, except in case of the university's right to reproduce the thesis for preservation purposes only. If a graduation thesis is based on the joint creative activity of two or more persons and the co-author(s) has/have not granted, by the set deadline, the student defending his/her graduation thesis consent to reproduce and publish the graduation thesis in compliance with clauses 1.1 and 1.2 of the nonexclusive licence, the non-exclusive license shall not be valid for the period.

Appendix 2 – Mapping of Investigated Incidents to the MITRE ATT&CK® Framework

The techniques listed below illustrate MITRE ATT&CK® mappings that were either a technical analysis of an incident simulated during the experiment or identified by SOC analysts during investigations. The techniques were gathered through MSC outputs and manual classification (Table 4).

Technique Name	Technique ID	Tactic Name
Application Layer Protocol	T1071	Command and Control
OS Credential Dumping	T1003	Credential Access
Masquerading	T1036	Defence Evasion
Match Legitimate Name or Location	T1036.005	Defence Evasion
Indicator Removal on Host	T1070	Defence Evasion
Impair Defences: Disable or Modify Tools	T1562.001	Defence Evasion
Modify Registry	T1112	Defence Evasion, Persistence
Password Guessing	T1110.001	Credential Access
Process Injection	T1055 (001,002)	Defence Evasion, Privilege Escalation
Network Service Discovery	T1046	Discovery
Command and Scripting Interpreter	T1059	Execution
PowerShell	T1059.001	Execution
Unix Shell	T1059.004	Execution
User Execution: Malicious File	T1204.002	Execution
Exploitation for Client Execution	T1203; T1566.001	Initial Access
Remote Services: Remote Desktop Protocol	T1021.001	Lateral Movement
SMB/Windows Admin Shares	T1021.002	Lateral Movement
Exploitation of Remote Services	T1210	Lateral Movement
Create Account	T1136	Persistence
Web Shell	T1505.003	Persistence
Create or Modify System Process	T1543.003	Persistence, Privilege Escalation
Account Manipulation	T1098	Credential Access
Establish Accounts	T1585	Resource Development
Compromise Accounts	T1586	Resource Development
Data Encrypted for Impact	T1486	Impact
Lateral Movement	TA0008	Lateral Movement
Privilege Escalation	TA0004	Privilege Escalation
Agent Tesla Malware Behaviour	T1059; T1105; T1041	Execution, Collection, Exfiltration, C2
Exploitation for Privilege Escalation; Valid Accounts: Local Accounts	T1068; T1078.003	Privilege Escalation
Exploit Public-Facing Application	T1190	Initial Access
Valid Accounts; Brute Force: Password Guessing	T1078; T1110.001	Defence Evasion, Persistence, Privilege Escalation, Initial Access, Credential Access
Pass the Hash	T1550.002	Command and Control, Execution, Lateral Movement
Enable or Modify Cloud Compute Infrastructure: Enable RDP Access	T1578.004	Lateral Movement, Defence Evasion

Table 4. MITRE ATT&CK® Technique Mapping of Investigated Incidents

Appendix 3 – Open-Ended Survey Responses

This appendix presents the full set of qualitative feedback provided by participants in response to open-ended survey questions. The responses include reflections on the usability, limitations, and effectiveness of the AI-assisted security tool MSC, as well as insights on the structured 5W1H investigation approach. These comments offer additional context to the structured survey data and contribute to a deeper understanding of participant experiences during the experiment (Table 5)

Participant Reflections on AI-based Security Tools, MSC, and 5W1H Approach

I would say it is a very effective tool for structuring investigations. It serves as a simple but comprehensive framework for gathering and analysing information, ensuring that all angles are considered.
The template was highly effective in guiding my work through a structured and focused process. By presenting the material in a clear, step-by-step format, it allowed me to move through each section systematically without becoming overwhelmed. The question-based approach was particularly useful, as it helped limit the scope of each step and ensured that I stayed aligned with the key objectives of the task. Rather than having to guess what to include or how much detail to provide, the template prompted me with targeted questions that clarified my thinking and direction. An additional strength was the use of scales and options when making choices. This reduced ambiguity, as I didn't need to overthink or justify selections—choosing between predefined options made the process more efficient and less mentally taxing. Furthermore, the sub-questions that accompanied each main prompt helped me drill down into the specifics of what needed to be addressed. This layered structure gave me the confidence that I was asking the right questions and capturing the most relevant information. Overall, the template supported both clarity and focus, which significantly improved the quality and coherence of my work. It not only made the process more manageable but also more reflective, helping me critically assess each component before moving forward.
I found the 5W1H incident analysis approach is very effective. It ensures that each step is followed properly and provides a detailed overview of every aspect of the incident, allowing for proper data collection. It enables quick investigation within a short time frame, gathering data through Microsoft Security Copilot.
It would truly improve the day-to-day life of the SOC analyst. Firstly, an analyst (especially junior/inexperienced) could handle way more incidents in a day with the help of an AI tool, especially taking into account that usually there are a lot of false positives. I think the AI tool could really help with filtering out those cases and handling them way quicker! Secondly, it gives you more confidence, because as for me, since I am still going through the onboarding on my first SOC job, sometimes I am not sure where to look for the details I need, or if the details I got are enough for the assessment. But with an AI tool, it gives you all it can find, and you don't need to guess whether it is enough or whether you look in the correct place.
During my recent experience using Microsoft Security Copilot for incident investigations, I found it to be a helpful but still maturing tool. One of the strengths of the platform is its ability to quickly gather and present relevant data across multiple sources. In many cases, it can outline what happened in an incident, which is valuable for SOC teams aiming to save time and reduce manual analysis.
However, I noticed that while it can provide the facts, it often falls short in fully understanding the broader scenario or context behind the incidents. For example, it could identify and list multiple possible outcomes, but it didn't clearly classify them, such as separating true positives from false positives. Ideally, I expected it to make a judgment based on global threat intelligence or behavioural statistics, indicating, for instance, that an action is likely malicious because it's uncommon for normal users.
This lack of contextual judgment means that while the tool reduces the initial workload, human oversight is still essential. SOC analysts still need to validate and interpret the output, especially in nuanced situations. That said, with the right input and guidance, Security Copilot can be a strong assistant that accelerates investigations.

Participant Reflections on AI-based Security Tools, MSC, and 5W1H Approach
Looking forward, improvements in contextual awareness, clearer scenario mapping, and more intuitive prompts could greatly enhance its value. If these areas are addressed—and if pricing is accessible—it has real potential to become an everyday asset for SOC teams.
My overall experience with Microsoft Security Copilot AI-assisted security tools has been largely positive, especially in the context of streamlining incident investigations and augmenting the capabilities of the SOC team. These tools significantly reduce the time it takes to triage alerts, analyse raw data, and generate reports. By automating repetitive tasks and providing contextual summaries, MSC allows analysts to focus more on decision-making and remediation rather than manual data correlation.
I wouldn't consider AI tools as a great way of assistance. Instead, experienced SOC analysts would achieve better results. AI and especially MSC lack the ability to comprehend the incident.
It could be very beneficial for newer analysts to get hang of what is expected of them, but also dangerous in that regard they would not get too dependent of it, while using it, it was already getting easy for me to become overly reliant on it, and it felt wrong to question the judgement of AI as it can process a lot more information and look for it faster than I can. AI is not always correct and therefore shouldn't rely on it too much, YET. But eventually, when it is trained a lot more, I think SOC teams can be reduced to smaller teams for incident handling if AI gets very good in future.
I believe this tool will become every SOC Analyst's future friend, much like ChatGPT. It significantly reduces the time required to handle incidents and provides clear summaries that make it easy to understand the context of each incident. It very well analyses malicious script behaviour; it often requires a huge amount of time and collaboration among many SOC analysts to understand the malicious context. This tool shows great promise and will undoubtedly become an integral part of every SOC. The natural language-based response requires a bit of improvement.
MSC or any other competent AI-assist tool would give the SOC team a significant advantage in terms of efficiency, improved accuracy, and scalability.
My experience with MSC-AI-assist security tools for incident investigations has been a mix of promising capabilities and identifiable limitations. On the positive side, the integration of AI-driven assistance within the SOC team has significantly accelerated the initial triage process, improved visibility into incident timelines, and reduced manual workload for repetitive tasks. This allows analysts to focus more on strategic decision-making rather than low-level data parsing.
However, one of the key challenges lies in the tool's contextual understanding. While MSC-AI can extract and organise data effectively, it sometimes struggles with the logical correlation of events or nuanced human behaviour, especially in complex multi-stage attacks. This can lead to incomplete root cause analysis or over-reliance on manual validation.
Another concern is in areas like Confidence Score and Impact assessment, where AI's judgment may fall short without sufficient contextual grounding or historical awareness. These aspects often require human expertise to verify and adjust.
For future enhancement, improving the AI's reasoning capabilities and its ability to interpret human-like logic would be a major step forward. Furthermore, adding features such as adaptive learning from analyst feedback, tighter integration with MITRE ATT&CK mapping, and contextual enrichment using threat intelligence feeds could significantly enhance investigation quality.
Overall, MSC-AI-assist tools show strong potential in supporting SOC teams, but they would benefit from refinement in understanding context and reducing dependence on human validation.
Microsoft Security Copilot is a powerful addition to the modern security toolkit. It empowers professionals with AI-driven insights and operational efficiency but could reach its full potential with further enhancements to its natural language processing capabilities.
The MSC we used to be not very successful in executing prompts, so I had to work quite a lot on refining them. Although it provided more accurate results compared to ChatGPT in some cases, it struggled with logical understanding. Therefore, I believe it could be further improved, especially in terms of understanding human intent. Additionally, since manual reviews are still necessary, MSC is weak in determining the Confidence Score and Impact. Nevertheless, I improved the 5W1H prompt used and was able to generate effective and understandable results despite these challenges.

Table 5. Participant Reflections on AI-based Security Tools, MSC, and 5W1H Approach