

# A Hybrid AI-Graph Engine for the All-in-One Automated Security Operations Center (SOC)

Ghaylan Muhammad Fatih, Ghazi Akmal Fauzan, Abdillah Ahmad, and John (Jong Uk) Choi

*MarkAny*

1 Ssanglim BLDG, 286 Toegye-Ro, Jung-Gu, Seoul, Korea, 04615

{ghaylan, ghazi, abid}@ganeshait.com, juchoi@markany.com

**Abstract**—The contemporary cybersecurity landscape is characterized by a fundamental asymmetry: highly automated offensive operations juxtaposed against defensive operations that remain heavily reliant on human analysts. This disparity introduces significant risk, particularly for under-resourced Small and Medium Enterprises (SMEs). To mitigate this challenge, this paper proposes an “All-in-One System for Automating SOC,” utilizing a four-layer architecture that integrates detection, interpretation, and response. This work presents the design and implementation of the system’s core interpretation engine: a hybrid AI-Graph API service. The engine receives raw telemetry from the Detection Layer and executes a two-stage analysis: (1) A Large Language Model (LLM) service analyzes raw data (e.g., process trees) to infer relevant MITRE ATT&CK Tactics, Techniques, and Procedures (TTPs). (2) The identified TTP is subsequently employed to query a Knowledge Graph (Neo4j), retrieving critical enrichment data (e.g., Threat Actors, Mitigations) and associated compliance mappings (e.g., NIST, CIS). The engine generates a structured JSON output, which serves as the designated input for the Playbook Synthesis Layer. By automating the manual interpretation phase, this engine transforms raw EDR alerts into actionable response playbooks, thereby bridging the critical gap between threat detection and automated response.

**Index Terms**—Automated SOC, Knowledge Graph, API, EDR, SOAR, FastAPI, LLM, Incident Response, MITRE ATT&CK, Compliance

## I. INTRODUCTION

A fundamental asymmetry defines the contemporary cybersecurity landscape. Offensive operations, including ransomware and credential theft, are increasingly automated and commoditized [6]. Conversely, defensive operations remain heavily reliant on manual human processes for analysis and response. This disparity results in an unsustainable operational posture for many organizations, particularly those lacking dedicated 24/7 security personnel and the requisite budget [5].

### A. The Need for Automation in Modern Security Operations

Personnel working in Security Operations Centers (SOCs) face significant operational challenges. They are frequently overwhelmed by high volumes of security telemetry without receiving the actionable intelligence required for rapid response. This influx of granular alerts from disparate security tools leads to analyst fatigue [5], increasing the probability that critical threats are overlooked. The primary issue is not a data deficit, but rather the system’s inability to correlate low-level computer events with broader adversarial

behaviors or organizational security policies. This inefficient process results in prolonged containment times—averaging 277 days—providing adversaries with substantial dwell time to achieve their objectives [4].

Security automation initially focused on centralized data collection via SIEM platforms and the use of Cybersecurity Knowledge Bases (KB) to assist manual interpretation [1]. However, to significantly reduce response times, defensive systems must transition to a model where detection is directly coupled with automated response. This requires shifting the initial stages of triage and interpretation from human analysts to automated computational processes.

### B. From Information to Action

The core objective of this paper is to transition security systems from passive information delivery to active remediation. Previous research focused on establishing database connections that provided analysts with threat context [1]. While beneficial, this approach resulted in static information for human consumption. This paper introduces a system that transforms this intelligence into an active component of the security infrastructure.

For automation to be effective, threat intelligence must be converted from human-readable reports into machine-executable commands. The architecture and Application Programming Interface (API) presented in this work provide the framework for this conversion. By establishing rigorous protocols for communication between the detection system and the knowledge base, a threat query is transformed into a direct remediation command, allowing the system to initiate self-defense protocols. This reduces the time-to-remediation from hours or days to mere seconds.

### C. Bridging Open-Source EDR and Automation

This research addresses a critical gap in security architecture: the integration of open-source Endpoint Detection and Response (EDR) tools with automated systems. Commercial automation platforms often rely on rigid, manually authored playbooks that lack the flexibility required for emerging threats [1]. Conversely, robust open-source EDR tools like OpenEDR provide granular telemetry but lack real-time integration with decision-making frameworks [12].

Typically, these tools export data to centralized storage (such as an ELK stack) for retrospective analysis. This batch-

processing approach is unsuitable for the low-latency, direct API calls required for immediate threat response. This paper presents a comprehensive design to resolve this connectivity gap, enabling EDR tools to function as integrated components of an automated defense ecosystem.

## II. BACKGROUND AND RELATED WORK

### A. The Threat-Compliance Knowledge Base

The foundation for this paper is our previous work, which presented a graph-based knowledge base designed to bridge the operational gap between threat intelligence and compliance management [1]. The core problem was that critical security domains operate in silos. Security teams track adversary behaviors using frameworks like MITRE ATT&CK [3], while governance teams manage defensive posture using frameworks like the NIST Cybersecurity Framework (CSF) [2].

Our solution was a Neo4j graph database that models both domains in a single, queryable resource. This was achieved through a generalized, self-referencing schema for Control nodes and an LLM-powered “semantic bridge” which automates the creation of [:MITIGATES] relationships.

### B. The EDR-to-SOAR Automation Gap

Modern security defense relies on a “detect and respond” paradigm, primarily driven by EDR and SOAR platforms. EDR tools are highly effective at monitoring endpoint activity and generating alerts mapped to TTPs from the ATT&CK framework [10]. However, a critical “automation gap” exists between EDR detection and effective SOAR-driven response. An EDR alert, such as “T1059.001 detected,” lacks essential business and security context. A SOAR playbook must answer: Is this part of a known campaign? What compliance mandates are at risk? This gap is currently filled by a human analyst, who is prone to alert fatigue [5].

### C. Knowledge Graphs and AI in Security Operations

The use of AI and knowledge graphs in cybersecurity is an active area of research [7]. Works such as the Unified Cyber Ontology (UCO) create a common vocabulary for disparate security domains [8]. Other research has demonstrated the power of graph databases for attack path analysis [9]. However, many existing approaches focus on static, offline analysis. Our work differentiates itself by operationalizing the knowledge graph in a real-time, event-driven architecture, combining an “AI Analyst” (LLM) and a “Knowledge Brain” (Graph KB).

## III. PROPOSED SYSTEM ARCHITECTURE

### A. End-to-End Operational Workflow

The proposed architecture establishes a seamless, automated workflow connecting threat detection to knowledge-based response. The process operates as follows:

- 1) **Detection:** The EDR Agent collects raw telemetry concerning system actions (e.g., process creation) [12].
- 2) **Threat Correlation:** The Correlation Engine analyzes the telemetry using heuristic rules to identify potential adversarial techniques.

- 3) **API Request:** The EDR Management Server initiates a direct API call to the API Bridge with a structured JSON payload.
- 4) **Knowledge Base Query:** The KB translates the request into a Cypher query for the Neo4j database.
- 5) **Response Identification:** The KB identifies associated “Control” nodes containing executable commands to remediate the threat.
- 6) **API Response:** The KB returns a structured JSON response containing contextual details and recommended countermeasures.
- 7) **Execution:** The EDR Agent executes the remediation command to neutralize the threat.

This entire sequence is designed for sub-minute execution. Figure 1 illustrates the interaction between these components.

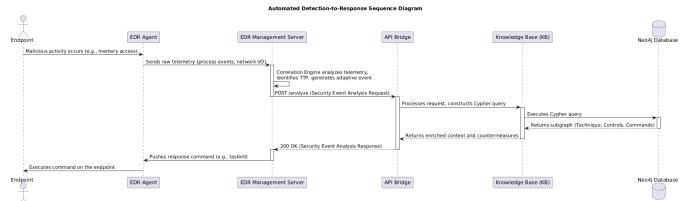


Fig. 1. Automated Detection-to-Response Sequence Diagram.

### B. Component Architecture

The system is composed of modular components. Figure 2 illustrates the system’s structural organization.

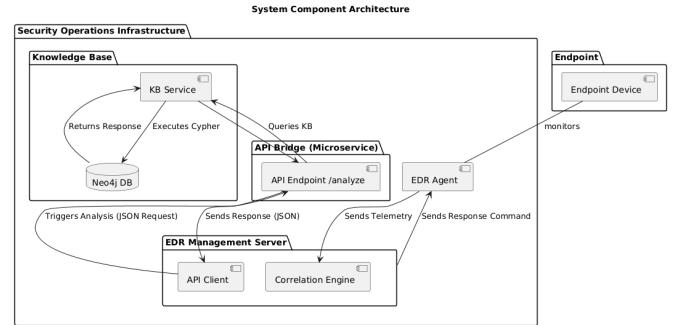


Fig. 2. System Component Architecture.

1) **The EDR System:** The architectural foundation is the EDR, built upon the OpenEDR framework [12]. The EDR Agent provides high-fidelity telemetry, while the Correlation Engine translates raw data into high-level security alerts (e.g., “Credential Stealing with Mimikatz”).

2) **The API Bridge:** The API Bridge serves as the orchestration hub, providing a secure interface between the EDR and the Knowledge Base. The Request Format is strictly structured (Table I), with the `process_parent_tree` providing critical forensic context for analysis.

The Response Format (Table II) is structured for both machine execution and human oversight, encompassing `technique_info`, `context`, and `countermeasures`.

TABLE I  
SECURITY EVENT ANALYSIS REQUEST DATA

| Field Name          | Data Type | Description  |
|---------------------|-----------|--|
| adaptive_event_type | String    | Threat classification determined by the EDR (e.g., “Credential Stealing”). |
| base_event_type     | String    | Primitive system event (e.g., “Virtual Memory Access”).                    |
| component           | String    | Source system generating the telemetry.                                    |
| process_hash        | String    | SHA1 hash of the subject binary.   |
| process_parent_tree | Array     | Hierarchical process lineage of the subject program.                       |

TABLE II  
SECURITY EVENT ANALYSIS RESPONSE DATA

| Field Name      | Data Type | Description  |
|-----------------|-----------|--|
| incident_id     | String    | Unique identifier for the security event.              |
| technique_info  | Object    | Metadata regarding the MITRE ATT&CK technique.         |
| context         | Object    | Enrichment data from the KB (Threat Groups, Software). |
| countermeasures | Array     | Executable actions/commands for remediation.           |

### C. The Knowledge Base as a Real-Time Decision Maker

The KB utilizes the dual-stage design established in prior work [1]. Upon receiving a request, the KB constructs a Cypher query for the Neo4j database. The `adaptive_event_type` is mapped to a MITRE Technique ID to initiate graph traversal.

```
// Information from the API request is used here
WITH "T1003.001" AS techniqueId
// Find the technique and the controls that stop it
MATCH (tech:Technique {id: techniqueId})
// The MITIGATES link connects threat and defense
MATCH (ctrl:Control)-[:MITIGATES]-(tech)
// Only get controls with executable commands
WHERE ctrl.command IS NOT NULL
// Get extra information
OPTIONAL MATCH (tech)<--(grp:Group)
OPTIONAL MATCH (tech)<--(sw:Software)
RETURN
  tech.id AS techniqueId,
  tech.name AS techniqueName,
  collect(DISTINCT grp.name) AS associatedGroups,
  collect(DISTINCT sw.name) AS associatedSoftware,
  collect({
    category: ctrl.category,
    action: ctrl.description,
    command: ctrl.command
  }) AS countermeasures
```

Listing 1. Cypher Query Example

### D. Core API Logic

The core logic of the system is visualized in the flowchart below. It orchestrates authentication, LLM inference, and graph traversal.

The responsibilities of each component are summarized in Table III.

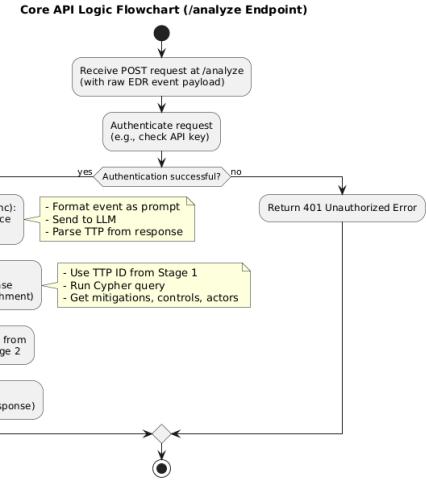


Fig. 3. Flowchart of the Core API Logic for the /analyze Endpoint.

TABLE III  
COMPONENT RESPONSIBILITIES

| Component          | Primary Role     | Key Activities                                   |
|--------------------|------------------|--|
| EDR Agent          | Data Collection  | Telemetry monitoring, command execution [12].    |
| Correlation Engine | Triage           | Behavioral analysis, machine learning [12].      |
| API Bridge         | Orchestration    | Microservice management, API gateway, OAuth 2.0. |
| Knowledge Base     | Decision Support | Neo4j graph traversal, MITRE ATT&CK mapping.     |

## IV. IMPLEMENTATION

### A. Technology Stack

The system is implemented as a containerized microservice using modern Python frameworks:

- FastAPI:** Utilized for native `async/await` support, which is critical for handling non-blocking I/O-bound calls to the LLM and Neo4j.
- Neo4j 5.x:** A graph database selected for modeling complex security relationships and performing high-speed multi-hop traversals [9].
- LLM Service:** Acts as the “AI Analyst,” serializing the `SecurityEventRequest` into a prompt to infer corresponding TTPs.
- Docker & NGINX:** Ensures environment reproducibility and provides reverse proxy services.

### B. Algorithms

The `/analyze` endpoint functions as an asynchronous orchestrator, as detailed in Algorithm 1.

The `Get_Graph_Enrichment` function executes a multi-hop traversal of the Neo4j graph, gathering relevant enrichment data in a single transaction (Algorithm 2).

## V. EVALUATION AND METRICS

To validate the efficacy and performance of the interpretation engine, we conducted tests evaluating two primary areas:

### Algorithm 1 Asynchronous Triage and Enrichment

```

Require: event: SecurityEventRequest, user: AuthToken
Ensure: response: SecurityEventResponse
1: Function ANALYZE_SECURITY_EVENT(event, user)
2:   // Stage 1: Call LLM to identify TTP from raw event
3:   inferred_ttp_data ← await Identify_TTP(event);
4:   ttp_id ← inferred_ttp_data.id;
5:   // Stage 2: Call Graph DB to enrich TTP
6:   enrichment_data ← await Get_Graph_Enrichment(ttp_id);
7:   // Stage 3: Combine and return response
8:   response ← Combine(inferred, enrichment);
9:   return response;
10: Function Identify_TTP(event)
11:   prompt ← Format_Event_As_Prompt(event);
12:   llm_result ← LLM_SERVICE.generate(prompt);
13:   ttp_data ← Parse_LLM_Result(llm_result);
14:   return ttp_data;

```

### Algorithm 2 Graph Enrichment Traversal

```

Require: ttp_id: String
Ensure: enrichment_data: Enrichment Data
1: Function Get_Graph_Enrichment(ttp_id)
2:   // Initialize empty result lists
3:   actors, software, mitigations, compliance ← []
4:   t ← GRAPH.FindNode(Technique, id = ttp_id)
5:   // Find related threat actors and software
6:   foreach g in GRAPH.FindNodes(Group, target = t)
    do
7:     ADD g TO actors;
8:   // Find strategic mitigations
9:   foreach m in GRAPH.FindNodes(Mitigation,
    source = t) do
10:    ADD m TO mitigations;
11:   // Find compliance controls (Semantic Bridge)
12:   foreach c in GRAPH.FindNodes(Control, target = t)
    do
13:     f ← GRAPH.FindNode(Framework, target = c);
14:     ADD {control : c, framework : f} TO
        compliance;
15:   enrichment ← {actors, software,
16:                 mitigations, compliance};
17:   return enrichment;

```

(1) a qualitative demonstration of the end-to-end analysis pipeline, and (2) a quantitative evaluation of latency against established industry benchmarks.

#### A. Use Case Demonstration: PowerShell TTP Inference

To demonstrate the system's capability, we submitted a real-world EDR event to the `/analyze` endpoint. This event simulates suspicious activity: a PowerShell process (`powershell.exe`) spawned from an unexpected parent (`svchost.exe`) under high-privilege (SYSTEM) credentials. The raw event payload is visualized in Figure 4.

```

Event 2
- Back to events list
Source Wazuh Event
{
  "timestamp": "2025-11-05T14:08:28.304+0700",
  "rule": "idle_timeout",
  "level": 5,
  "id": "68702",
  "parent": "1",
  "mail": false,
  "groups": [
    "windows",
    "windows_application"
  ],
  "agent": {
    "id": "00110000000000000000000000000000",
    "ip": "192.168.1.25",
    "name": "markiny-3-B658H-K"
  },
  "decoder": {
    "name": "windows_eventchannel"
  },
  "data": {
    "data": {
      "data": {
        "providerName": "VSS",
        "eventId": "8224",
        "version": "1.0",
        "level": "4",
        "task": "0",
        "opcode": "0",
        "keywordId": "8000000000000000",
        "sysEventId": "2025-11-05T04:21:09.37629652",
        "eventRecordId": "11151",
        "processHash": "2836",
        "threadId": "0",
        "channel": "Application",
        "computation": "DESKTOP-T0Q55TB",
        "severityValue": "Information",
        "message": "\u0022The VSS service is shutting down due to idle timeout.\u0022"
      }
    }
  }
}

LLM Decision Result
{
  "incident_id": "VSS-DESKTOP-T0Q55TB-1762326599",
  "technique_info": {
    "id": "T1059"
  }
}

```

Fig. 4. Raw Security Event Input Payload (PowerShell Suspicion).

The API successfully processed this event and returned a comprehensive SecurityEventResponse (Figure 5).

```

{
  "events": [
    {
      "binary": "base64-encoded-binary-data-here"
    }
  ],
  "location": "EventChannel"
}

LLM Decision Result
{
  "incident_id": "VSS-DESKTOP-T0Q55TB-1762326599",
  "technique_info": {
    "id": "T1059"
  }
}

```

Fig. 5. Enriched JSON Response Output.

The output confirms the successful execution of the two-stage pipeline:

- 1) **Stage 1 (LLM Inference):** The engine analyzed the input and correctly identified the threat as TTP T1059.001 (PowerShell).
- 2) **Stage 2 (Graph Enrichment):** The engine enriched the TTP, identifying "FIN7" as the associated actor, "Execution Prevention" as a mitigation strategy, and "CIS Controls 10.7" as the compliance mapping.
- 3) **Response Generation:** The engine generated a `recommended_playbook` containing immediate remediation steps.

### B. Performance Evaluation: Latency

We conducted 100 stress tests against the `/analyze` endpoint to measure the system's end-to-end interpretation time. The average response time for the complete two-stage analysis was **24.97 seconds**.

To contextualize this performance, we benchmarked the system against the industry-standard “1/10/60 Challenge” framework [11]. This framework dictates that organizations should detect an intrusion within 1 minute, **investigate** within 10 minutes, and remediate within 60 minutes.

Our engine specifically targets the “10-minute” investigation window—the phase where an analyst must interpret the alert’s context. As shown in Table IV, our engine reduces investigation time from the 10-minute (600-second) benchmark to approximately 25 seconds. This represents an improvement of over 95.8%, enabling investigations to occur at machine speed.

TABLE IV

PERFORMANCE COMPARISON: 1/10/60 BENCHMARK VS. AUTOMATED

| Process            | Industry Benchmark (1/10/60) | Automated Engine     |
|--------------------|------------------------------|----------------------|
| Phase              | Investigation (The “10”)     | Stage 1 + Stage 2    |
| Time               | 10 Minutes (600s) [11]       | <b>24.97 seconds</b> |
| Output             | Analyst Understanding        | Structured JSON      |
| <b>Improvement</b> | <b>&gt;95.8% Reduction</b>   |                      |

## VI. CONCLUSION AND FUTURE WORK

This paper has detailed the design and implementation of a hybrid AI-graph interpretation engine, the core of the “All-in-One Automated SOC” architecture. By utilizing a two-stage hybrid AI architecture, this work automates the manual analyst workflow. The FastAPI service orchestrates this analysis in sub-minute time, providing essential interconnectivity.

Future work will focus on implementing the remaining layers:

- **Playbook Synthesis Layer (Layer 3):** Developing a service to consume JSON output and generate standard CACAO playbooks.
- **Playbook Operations Layer (Layer 4):** Creating a user interface to allow non-expert operators to execute playbooks.
- **Bi-Directional API:** Extending the engine to allow the Operations Layer to update the graph, enabling a continuous learning loop.

Additionally, future optimizations will address the methodology of Stage 1. Instead of relying solely on an LLM to determine the corresponding ATT&CK TTP ID, future iterations will explore using a vector database to create a search-based retrieval method, aiming for higher accuracy, consistency, and traceability.

## REFERENCES

- [1] G. M. Fatih, G. A. Fauzan, and J. Choi, “A Graph-Based Knowledge Base for Integrating Adversary Behavior with Compliance Frameworks,” 2025.
- [2] National Institute of Standards and Technology, “The NIST Cybersecurity Framework (CSF) 2.0,” U.S. Department of Commerce, Feb. 2024.
- [3] The MITRE Corporation, “MITRE ATT&CK,” [Online]. Available: <https://attack.mitre.org>. [Accessed: Sep. 15, 2025].
- [4] IBM, “Cost of a Data Breach Report 2025,” Ponemon Institute, 2025.
- [5] T. Olsik, “The Life and Times of Cybersecurity Professionals 2023,” Enterprise Strategy Group (ESG), 2023.
- [6] Verizon, “2025 Data Breach Investigations Report,” Verizon Enterprise, 2025.
- [7] L. F. Sikos, *AI in Cybersecurity*. Springer, 2020.
- [8] A. S. Boddy, U. Jaimini, N. L. Skarlat, and J. R. Santos, “The Unified Cyber Ontology,” The MITRE Corporation, 2016.
- [9] S. Noel and S. Jajodia, “Understanding complex network attack graphs,” in *Graph Theoretic and Topological Approaches to Security, Trust, and Resilience*, Cham: Springer, 2014, pp. 1–26.
- [10] B. E. Strom et al., “MITRE ATT&CK design and philosophy,” The MITRE Corporation, 2020.
- [11] CrowdStrike, “The 1-10-60 Minute Challenge: A Framework for Stopping Breaches Faster,” [Online]. Available: <https://www.crowdstrike.com/en-us/resources/crowdcasts/the-1-10-60-minute-challenge-a-framework-for-stopping-breaches-faster/>. [Accessed: Nov. 27, 2025].
- [12] OpenEDR, “Open Source Endpoint Detection and Response (EDR),” [Online]. Available: <https://www.openedr.com/>. [Accessed: Oct. 30, 2025].