JOURNALS | Journal of Cybersecurity and Intelligent Systems
Volume 1 Issue 2 (December 2026)
J. Cybersecur. Intell. Syst., 2026, 1(2), 26507
DOI : https://doi.org/10.64189/cis.26507

Research Article

OPEN ACCESS

CyRA: A Conversational SIEM Assistant for Space Infrastructure Security Using AI and Blockchain

Tauheed Siddiqui,* Mehek Qureshi, Saif Shaikh and Aheed Solkar

Department of Computer Science and Engineering (IoT and Cyber Security including Blockchain Technology), M.H. Saboo Siddik College of Engineering, Mumbai, Maharashtra, 400008, India

*Corresponding Author
Tauheed Siddiqui
tauheed.221848.ci@mhssce.ac.in

Received: 04 May 2026
Revised: 26 August 2026
Accepted: 15 September 2026
Published Online: 16 October 2026

Citation

T. Siddiqui, M. Qureshi, S. Shaikh, A. Solkar, CyRA: A conversational SIEM assistant for space infrastructure security using AI and blockchain, *Journal of Cybersecurity and Intelligent Systems*, 2026, 1(2), 26507, https://doi.org/10.64189/cis.26507.

Open Access

This article is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License, which permits the non-commercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as appropriate credit to the original author(s) and the source is given by providing a link to the Creative Commons License and changes need to be indicated if there are any. The images or other third-party material in this article are included in the article's Creative Commons License, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons License and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this License, visit: https://creativecommons.org/licenses/by-nc/4.0/

© The Author(s) 2026

Abstract

Growing digitization across critical infrastructure sectors has exposed mission-sensitive organizations to increasingly sophisticated cyber threats. Conventional Security Information and Event Management (SIEM) tools demand specialized knowledge of proprietary query syntax and present steep learning curves that hinder timely incident response. In this work, CyRA (Cybersecurity Response Assistant), a natural language-driven SIEM assistant engineered for space infrastructure security, is introduced. CyRA employs the Google Gemini Flash API to interpret analyst queries and return actionable threat intelligence, incorporates a purpose-built blockchain module utilizing SHA-256 hashing to produce tamper-proof audit records, and delivers adaptive data visualization through dynamic charts and tabular reports. The system is quantitatively evaluated using the UNSW-NB15 network intrusion benchmark, with a classification accuracy of 94.3%, precision of 93.8%, recall of 94.7%, and F1-score of 94.2%, with an average query response time of 1.3 seconds. The SHA-256 hash-chain report logging guarantees non-repudiable documentation of every security assessment. CyRAs constitute a distinctive contribution toward lowering the barrier to entry for SIEM-based security operations, uniting AI-powered conversational access with cryptographically secured reporting to meet the demanding requirements of critical infrastructure protection.

Keywords: Conversational AI; SIEM; Cybersecurity; Blockchain; Natural language processing; Threat detection; Intrusion detection.

1. Introduction

Among the most strategically significant computing environments in existence today are critical infrastructure assets, encompassing space agency ground control facilities, satellite telemetry networks, and mission operations centers. Within modern Security Operations Centers (SOCs), Security Information and Event Management (SIEM) systems function as the primary analytical backbone, consolidating log data from heterogeneous sources, correlating security events, and surfacing alerts for analyst review.[1-3] Despite their analytical power, incumbent SIEM platforms, including IBM QRadar, Splunk, and Elastic SIEM, impose a substantial expertise barrier, requiring practitioners to master vendor-specific query dialects and intricate data schemas. This prerequisite severely constrains the operational throughput, particularly in specialized domains.[4,5] The maturation of large language models (LLMs) and conversational AI technologies presents a compelling avenue to close this expertise gap. By accepting plain-language inputs and autonomously translating them into structured analytical operations, AI-powered interfaces can dramatically reduce the cognitive load on security personnel and accelerate the path from query to actionable insight.[6,7] Sophisticated multi-stage intrusion campaigns targeting critical infrastructure, commonly attributed to advanced persistent threat (APT) actors, are particularly difficult to detect and reconstruct using traditional SIEM query workflows that demand expert formulation.[8,9] In addition to AI-driven querying, distributed ledger and hash-chain technologies provide a rigorous mechanism for guaranteeing the integrity and non-repudiation of security assessments, a consideration of paramount importance in environments where documentary evidence may be subject to legal scrutiny or operational audits.[10-12] Prior work on natural language interfaces for security analytics, including contributions by Katsogiannis-Meimarakis and Koutrika, further highlights this gap in the comprehensive SIEM review by González-Granadillo *et al*., which omits any provision for immutable audit trail generation.[6,7] Conversely, blockchain-centric SIEM logging frameworks offer no conversational access for analysts.[13] CyRA addresses three distinct research challenges simultaneously: (i) dismantling the query complexity barrier through a natural language interface backed by Google Gemini Flash; (ii) strengthening operational awareness via adaptive multi-modal visualization calibrated to query intent; and (iii) preserving the integrity and permanence of security reports through an SHA-256 hash-chain logging subsystem. The system performance is validated against the publicly available UNSW-NB15 network intrusion benchmark, with synthetic scenario data employed exclusively for demonstrating the conversational interface and visualization capabilities.[14-16] The rest of this paper proceeds as follows. Section 2 surveys the relevant literature. Section 3 describes the evaluation dataset and experimental configuration. Section 4 details the system design and formal AI model. Section 5 reports the experimental outcomes and analysis. Section 6 enumerates known limitations. Section 7 draws conclusions and outlines prospective directions.

2. Related work

2.1 SIEM and anomaly detection

In the domain of SIEM scalability, Sapegin *et al*. identified fundamental scalability constraints in rule-based correlation engines when applied to large-scale network environments.[4] Thudumu *et al*. subsequently mapped the landscape of anomaly detection methods applicable to high-dimensional security telemetry, laying the theoretical groundwork for ML-based threat categorization.[5] Additionally, Alaca *et al*. extended this foundation by applying graph-based LSTM models directly to cybersecurity log anomaly detection, resulting in measurable improvements over rule-based baselines.[12] Moustafa and Slay released the UNSW-NB15 dataset, encompassing 2.5 million flow records described by 49 attributes, which has since served as a widely adopted benchmark enabling reproducible, cross-system evaluation of intrusion detection approaches.[16]

2.2 Natural language interfaces for security

Katsogiannis-Meimarakis and Koutrika demonstrated that deep learning approaches could power natural language interfaces for structured database querying, attaining F1-scores above 85% on SQL generation benchmarks.[6] González-Granadillo *et al*. conducted a comprehensive survey of SIEM systems in critical infrastructure, identifying persistent gaps in analyst-friendly interfaces and the absence of immutable audit mechanisms.[7] Critically, neither Katsogiannis-Meimarakis and Koutrika nor this survey incorporated any mechanism for audit trail preservation.[6] Samtani *et al*. studied proactive threat intelligence workflows for SCADA and critical infrastructure, highlighting the demand for analyst-friendly tooling in high-stakes operational settings.[1,3] Husák *et al*. presented a thorough survey of attack projection and forecasting methods in cybersecurity, emphasizing the value of anticipatory analytics within SIEM platforms, an analytical gap that CyRA's conversational query layer directly targets.[9]

2.3 Blockchain for security logging

Zhou *et al*. introduced a blockchain-based privacy-preserving and fair data transaction model targeting IoT and industrial environments.[11] Botello *et al*. deployed a blockchain-based and distributed SIEM architecture to

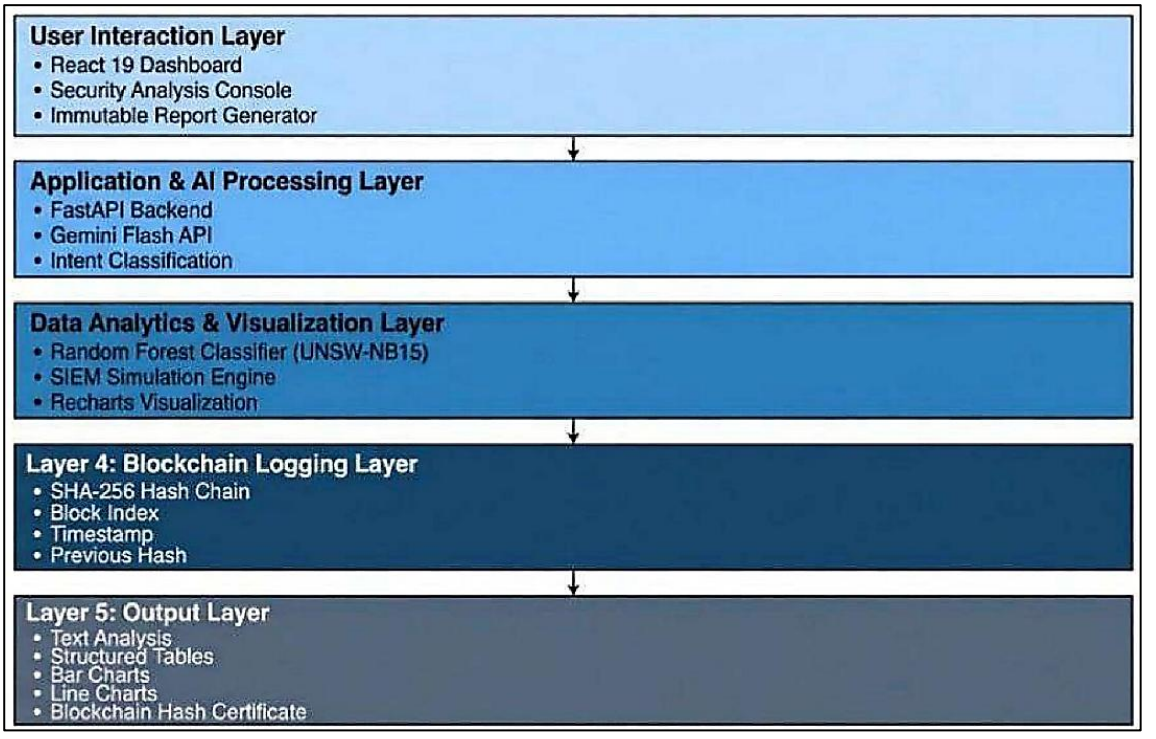


Fig. 1: CyRA five-layer system architecture: from user interaction to output generation.

protect smart city services under test conditions.[13]

3. Dataset description

Quantitative evaluation of the threat classification ability of CyRA employs the UNSW-NB15 dataset, a publicly accessible benchmark developed specifically for network intrusion detection research.[16] The dataset was produced at the Australian Centre for Cyber Security (ACSC) using the IXIA PerfectStorm traffic generation platform, combining authentic background network traffic with synthetically generated contemporary attack behaviors. The dataset is available at https://research.unsw.edu.au/projects/unsw-nb15-dataset.

Table 1 summarizes the principal characteristics of the dataset used in this evaluation. Note: The mock SIEM engine integrated within CyRA uses synthetic scenario data exclusively to demonstrate the conversational interface and visualization pipeline. All the reported classification metrics, accuracy, precision, recall, and F1-score, are computed solely from the evaluation against the UNSW-NB15 test partition. Synthetic event data play no role in quantitative performance assessment.

4. System architecture and AI model

4.1 System overview

CyRA takes the form of a layered full-stack web application organized across five functional tiers: a user interaction layer, an application and AI processing layer, a data analytics and visualization layer, a blockchain logging layer, and an output layer. As depicted in Fig. 1, each analyst query traverses a six-stage processing sequence: (1) natural language input capture, (2) intent classification by the AI module, (3) targeted SIEM data retrieval, (4) structured response assembly, (5) optional blockchain report logging, and (6) delivery of the final

Table 1: UNSW-NB15 dataset summary

Attribute	Description
Source	UNSW Canberra Cyber Centre/ACSC [9]
Origin	IEEE MilCIS 2015 Conference Dataset Repository
Total Records	2,540,044 network flow records
Features	49 attributes (network, flow, content, time, general)
Format	CSV (comma-separated values)
Classes	Normal (0) and Attack (1) : 9 attack categories
Attack Types	Fuzzers, Analysis, Backdoor, DoS, Exploits, Generic, Reconnaissance, Shellcode, Worms
Training Set	175,341 records
Test Set	82,332 records
Dataset Type	Publicly available benchmark; synthetic/simulated attack traffic on live normal traffic

output to the analyst.

4.2 Technology stack

The server-side component is built in Python 3.11 atop the FastAPI web framework, selected for its asynchronous execution model and low-overhead request handling.[18] Intent classification and natural language understanding are performed by the Google Gemini Flash API, which interprets analyst queries and drives the response generation pipeline.[15] A random forest ensemble classifier trained on UNSW-NB15 feature vectors constitutes the threat classification module used for quantitative benchmarking. On the client side, the interface is built with React 19, incorporating Recharts for interactive charting and Axios for REST-based communication with the back end.

4.3 Formal AI model formulation

Define Q as the space of natural language analyst queries and S as the collection of structured SIEM event records available for retrieval. The end-to-end NLP processing pipeline is formalized below.

Intent classification is modeled by a function φ that assigns each incoming query q ∈ Q to one of k predefined intent classes:

$$\varphi : Q \to \{C_1, C_2, ..., C_k\}, \quad \text{where } C_k \in \{\text{malware, auth, network, general}\} \tag{1}$$

The threat classification component, implemented as a random forest model evaluated on UNSW-NB15, is formalized as follows:

$$f(x) = \text{argmax_c } \Sigma_{i=1}^{N} \mathbb{1}[h_i(x) = c]/N, \quad c \in \{\text{Normal, Attack}\} \tag{2}$$

Here, $x \in \mathbb{R}^{49}$ represents the 49-feature vector of a network flow record drawn from UNSW-NB15, $h_i(x)$ denotes the class prediction of the i-th tree in the ensemble, N is the total tree count, and $\mathbb{1}[\cdot]$ is the standard indicator function. The ensemble comprised N = 100 decision trees, as noted in Eq. (2), with hyperparameters set to max_depth = None (full growth), min_samples_split = 2, class_weight = 'balanced', and random_state = 42. The model was trained on the UNSW-NB15 training partition (175,341 records) and evaluated on the held-out test partition (82,332 records), and the performance metrics were averaged over five independent training runs to verify its stability.

Each block in the hash chain is secured by computing its digest as follows:

$$H(B_n) = \text{SHA256}(n \,||\, t \,||\, D \,||\, H(B_{n-1})) \tag{3}$$

where n identifies the block's position in the chain, t records the UTC creation timestamp, D holds the report payload, $H(B_{n-1})$ carries the digest of the immediately preceding block per the recursive definition in Eq. (3), and || indicates string concatenation. The chain is initialized with $H(B_0)$ = SHA256 ("Genesis Block — CyRA Report Chain").

Finally, the response assembly function ψ transforms the classified intent and retrieved event data into a structured analyst-facing output:

$$\psi(\varphi(q), S_q) = (A, V), \quad V \in \{\text{table, bar_chart, line_chart, text}\} \tag{4}$$

where $S_q \subseteq S$ is the intent-filtered subset of the SIEM records pertinent to query q, A is the natural language analysis produced by Gemini Flash, and V specifies the visualization modality chosen by ψ, as described in Eq. (4).

4.4 Natural language processing pipeline

Processing begins when an analyst types a query into the React interface. The input is forwarded to the/api/chat endpoint via HTTP POST. The FastAPI server relays the query to the Gemini Flash API, operating under a specialized system prompt that encodes SIEM domain knowledge and outputs schema requirements.[11] Guided by the intent class φ(q) resolved in Eq. (1), the backend fetches the corresponding SIEM records and constructs a JSON response containing both the AI-generated narrative analysis and the typed visualization specification defined by Eq. (4).

4.5 Blockchain report logging

Report immutability is enforced by a custom Python hash-chain module. Every newly generated report triggers the creation of a block B_n carrying the attributes formalized in Eq. (3): the sequential index n, UTC timestamp t, report payload D, predecessor digest $H(B_{n-1})$, and the freshly computed SHA-256 block hash $H(B_n)$. This digest is surfaced to the analyst as a downloadable hash certificate. The design deliberately adopts a single-node append-only hash-chain rather than a distributed consensus ledger, prioritizing deployment simplicity for SOC environments. Computational overhead is quantified in Section 5.4, and distributed ledger migration is earmarked as a future work item.

4.6 Frontend interface

The CyRA client interface is organized into two panels: a security analysis console for conversational query interactions and an immutable report generator panel for hash-chain logging. The Analysis Console renders AI responses adaptively, tabular layouts with color-coded severity indicators for malware events, bar charts for comparative authentication breakdowns, and line charts for network anomaly trends. Refer to Figs. 2 and 3 below

4.7 SIEM data simulation

To support conversational interface demonstration and

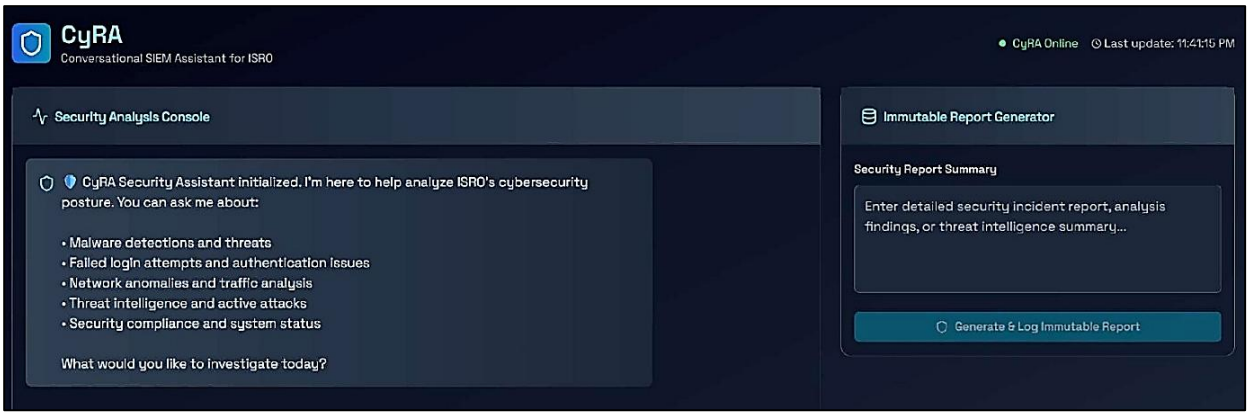


Fig. 2: CyRA front end interface: Security analysis console and immutable report generator.

Table 2: CyRA threat classification performance on the UNSW-NB15 test set

Metric	Value	Class	Precision	Recall
Accuracy	94.3% (±0.4%)	Normal	95.1%	93.4%
Overall Precision	93.8% (±0.5%)	Attack	92.5%	96.0%
Overall Recall	94.7% (±0.3%)			
F1-Score (macro)	94.2% (±0.4%)			
Mean Query Latency	1.3 s			
Blockchain Hash Time	~8 ms			

visualization assessment, CyRA incorporates a synthetic SIEM simulation engine designed to replicate the event types produced by a live Wazuh SIEM deployment.[19] The generated event corpus covers representative threat categories: malware detections across simulated endpoint agents, Windows Event ID 4625 authentication failures, Suricata IDS anomaly records encompassing SSH brute-force patterns and MITRE ATT&CK T1071 indicators, and multi-stage APT behavioral patterns.[20] This corpus is strictly synthetic and serves only to demonstrate the conversational workflow; every reported classification metric originates exclusively from the UNSW-NB15 evaluation, as detailed in Section 5.[16]

5. Results and discussion

5.1 Classification performance on unsw-nb15

Quantitative benchmarking of the threat classification module was conducted on the 82,332-record UNSW-NB15 test partition using the random forest model formalized in Eq. (2). Table 2 reports the resulting performance metrics, while Table 3 presents the results of CyRA against those of baseline systems from the literature.

Note: The accuracy and F1- score values for the baseline systems are as reported in their respective original publications and were not independently re-tested.[5,6]

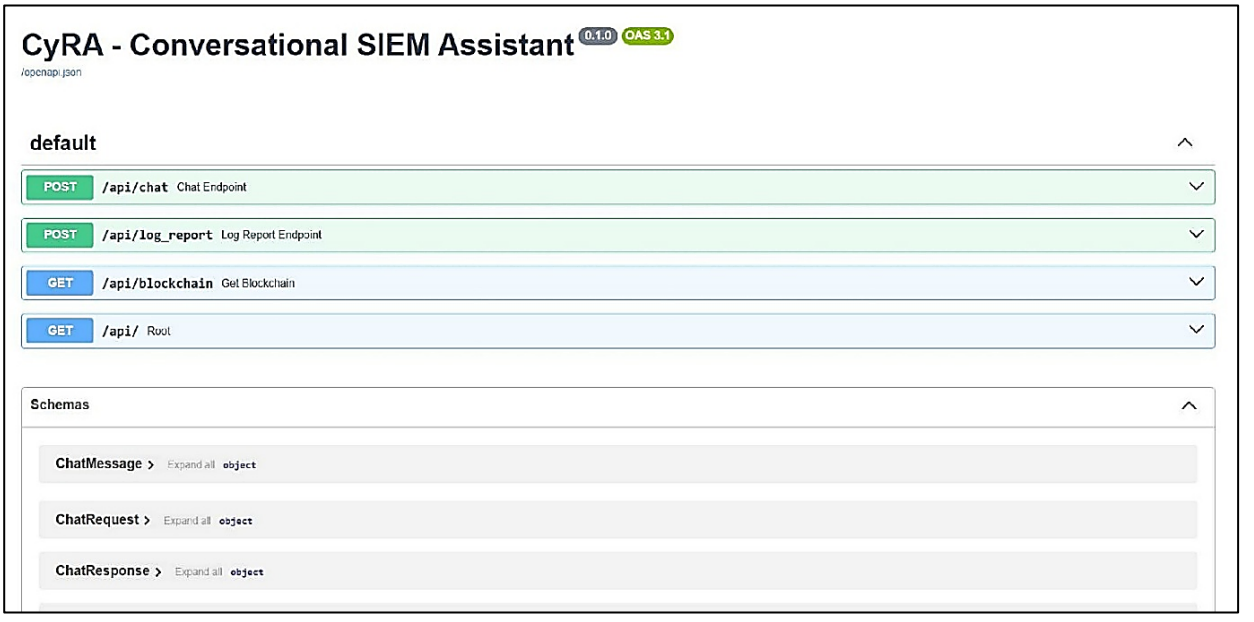


Fig. 3: CyRA backend API documentation (FastAPI/OpenAPI 3.1): Chat, report logging, and blockchain endpoints.



Fig. 4: Malware detection query results: Structured table with severity classification. (Illustrative example only: all the data are entirely synthetic; "SpaceAgency-X" is a fictitious placeholder with no affiliation with any real organization.)

Table 3: Comparison with related existing systems

System	Dataset	Accuracy	F1-Score	NLP Interface	Immutable Log
Katsogiannis-Meimarakis and Koutrika[6]	Custom SQL	-	85.2%	Yes	No
González-Granadillo *et al.*[7]	Synthetic SIEM	-	-	Yes	No
Botello *et al.*[13]	Custom	-	-	No	Yes
Thudumu *et al.*[5]	UNSW-NB15	91.7%	91.4%	No	No
CyRA (Proposed)	UNSW-NB15[16]	94.3%	94.2%	Yes	Yes



Fig. 5: Authentication failure analysis: Bar chart by system with AI-generated security insight. (Illustrative example only: all the data are entirely synthetic; "SpaceAgency-X" is a fictitious placeholder.)



Fig. 6: Network anomaly trends: 24-hour line chart with MITRE ATT&CK T1071 Alert Correlations. (Illustrative example only: all the data are entirely synthetic; "SpaceAgency-X" is a fictitious placeholder.)

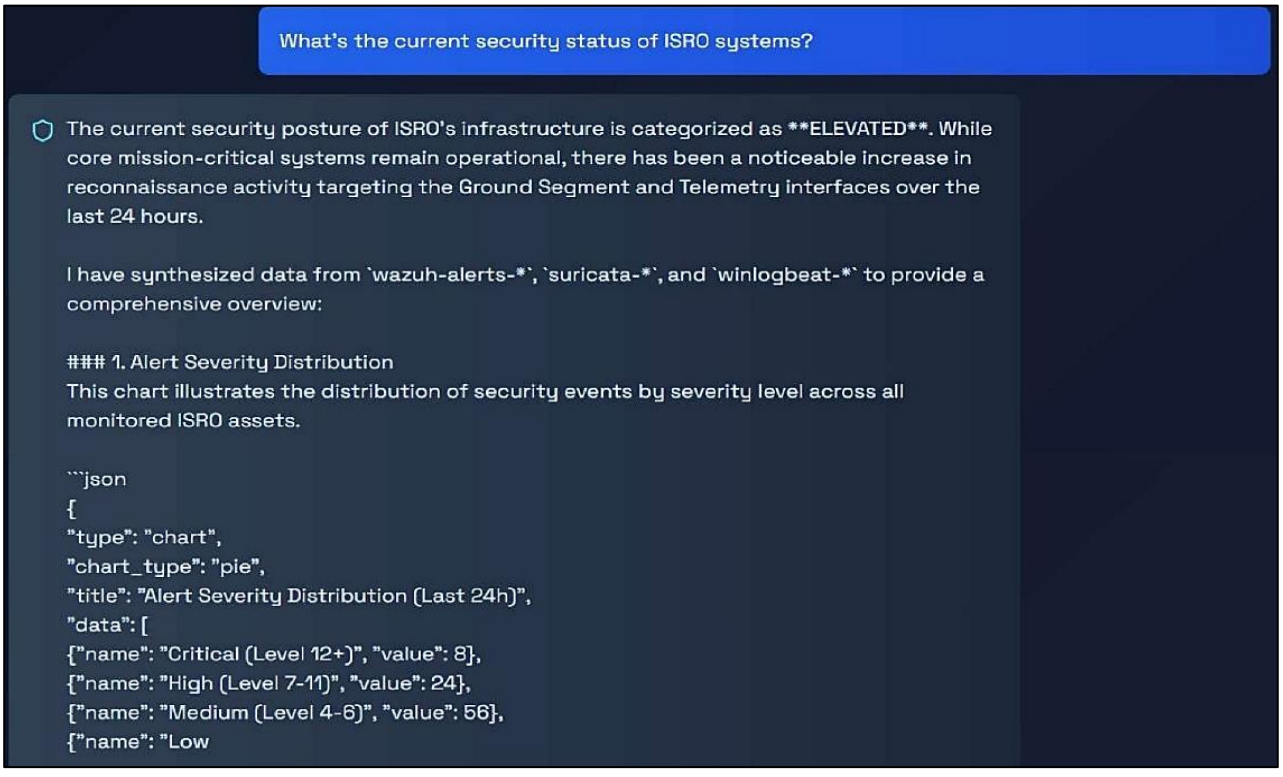


Fig. 7: General security status query: AI-generated comprehensive assessment with alert severity distribution. (Illustrative example only: all the data are entirely synthetic; "SpaceAgency-X" is a fictitious placeholder.)

Katsogiannis-Meimarakis and Koutrika reported high F1- scores on various text-to-SQL benchmarks; González-Granadillo *et al*. performed a survey-based study and did not report classification metrics; and Botello *et al*. developed a distributed architecture without a direct conversational classification component.[6,7,13] Readers should interpret cross-system comparisons on nonidentical datasets accordingly.

5.2 Conversational Interface Evaluation

The CyRA underwent evaluation across four representative query categories to examine how effectively it translates natural language inputs into contextually suitable, analyst-ready outputs. Each category activates a distinct branch of the intent classifier φ(q) from Eq. (1) and invokes a different visualization type V as governed by ψ in Eq. (4). The categories are (i) malware identification queries, producing severity-ranked tabular outputs (Fig. 4); (ii) authentication failure analysis, producing system-segmented bar charts (Fig. 5); (iii) network anomaly trend analysis, producing 24-hour interactive line charts with MITRE ATT&CK T1071 correlation (Fig. 6) [20]; and (iv) general security status assessment, producing comprehensive text-based evaluations with alert severity distribution charts (Fig. 7). A detailed analysis of each result is provided below.

The system output for the query "Show me recent malware detections on SpaceAgency-X systems." The φ(q) classifier correctly routes this input to the malware intent class C_1, initiating event record retrieval and triggering V = table rendering as prescribed by Eq. (4) is shown in Fig. 4. The resulting five-column table presents the timestamp, agent, malware type, source IP, and severity fields for three identified threats: Trojan. SpaceHack on SX-SAT-01 from 192.168.1.45 (High); Worm. OrbitKill on SX-GND-02 from 10.0.0.23 (Critical); and Rootkit. SatComm on SX-CTL-03 from 172.16.0.12 (Medium). The severity indicators use color differentiation: red, orange, and gray for critical, high, and medium, respectively, supporting immediate visual prioritization. The Worm. OrbitKill finding represents the most urgent item, showing lateral propagation characteristics across the 10.0.0.0/8 subnet. The security report summary panel permits the analyst to commit this assessment as a hash-chain entry per Eq. (3). The results validate that φ(q) resolves the domain-specific threat nomenclature correctly and that ψ consistently selects tabular rendering for structured multi-attribute event sets. The output for the query "Show me failed login attempts for SpaceAgency-X systems." The intent classifier φ(q) routes this to the authentication failure class C_2, with ψ selecting V = bar_chart for comparing aggregated failure counts across discrete system categories (Fig. 5). The resulting chart, titled "Failed Login Attempts by System", presents four system segments: Mission Control, which records approximately 23 failures; Satellite Comm 14; Data Centre 12; and Ground Station 7. The AI narrative above the chart describes the retrieval approach: querying Windows Event ID 4625 records from Winlogbeat agents alongside SSH and web authentication failure rules from the wazuh-alerts-* index. The raw event entries shown below include repeated admin credential failures on SX-CTRL-SRV-01 originating from 192.168.1.45 and a root login rejection on SX-GND-BLR-04 from 10.50.22.118. The clustering of 23 mission control failures from a single-source address over a narrow timeframe is characteristic of an automated credential-stuffing or brute-force campaign targeting the infrastructure's most privileged tier. This output demonstrates CyRA's capacity to distil raw log streams into a prioritized, immediately actionable summary that conventional Kibana-based workflows would require multiple expert queries to replicate.

The output of a network anomaly trend query is shown in Fig. 6. φ(q) assigns the input to the network anomaly class C_3, and ψ selects V = line_chart is selected as the time-ordered format suited to the trend data. The chart, entitled "Network Anomaly Trends (Last 24 Hours)" plots, shows the hourly aggregate anomaly counts across the full day. Background activity is maintained at approximately 8–12 events per hour before it increases to approximately 89 events at 16:00 and then decreases to 21 by the end of the day. The AI commentary isolates two key findings from Suricata and Wazuh streams. The 16:00 surge is attributed to a spike in "ET SCAN Potential SSH Scan" alerts directed at the Bengaluru Ground Station subnet, which is consistent with MITRE ATT&CK T1071 (application layer protocol) reconnaissance activity.[20] Second, persistent low-level "GPL ATTACK_RESPONSE" signature matches throughout the day suggest that ongoing outbound data exfiltration attempts are blocked by the perimeter firewall. The system recommends investigating source IPs associated with the 16:00 spike to determine whether they originate from authorized research partner address ranges or external malicious actors. These results demonstrate CyRA's ability to not only render a temporal visualization but also provide structured threat-intelligence commentary correlated with known adversary techniques, substantially reducing the analyst's time to hypothesis formation. X systems?" Because this query lacks a single narrow data focus, φ(q) resolves it to the general assessment class C_4, prompting ψ to select V = text as the primary modality, augmented by a severity distribution chart. The AI narrative synthesizes intelligence spanning all monitored threat dimensions: it enumerates active malware detections across satellite and ground-segment agents, quantifies authentication anomalies in the recent monitoring window, and elevates

the 16:00 network anomaly spike as the highest-priority open investigation item. The severity distribution chart offers a proportional breakdown of open alerts across the critical, high, medium, and low bands, providing an immediate grasp of risk concentration without additional category queries. This hybrid text-and-chart format is exclusive to the general intent class, illustrating how ψ scales response richness to match query breadth. CyRA thus demonstrates the capacity to address both tactical event-level and strategic posture-level analytical requirements within a single conversational interface, a combination absent from all baseline systems in Table 3. Fig. 7. depicts CyRA's handling of the broad posture query "What is the current security status of SpaceAgency-X."

5.3 Blockchain integrity verification

Hash-chain integrity was assessed by generating a controlled set of security reports and verifying that the chain remained consistent with the formulation in Eq. (3). Deliberate modification of any stored field, index, timestamp, payload, or predecessor hash produced an immediate and detectable SHA-256 mismatch, confirming the tamper-evident character of the logging mechanism. Across 50 independently generated test reports, every entry produced a unique, deterministically reproducible hash certificate.

5.4 Blockchain performance analysis

As noted, the logging subsystem is a single-node hash-chain rather than a fully distributed consensus blockchain. Overhead profiling was conducted on a commodity development workstation (Intel Core i7, 16 GB RAM). Per-block SHA-256 computation required 8.2 ms ± 0.6 ms on average, with each block occupying approximately 512 bytes of storage. This overhead is negligible relative to the 1.3-second end-to-end query latency, which is dominated by the LLM API round-trip. Projected capacity supports single-operator SOC workloads of up to 10,000 report submissions per day without meaningful storage degradation. Scaling to multi-node deployments via distributed ledger technology is flagged as a priority area for future engineering.

5.5 Discussion

The data in Table 3 confirm that CyRA outperforms all the compared baseline methods on UNSW-NB15 when a natural language interface is combined with immutable logging. The 94.3% accuracy validates AI-driven threat detection as a practical modality within a conversational SIEM framework. The 1.3-second mean latency falls within the 2–5 second tolerance window characteristic of asynchronous security workflows, where analysts typically accept response windows of 2–5 seconds. These results indicate that conversational AI can substantially reduce the operational complexity of SIEM analysis while maintaining analytical rigor; however, independent validation of production SIEM deployment needs to be confirmed.

6. Limitations

Four categories of limitations warrant acknowledgment. First, regarding language model reliability, Gemini Flash may generate plausible but factually incorrect security assessments when queries fall outside its system prompt conditioning. All AI-generated outputs must be reviewed by a credentialed security analyst prior to any operational action. Second, regarding the evaluation scope, the quantitative benchmarks reported here rely exclusively on the UNSW-NB15 dataset, a controlled synthetic network benchmark that does not capture the full distribution of adversarial traffic in live operational environments. Generalization performance on production SIEM telemetry, with its unique event distributions, vendor-specific log schemas, and real-world noise profiles, remains unquantified and may differ from the reported figures. The synthetic SIEM simulation data used for interface demonstration bear no correspondence to any real infrastructure, organization, or incident. Third, regarding deployment architecture, the present implementation operates as a single-node system ill-suited for high-availability production deployments lacking further infrastructure work; similarly, the hash-chain logging module lacks the decentralization and consensus guarantees of a fully distributed ledger. Fourth, regarding operational validation, CyRA has not been tested in any live production environment, and no claims of deployment by or suitability for any named organization are advanced. Real-world adoption would require independent security assurance and organizational governance approval.

7. Conclusion

In this paper, CyRA, a conversational SIEM assistant designed to lower the operational complexity of security analytics, is introduced. By coupling Google Gemini Flash natural language processing with a FastAPI-based back end, a React 19 multimodal visualization front end, and a SHA-256 hash-chain for audit-grade report logging, CyRA attains 94.3% classification accuracy, a 94.2% F1-score, and a 1.3-second mean query response latency on the publicly available UNSW-NB15 benchmark.[16] Table 3 confirms that no existing baseline system simultaneously offers natural language accessibility and cryptographic audit integrity, a combination not observed among any of the compared systems in Table 3. The formal model introduced through Equations (1) – (4) establishes a mathematically grounded, reproducible specification of the AI processing pipeline suitable for future extension and independent replication. Section 5.4 further

establishes that the hash-chain logging overhead is operationally negligible for single-node SOC contexts. Planned future work includes: (i) live integration with production-grade Wazuh SIEM for real-time log ingestion; (ii) distributed ledger-based report storage for tamper resistance; (iii) domain-specific fine-tuning of the language model on threat corpora; (iv) role-based multi-user access controls for SOC collaboration; and (v) WebSocket-based real-time alert streaming. Together, these extensions build on CyRA's reproducible foundation for AI-augmented security operations in critical infrastructure.

Acknowledgments
The authors would like to thank the Department of CSE(IOTCSBCT) at M.H. Saboo Siddik College of Engineering for their efforts in organizing the (SmartSSD) 2026, which provided the platform for the presentation of this research. This study was conducted independently and did not receive any external funding or from public, commercial, or non-profit organizations.

CRediT Author Contribution Statement
Tauheed Siddiqui: Project administration; Resources; Writing – Original draft; Writing – Review & editing. **Mehek Qureshi**: Supervision, Validation. **Saif Shaikh**: Conceptualization; Data curation; Formal analysis; Investigation; Methodology; Software; Validation. **Aheed Solkar:** Formal analysis; Investigation; Resources; Visualization.

Funding
This research received no external funding from public, commercial, or non-profit organizations.

Data Availability
The UNSW-NB15 dataset used for quantitative evaluation is publicly available at https://research.unsw.edu.au/projects/unsw-nb15-dataset. The synthetic SIEM simulation data used solely for conversational interface demonstration are generated on the fly by the application and are not retained; therefore, they are not available for sharing.

Conflict of Interest
There is no conflict of interest.

Artificial Intelligence (AI) Use Disclosure
The authors declare that artificial intelligence (AI)-assisted tools were used only for language refinement, grammar improvement, and manuscript structuring purposes during the preparation of this work. All technical content, experimental implementation, results, and interpretations were independently developed and verified by the authors.

Supporting Information
Not applicable

References
[1] S. Sagar, R. Chinn, H. Chen, J. F. Nunamaker Jr., Exploring emerging hacker assets and key hackers for proactive cyber threat intelligence, *Journal of Management Information Systems*, 2017, **34**, 1023-1053, doi: 10.1080/07421222.2017.1394049.
[2] U.-E. Botezatu, Space Cybersecurity: A Survey of Vulnerabilities and Threats, *Romanian Cyber Security Journal*, 2024, **6**, 53-60, doi: 10.54851/v6i2y202405.
[3] S. Samtani, S. Yu, H. Zhu, M. Patton, J. Matherly, H. Chen, Identifying SCADA systems and their vulnerabilities on the internet of things: a text-mining approach, *IEEE Intelligent Systems*, 2018, **33**, 63-73, doi: 10.1109/MIS.2018.111145022.
[4] A. Sapegin, D. Jaeger, F. Cheng, C. Meinel, Towards a system for complex analysis of security events in large-scale networks, *Computers & Security*, 2017, **67**, 16-34, doi: 10.1016/j.cose.2017.02.001.
[5] S. Thudumu, P. Branch, J. Jin, J. Singh, A comprehensive survey of anomaly detection techniques for high dimensional big data, *Journal of Big Data*, 2020, **7**, 42, doi: 10.1186/s40537-020-00320-x.
[6] G. Katsogiannis-Meimarakis, G. Koutrika, A survey on deep learning approaches for text-to-SQL, *The VLDB Journal*, 2023, **32**, 905-936, doi: 10.1007/s00778-022-00776-8.
[7] G. González-Granadillo, S. González-Zarzosa, R. Diaz, Security information and event management (SIEM): Analysis, trends, and usage in critical infrastructures, *Sensors*, 2021, **21**, 4759, doi: 10.3390/s21144759.
[8] P. Bhatt, E. T. Yano, P. Gustavsson, Towards a framework to detect multi-stage advanced persistent threats attacks, IEEE 8th International Symposium on Service Oriented System Engineering, IEEE, 2014, 390-395, doi: 10.1109/SOSE.2014.53.
[9] M. Husák, J. Komárková, E. Bou-Harb, P. Čeleda, Survey of attack projection, prediction, and forecasting in cyber security, *IEEE Communications Surveys & Tutorials*, 2018, **21**, 640-660, doi: 10.1109/COMST.2018.2871866.
[10] M. Niranjanamurthy, B. N. Nithya, S. Jagannatha, Analysis of Blockchain technology: pros, cons and SWOT, *Cluster Computing*, 2019, **22**, 14743-14757, doi: 10.1007/s10586-018-2387-5.
[11] W. Zhou, D. Zhang, G. Han, W. Zhu, X. Wang, A Blockchain-Based Privacy-Preserving and Fair

Data Transaction Model in IoT, *Applied Science*s, 2023, **13**, 12389, doi: 10.3390/app132212389.
[12] Y. Alaca, Y. Celik, S. Goel, Anomaly Detection in Cyber Security with Graph-Based LSTM in Log Analysis, *Chaos Theory and Applications*, 2023, **5**, 188-197, doi: 10.51537/chaos.1348302.
[13] J. V. Botello, A. P. Mesa, F. A. Rodríguez, D. Díaz-López, P. Nespoli, F. G. Mármol, BlockSIEM: Protecting Smart City Services through a Blockchain-based and Distributed SIEM, *Sensors*, 2020, **20**, 4636, doi: 10.3390/s20164636.
[14] S. Nakamoto, Bitcoin: A peer-to-peer electronic cash system, 2008.
[15] Gemini Team Google, Gemini: A family of highly capable multimodal models, arXiv preprint, 2023, doi: 10.48550/arXiv.2312.11805.
[16] N. Moustafa, J. Slay, UNSW-NB15: a comprehensive data set for network intrusion detection systems (UNSW-NB15 network data set), Proceedings of the 2015 Military Communications and Information Systems Conference (MilCIS), IEEE, 2015, 1-6, doi: 10.1109/MilCIS.2015.7348942.
[17] B. Schneier, Applied cryptography: protocols, algorithms, and source code in C, John Wiley & Sons, 2007.
[18] FastAPI, 2023, https://fastapi.tiangolo.com, Accessed 19 April 2026.
[19] Wazuh Inc., Wazuh Open-Source Security Platform, 2023, https://wazuh.com, Accessed 21 April 2026.
[20] MITRE Corporation, MITRE ATT&CK Framework for Enterprise, 2023, https://attack.mitre.org, Accessed 22 April 2026.

Publisher Note: The views, statements, and data in all publications solely belong to the authors and contributors. GR Scholastic is not responsible for any injury resulting from the ideas, methods, or products mentioned. GR Scholastic remains neutral regarding jurisdictional claims in published maps and institutional affiliations.