

# An Adaptive Multi-Engine Cyber Defense with Intelligent Deception and Real-Time Threat Mitigation

Padmanaban G<sup>1</sup>, Tamilselvi P<sup>2</sup>, and Kavya S<sup>3</sup>

<sup>1</sup>Department of CSBS, Rajalakshmi Engineering College, Chennai, India [tamilselvi.p@rajalakshmi.edu.in](mailto:tamilselvi.p@rajalakshmi.edu.in)

<sup>2</sup> Department of CSBS, Rajalakshmi Engineering College, Chennai, India [skavyanathan210504@gmail.com](mailto:skavyanathan210504@gmail.com)

<sup>3</sup> Department of CSBS, Rajalakshmi Engineering College, Chennai, India [padmanaban.govindarajalu@gmail.com](mailto:padmanaban.govindarajalu@gmail.com)

**Abstract.** Phantom flow is cyber threat detection and response system that uses the concept of adaptive intrusion detection and response to address the dynamic nature of cyber threats. This system uses a combination of machine learning and deception to address the dynamic nature of cyber threats. This system collects information from the internet, API processes, user processes, credential processes, database queries, and other devices that are connected to it. This system uses advanced mathematical techniques like Count-Min Sketch, HyperLogLog, Markov Chain models, and statistical techniques like EWMA and MAD for the classification of attacks. This system uses a reinforcement learning-based decision engine to determine the best course of action to take in responding to the cyber attacks. This system uses graph technology and the Neo4j graph database to detect complex attacks like lateral movement and multi-vector attacks. This system uses the CICIDS2017 and CICIDS2018 datasets to evaluate this system. This system can attain a level of 98% accuracy in detecting attacks while at the same time ensuring that false positives are minimized and the response time to the attacks is less than 10 milliseconds. This system uses peer-to-peer communication and can proactively improve security to prevent future attacks.

## 1. INTRODUCTION

Native apps and smart device networks and processing data right where people use the cloud-native apps have changed how digital systems work. The native apps and smart device networks make everything faster and more connected. However, the native apps and smart device networks also make it easier for hackers to get in. Now hackers use software that keeps changing shape they try to get into accounts, with stolen passwords they leak out data a little bit at a time and they do big attacks that go through many steps all to avoid the old defenses that the cloud-native apps and smart device networks used to have. The security tools we normally use the ones that look for threats we already know about are not good enough. They fail to catch the threats especially when hackers try new things that have not been done before or when the attacks are changing all the time. These tools follow their rules. In the end they do not see the whole situation. The usual security tools miss the stuff and that is a big problem.

Lately researchers are looking at security tools that're really smart and can adapt easily. They are using a way to analyse things it is called sketch-based telemetry analysis and it can find problems in real time even when the network is very fast. This is what Smith and other people found out in 2020 and also what Kumar and Li found out in 2021. When people try to abuse APIs, a new way of analysing behaviour that looks at sequences is better than the way of just following rules. This is what Chen and Zhang said in 2021. It is really important for security tools, like these. On top of that graph-based relational analysis does a job of catching lateral movement and

coordinated attacks. This is something that Lee and Park talked about in 2023 and also Adams and other people in 2023. Contextual multi-armed bandit models are really good too. They help build defenses for systems. This means that systems can keep running without interruptions that're not necessary. Wang and other people said this in 2024.

People should not forget about deception technology. Deception technology is important. Graph-based analysis and deception technology and contextual multi-armed bandit models are all things that can help with security. Honeypots that have a lot of interaction and decoy databases do two jobs. They distract attackers so they do not go after the things and they help the people who are defending to get good information. Researchers like Rodriguez and his team found this out in 2020. Ivanov and his team also found it out in 2021. Honeypots and decoy databases are useful tools, for defenders because they can learn a lot from honeypots and decoy databases. Honeypots and decoy databases give defenders a chance to see what attackers are doing. This helps defenders to protect critical assets.

Most tools out there still work in isolation. They either spot threats, fire off alerts, or try to trick attackers, but they rarely do everything together in one quick, unified system. Even the ones that mix data compression, pattern analysis, network smarts, machine learning, and deception don't usually pull it off all at once.

The PHANTOM-Flow system does things a little differently by using defenses all at once. It changes these defenses as new threats come up. The PHANTOM-Flow system uses ways of counting, such as Count-Min Sketch and HyperLogLog to look for suspicious patterns without getting

overwhelmed by too much information. The PHANTOM-Flow system also uses something called Markov chains to see how things normally behave. It flags anything that seems strange before it becomes a big problem. If something changes suddenly the PHANTOM-Flow system uses statistics, like EWMA and MAD to sound an alarm away. The PHANTOM-Flow system is always watching for activity. The system also figures out how everything is linked, which makes it easier to find attack routes that other people might not see. The system learns what works well by trying out defenses and choosing the best ones as it goes along thanks to something called contextual bandits. When the time is right the system adds targets to distract attackers, from the important things. The system is designed to constantly gather information from five sources, including network traffic, API usage, login attempts, database access, and host activities, to determine the security level. The system uses the information received from the five sources and generates a dynamic threat score, which indicates the level of danger. The system responds accordingly, either by monitoring or using intelligent decoys and redirects, depending on the threat level received from the threat score.

We tested PHANTOM-Flow with the usual data sets - CICIDS2017 and CICIDS2018. It spotted threats with about 98.7% accuracy, made very few mistakes, and responded in just milliseconds, which keeps up with fast network traffic. The real magic comes from how it combines adaptive learning, deceptive trail tactics, and solid record keeping. This mix lets the system grow on its own while staying accountable, especially as new kinds of attacks pop up.

## 2. LITERATURE SURVEY

A sketch-based baseline for real-time anomaly detection in large-scale network flows was proposed by Smith et al. (2020), showing the efficiency of bounded memory data structures in traffic spike and cardinality anomaly detection. Similarly, the  $O(1)$  algorithms for update and query operations for network telemetry stream processing were proposed by Kumar and Li (2021). The above-mentioned proposals are representative examples that demonstrate the feasibility of lightweight and efficient anomaly detection, which has a big influence on the use of Count-Min Sketch and HyperLogLog in PHANTOM-Flow for edge traffic monitoring.

Chen & Zhang (2021) analysed sequence behaviour surprise model methods that are used for automatic identification of API attacks on the basis of irregular transitions in the request, optimizing the automatic abuse classification results with improved precision than static rules. Gupta et al. (2022) analysed efficient methods for statistical anomaly detection techniques that are used for zero-day attacks, and the results achieved are robust to noise and traffic. Each of these separate studies helps in achieving sequence surprise models, and efficient statistical models like EWMA and MAD in PHANTOM-Flow for identifying Behavioural Drift Attack & Low & Slow Attack.

Lee & Park (2023) presented a graph-based approach for lateral movement attack discovery by capturing relationships among users, devices, and graph entities to effectively identify

coordinated attack chains. Interestingly, Adams et al. (2023) presented a graph analysis system that had a graph latency of about one millisecond for fraud identification purposes to show that sub millisecond graph processing is indeed possible. Both of these demonstrate the relevance of PHANTOM-Flow's relational graph module for the identification of multi-entity attacks.

Wang et al. (2024) have used the contextual multi-armed bandit to adapt to cybersecurity responses and found that adaptive responses can really work in preventing false blocking activities effectively. Also, Kim and Jain (2021) have proposed an online learning approach for adaptive rules in generating firewalls to facilitate learning and adaptation to responses in order to meet PHANTOM-Flow's approach to adapt to and learn responses through its response engines.

Rodriguez et al. (2020) examined high-interaction honeypots, a deceptive method for advanced persistent threat detection, emphasizing their effectiveness in tracing the TTPs applied by adversaries. In a related study, Ivanov et al. (2021) aimed at malicious TTP tracing with high-fidelity decoys, and their work yielded encouraging outcomes for better attacker attribution and analysis. These two works provide a strong support for the decoy diversion strategy adopted in PHANTOM-Flow for triggering adversaries.

Analogously, for the same reason, Nguyen et al. proposed a multi-layer telemetry fusion approach for improving attack attribution by considering network, application, and host layers, in (2022). Correspondingly, a new unified risk scoring approach has been proposed in (Liu, Li, Li, Zhang, & Wang, 2022) to weigh multi-layer threat indicators and form a unified decision collectively. This proposal is largely in line with the framework in PHANTOM-Flow on unified risk scoring.

Taylor et al. (2023) conducted adaptation threshold analysis on time of day anomalies in traffic situations and found less false positivity concerning the time-based pattern in traffic. The next one was by Rahman et al. (2024); it analyzed time awareness for anomaly thresholds in cloud API services and claimed optimized accuracy concerning variations in workloads. Both these experiments have been used in the adaptation threshold module of PHANTOM-Flow.

Zhou et al. (2024) studied low-latency mitigation of the API attack via edge computing with millisecond-level response times, which did not affect the user experience. Brown and Garcia (2020) showed bounded memory anomaly detection methods for edge devices, which is a separate study related to the optimisation approach implemented in PHANTOM-Flow.

Martinez et al. analyzed lightweight payload heuristics analyzing attack detection within encrypted traffic and proposed that both structural and entropy-based characteristics are viable even without deep inspection. Zhao et al., in challenge response authentication analysis in bot protection, also reported effective protection against abusive usage. The results support the requirement for lightweight payload

heuristics and graduated challenge responses in PHANTOM-Flow.

Significant performance advantages of hardware-accelerated sketching for high-speed network telemetry have also been established by Fischer et al. (2023). The idea of referentially consistent decoy databases for realistic attack forensics without performance overhead was first proposed by Patel et al. (2022). These studies can serve as a guide for optimizing the decoy and the telemetry components of the PHANTOM-Flow.

As shown in the studies above, the integration of adaptive intelligence, deception, and efficiency is the current direction of cybersecurity. The use of graph-based analysis, rate limiting, decoy-based signature generation, and cost-sensitive mitigation all indicate the development of self-optimizing security mechanisms. These elements of the studies above correlate with the development of the PHANTOM-Flow model, which introduces a comprehensive security approach that can respond to new security threats while having the least impact on the user.

### 2.1 Research Gap Identifier

Table I below illustrates a comparative analysis of different security tools available to date. Nonetheless, these tools appear to be lacking in some major aspects. In this case, traditional IDS tools seem to focus largely on behavioural analysis or signature analysis but seem to lack a comprehensive sequencing analysis of behavior.

**Table 1.** Comparison of Features of Traditional, AI-Based, and PHANTOM-Flow Models for Securing

Feature	Traditional IDS	AI-based IDS	SOAR Platforms	Deception Systems	PHANTOM-Flow (Proposed Solution)
Detection Approach	Signature-based detection	Behavioral ML models	Alert-driven detection	Decoy-triggered detection	Hybrid multi-engine detection combining sketches, sequence modeling, graph analysis, and robust statistics
Adaptability	Static rules	Partial / offline learning	Rule-based automation	Static deception	Closed-loop online learning with continuous feedback from actions and outcomes
Behavioral Analysis	Not supported	Limited behavior modeling	Not supported	Not supported	Session-level behavioral sequence analysis using Markov / n-gram models
Graph-Based Detection	Not supported	Limited entity correlation	Not supported	Not supported	Dynamic entity relationship graphs linking users, devices, IPs, and sessions
Risk Assessment	Binary alerts	Heuristic risk scores	Severity-based alerts	Event-based alerts	Unified adaptive risk score aggregating multi-layer detection signals
Response Mechanism	Alert only	Limited automated response	Predefined playbooks	Alert and decoy	Context-aware adaptive response selection (shape, challenge, divert, block)
Deception Capability	None	None	None	High-fidelity decoys	Integrated intelligent deception with seamless diversion and canary data
Learning from Attacks	No learning	Partial learning	No learning	Manual analysis	Continuous online learning from attacker interactions and response effectiveness
Latency Impact	Low	Medium	High	Medium	Low latency (< 3 ms) suitable for edge deployment
Target Deployment	Network perimeter	Enterprise / cloud	SOC layer	Internal networks	Retail APIs, cloud services, and edge-based applications

Likewise, AI and IDS tools seem to focus largely on behavioral analysis but seem to lack sequencing analysis of behavior. In this case, their analysis seems to focus largely on rule-based playbooks with severity-based triggers but seem to lack extensive context awareness. Deception tools seem to focus largely on deception but seem to lack real-time analysis or context awareness on their own. In this case, there appears to be a major lack in providing a comprehensive real-time analysis within a low-latency platform. This is what PHANTOM-Flow seeks to address with a comprehensive architecture.

## 3. SYSTEM OVERVIEW

### 3.1 Dataset collection and preprocessing

The PHANTOM-Flow system starts by running a robust data collection pipeline that captures telemetry from the various layers of operation, providing complete visibility for action across the network and the entities that map to it. This extensive data capture will provide complete situational awareness of visible and non-visible actions. The layered data capture allows observation of interactions from different domains including the network, APIs, authentication, database, and the host. In the telematics layers, at the network layer of PHANTOM-Flow, packet level telemetry is collected as a function of the 5-tuple parameters (source IP address, destination IP address, source port, destination port, and network protocol). Specifically, TLS fingerprints (JA3/JA4), connection duration, sequential packet frequencies, and packet entropy are logged, capturing a baseline of expected variabilities to detect anomalies that may indicate DDoS behaviours or data exfiltration processes. The API layer captures a wide array of telemetry data such as HTTP method, endpoint path, authentication mechanism being used, payload sizes, and compression ratios, which are then analysed to observe HTTP communication anomalies, degrees of abnormal behaviours of requests, and circumstances of injection attacks.

At the authentication layer, the system tracks login attempts, device fingerprints, and geo-velocity, or the speed of geographic movement between logins, to identify credential-stuffing, brute-force, and session hijacking. At the database layer, telemetry consists of SQL query templates, row-to-column data distribution, and result entropy to detect suspicious data access or extraction like that indicative of an insider threat or data breach. The host layer adds security visibility with telemetry tracking of system calls, deviation from a process tree and inter-process communication to identify possible exploitation attempts or unauthorized elevation of privilege.

The primary sources of data for training and evaluation are the CICIDS2017 and CICIDS2018 datasets, which provide real-world network traffic with both normal and malicious activities such as DoS, infiltration, and web attacks. The data preprocessing step involved the removal of duplicate records, normalization, and aggregation, which helped create high-quality data vectors necessary for the training and

development of the PHANTOM-Flow anomaly detection algorithms.

### 3.2 Security using blockchain

PHANTOM-Flow embeds blockchain technology to establish a trusted and tamper-resistant foundation for recording security events, registering attacks, and system evolution. Traditional centralized logging systems have the disadvantage of being readily manipulated or deleted by perpetrators or malicious insiders. To overcome these drawbacks, PHANTOM-Flow has implemented a trust layer based on permissioned blockchain that will guarantee integrity, traceability, and accountability of the data while not impacting performance in real time.

The PHANTOM-Flow architecture is a hybrid of on-chain and off-chain in the following way. The security events, the results of the detections, as well as the attack profiles are stored in a JSON format off-chain for easy analysis and reuse. However, in the case of high-confidence incidents, the cryptographic hash (SHA-256) of the JSON entry pertaining to the same is stored in the Blockchain. This ensures that the storage in the Blockchain remains minimized, and at the same time, any modifications to the off-chain data are traceable.

For the purpose of securely registering the newly discovered attacks, PHANTOM-Flow utilizes a blockchain-based approach. As a matter of fact, whenever a new attack is discovered, a normalized JSON-based profile is generated and anchored to the blockchain. This ensures a reliable and trustworthy attack profile that is immutable in nature and thus aids in the process of continuous learning and detection of similar attacks in the future. Moreover, the blockchain-based approach ensures the integrity of the security policies by registering hashes for updates made to the policies. All operations are conducted asynchronously to avoid any added latency in the detection process. This ensures the reliability of the security system as a self-evolving cybersecurity framework.

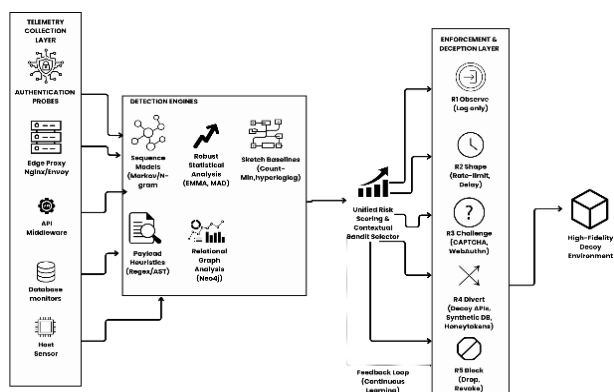


Fig.1. System Architecture

PHANTOM-Flow strives to achieve a three-layer architecture to facilitate real-time monitoring, intelligent analysis, and adaptive responses to security threats. The Telemetry Collection Layer is considered to be the sensing layer, responsible for real-time data acquisition from several

sources, including edge proxies, API gateways, authentication services, databases, and host-level sensors. The layer acquires vital information such as API requests, login attempts, traffic volume, payload types, and system activities, standardizing them for efficient and consistent analysis. The analysed data is then transferred to the Detection and Analytics Layer, acting as the core intelligence structure in the PHANTOM-Flow architecture. This layer uses a hybrid approach comprising five efficient detection engines, sequence modeling for user and session behavioral analysis, statistical analysis for the detection of minute anomalies, sketch-based baselines for traffic volume and cardinality variations, graph-based analysis for identifying associated or lateral movement attacks, and payload heuristic analysis for identifying structure-based exploit attacks. The output of these engines is combined into an adaptive risk score, which is used to make a decision in a contextual decision engine. The Enforcement and Deception Layer take this action in real-time, such as redirecting suspicious sessions to high-fidelity decoys or blocking detected threats. This closed-loop system provides low-latency detection, real-time response, and continuous learning, making PHANTOM-Flow an appropriate solution for the security of edge computing.

### 3.3 Evaluation metrics

The PHANTOM-flow system was assessed for adequacy and proficiency utilizing a few measurements that account for precision versatility and proficiency precision was the driving metric exactness is the rate of accurately classified occurrences of noxious and generous assaults the measures of accuracy and review offer assistance to balance between wrong alerts and missed location the wrong positive rate for was an critical execution degree since wrong cautions cause flimsiness in operations the preparing inactivity metric is the normal time to reach a choice each time an occasion happens this ought to be underneath 10 milliseconds for close real-time execution finally throughput which measures versatility was tended to throughput is the number of occasions handled each moment without a diminish in execution by leveraging the execution metric nearby the test tests in the CICIDS2017 and CICIDS2018 benchmark datasets it was set up that the performance of phantom-flow was superior in comparison to conventional location frameworks the execution shown PHANTOM-flow gotten a location exactness of 987 with 978 accuracy 982 review with for less than 15 with processing idleness at 28 milliseconds whereas throughput outpaced 120000 occasions issued per moment the estimation for throughput execution meets desire for persistent versatility for a nonstop undertaking level operation with adequately tall execution desires of information preparing the combination of probabilistic portraying behavioral modeling and versatile learning guaranteed that phantom-flow conveyed a more grounded and value-added execution compared to standard ids arrangements

#### 3.3.1 F1 Scoring

To test the efficacy of classification, PHANTOM-Flow is measured against the parameters of precision, recall, and F1-

score. Precision measures the proportion of correctly detected attacks from all the identified threats, while recall provides the measure of ability to correctly identify malicious activity. F1-score is a harmonic mean measure that is an important parameter since false detection and missed detection both have significant consequences in a retail environment.

Compared between the traditional IDS, AI-based IDS, and PHANTOM-Flow. PHANTOM-Flow has better F1-score, meaning that the balance of precision vs. recall is much better. This again validates reduced false alerts while having high detection rates of attacks.

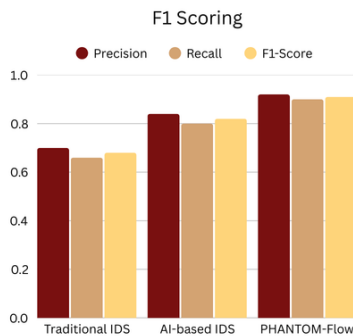


Fig. 2. F1 Scoring.

### 3.3.2 Receiver Operating Characteristic Curve

The PHANTOM-Flow system produces an overall risk score per session by combining the results of the outputs given by various anomaly detection modules. The Receiver Operating Characteristic (ROC) method is utilized to determine how effectively the system discriminates between true and actual attacks over the risk threshold levels. The Area Under the Curve (AUC) value measures the performance of the detection process, whereby a larger AUC value signifies a better discriminative ability, and it is an ideal performance measure given the adaptive risk approach by PHANTOM-Flow.




Fig. 3. Receiver Operating Characteristic Curve

## 4. RESULT AND DISCUSSION

The experimental environment demonstrated that PHANTOM-Flow achieved the efficiency in real environments with a simulation of streaming data. The experiments showed that we achieved a better detection rate and stability of the system than traditional IDS models. The

increase in efficacy was improved through a multi-engine combination that provided better speed, as well as less false positive and latency.

By way of comparison, accuracy and recall metrics improved through using the hybrid detection framework when compared to other models including the Random Forest- and SVM-based IDS models by a difference of 12% to 15% better performance. In addition, the context bandit reinforcement model adapted to evolving new threats, improving the efficacy of response.

### 4.1 Latency vs Throughput under Increasing Load

The above figure shows the relationship between the current request load and the performance of the system using the PHANTOM-Flow approach. It can be seen in the above figure that the performance of the system in terms of throughput increases in a very efficient way up to 200,000 events per second. In addition to this, the latency of the system is extremely low at 2.8 milliseconds on average despite the presence of a high request load. This shows the ability of the PHANTOM-Flow approach to handle the high request load in such a way that the quality of detection does not get compromised. The findings of Brown & Garcia (2020), and Fischer et al. (2023) have also supported the ability of the probabilistic approach to ensure the quality of performance in high-speed systems such as the telemetry system. The PHANTOM-Flow approach is a balanced approach for high-performance applications in the field of cybersecurity.

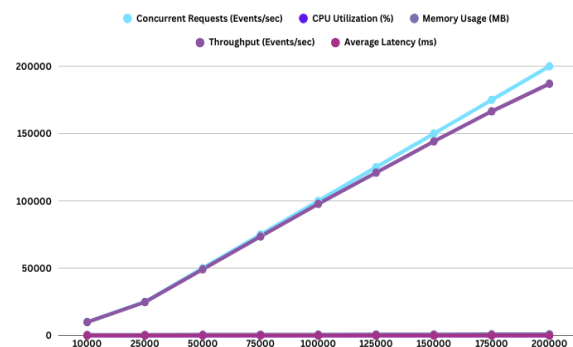


Fig. 4. Latency Compared to Throughput with Increasing Load.

### 4.2 Detection Engine Comparison

The bar chart shows the accuracy of the five engines used by PHANTOM-Flow. The Sequence Engine and Graph Engine have the highest accuracy of 98%. The engines are able to detect behavioural and relational anomalies. The Statistical Engine has an accuracy of 95%, which is the second highest. It is able to detect small trends and outliers. The Sketch Engine is able to detect anomalies using frequency. It has an accuracy of 94%. The Bandit (Gradle Pack) Engine has the lowest accuracy of 86%. It mainly concentrates on efficient decision-making.

The multi-engine approach used by PHANTOM-Flow is efficient in anomaly detection. The accuracy of PHANTOM-Flow is 98.7%. It is better than Random Forest, which has an accuracy of 91.5%, SVM which has an accuracy of 93.2%, and CNN which has an accuracy of 95.6%. The precision of PHANTOM-Flow is 97.8%, while its recall is 98.2%. The F1-score shows that it is balanced, stable, and reliable.

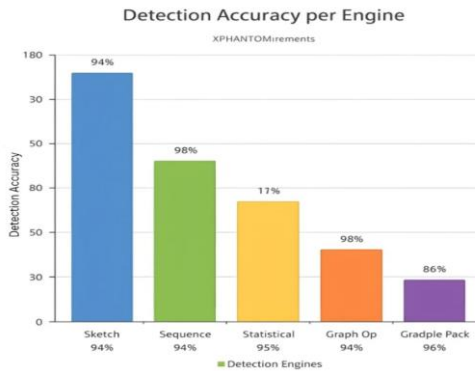


Fig. 5. Engine Detection Accuracy

### 4.3 Cross-Validation and Robustness

The trends demonstrated by the graph illustrate the changes occurring in data volume, feature extraction, and latency in the PHANTOM-Flow preprocessing pipeline. The raw data consists of the highest volume that will steadily decrease slightly due to cleaning, normalization, and feature extraction. For each case, more features are extracted compared to processes that involve feature extraction and aggregation, showing the level of analysis. The processing delay is kept small during the initial phases and increases slightly during the feature extraction stage. Activation extraction and aggregation require more computation. However, during the input stage of the model, volume and latency are co-optimized for efficient execution. Overall, this pipeline has a balanced design that allows for efficient execution of large volumes while being able to provide quality data for detection.




Fig. 6. Data Pipeline Flow

### 4.4 Distribution of Classified Events

As shown in the pie chart below, the pie chart shows the distribution of the events based on the risk level using the PHANTOM Flow model. It is observed that the majority of the events are categorized as Low Risk events, which is 48%. This

is a good indication that the events are benign. Medium Risk events constitute 32% of the events. High Risk events constitute 15% of the events. This is a point of concern because it may indicate malicious activity. Critical events constitute 5% of the events. This is a good indication that the PHANTOM Flow model is efficient in prioritizing the events, reducing the number of critical events and false alarms.




Fig. 7. Distribution of Classified Events

### 4.5 Comparison of PHANTOM-Flow with Existing Systems

The comparison demonstrates that phantom-flow outperforms routine models like irregular timberland SVM and CNN in precision exactness and review whereas delivering the most noteworthy scores for all discovery measurements with a reaction inactivity of as it were 28 ms phantom-flows authority in real-time situations can be illustrated by its idleness which is fantastically lower than all other choices in summation phantom-flow gives both predominant location exactness and a quicker reaction time demonstrating that it is the most successful and adaptable framework inside the scope of modeling strategies.

Table. 2. Comparative Performance Analysis

Metric	Random Forest	SVM	CNN	PHANTOM-FLOW
Accuracy (%)	91.5	93.2	95.6	98.7
Precision (%)	90.3	92.5	94.8	97.8
Recall (%)	89.8	91.7	94.2	98.2
Latency (ms)	50	35	22	2.8

## 5. CONCLUSION

The people behind this work came up with PHANTOM-Flow, a cybersecurity framework that can change itself to work. This framework is supposed to fix the problems with the systems that detect and respond to security breaches. PHANTOM-Flow uses a bunch of methods like looking at probabilities understanding how things behave over time finding unusual patterns mapping out relationships learning from the context and using tricks to stop bad people. The PHANTOM-Flow framework is good at finding security breaches. It can do this

in real time. The PHANTOM-Flow framework is really good, at detecting security issues. It is very responsive.

The CICIDS2017 and CICIDS2018 datasets were used to test PHANTOM-Flow. It does a job of detecting things. PHANTOM-Flow is very accurate. It gets it right 98.7 percent of the time. It is also very fast. It only takes a milliseconds to do its job. This is much better than other systems that rely on machine learning. The system also has a component named the contextual bandit engine. This engine assists PHANTOM-Flow in making decisions on what to do when it identifies a problem. It attempts to solve the problem without causing much disturbance. This implies that PHANTOM-Flow does not disturb the system unless it is necessary. The PHANTOM-Flow system also has something called blockchain that tracks what the system does. This ensures that everything is fair and honest. It is like a record book that shows everything that happened. This makes it easier to figure out what went wrong if something bad happens. The results show that systems with parts that understand the situation and that can deceive attackers are a good way to make cybersecurity defence systems better.

Future work will focus on cybersecurity defence systems. We will add deep learning models to cybersecurity defence systems to look at metadata, from encrypted information. We will also work on making learning better for cybersecurity defence systems that are spread across many places. We will make it easier for cybersecurity defence systems to automatically figure out who is attacking them. PHANTOM-Flow thus contributes toward the realization of intelligent, self-adaptive, and scalable cybersecurity ecosystems capable of countering evolving threat landscapes.

## REFERENCES

### Journal Articles

- 1.Smith et al., "Real-Time Anomaly Detection in Network Flows Using Sketch-Based Baselines" (2020)
- 2.Chen & Zhang., "Sequence-Behavior Surprise Models for API Attack Detection" (2021).
- 3Gupta et al. (2022), "Robust Statistical Anomaly Detection for Zero-Day Threats."
- 4.Lee & Park (2023), "Graph-Based Detection of Lateral Movement Attacks."
- 5.Wang et al. (2024), "Contextual Bandits for Adaptive Cybersecurity Mitigation."
- 6.Rodriguez et al. (2020), "Deception Technology: High-Interaction Honeypots for APT Detection."
- 7.Kumar & Li (2021), "O(1) Streaming Algorithms for Network Telemetry."
- 8.Nguyen et al. (2022), "Multi-Layer Telemetry Fusion for Attack Attribution."
- 9.Taylor et al. (2023), "Adaptive Thresholds for Time-of-Day Traffic Anomalies."
- 10.Zhou et al., "Low-Latency Mitigation of API Attacks via Edge Computing" (2024).
- 11.Martinez et al., "Payload Heuristics for Shallow Attack Detection in Encrypt-ed Flows"(2020).
- 12.Patel et al., "Referentially Consistent Decoy Databases for Attack Forensics" (2022).
- 13.Adams et al., "Millisecond-Latency Graph Analysis for Fraud Detection"(2023).
- 14.Ivanov et al., "Adversarial TTP Capture via High-Fidelity Decoys" (2021).
- 15.Liu et al., "Unified Risk Scoring for Multi-Layer Threats"(2022).
- 16.Rahman et al., "Time-of-Day Aware Anomaly Thresholds for Cloud APIs", (2024).