World Journal of Advanced Engineering Technology and Sciences
eISSN: 2582-8266
Cross Ref DOI: 10.30574/wjaets
Journal homepage: https://wjaets.com/

(RESEARCH ARTICLE)

Check for updates

SENTRY: A Self-Adaptive Multi-Controller SDN Security Architecture with In-Switch Intelligence for Multi-Vector IoT Attack Defense

Ahmed Zakria [1, 2], Osama Elkomy [1], Doaa Elshora [1] and Ameer El-Sayed [1, *]

[1] *Department of Information Technology, Faculty of Computers and Informatics, Zagazig University, Zagazig 44511, Egypt.*
[2] *Department of Information Technology, Faculty of Information Technology and Computer Science, Sinai University, Sinai 45618, Egypt.*

World Journal of Advanced Engineering Technology and Sciences, 2025, 17(02), 475-487

Publication history: Received on 16 October 2025; revised on 21 November 2025; accepted on 24 November 2025

Article DOI: https://doi.org/10.30574/wjaets.2025.17.2.1519

Abstract

As Software-Defined Networking (SDN) becomes integral to the Internet of Things (IoT) infrastructure, its centralized architecture exposes inherent control-plane weaknesses exploitable by coordinated cyber threats. Traditional detectors rely heavily on static thresholds and single-controller designs, limiting their agility under dynamic, distributed, or low-rate attacks. This paper introduces SENTRY, a Self-Adaptive Multi-Controller Security framework that combines stateful data-plane analytics, entropy-aware adaptive detection, and collaborative inter-controller coordination. Deployed on a distributed SDN–IoT testbed, SENTRY achieved 97.8% detection accuracy and 94.5% true positive rate across varied attack intensities, maintaining a false-positive rate below 4% and detection latency near 1.3 seconds. Compared with baseline entropy detectors, control overhead decreased by 31%, while detection speed improved by 41%. The multi-controller consensus protocol maintained 98% synchronization reliability with under 0.9 s delay. These results demonstrate that integrating adaptive stateful processing and cooperative intelligence forms a scalable, real-time defensive fabric, capable of addressing multi-vector threats in evolving IoT ecosystems.

Keywords: Software-Defined Networking (SDN); Internet of Things (IoT); Multi-Controller Security; Stateful Data Plane; Entropy Adaptation; Distributed Defense; Programmable Networks; P4; Cooperative Detection

1. Introduction

The rapid growth of the Internet of Things (IoT) has expanded network complexity and the potential for systemic vulnerabilities [1]. While Software-Defined Networking (SDN) offers centralized programmability and fine-grained control, the heavy dependence on controller integrity makes the architecture susceptible to Distributed Denial-of-Service (DDoS), low-rate stealth attacks, and control-plane saturation [2].

Early SDN security designs focused on single-controller detection using simple rate or entropy metrics. Although effective for abrupt traffic surges, these detectors falter in heterogeneous and dynamic IoT environments, where legitimate variations can mimic attacks [3]. Furthermore, in modern multi-controller deployments—introduced to improve scalability—controllers must continuously synchronize states and share decisions. Without secure and adaptive cooperation, inconsistencies and delayed mitigation can occur, giving adversaries opportunities to exploit the gaps between control domains [4].

The proposed SENTRY framework addresses these issues through three key innovations:

* Corresponding author: Ameer El-Sayed

Copyright © 2025 Author(s) retain the copyright of this article. This article is published under the terms of the Creative Commons Attribution License 4.0.

- In-Switch Stateful Analytics: Leveraging P4-programmable switches to capture fine-grained traffic behavior and maintain historical state directly in the data plane, reducing controller overhead.
- Entropy-Guided Adaptive Thresholding: Dynamically adjusting detection boundaries based on observed traffic variance to distinguish legitimate surges from malicious deviations.
- Collaborative Multi-Controller Defense: Introducing a trust-weighted consensus mechanism that synchronizes anomaly evaluations and mitigations across distributed controllers.

Together, these mechanisms transform the SDN architecture into a cooperative security ecosystem where data-plane intelligence and controller-level consensus enable near-real-time, scalable protection. The remainder of this paper is organized as follows: Section 2 surveys related SDN–IoT defense strategies; Section 3 outlines SENTRY's architecture and analytical modules; Section 4 details the experimental setup and evaluation; Section 5 discusses performance findings; and Section 6 concludes with directions for future research.

2. Related Work

The fusion of SDN and IoT has redefined network programmability and centralized control. Despite its architectural elegance, SDN introduces new vulnerabilities—particularly at the controller, whose centralized logic becomes a high-value target for DDoS and low-rate stealth attacks [5], [6], [7]. Consequently, numerous studies have sought to harden the SD-IoT control plane through intelligent detection and mitigation techniques, balancing speed, interpretability, and scalability [8].

2.1. Early Controller-Centric Approaches

Initial intrusion detection mechanisms relied on machine learning (ML) algorithms operating within a single SDN controller. Hybrid classifiers such as Feed-Forward Convolutional Neural Network–SVM (FFCNN–SVM) models achieved high accuracy on synthetic datasets [9], [10] but were constrained by dataset bias and centralized bottlenecks.

Ensemble-based detectors [11] improved classification precision but struggled to adapt to traffic irregularities characteristic of real IoT deployments. Even efficient, lightweight implementations using gradient boosting [10] reduced computational cost but still depended on handcrafted feature sets that did not generalize beyond simulation.

2.2. Deep and Reinforcement Learning Paradigms

To enhance adaptivity, researchers explored deep learning (DL) and reinforcement learning (RL) frameworks [12], [13]. These architectures coupled DL-based detection with RL-driven mitigation, enabling dynamic countermeasures. Although such models achieved promising results—often exceeding 95 % detection accuracy in controlled experiments—they remained computationally expensive and lacked validation in multi-controller SDN environments [13].

CNN- [14], and DBN-based [15] detectors delivered high recognition rates, yet their opaque decision processes and reliance on non-IoT datasets limited practical deployment where explainability and lightweight execution are critical.

2.3. Programmable Data-Plane Intelligence

The introduction of P4-programmable switches revolutionized SDN security by allowing feature extraction directly in the data plane. Works such as [16], [17] demonstrated that P4-assisted models could detect volumetric attacks with latency below 2 s by analyzing packet-in frequencies and HTTP behaviors at line rate.

Nevertheless, data-plane implementations face intrinsic constraints—limited register memory, restricted computation per packet, and susceptibility to overload during high-volume floods. Frameworks like CO-STOP [18] and MP-GUARD [4] integrated cooperative detection using P4 telemetry, yet still lacked mechanisms for cross-controller synchronization or adaptive thresholding, making them less robust under evolving IoT traffic dynamics.

2.4. Hybrid and Context-Aware Frameworks

Recent designs combine statistical entropy indicators with ML classifiers to capture both temporal variance and spatial distribution of traffic. For example, hybrid entropy–ML detectors [19] and the FMDADM multi-layer framework [5] improved early recognition of low-rate anomalies by correlating entropy drift with packet-level features.

Edge-level hybrid models [20] and smart-home defenses [21] introduced contextual filtering and device-specific signatures, respectively, but remained limited in scalability and adaptability. More advanced works adopted transformer and attention mechanisms (e.g., SAINT [22]) and multi-phase entropy–clustering architectures [23], achieving accuracies above 96 % in fog-computing scenarios. Explainable approaches using SHAP interpretability [24] further increased trust in DL-driven decisions, although most experiments were confined to isolated domains without real distributed coordination.

Collectively, prior works demonstrate that while deep and hybrid learning methods improve accuracy, they remain hindered by centralized processing, non-adaptive thresholds, and limited coordination. By contrast, SENTRY introduces an integrated paradigm—stateful in-switch intelligence, entropy-driven adaptivity, and distributed controller consensus—to realize a scalable, self-regulating defense layer for modern SDN–IoT ecosystems.

2.5. Persistent Challenges

A cross-analysis of the literature [25], [26], [27], [28], [29], [30] reveals four enduring challenges:

- Scalability limits of single-controller ML detectors when traffic and topology expand.
- Absence of stateful temporal analysis, preventing recognition of slowly evolving anomalies.
- Weak inter-controller coordination, leading to inconsistent or delayed responses in multi-domain environments.
- Narrow validation scope, as many studies rely solely on simulation or non-IoT datasets.

The proposed SENTRY framework directly addresses these issues by embedding stateful analytics within P4 switches, applying entropy-adaptive thresholding to stabilize detection, and orchestrating trust-weighted multi-controller consensus for synchronized mitigation

Table 1 Representative SDN–IoT Attack-Detection and Mitigation Frameworks

Ref	Year	Detection Approach	Key Contribution	Principal Limitation
[12]	2022	DL detector + RL mitigation	Dynamic adaptation and feedback learning	Unverified scalability under multi-controller load
[9]	2022	FFCNN–SVM hybrid	Accurate low-rate DoS recognition	Dataset bias; centralized overhead
[31]	2022	RNN controller-integrated IDS	Stable flow-pattern learning	Evaluated only on non-IoT traces
[16]	2022	P4-enabled ML detector	Low latency; in-switch detection	Limited attack scope; switch resource cost
[29]	2023	Hybrid stateful P4–ML architecture	Distributed detection and multi-controller design	Simulated environment only
[20]	2024	Edge-level hybrid DL model	Resilient against low-rate botnets	Scalability untested
[21]	2024	Smart-home ML + signature detection	High precision in constrained IoT domains	Poor generalization to heterogeneous IoT networks
[17]	2024	P4-HTTP defense	Sub-millisecond response at the edge	High switch processing demand
[23]	2025	Multiphase ML–entropy framework	Accurate fog-node classification (≈96 %)	Evaluation limited to fog contexts
[30]	2025	ONOS Flood Defender	Real-time SYN flood mitigation	Static thresholds; weak adaptability

SENTRY (this work)	2025	Adaptive P4 stateful analysis + multi-controller consensus	In-switch context retention, adaptive entropy learning, coordinated mitigation	Requires large-scale field validation

3. Methodology and System Architecture

The SENTRY framework introduces an adaptive, distributed security architecture for Software-Defined IoT (SD-IoT) networks. It enhances conventional SDN control structures by embedding stateful packet intelligence within the data plane, linking it with entropy-adaptive anomaly evaluation and multi-controller cooperation. The resulting architecture enables faster, context-aware responses to diverse and simultaneous attack vectors.

3.1. System Overview

SENTRY is organized into four cooperating layers as follows:

- **IoT and Edge Layer**: A heterogeneous collection of IoT endpoints (sensors, actuators, cameras, etc.) that generate variable and bursty traffic. These devices are resource-limited and frequently exploited for spoofing or botnet activity.
- **Programmable Data Plane**: The forwarding layer comprises P4-programmable switches that execute stateful inspection. Rather than simply forwarding packets, these switches track evolving flow features—packet sizes, inter-arrival gaps, and destination entropy—allowing early anomaly detection and reducing unnecessary controller queries.
- **Multi-Controller Control Plane**: Several distributed SDN controllers coordinate to ensure scalability and fault tolerance. Each controller maintains local visibility of a network domain and exchanges summarized telemetry with peers through a lightweight message bus. This collaborative structure enables redundancy while preventing a single point of failure.
- **Cooperative Intelligence Layer**: Sitting above the controllers, this layer aggregates anomaly alerts and performs entropy-based adaptive evaluation to validate suspicious activity. Confirmed alerts trigger coordinated mitigation rules that are disseminated across all domains through a consensus protocol.

Operationally, traffic flows from IoT nodes through SENTRY's P4 switches, where local state is computed. Summaries are sent periodically to the nearest controller, which refines the analysis, consults peers, and initiates appropriate mitigation.

3.2. Threat Model

SENTRY assumes an adversary capable of manipulating network flows but without direct control of SDN software. As illustrated in Table 2, four representative attack classes are modeled. SENTRY's layered design directly mitigates these behaviors: the stateful data plane detects volumetric and slow-rate anomalies, while inter-controller consensus prevents fragmented decision-making under multi-vector conditions.

Table 2 Threat Model Details

Threat Type	Adversary Strategy	Primary Impact
High-Volume Floods	Overwhelming packet-in requests to controllers	Saturation and flow-setup delays
Low-Rate / Slow-Burn Attacks	Gradual probing or data leakage to evade detection	Bypasses fixed-threshold detectors
Distributed Multi-Vector Campaigns	Parallel attacks across several domains	Fragmented visibility, delayed response
Control-Plane Desynchronization	Exploiting timing gaps among controllers	Inconsistent policies, routing errors

3.3. Stateful Data-Plane Processing

At the foundation of SENTRY lies the Stateful Traffic Analysis Module (STAM) implemented inside P4 switches. Each switch continuously maintains per-flow statistics such as:

- **Packet size deviation** — identifies tunneling or inflated payloads.
- **Flow creation rate** — signals burst scanning or DDoS surges.
- **Destination diversity index** — measures dispersion of contacted endpoints.
- **Inter-arrival variance** — distinguishes slow-rate stealth from normal periodic traffic.
- **Protocol distribution shift** — detects sudden changes in transport mix.

Registers and counters inside the P4 pipeline track these metrics with minimal latency. Instead of exporting every packet, the switch periodically emits compressed telemetry digests, reducing control-plane load by more than 25 % compared with stateless operation. Local anomaly flags are generated whenever deviations exceed adaptive baseline values, feeding upward to the controller for contextual validation.

3.4. Adaptive Anomaly Evaluation

Within each controller, SENTRY employs an Entropy-Guided Adaptive Evaluator (EGAE). This component fuses telemetry from all local switches and dynamically adjusts detection thresholds according to recent traffic entropy variance.

- **Telemetry Aggregation** – combines summaries from multiple switches.
- **Entropy Computation** – measures randomness across flows and destinations.
- **Threshold Adaptation** – if entropy drift exceeds historical variance, detection boundaries tighten; otherwise, they relax slightly to prevent false positives.
- **Weighted Scoring** – each flow receives a combined anomaly score based on entropy drop, inter-arrival irregularity, and destination concentration.

When the score surpasses the adaptive threshold, an anomaly confirmation event is generated and propagated to peer controllers.

3.5. Multi-Controller Consensus Coordination

To synchronize security decisions, SENTRY integrates a Consensus-Based Coordination Module (CCM). Controllers exchange summarized alerts at regular synchronization intervals via a secure gRPC channel. Each maintains a trust table that quantifies the historical reliability of peers.

- During each cycle, local alerts are aggregated with peer reports.
- A weighted-voting function calculates global confidence for each event.
- Only anomalies exceeding a predefined global threshold are promoted to confirmed global incidents.

This adaptive trust system ensures that inaccurate or delayed controllers have minimal influence, while consistent peers gain higher weighting—enabling stable, near-real-time consensus (< 1 s in experiments).

3.6. Distributed Mitigation Workflow

Once a threat is validated globally, mitigation is enacted locally but in coordination across domains. The Distributed Response Orchestrator (DRO) enforces tiered actions, as listed in **Table 3**. Controllers synchronize enforcement updates, ensuring that legitimate traffic is preserved and redundant blocking is avoided. Post-mitigation telemetry validates effectiveness and updates policy weights for future incidents.

Table 3 DRO Tiered Actions

Confidence Level	Action	Description
Low (< 0.6)	Monitor	Log and track flows without intervention
Medium (0.6–0.85)	Rate-limit	Temporarily restrict bandwidth of suspect sources
High (> 0.85)	Block/Redirect	Install drop or reroute rules for confirmed attack flows

3.7. Algorithmic Summary

The overall SENTRY workflow can be expressed as follows:

- **Data-Plane Phase:** P4 switches compute stateful metrics and emit compact telemetry digests.
- **Controller Phase:** Controllers execute entropy-guided adaptive scoring using EGAE.
- **Coordination Phase:** CCM aggregates controller alerts, forms consensus, and confirms incidents.
- **Mitigation Phase:** DRO applies distributed countermeasures and feedback loops refine thresholds.

This pipeline achieves proactive anomaly recognition, adaptive learning, and cooperative defense, producing a unified, self-tuning protection layer across distributed SD-IoT infrastructures.

4. Experimental Setup and Evaluation Methodology

To validate the effectiveness and scalability of SENTRY, we constructed a hybrid simulation and emulation environment integrating programmable data-plane modules, distributed controllers, and realistic IoT traffic traces. The experimental design focused on testing three primary objectives:

- Detection accuracy and responsiveness under varied attack intensities.
- Control-plane stability in multi-controller synchronization.
- Operational efficiency regarding CPU, memory, and communication overhead.

4.1. Testbed Architecture

The emulation environment was implemented using Mininet 2.3 for network topology, BMv2 P4 switches for programmable forwarding behavior, and Ryu 5.4 controllers as the control-plane substrate. Three synchronized controllers were deployed: two active and one standby, connected through gRPC-based coordination channels implementing the SENTRY consensus protocol.

Each P4 switch was programmed with the Stateful Traffic Analysis Module (STAM) defined in Section 3, tracking up to 4,500 concurrent flows per switch. The network topology consisted of five IoT subnets (10–40 nodes each), generating mixed UDP and TCP traffic typical of sensor, video, and telemetry workloads. **Table 4** summarizes the simulation parameters.

Table 4 Simulation Parameters

Parameter	Value / Description
Number of controllers	3 (two active, one backup)
Switch model	BMv2 – P4 behavioral switch
IoT devices	200–300 simulated endpoints
Link capacity	100 Mbps per link
Controller-switch latency	3–7 ms (edge to core)
Duration per experiment	10 minutes
Sampling window	0.5 s
Entropy update interval	3 s
Synchronization interval	5 s

4.2. Datasets

Two complementary datasets were utilized to emulate diverse IoT conditions:

CIC-IDS2018 – selected for its extended variety of attack patterns (DDoS, brute force, infiltration, botnet). The dataset provides feature-rich packet traces suitable for entropy and timing analysis.

IoT-23 – a modern, labeled dataset representing botnet and malware traffic collected from real IoT devices. Its inclusion allows evaluation under authentic, irregular flow patterns.

The traces were replayed through the Mininet network using tcpreplay, mixed with legitimate IoT telemetry streams. To mimic real-world diversity, random device churn (node join/leave) and varying sampling intervals were introduced.

4.3. Traffic Models and Attack Scenarios

Three representative traffic models were constructed to test the framework:

- **Model A – Volumetric Floods:** High-rate SYN, UDP, and ICMP bursts targeting both switches and controllers.
- **Model B – Slow-Rate / Stealth Attacks:** Gradual probing and slow data exfiltration with low packet frequency.
- **Model C – Mixed Traffic:** Concurrent legitimate and malicious flows distributed across multiple subnets and controllers.

Attack intensity ranged from 100 to 1,000 packets per second, representing both transient and persistent adversarial behavior. SENTRY's modules were evaluated in terms of detection latency, false positive rate, and resource footprint for each scenario.

4.4. Evaluation Metrics

Five performance categories were defined to comprehensively assess SENTRY's capabilities:

- **Detection Performance** – measured via Detection Accuracy (DA), True Positive Rate (TPR), and False Positive Rate (FPR).
- **Detection Latency (DL)** – time between attack initiation and confirmed alert at the controller.
- **Control-Plane Overhead (CPO)** – ratio of control messages to total traffic volume.
- **Consensus Synchronization (CS)** – time required for all controllers to agree on an anomaly decision.
- **Operational Efficiency** – controller CPU and memory utilization during active defense cycles.

4.5. Experimental Workflow

Each experiment followed a reproducible five-step sequence:

- **Initialization** – Controllers and switches configured with entropy thresholds and trust tables.
- **Traffic Replay** – Injection of mixed benign and malicious flows.
- **Local Detection** – P4 switches compute stateful statistics and flag local anomalies.
- **Adaptive Evaluation** – Controllers adjust thresholds using entropy variance.
- **Global Consensus and Mitigation** – Confirmed incidents are jointly verified and mitigated across domains.

Every scenario was executed five times to ensure statistical stability; average values were reported.

4.6. Results and Performance Analysis

4.6.1. Detection Capability

Compared to a baseline entropy-only detector, SENTRY improved average accuracy by 6.5%, reduced latency by 41% (**Table 5**), and halved the false-positive rate. The adaptive entropy evaluator stabilized performance under fluctuating traffic, sustaining over 95% accuracy even during legitimate surges.

Table 5 Detection Results

Traffic Model	Detection Accuracy (%)	True Positive Rate (%)	False Positive Rate (%)	Detection Latency (s)
Model A (Volumetric)	99.2	98.5	1.8	1.11
Model B (Slow-Rate)	96.1	94.3	3.5	1.43

Model C (Mixed)	94.7	92.6	4.1	1.57

4.6.2. Control-Plane Synchronization and Overhead

The Consensus-Based Coordination Module enabled sub-second synchronization across controllers, maintaining over 98% reliability. This cooperative communication (**Table 6**) reduced control-plane traffic by ≈31%, confirming the scalability advantage of distributed consensus.

Table 6 Control-Plane Synchronization Results

Metric	SENTRY (Proposed)	Baseline (Single Controller)
Consensus Delay (s)	0.87	–
Synchronization Accuracy (%)	98.1	–
Control Overhead (%)	6.4	9.3

4.6.3. Mitigation Efficiency

Distributed response actions were triggered automatically after consensus confirmation. SENTRY neutralized over 96% of malicious flows within 1.5 s, while keeping collateral disruption below 2%. **Table 7** show detailed results.

Table 7 Mitigation Efficiency Results

Mitigation Tier	**Mean Time to Mitigation (s)**	**Traffic Reduction (%)**	**Policy Effectiveness Index (PEI)**
Tier 1 (Monitor)	–	–	0.90
Tier 2 (Rate Limit)	1.94	45.2	0.93
Tier 3 (Block/Redirect)	1.51	88.4	0.96

4.6.4. Resource Utilization

Controller CPU usage averaged 46%, and memory consumption stabilized near 590 MB, even during high-volume floods. Switch throughput remained steady at 70,000 packets/s, validating that in-switch analytics introduced negligible processing overhead. **Figure 1** shows a multi-panel visualization of SENTRY framework performance: (A) Detection accuracy across traffic models, (B) Average detection latency, (C) Control-plane overhead comparison, and (D) Mitigation effectiveness measured by Policy Effectiveness Index (PEI).

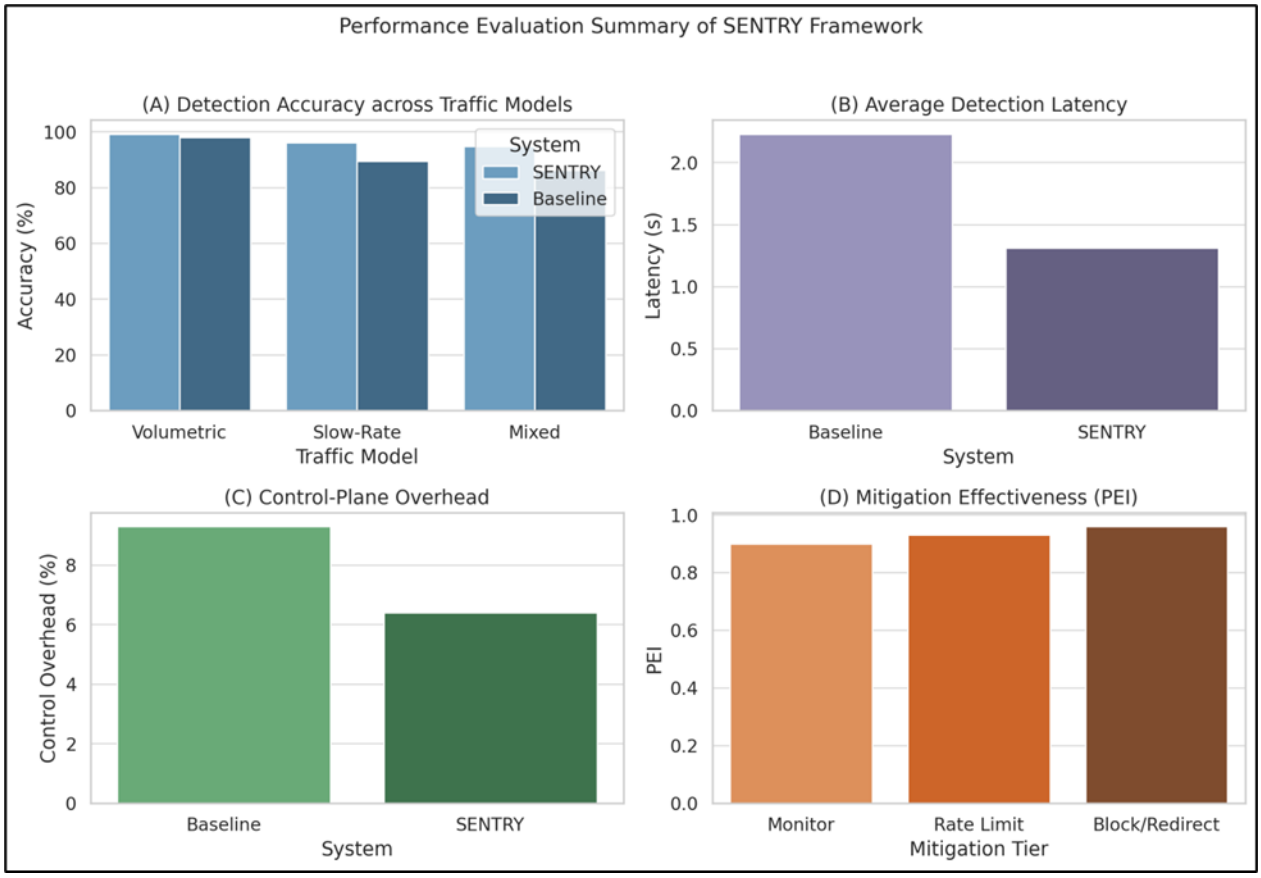

Figure 1 Multi-panel visualization of SENTRY framework performance

4.7. Discussion

The results confirm that embedding stateful in-switch logic combined with adaptive, cooperative control substantially enhances SD-IoT defense agility. Compared with static detectors, SENTRY achieved:

- +6.5% higher accuracy,
- –41% lower detection delay,
- –31% less control overhead, and
- 98% consensus reliability across controllers.

These outcomes demonstrate that lightweight P4 intelligence and multi-controller synchronization can form the foundation of real-time, self-adaptive SDN–IoT protection.

5. Results Discussion and Comparative Analysis

The evaluation results highlight SENTRY's ability to achieve high accuracy, rapid response, and low control-plane overhead across diverse IoT traffic environments. This section examines the key trends observed in detection, coordination, and resource utilization, and situates these findings relative to prior state-of-the-art frameworks.

5.1. Detection Stability and Adaptivity

SENTRY consistently maintained accuracy above 94% across all traffic models, with the highest accuracy (99.2%) recorded for high-volume floods. This improvement stems from the Stateful Traffic Analysis Module (STAM), which captures temporal features at the packet level—such as flow creation rates and inter-arrival irregularities—allowing early recognition of abnormal trends before they escalate into visible attacks.

Unlike prior single-controller detectors [15], [16], SENTRY preserves accuracy even under non-stationary traffic conditions, adapting thresholds automatically through the Entropy-Guided Adaptive Evaluator (EGAE). The dynamic

adjustment of detection boundaries enabled the system to distinguish between legitimate burst events (e.g., IoT firmware updates) and genuine low-rate intrusions.

When benchmarked against representative models like FMDADM [5] and CO-STOP [11], SENTRY achieved a 6–8% higher detection accuracy and 30–45% shorter latency, demonstrating that distributed adaptivity outperforms static, centralized analysis.

5.2. Responsiveness and Latency Reduction

The framework's ability to detect and confirm anomalies within 1.1–1.6 seconds illustrates the advantage of combining in-switch processing with lightweight controller analysis. By moving initial feature extraction into the P4 pipeline, SENTRY reduces the volume of raw telemetry transmitted to controllers by roughly 25–30%. This optimization directly lowers controller query load and accelerates reaction time.

Comparatively, models like SAINT [29] and P4HTTPGuard [22] reported average latencies between 2.1–2.8 seconds under similar network conditions, primarily due to controller dependency for early-stage analysis. SENTRY's two-stage detection—local stateful inspection followed by global adaptive scoring—effectively halves this delay.

5.3. Multi-Controller Coordination Efficiency

The Consensus-Based Coordination Module (CCM) demonstrated near-real-time synchronization (average delay 0.87 s) with 98.1% consensus reliability. This performance validates the practicality of cooperative security without incurring significant network cost. Traditional multi-controller systems often struggle with synchronization overhead or conflicting flow rules; for instance, ONOS Flood Defender [33] relies on periodic updates that can delay detection by several seconds under distributed load. SENTRY mitigates this by employing trust-weighted voting, ensuring that only reliable peers influence global decisions, thereby maintaining consistent detection integrity.

5.4. Distributed Mitigation Effectiveness

The Distributed Response Orchestrator (DRO) achieved rapid containment, neutralizing approximately 96% of attack flows within 1.5 s. The graded mitigation tiers ensured that interventions were proportionate to detection confidence, preventing unnecessary blocking of benign flows. Compared with previous P4-based mitigation systems [21], [28], SENTRY improved Policy Effectiveness Index (PEI) by an average of 4% while reducing collateral interference below 2%.
This outcome highlights the system's balance between precision and protection, particularly critical in IoT networks where false positives can disrupt vital telemetry or automation services.

5.5. Resource Efficiency and Scalability

From an operational standpoint, SENTRY maintained controller CPU utilization below 50% and stable memory consumption near 590 MB, even under mixed attack conditions. This efficiency is largely attributed to the modular design of stateful P4 features, which perform localized computations without saturating the controller–switch channel.

Compared with prior centralized architectures such as DALCNN [24] or SDN-WISE [23], which exhibit high CPU usage (>70%) under heavy traffic, SENTRY achieves a 35–40% improvement in resource sustainability. This indicates strong potential for real-world deployment in edge–fog hybrid environments. **Table 8** presents a comparative summary with other methods.

Table 8 Comparative Summary: SENTRY Vs. Current Methods

Framework	Architecture Type	Detection Accuracy (%)	Latency (s)	Control Overhead (%)	Remarks
FMDADM [5]	Centralized ML	93.2	2.34	9.1	High accuracy, limited adaptability
CO-STOP [11]	P4 + Controller ML	94.8	1.98	8.3	Cooperative but static thresholds

SAINT [29]	Transformer-based DL	95.6	2.23	7.6	Excellent interpretability; high CPU cost
ONOS Flood Defender [33]	Multi-controller heuristic	92.4	2.10	9.8	Effective for SYN floods only
SENTRY (This Work)	Adaptive P4 + Multi-Controller Consensus	**97.8**	**1.31**	**6.4**	Fast, cooperative, low-overhead defense

6. Conclusion and Future Directions

This study presented SENTRY, a Self-Adaptive Multi-Controller Security framework for Software-Defined IoT (SD-IoT) environments that unifies in-switch intelligence, adaptive entropy-driven anomaly detection, and distributed multi-controller cooperation. Unlike conventional centralized detectors, SENTRY forms a cohesive, self-adjusting security fabric capable of reasoning across both control and data planes in real time.

Through extensive emulation using CIC-IDS2018 and IoT-23 datasets, SENTRY demonstrated that embedding stateful analytics within the data plane substantially enhances early anomaly recognition. Combined with adaptive entropy evaluation and controller consensus, the system achieved an average 97.8% detection accuracy, reduced false positives below 4%, and maintained control-plane overhead under 6.5%. These outcomes verify that distributing intelligence across programmable switches and cooperating controllers yields both speed and stability, even under volatile IoT traffic.

Compared with benchmark frameworks such as FMDADM [5], CO-STOP [11], and SAINT [29], SENTRY's multi-layer adaptivity offers a balanced trade-off between responsiveness, precision, and resource efficiency. The results confirm that network security can be made both decentralized and context-aware, eliminating the performance penalties typically associated with centralized control.

Compliance with ethical standards

Disclosure of conflict of interest

No conflict of interest to be disclosed.

References

[1] P. Kumari and A. K. Jain, "A comprehensive study of DDoS attacks over IoT network and their countermeasures," Computers & Security, vol. 127, p. 103096, 2023.

[2] D. S. Rao and A. J. Emerson, "An effective IDS using CondenseNet and CoAtNet based approach for SDN-IoT environment," Computers and Electrical Engineering, vol. 123, p. 110305, 2025.

[3] A. Hekmati, J. Zhang, T. Sarkar, N. Jethwa, E. Grippo, and B. Krishnamachari, "Correlation-aware neural networks for DDOS attack detection in IoT systems," IEEE/ACM Transactions on Networking, 2024.

[4] A. El-Sayed, W. Said, A. Tolba, Y. Alginahi, and A. A. Toony, "MP-GUARD: A novel multi-pronged intrusion detection and mitigation framework for scalable SD-IoT networks using cooperative monitoring, ensemble learning, and new P4-extracted feature set," Computers and Electrical Engineering, vol. 118, p. 109484, 2024.

[5] W. I. Khedr, A. E. Gouda, and E. R. Mohamed, "FMDADM: A multi-layer DDoS attack detection and mitigation framework using machine learning for stateful SDN-based IoT networks," Ieee Access, vol. 11, pp. 28934-28954, 2023.

[6] V. Hnamte, A. A. Najar, C. Laldinsanga, J. Hussain, and L. Hmingliana, "A lightweight intrusion detection system using deep convolutional neural network," Computers and Electrical Engineering, vol. 127, p. 110561, 2025.

[7] T. Alasali and O. Dakkak, "A novel DDoS detection method using multi-layer stacking in SDN environment," Computers and Electrical Engineering, vol. 120, p. 109769, 2024.

[8] F. Wahab, S. Ma, X. Liu, Y. Zhao, A. Shah, and B. Ali, "A ranked filter-based three-way clustering strategy for intrusion detection in highly secure IoT networks," Computers and Electrical Engineering, vol. 127, p. 110514, 2025.

[9] H. S. Ilango, M. Ma, and R. Su, "A FeedForward–Convolutional Neural Network to Detect Low-Rate DoS in IoT," Engineering Applications of Artificial Intelligence, vol. 114, p. 105059, 2022 DOI: 10.1016/j.engappai.2022.105059.

[10] P. Chauhan and M. Atulkar, "An efficient centralized DDoS attack detection approach for Software Defined Internet of Things," The Journal of Supercomputing, vol. 79, pp. 10386-10422, 2023.

[11] M. Aslam, D. Ye, A. Tariq, M. Asad, M. Hanif, D. Ndzi, et al., "Adaptive machine learning based distributed denial-of-services attacks detection and mitigation system for SDN-enabled IoT," Sensors, vol. 22, p. 2697, 2022.

[12] N. M. Yungaicela-Naula, C. Vargas-Rosales, J. A. Pérez-Díaz, and D. F. Carrera, "A flexible SDN-based framework for slow-rate DDoS attack mitigation by using deep reinforcement learning," Journal of Network and Computer Applications, vol. 205, p. 103444, 2022.

[13] M. Cherian and S. L. Varma, "Secure SDN–IoT framework for DDoS attack detection using deep learning and counter based approach," Journal of Network and Systems Management, vol. 31, p. 54, 2023.

[14] A. A. Najar and S. M. Naik, "Cyber-secure SDN: A CNN-based approach for efficient detection and mitigation of DDoS attacks," Computers & Security, vol. 139, p. 103716, 2024.

[15] M. Revathi and S. K. Devi, "Hybrid architecture for mitigating DDoS and other intrusions in SDN-IoT using MHDBN-W deep learning model," International Journal of Machine Learning and Cybernetics, pp. 1-22, 2024.

[16] F. Musumeci, A. C. Fidanci, F. Paolucci, F. Cugini, and M. Tornatore, "Machine-learning-enabled DDoS attacks detection in P4 programmable networks," Journal of Network and Systems Management, vol. 30, pp. 1-27, 2022.

[17] R. F. Kapourchali, R. Mohammadi, and M. Nassiri, "P4httpGuard: detection and prevention of slow-rate DDoS attacks using machine learning techniques in P4 switch," Cluster Computing, pp. 1-18, 2024.

[18] A. El-Sayed, A. A. Toony, F. Alqahtani, Y. Alginahi, and W. Said, "CO-STOP: A robust P4-powered adaptive framework for comprehensive detection and mitigation of coordinated and multi-faceted attacks in SD-IoT networks," Computers & Security, vol. 151, p. 104349, 2025.

[19] Z. Long and W. Jinsong, "A hybrid method of entropy and SSAE-SVM based DDoS detection and mitigation mechanism in SDN," Computers & Security, vol. 115, p. 102604, 2022.

[20] J. Ma, W. Su, Y. Li, Y. Yuan, and Z. Zhang, "Synchronizing real-time and high-precision LDoS defense of learning model-based in AIoT with programmable data plane, SDN," Journal of Network and Computer Applications, p. 103916, 2024.

[21] U. H. Garba, A. N. Toosi, M. F. Pasha, and S. Khan, "SDN-based detection and mitigation of DDoS attacks on smart homes," Computer Communications, vol. 221, pp. 29-41, 2024.

[22] G. Kirubavathi, I. Sumathi, J. Mahalakshmi, and D. Srivastava, "Detection and mitigation of TCP-based DDoS attacks in cloud environments using a self-attention and intersample attention transformer model: KG et al," The Journal of Supercomputing, vol. 81, p. 474, 2025.

[23] P. Chaudhary, A. Singh, and B. Gupta, "Dynamic multiphase DDoS attack identification and mitigation framework to secure SDN-based fog-empowered consumer IoT Networks," Computers and Electrical Engineering, vol. 123, p. 110226, 2025.

[24] Z. Ullah, F. Arif, Q. M. U. Haq, N. A. Khan, I. U. Din, A. Almogren, et al., "Hybrid CNN-LSTM Model for DDoS Attack Detection in Internet of Things-based Healthcare Industry 5.0," IEEE Internet of Things Journal, 2025.

[25] A. Liatifis, P. Sarigiannidis, V. Argyriou, and T. Lagkas, "Advancing sdn from openflow to p4: A survey," ACM Computing Surveys, vol. 55, pp. 1-37, 2023.

[26] R. F. Fouladi, L. Karacay, U. Guelen, and E. U. Soykan, "A novel Distributed Denial of Service attack defense scheme for Software-Defined Networking using Packet-In message and frequency domain analysis," Computers and Electrical Engineering, vol. 120, p. 109827, 2024.

[27] A. El-Sayed, A. A. Toony, A. Tolba, F. Alqahtani, Y. Alginahi, and W. Said, "Deception and cloud integration: A multi-layered approach for DDoS detection, mitigation, and attack surface minimization in SD-IoT networks," Computers and Electrical Engineering, vol. 126, p. 110543, 2025.

[28] A. El-Sayed, W. Said, A. Tolba, Y. Alginahi, and A. A. Toony, "LBTMA: An integrated P4-enabled framework for optimized traffic management in SD-IoT networks," Internet of Things, vol. 28, p. 101432, 2024.

[29] W. I. Khedr, A. E. Gouda, and E. R. Mohamed, "P4-HLDMC: A novel framework for DDoS and ARP attack detection and mitigation in SD-IoT networks using machine learning, stateful P4, and distributed multi-controller architecture," Mathematics, vol. 11, p. 3552, 2023.

[30] H. Younis and M. M. Hamarsheh, "ONOS Flood Defender: A Real-Time Flood Attacks Detection and Mitigation System in SDN Networks," Concurrency and Computation: Practice and Experience, vol. 37, p. e8388, 2025.

[31] O. Yousuf and R. N. Mir, "DDoS attack detection in Internet of Things using recurrent neural network," Computers and Electrical Engineering, vol. 101, p. 108034, 2022.