



1 INTRODUCTION

Modern networks face increasingly sophisticated attacks, including zero-day exploits unknown to security vendors. Traditional signature-based IDS such as Snort rely on pre-defined rules and cannot detect previously unseen threats.

Anomaly-based detection using ML can flag unseen threats but suffers from high false-positive rates in isolation. This project asks whether a **hybrid** ML approach can deliver both high accuracy and low false alarms in a deployable form.

2 AIM & OBJECTIVES

Aim: Design, build and evaluate a hybrid ML-IDS that outperforms a signature-based baseline (Snort) in both detection accuracy and false-positive rate.

- Train an XGBoost classifier on labelled CIC-IDS2017 traffic.
- Train a deep autoencoder on benign-only traffic for anomaly scoring.
- Fuse the two via weighted voting and tune for low false-positive rate.
- Benchmark against Snort 2.9 on identical traffic.
- Deploy a live prototype on Raspberry Pi 4 via SPAN port mirroring.

3 METHODOLOGY

Dataset. CIC-IDS2017 (2.83M flows, 78 features, 14 attack classes including DDoS, Brute-Force, Botnet and Web Attacks). Stratified 70/15/15 split.

Model A — XGBoost. Supervised gradient-boosted trees (300 estimators, depth 8) trained on labelled flows. Outputs probability of attack.

Model B — Deep Autoencoder. Unsupervised 78–48–24–12–24–48–78 network trained on benign flows only; reconstruction error scores anomalies.

Fusion. Weighted voting $w_1 \cdot P(\text{attack}) + w_2 \cdot A(\text{anomaly})$, with weights tuned on validation to minimise false-positive rate at fixed recall.

4 SYSTEM ARCHITECTURE





Fig. 1: Hybrid ML-IDS pipeline. Both models score the same flow; weighted voting balances supervised confidence with anomaly score.

5 LIVE PROTOTYPE

The trained models were exported and deployed to a **Raspberry Pi 4 (4GB)** running a Python detection daemon.

- Cisco Catalyst 2960 SPAN port mirrors live traffic.
- CICFlowMeter extracts 78 features per flow.
- Real-time inference at ~120 flows/sec; alerts logged via syslog.



6 KEY RESULTS

99.4%	98.9%
Accuracy	Recall (attack)
0.6%	8.4 ms
False-positive rate	Avg. inference latency

7 BENCHMARK vs. SNORT

Attack class	Snort 2.9	Hybrid	Δ
DDoS	92.1%	99.7%	+7.6
Brute-Force	71.4%	97.2%	+25.8
Zero-day (held-out)	0%	84.3%	+84.3

8 CONCLUSIONS & FUTURE WORK

The hybrid system combined XGBoost's precision on known attacks with the autoencoder's ability to flag unfamiliar traffic, achieving **99.4% accuracy at a 0.6% false-positive rate** — a meaningful improvement over Snort, particularly on the held-out zero-day class where signature matching scored 0%.

Real-time deployment on a Raspberry Pi 4 confirms the approach is viable on commodity hardware at small-network scale.

Future work

- Online learning to adapt to traffic drift.
- Encrypted-traffic features (TLS handshake metadata).
- Adversarial robustness testing against evasion attacks.
- Federated training across multiple sites.

Scan to read
the full report

