



Gannon University
Erie, PA 16541

Towards a Serverless Intelligent Firewall: AI-Driven Security, and Zero-Trust Architectures

Md Anisur Rahman Chowdhury

Student ID: 3195493 · chowdhur014@gannon.edu · Master's of Information Technology

Professor: Ronny C. Bazan-Antequera, Ph.D. · bazanant001@gannon.edu · Department of Computer and Information Science, College of Engineering & Business, Gannon University

Aishwarya

Student ID: 3214245 · aishwary001@gannon.edu · Master's of Software Engineering

CG 2026

Gannon
University

ABSTRACT

The ephemeral, stateless nature of serverless computing presents unprecedented challenges to traditional network security. Conventional perimeter-based defenses and rule-based IDS are ill-equipped to protect transient, event-driven functions lacking persistent state.

This research presents a **Serverless Intelligent Firewall (SIF)** combining **LSTM-based deep learning** intrusion detection with **Zero-Trust Architecture (ZTA)** enforcement for adaptive, real-time threat detection in cloud-native environments.

The 3-layer LSTM achieved **98% accuracy, precision, recall, and F1-score** on CICIDS2017, outperforming SVM (88.4%), Decision Tree (90.2%), and CNN (93%).

Keywords: Serverless · LSTM · Zero-Trust · IDS · CICIDS2017 · AWS Lambda

RESEARCH OBJECTIVES / PURPOSE OF THE STUDY

- Design a **deep learning IDS** optimized for stateless serverless cloud environments
- Develop an **LSTM architecture** capturing temporal flow dependencies for coordinated attack detection
- Implement **strategic undersampling** to resolve class imbalance across 2.8M CICIDS2017 records
- Benchmark against **SVM, Decision Tree, and CNN** using standardized evaluation metrics
- Integrate **NIST SP 800-207 Zero-Trust** — continuous verification, least-privilege, breach assumption
- Deploy on **AWS Lambda** demonstrating <100 ms real-time inference latency

RESEARCH QUESTIONS

- RQ1:** Can serverless architectures maintain <100 ms inference latency for real-time IDS?
- RQ2:** Does LSTM provide superior accuracy over CNN/SVM/DT on CICIDS2017?
- RQ3:** Can ZTA be integrated without significant operational overhead?

All three answered affirmatively — statistically significant ($p < 0.05$).

DATASET — CICIDS2017

Total Records	2,830,540 flow records
Features	78 numerical traffic features
Classes	BENIGN · DDoS · DoS · PortScan · Other
Balancing	Undersampled to 50,000 / class
Split	80% train / 20% test (stratified)

METHODOLOGY (AN OVERVIEW)

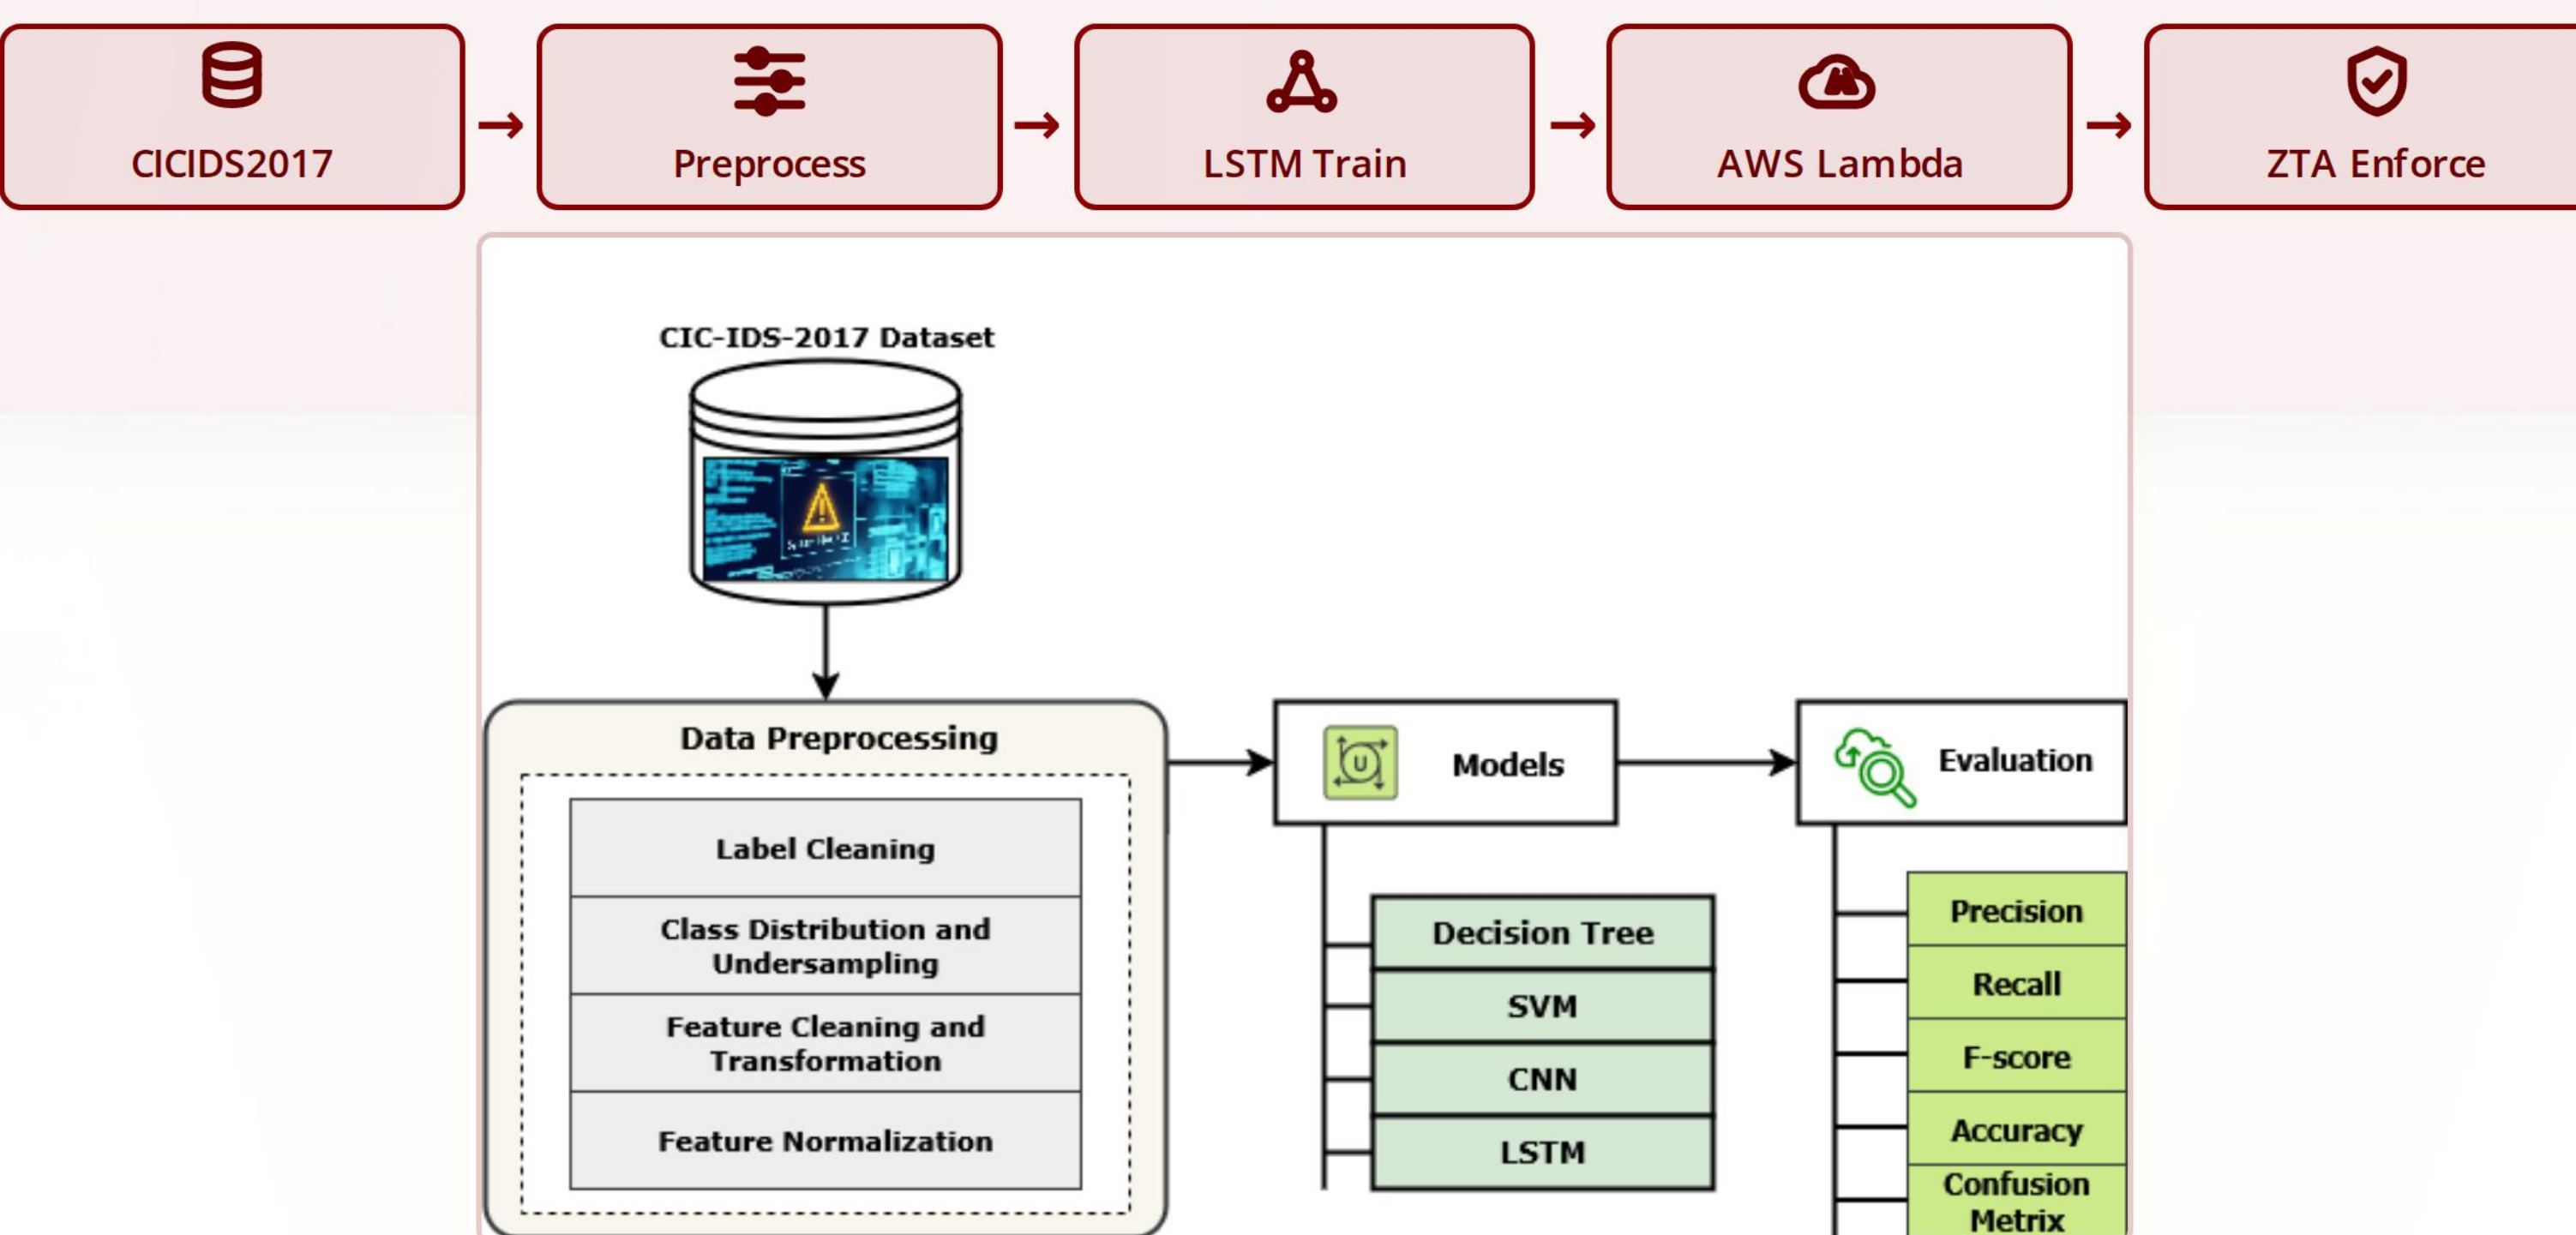




Fig. 1 — Serverless Intelligent Firewall system architecture

LSTM Parameter	Value
Input features	78 numerical flow features
Hidden units	3 layers: 128 → 64 → 32
Dropout	0.3 after each LSTM layer
Optimizer / LR	Adam · 0.001
Epochs / Patience	120 epochs · Early stop = 10
Lambda latency	<15 ms warm · <100 ms cold



RESULTS & EVALUATION

	98% Accuracy	98% Precision	98% Recall	98% F1-Score
SVM	88.4%			
Dec.Tree	90.2%			
CNN	93.0%			
LSTM ★	98.0%			

Model	Prec.	Recall	F1
SVM	84.1%	77.8%	80.8%
Dec.Tree	87.6%	81.3%	84.3%
CNN	95.1%	85.4%	89.9%
LSTM	98.0%	98.0%	98.0%




Fig. 4 — Confusion Matrix




Fig. 5 — Class-wise Performance

Class	Precision	Recall	F1	Correct/Total
BENIGN	0.99	0.94	0.96	9,446 / ~10,000
DDoS	0.99	1.00	0.99	9,989 / 10,000
DoS	0.98	0.99	0.98	9,909 / ~10,000
PortScan	0.99	1.00	0.99	9,982 / 10,000
Other	0.94	0.99	0.96	3,556 / 3,606
Macro Avg	0.978	0.984	0.976	—

RESEARCH CONTRIBUTIONS

- Novel SIF Framework:** First integration of LSTM-based IDS with serverless AWS Lambda deployment enforcing NIST SP 800-207 ZTA in a unified, production-viable architecture
- State-of-the-Art Accuracy:** 98% across all metrics, surpassing federated IDS (97%), hybrid CNN-LSTM (95%), and all classical baselines
- Scalable & Zero-Cost Idle:** Lambda scales to zero — eliminates persistent attack surfaces; cost-per-invocation model superior to always-on EC2
- Temporal Feature Exploitation:** LSTM memory gates capture flow evolution patterns inaccessible to CNN convolutions or SVM kernels
- Sub-100ms Latency:** 15 ms warm execution; 1.1 s cold-start via TorchScript + provisioned concurrency
- ZTA with Minimal Overhead:** mTLS + IAM adds ~8 ms (5%) while providing per-flow continuous verification
- Open Reproducibility:** Full code, trained model (.pth), Docker config, Lambda scripts on GitHub

ZERO-TRUST ARCHITECTURE

- **Identity:** mTLS via AWS Cognito on every flow
- **Classify:** LSTM inference · confidence >0.85
- **Policy:** IAM least-privilege condition evaluation
- **Enforce:** ALLOW (BENIGN) / BLOCK + SNS alert
- **Re-verify:** Every 5 min or 500 flows

~8 ms overhead per decision (5% of total execution time)

CONCLUSION & FUTURE WORK

The SIF framework proves **serverless + LSTM + ZTA are synergistic**: serverless eliminates persistent attack surfaces; LSTM captures temporal threat patterns; ZTA converts ML predictions into auditable security decisions.

Future directions:

- Federated learning across multi-tenant environments
- Adversarial robustness evaluation & hardening
- Transformer-based traffic classifiers
- Encrypted TLS traffic analysis
- XAI (SHAP/LIME) for analyst dashboards

SOTA COMPARISON

Reference	Model	Accuracy
Proposed (2025)	3-Layer LSTM	98.0%
Neto et al. 2022	Federated IDS	97.0%
Bamber et al. 2025	CNN-LSTM	95.0%
Altunay et al. 2023	CNN+LSTM	93.2%

REFERENCES

- Sharafaldin et al. (2018). CICIDS2017. *ICISSP*
- Rose et al. (2020). NIST SP 800-207 Zero Trust Architecture
- Hochreiter & Schmidhuber (1997). LSTM. *Neural Computation*
- Altunay & Albayrak (2023). CNN+LSTM. *ESTIJ*
- Bamber et al. (2025). CNN-LSTM IDS. *arXiv*
- Kingma & Ba (2015). Adam Optimizer. *ICLR*
- AWS (2024). Lambda Developer Guide