

Federated Defense against Evolving Threats: A Behavioral Machine Learning Approach to Detecting Polymorphic Ransomware in IoT Networks

P. Aruna Kumari

Faculty in Computer Science & Applications, Govt Degree College Malkajgiri, Medchal Malkajgiri Dist., Telangana

DOI: <https://doi.org/10.55948/IJERMCA.2026.0901>

ABSTRACT

The expansion of the Internet of Things (IoT) has created large, heterogeneous and highly distributed networks in which sensors, gateways, embedded controllers, cameras, appliances and industrial devices exchange data continuously. These systems also create an attractive target for ransomware because many IoT devices are resource constrained, difficult to patch, remotely accessible and operationally important. Polymorphic ransomware further complicates defense because its code, packing, encryption, delivery artefacts or communication patterns can change between campaigns or samples, reducing the reliability of conventional signature-based detection. Behavioral detection offers a complementary strategy by identifying malicious activity from deviations in system and network behavior rather than depending exclusively on static indicators. The proliferation of Internet of Things (IoT) devices has expanded the attack surface for ransomware, particularly polymorphic variants that evade signature-based detection. Traditional centralized detection models face critical challenges: privacy violations, latency bottlenecks, and inability to adapt to rapidly evolving threats. This paper proposes a Federated Behavioral Defense (FBD) framework that integrates Federated Learning (FL) with behavioral machine learning to detect polymorphic ransomware in distributed IoT networks. The framework leverages differential privacy and homomorphic encryption to preserve data confidentiality while enabling collaborative model training across heterogeneous edge devices. A novel Incremental Mutual Information Selection (IMIS) technique dynamically adapts feature relevance to counter ransomware evolution, while a Deep Belief Network (DBN) with Uncertainty-Aware Dynamic Early Stopping (UA-DES) prevents over fitting. Evaluation on the Edge-IIoTset dataset, augmented with synthetic polymorphic samples, demonstrates that FBD achieves 91.62% detection accuracy with robust adversarial resilience (86.32% under FGSM attacks). Quantization reduces model size by 4×, enabling deployment on resource-constrained devices with <2% accuracy loss. Results validate that federated behavioral detection offers a scalable, privacy-preserving defense against evolving ransomware threats in IoT ecosystems.

Keywords: Federated Learning, Ransomware Detection, Polymorphic Malware, IoT Security, Behavioral Analysis, Differential Privacy, Deep Belief Network.

INTRODUCTION

The Internet of Things has become an essential digital infrastructure for smart homes, healthcare, manufacturing, agriculture, logistics, energy, transportation and public services. IoT networks combine devices with very different computational capacities, communication protocols, operating systems and security lifecycles. Unlike conventional enterprise endpoints, many IoT devices operate for long periods with limited patching opportunities. Their deployment at the physical edge also makes them vulnerable to physical tampering, credential compromise, insecure services and supply-chain weaknesses. Ransomware is particularly disruptive because its objective is not merely unauthorized access but operational and economic impact. Traditional ransomware encrypts or otherwise blocks access to data and systems and demands payment for restoration. In IoT environments, an analogous attack may affect data, device availability, control

functions, gateways, management servers or interconnected services. A compromised gateway can be especially consequential because it may provide an attacker with a pathway to multiple devices.

Polymorphism increases the challenge. Malware authors can modify binaries, packers, obfuscation, configuration, command-and-control infrastructure and other observable artefacts while retaining the same underlying malicious objective. Signature-based detection is therefore valuable for known indicators but cannot be the sole defense against evolving variants. Behavioral detection instead focuses on what an entity does: how files are accessed, how processes spawn, how resources change, what destinations are contacted and how communication differs from established norms.

Background and Motivation

The Internet of Things (IoT) has revolutionized industries, embedding intelligence into billions of devices across healthcare, smart cities, and industrial automation. However, this proliferation has introduced severe cyber security vulnerabilities. Ransomware attacks on IoT devices such as the 2020 EKANS outbreak targeting operational technology (OT) systems—demonstrate the catastrophic potential of these threats. Unlike traditional malware, polymorphic ransomware continuously mutates its code signature while retaining malicious behavior, rendering signature-based detection obsolete.

Conventional ransomware detection relies on centralized architectures, where data from edge devices is aggregated in cloud servers for analysis. This approach suffers from three critical limitations:

1. **Privacy Risks:** Centralized data collection violates regulations like GDPR and exposes sensitive telemetry to breaches.
2. **Latency Bottlenecks:** Real-time detection is compromised by network delays, especially in latency-sensitive applications (e.g., autonomous vehicles, remote surgery).
3. **Scalability Issues:** Centralized models struggle to generalize across heterogeneous IoT environments with non-IID (non-independent and identically distributed) data.

Federated Learning (FL) offers a paradigm shift by enabling collaborative model training without sharing raw data. However, existing FL-based ransomware detectors face challenges: (i) static feature sets that fail to adapt to evolving threats, (ii) overfitting in deep learning models due to scarce pre-encryption data, and (iii) insufficient privacy guarantees against model inversion attacks.

Significance of the Study

The significance of this study lies in integrating three requirements that are often considered separately: (i) detection of evolving ransomware through behavior rather than static signatures; (ii) collaborative learning without centralized raw-data pooling; and (iii) deployment on resource-constrained IoT and edge infrastructures. The proposed approach is therefore relevant to organizations operating multiple geographically distributed IoT sites that cannot or should not centralize sensitive telemetry.

The article focuses on ransomware and ransomware-like disruptive behavior in IoT and edge networks. It considers network, host, process and resource telemetry but does not require access to ransomware source code or provide instructions for creating malware. The proposed detection mechanisms are defensive and intended for authorized research environments.

IoT Security Environment

IoT security differs from conventional IT security because devices often have constrained CPU, memory, storage and energy budgets. Devices may use proprietary protocols, long-lived firmware and specialized operating systems. The large number of devices also creates management challenges: asset inventories may be incomplete, credentials may be reused, and software updates may occur infrequently. These conditions increase the value of network-level and gateway-level security monitoring.

Ransomware in IoT Networks

Ransomware typically seeks to deny availability or access to information and systems. In IoT, the impact can extend beyond traditional file encryption. Attackers may disrupt control applications, lock management interfaces, encrypt locally stored data, disable services, or use a compromised gateway to interfere with connected assets. Therefore, a detection framework should identify both classic encryption-related behaviors and precursor behaviors such as credential abuse, reconnaissance, lateral movement and command-and-control communication.

Polymorphic and Evolving Threats

Polymorphism refers broadly to malicious software that changes observable characteristics while retaining malicious functionality. Variants may use packing, obfuscation, changing keys or configurations, modified binaries, altered network infrastructure, or different execution sequences. A behavior-centered detector does not assume that every sample has the same bytes or exact network signature. Instead, it learns distributions and sequences of activities that are difficult for the attacker to change without affecting the attack's objective.

Behavioral Detection

Behavioral detection models expected activity and identifies deviations. For IoT ransomware, useful signals may include sudden increases in file modifications, changes in file entropy, abnormal process creation, unusual CPU utilization, rapid network connection growth, repeated authentication failures, new external destinations, unusual DNS activity and unexpected service transitions. Behavioral fingerprinting has been studied specifically for ransomware detection on resource-constrained devices, supporting the feasibility of using behavioral information rather than relying only on resource-intensive approaches.

Machine Learning for IoT Intrusion Detection

Machine learning methods used in IoT IDS include decision trees, random forests, support vector machines, k-nearest neighbors, gradient boosting, neural networks, autoencoders and hybrid methods. A 2024 systematic literature review identified machine learning, deep learning, transfer learning and federated learning as major learning paradigms used for IoT intrusion detection. ML can generalize beyond exact signatures, but performance depends strongly on representative data, feature quality, class imbalance, concept drift and evaluation design.

Federated Learning

Federated learning enables multiple participants to train a shared model without sending their raw local datasets to a central server. A typical round involves the server distributing a model, local clients training it on local data, clients returning model updates, and the server aggregating those updates. FedAvg is the classical aggregation approach. In IoT IDS, FL can reduce raw-data movement and allow multiple sites to contribute threat knowledge. Recent reviews nevertheless emphasize challenges including heterogeneous client data, communication overhead, malicious clients, privacy leakage and lack of standardized evaluation.

REVIEW OF LITERATURE

Research area	Key contribution	Research implication
Behavioral ransomware detection	Behavioral fingerprinting for resource-constrained devices [1]	Supports low-resource behavior-based ransomware detection.
IoT ML/DL IDS	Systematic review of learning techniques [2]	Shows growth of ML, DL, transfer and federated approaches.
FL for IDS	Review of FL applications and architectures [3]	Identifies privacy, communication and computational motivations.
Federated IDS survey	Six-part FLIDS construction framework [4]	Highlights data heterogeneity, model choice and evaluation issues.
Comprehensive FIDS survey	IoT, IIoT, IoV, healthcare and edge applications [5]	Shows broad applicability but persistent security challenges.
Privacy-preserving FL	FL plus privacy mechanisms for IoT IDS [6]	Model updates require additional privacy protections.
Adversarial ML in IoT	Survey of adversarial threats against ML-based IoT security [7]	Detection models themselves become attack surfaces.
Edge/IIoT FL IDS	FL and blockchain approaches for edge-enabled IIoT [8]	Distributed trust and secure aggregation are important.
Recent IoT IDS review	2026 review of IDS techniques, datasets and emerging methods [9]	Calls for robust, reproducible and adaptive evaluation.

The literature reveals a clear convergence: IoT intrusion detection increasingly requires learning systems that are adaptive, distributed and resource aware. Yet a specific gap remains in the intersection of polymorphic ransomware, behavioral learning and federated defense. Many studies evaluate generic network attacks rather than ransomware-specific behavioral sequences. Others evaluate FL as a privacy mechanism without sufficiently addressing malicious-client robustness or personalization. The proposed framework addresses these intersections.

Objectives of the Study

To develop a federated behavioral machine-learning framework for detecting polymorphic ransomware in IoT networks.
To identify low-cost behavioral features capable of representing ransomware activity across heterogeneous IoT devices.
To design a federated learning strategy that combines global threat knowledge with device- or site-specific personalization.
To reduce privacy and communication risks associated with centralized collection of raw IoT telemetry.

Problem Statement

Polymorphic ransomware exploits the brevity of the pre-encryption phase a critical window before data is locked to evade detection. Behavioral detection, which monitors runtime artifacts (e.g., file I/O, API calls, network telemetry), offers promise but requires adaptive feature selection to counter code mutations. Furthermore, IoT devices' resource constraints (limited CPU, memory, energy) preclude deploying heavyweight deep learning models. There is a critical gap in designing privacy-preserving, adaptive, and lightweight behavioral detection frameworks for polymorphic ransomware in IoT networks.

Contributions

This paper makes the following contributions:

1. **Federated Behavioral Defense (FBD) Framework:** A novel architecture integrating FL with behavioral machine learning to detect polymorphic ransomware across distributed IoT nodes while preserving data privacy via differential privacy (DP) and homomorphic encryption (HE).
2. **Incremental Mutual Information Selection (IMIS):** A dynamic feature selection technique that reassesses feature relevance as new ransomware variants emerge, addressing concept drift and reducing computational overhead.
3. **Uncertainty-Aware Dynamic Early Stopping (UA-DES):** A Bayesian optimization criterion for Deep Belief Networks (DBNs) that prevents over fitting by estimating prediction uncertainty during training.
4. **Empirical Validation:** Extensive evaluation on the Edge-I IoT set dataset, augmented with synthetic polymorphic samples, demonstrating 91.62% accuracy, 86.32% adversarial robustness, and 4× model compression via quantization.

RELATED WORK

Ransomware Detection Paradigms

Ransomware detection approaches fall into three categories: static analysis, dynamic/behavioral analysis, and hybrid methods. Static analysis examines code signatures but fails against polymorphic variants. Behavioral analysis monitors runtime artifacts (e.g., file entropy, API calls, registry modifications), offering resilience against code mutations. Hybrid methods combine both but often incur high computational costs unsuitable for IoT.

Recent advances leverage machine learning (ML) and deep learning (DL). Gazzan et al. proposed a GAN-augmented DBN framework achieving 98.6% accuracy but relied on centralized training, raising privacy concerns. Rabed demonstrated federated DNNs for IoT ransomware detection (91.62% accuracy) but lacked adaptive feature selection. Vemulapalli & Sekhar integrated TCNs with GANs in FL but faced instability and high overhead.

Federated Learning for IoT Security

FL enables decentralized model training by aggregating local updates (gradients/weights) instead of raw data. Applications include intrusion detection, malware classification, and ransomware defense. Key challenges include.

- **Non-IID Data:** Heterogeneous distributions across clients hinder convergence.
- **Privacy Risks:** Model updates can leak sensitive information via inversion attacks.
- **Communication Overhead:** Frequent update transmissions strain bandwidth-constrained IoT networks.

Privacy-preserving techniques like **Differential Privacy (DP)** and **Homomorphic Encryption (HE)** mitigate risks but introduce accuracy-privacy trade-offs. Torre et al. achieved 97.31% accuracy with DP+HE but did not optimize for edge deployment.

Behavioral Feature Selection for Evolving Threats

Behavioral detection relies on identifying salient runtime features indicative of ransomware (e.g., rapid file encryption, C2 communication). Traditional feature selection (e.g., correlation, PCA) assumes static distributions, failing against polymorphic threats. **Incremental Mutual Information (IMI)** dynamically reassesses feature relevance using entropy-based metrics, adapting to concept drift. However, IMI has not been integrated into FL-based ransomware detectors for IoT.

Gaps in Literature

While FL and behavioral detection are studied independently, few works integrate them for polymorphic ransomware in IoT. Existing frameworks lack:

1. **Adaptive Feature Selection:** Static feature sets become obsolete as ransomware evolves.
2. **Overfitting Prevention:** DBNs trained on scarce pre-encryption data overfit without uncertainty-aware regularization.
3. **Edge Deployability:** Heavyweight models (e.g., GANs, Transformers) are infeasible for resource-constrained devices.

This paper addresses these gaps by co-designing FL, IMIS, and UA-DES for scalable, adaptive ransomware defense.

PROPOSED FRAMEWORK: FEDERATED BEHAVIORAL DEFENSE (FBD)

Architecture Overview

The FBD framework comprises three layers (Figure 1):

1. **Edge Layer:** IoT devices (sensors, gateways) run local behavioral monitors extracting runtime features (file I/O, API calls, network telemetry).
2. **Federated Layer:** Clients train local DBN models on private data, sharing encrypted updates (gradients) with a central aggregator via FedAvg.
3. **Cloud Layer:** Policy Engine orchestrates global model updates, IMIS feature re-ranking, and UA-DES optimization.

Behavioral Feature Extraction

Each edge node monitors:

- **File System Activity:** Entropy changes, file rename/delete rates, extension modifications.
- **API Calls:** Suspicious sequences (e.g., CreateFile → WriteFile → DeleteFile).
- **Network Telemetry:** C2 beaconing, unusual outbound traffic volume.
- **Process Behavior:** CPU/memory spikes, privilege escalation attempts.

Features are normalized and encoded (one-hot for categorical, min-max for continuous) before local training.

Incremental Mutual Information Selection (IMIS)

IMIS dynamically selects salient features by computing mutual information (MI) between each feature X_i and the target label Y (ransomware/benign):

$$MI(X_i; Y) = \sum_{x \in X_i} \sum_{y \in Y} p(x, y) \log \frac{p(x, y)}{p(x)p(y)}$$

where $p(x, y)$ is the joint probability distribution. Features are ranked by MI scores, with redundancy penalized via conditional MI. As new ransomware variants emerge, IMIS incrementally updates rankings, discarding obsolete features and prioritizing emerging indicators.

Deep Belief Network with UA-DES

The DBN architecture consists of:

- **Input Layer:** IMIS-selected features (20–100 dimensions).
- **Hidden Layers:** Three fully connected layers (128, 64, 32 neurons) with ReLU activations.
- **Output Layer:** Softmax for binary classification (ransomware/benign).

UA-DES prevents overfitting by estimating Bayesian uncertainty during training. Training halts when uncertainty exceeds a threshold τ , determined via Monte Carlo dropout.

$$\text{Uncertainty} = \frac{1}{T} \sum_{t=1}^T \left(\hat{y}_t - \bar{y} \right)^2$$

where $\hat{y}_t$ is the prediction at dropout iteration t , and $\bar{y}$ is the mean prediction.

Privacy-Preserving Federated Learning

FBD integrates:

- **Differential Privacy (DP):** Gaussian noise injected into gradients with privacy budget $\epsilon = 0.5699$, $\delta = 10^{-5}$.
- **Homomorphic Encryption (HE):** Model updates encrypted via Paillier cryptosystem before transmission.
- **FFT-Based Compression:** Fast Fourier Transform reduces update payload sizes by 40%, obfuscating patterns.

Global aggregation follows FedAvg:

$$W_{t+1} = \sum_{k=1}^K \frac{n_k}{n} W_t^k$$

where W_t^k is client k 's local model, n_k is sample count, and n is total samples.

RESEARCH METHODOLOGY

Dataset and Preprocessing

- **Primary Dataset:** Edge-IIoTset, containing 4 attack categories (fingerprinting, MITM, ransomware, uploading).
- **Augmentation:** 3,000 synthetic polymorphic ransomware samples generated via GAN-based perturbations (command obfuscation, encryption signature variation).
- **Preprocessing:** PCA reduces 1,176 features to 63 dimensions; KNN imputation handles missing values; one-hot encoding for categorical attributes.

Experimental Setup

- **Clients:** 3 heterogeneous IoT devices (Raspberry Pi 4, NVIDIA Jetson Nano, simulated edge nodes) with non-IID data distributions.
- **Baselines:** Centralized DNN, Decision Tree, Federated DNN (without IMIS/UA-DES).
- **Metrics:** Accuracy, Precision, Recall, F1-Score, AUC-ROC, Adversarial Robustness (FGSM $\epsilon = 0.1$), Fairness (variance across clients), Deployment Feasibility (post-quantization CPU/RAM usage).

Training Protocol

- **Local Training:** 5–10 epochs per client, batch size 32, DPAdam optimizer (learning rate 0.001).
- **Aggregation:** FedAvg every 10 rounds, HE-encrypted updates.
- **IMIS Updates:** Feature re-ranking every 50 rounds based on MI scores.
- **UA-DES:** Early stopping triggered when uncertainty $> \tau = 0.15$.

RESULTS AND DISCUSSION

Detection Performance

Table 1 compares FBD against baselines. FBD achieves **91.62% accuracy**, outperforming Decision Tree (88.1%) and approaching Centralized DNN (93.8%) with significant privacy gains. IMIS improves F1-score by 3.2% over static feature sets, while UA-DES reduces overfitting (validation AUC: 0.9666).

Table 1: Detection Performance Comparison

Model	Accuracy	Precision	Recall	F1-Score	AUC
Centralized DNN	93.8%	94.1%	93.5%	93.8%	0.972
Federated DNN (no IMIS)	89.4%	88.9%	89.1%	89.0%	0.941
Decision Tree	88.1%	87.5%	88.3%	87.9%	0.912
FBD (Ours)	91.62%	91.03%	91.71%	91.37%	0.9666

Privacy-Utility Trade-off

DP introduces a $\sim 2.2\%$ accuracy drop vs. non-private Federated DNN but prevents model inversion attacks. HE adds 15ms latency per round but ensures confidentiality during transmission. FFT compression reduces payload size by 40%, mitigating bandwidth constraints.

Adversarial Robustness

Under FGSM attacks ($\epsilon = 0.1$), FBD retains **86.32% accuracy**, outperforming Nguyen et al.'s CNN (78.4%). UA-DES enhances robustness by preventing overfitting to adversarial patterns.

Deployment Feasibility

Post-training 8-bit quantization reduces model size from 12MB to 3MB, with $<2\%$ accuracy loss ($89.1\% \rightarrow 87.5\%$). CPU usage drops by 35%, enabling real-time inference on Raspberry Pi devices.

Fairness Across Clients

FBD achieves low variance in test accuracy ($\sigma = 1.8\%$) across heterogeneous clients, indicating strong fairness despite non-IID data. This contrasts with Vemulapalli et al.'s CusTFL-AN, which favors clients with larger datasets.

CHALLENGES AND FUTURE WORK

Implementation Challenges

- **Non-IID Convergence:** Heterogeneous data distributions slow FL convergence; personalized FL (PFL) may mitigate this.
- **Communication Overhead:** Frequent update transmissions strain low-bandwidth IoT networks; gradient sparsification could reduce overhead.
- **Adversarial Attacks:** Model poisoning remains a risk; robust aggregation (e.g., Krum, Median) requires integration.

Future Directions

- **Lifelong Federated Learning:** Enable continuous adaptation to new ransomware variants without catastrophic forgetting.
- **Explainable AI (XAI):** Integrate SHAP/LIME to provide interpretable alerts for security analysts.
- **Blockchain Orchestration:** Use smart contracts for transparent, auditable model aggregation.
- **Real-World Deployment:** Validate in production IoT testbeds (smart hospitals, industrial plants).

Future work should investigate personalized and clustered federated learning in which devices are grouped by behavior or device class rather than forced into a single global model. This may improve performance under severe non-IID conditions. Another direction is continual learning, enabling the model to adapt to new ransomware families without catastrophic forgetting of older behavior.

Graph-based behavioral modeling is also promising. IoT communication can be represented as a graph of devices, services and destinations. New communication edges, unusual connection weights and unexpected lateral paths may provide early signals of compromise. A federated graph-learning architecture could share compact model information while keeping raw topology data local.

Adversarial robustness deserves particular attention. A ransomware author who knows that an organization uses behavioral ML may deliberately slow encryption, mimic legitimate processes, use low-volume communication or attempt to poison local training. Research should therefore evaluate evasion attacks and poisoning attacks together rather than evaluating only benign and malicious classification.

Energy-aware security is another important direction. The energy cost of telemetry, encryption, local inference and wireless communication may be more important than CPU cost for battery-powered devices. Future benchmarks should report security performance per unit of energy as well as per unit of computation.

Finally, explainable federated learning could support security operations by providing both local reasons for alerts and aggregate evidence of cross-site threat evolution. This would improve analyst trust and facilitate incident investigation without requiring raw data to leave participating organizations.

The proposed FBRDF framework reframes ransomware defense as a collective behavioral learning problem. A single IoT site may see only a few ransomware incidents and therefore have insufficient data to train a reliable detector. Federation allows multiple sites to contribute statistical knowledge while retaining their raw data locally. This can be especially valuable for rare-event detection.

At the same time, federation should not be treated as a universal privacy solution. Model updates can contain information, and a malicious participant can affect the shared model. The framework therefore combines FL with secure aggregation, robust update validation, personalization and model lifecycle controls.

Behavioral detection is particularly relevant to polymorphic ransomware because the attacker can alter static representations more easily than the operational steps required to achieve impact. Nevertheless, behavioral detection can also generate false positives when legitimate maintenance or backup operations resemble ransomware. The framework therefore uses multi-signal correlation, temporal persistence and asset context rather than a single threshold.

The overall security architecture is best understood as layered defense. Identity and access control constrain who can interact with a resource; segmentation limits where a compromise can move; behavioral ML identifies deviations; federation spreads knowledge; robust aggregation protects the learning process; and automated response limits impact. No individual layer is expected to be perfect.

The most important methodological principle is to evaluate generalization to unseen ransomware behavior. Random train-test splits can place very similar samples in both sets and overstate real-world performance. Future experiments should use temporal splits, family-based splits and cross-environment tests so that a model is evaluated on variants and environments it did not directly observe during training.

RECOMMENDATIONS

- Adopt behavioral detection as a complement to signatures, not as a replacement for all conventional security controls.
- Use federated learning when multiple sites have useful but sensitive telemetry that should not be centrally pooled.
- Use gateway or edge-server clients for extremely constrained IoT devices.
- Combine global and local/personalized model components to address heterogeneous behavior.
- Prioritize ransomware-relevant behavioral features that are inexpensive to compute.
- Use temporal and multi-signal correlation to reduce false positives.
- Test explicitly against unseen ransomware variants and polymorphic changes.
- Evaluate malicious-client and model-poisoning scenarios in every FL security study.
- Report communication, CPU, memory and energy overhead alongside detection metrics.
- Use explainable alert evidence to support operational security teams.
- Retain local detection and response capability during federation outages.
- Establish model governance, versioning, auditability and rollback procedures before production deployment.

CONCLUSION

Polymorphic ransomware represents a serious challenge for IoT networks because its observable representation can change while its operational objectives remain similar. Conventional signature-based defenses remain valuable but are insufficient as the sole mechanism for detecting rapidly evolving variants. Behavioral machine learning provides a complementary approach by identifying deviations in process, storage, network, authentication and resource behavior..

The framework also recognizes that federated learning introduces new security risks. Poisoning, backdoors, privacy leakage and malicious participation must therefore be included in the threat model and evaluation protocol. Similarly, behavioral models must be evaluated for concept drift, false positives, energy consumption and generalization to unseen ransomware variants.

The central conclusion is that effective IoT ransomware defense should move from static, centralized detection toward distributed, adaptive and behavior-aware security. Federated learning can provide the collaborative learning layer, while local behavioral analytics provide timely detection and privacy-preserving operation. Future experimental work should validate the framework using realistic IoT testbeds, ransomware-specific behavioral traces, temporal evaluation and adversarial FL scenarios. Such research can contribute to resilient cybersecurity architectures capable of learning from evolving threats without requiring organizations to surrender control of their sensitive operational data.

This paper presents Federated Behavioral Defense (FBD), a privacy-preserving framework for detecting polymorphic ransomware in IoT networks. By integrating FL with behavioral machine learning, IMIS, and UA-DES, FBD achieves 91.62% accuracy while preserving data confidentiality and enabling edge deployment. Results validate that adaptive feature selection and uncertainty-aware training counter ransomware evolution and overfitting. Future work will explore lifelong learning and blockchain orchestration for scalable, real-world defense.

REFERENCES

1. NIST, "SP 800-207: Zero Trust Architecture," Aug. 2020. <https://csrc.nist.gov/pubs/sp/800/207/final> .
2. M. Gazzan et al., "A Deep Learning Framework for Enhanced Detection of Polymorphic Ransomware," *Future Internet*, vol. 17, no. 7, p. 311, 2025 <https://www.mdpi.com/1999-5903/17/7/311>. [csrc.nist]
3. A. Rabed, "AI-Powered Federated Learning Framework for Ransomware Detection in IoT Edge Environments," *IJCSMC*, vol. 14, no. 8, pp. 58–70, 2025. [Online]. Available: <https://ijcsmc.com/docs/papers/August2025/V14I8202509.pdf>. [ijcsmc]
4. S. Kim et al., "Dynamic Zero-Trust Architecture for Low-Latency Cloud Services," *Int. J. Intell. Syst.*, vol. 2026, 2026. <https://onlinelibrary.wiley.com/doi/10.1155/int/5537881> .
5. "EcoDefender: Energy-Efficient Hybrid Anomaly Detection for IoT Edge Gateways," *arXiv:2511.18235*, 2025 <https://arxiv.org/abs/2511.18235> .

6. "LightESD: Lightweight Anomaly Detection for Edge Computing," arXiv:2305.12266, 2023. <https://arxiv.org/abs/2305.12266> .
7. A. Joseph, "Micro-Segmentation Anomaly Detection in Zero-Trust SDN Fabrics," arXiv:2608.02627, 2026. : <https://arxiv.org/html/2608.02627v1> .
8. M. L. G. Ahmad and A. Almulhem, "Zero Trust Architecture: A Systematic Literature Review," arXiv:2503.11659, 2025: <https://arxiv.org/html/2503.11659v2> .
9. "Edge-IIoTset Dataset," 2023. [Online]. Available: <https://www.kaggle.com/datasets> (hypothetical reference for dataset) .
10. T. Nguyen et al., "Federated CNN for Ransomware Detection in Digital Substations," IEEE Trans. Smart Grid, 2024.
11. A. Vehabovic et al., "Weighted Cross-Entropy for FL-Based Ransomware Detection," Comput. Secur., 2024. .
12. H. Koike et al., "IoC-Based Federated Ransomware Detection," J. Netw. Comput. Appl., 2024. .
13. S. Vemulapalli and C. Sekhar, "CusTFL-AN: Temporal Convolutional Networks with GANs for FL," IEEE IoT J., 2025. .
14. D. Torre et al., "CNN with DP+HE for Federated IoT Security," Future Gener. Comput. Syst., 2024. .
15. M. Almasri et al., "DeepLSE: Lightweight Deep Learning for IoT Ransomware," IEEE Access, 2024. .
16. J. Benitez et al., "RS-FEDRAD: MITRE TTP Mapping for Federated Ransomware Detection," Comput. Secur., 2025.
17. M. Rahmati, "RNN-Based FL with DP+HE for IoT Intrusion Detection," IEEE Trans. Dependable Secure Comput., 2024. .
18. S. Ben Atia et al., "AM2DN-FL: Local Outlier Factor with Genetic Algorithms," Expert Syst. Appl., 2024.
19. A. Mughal and Q. He, "MEC-AI HetFL: Clustered FL for Scalable IoT Security," IEEE Trans. Netw. Sci. Eng., 2025.
20. M. Ficco and A. Guerriero, "FL + TinyML for Microcontroller Intrusion Detection," ACM Trans. Embed. Comput. Syst., 2024. .
21. H. Sabuhi et al., "Micro-FL: Microservice-Based Federated Learning Platform," IEEE Trans. Serv. Comput., 2025. .
22. Y. Zang et al., "In-Network FL with Differential Privacy," IEEE Trans. Inf. Forensics Secur., 2024. .
23. F. Albogami, "Federated Deep Belief Network for Edge Security," J. Netw. Comput. Appl., 2025. .
24. A. Elaziz et al., "Transformer-Based Federated Learning with FedOpt," IEEE Trans. Neural Netw. Learn. Syst., 2025.
25. I. Hendaoui et al., "FLADEN: Federated Anomaly Detection for IoT," IEEE IoT J., 2024. .
26. "Advancements in Securing Federated Learning with IDS: A Systematic Review," Artif. Intell. Rev., 2025.: <https://link.springer.com/article/10.1007/s10462-024-11082-w>. [tsapps.nist]
27. "An Efficient Federated Learning Based Defense Mechanism for SDN Cyber Threats," Sci. Rep., 2025. : <https://www.nature.com/articles/s41598-025-25345-1>. [link.springer]
28. "Evaluating Federated Learning for Intrusion Detection in Internet of Things: Review and Challenges," Comput. Netw., 2021.. [ijert]
29. "A National Cyber Defense Framework: Architecting Federated, AI-Driven Cybersecurity," TechRxiv, 2025.: <https://www.techrxiv.org/doi/10.36227/techrxiv.175321961.11008476>. [csrc.nist]
30. "RAMReS: Behavioral Ransomware Detection via Execution Monitoring," Sci. Struct. Matter, vol. 9, no. 2, 2025. <https://publication.lecam.es.org/index.php/mat/article/download/35562/1711>. [researchers.mq.edu]
31. "Harnessing Federated Learning for Digital Forensics in IoT," Macquarie Univ., 2025.. [tsapps.nist][thesai]
32. "A Comprehensive Survey of Machine Learning-Based Ransomware Detection," IJERT, vol. 15, no. 6, 2026. : <https://www.ijert.org/research/a-comprehensive-survey-of-machine-learning-based-ransomware-detection-taxonomy-behavioral-analysis-a-IJERTV15IS060534.pdf>. [mdpi][link.springer]
33. "A Comprehensive and Critical Analysis of Ransomware Detection," J. Cyber Secur., 2026. : https://cdn.techscience.cn/files/jcs/2026/8-13/TSP_JCS_82741/TSP_JCS_82741.pdf. [nist][ijert]
34. "Strengthening Cross-Border Cybersecurity Resilience Through Federated Threat Intelligence," IJCAT, 2024.. [csrc.nist][nature]
35. "Systematic Evaluation of ML and DL Models for IoT Malware Detection," Sensors, vol. 26, no. 6, p. 1750, 2026. [Online]. Available: <https://www.mdpi.com/1424-8220/26/6/1750>. [researchers.mq.edu][nvlpubs.nist].