



B-Prim: A Multidisciplinary Framework for IoT–Cloud Ecosystems Privacy – Preserving Real-Time Intrusion Detection and Isolation

Rajul Ravi*¹ Sandeep Kumar² and
Aashutosh K A Prashant³

*¹Tula's Institute, Dhoolkot, Dehradun, 248007, UK, India

²Tula's Institute, Dhoolkot, Dehradun, 248007, UK, India

³Bahra University, Wakhnaghat, 173234, HP, India
rajulravi223@gmail.com

Abstract: The worldwide network perimeter is now a pathwork of exploitable gaps due to billions of unprotected IoT devices. Current intrusion detection systems are confronted with a recurring also breaches the privacy regulations that contemporary deployments are required by law. This study presents B-PriM, a decentralized IoT security architecture that provides autonomous node isolation and real-time intrusion detection without ever disclosing raw traffic. B-PriM by combining a blockchain-backed decentralized reputation score system with a lightweight Edge-GRU with local differential privacy. We prove (epsilon delta) –differential privacy for each communication round and formalize detection-isolation as a combined optimization over privacy loss and inference delay. F1 of 95.8% +/- 0.5% at operating configuration $\epsilon = 1.0(\delta = 1e-5)$, sub- 3.8 ms latency, and 91.0% accuracy at 30% Byzantine fraction are obtained from experiments on the Edge-IIoTset and CICIoT2023 Datasets, which significantly exceed four federated baselines ($P < 0.01$).

Keywords: Blockchain security, Internet of Things, edge computing, local differential privacy, intrusion detection, privacy- preserving machine learning, and Byzantine resilience.

1. Introduction

Consider the global cost of cybercrime as a percentage of GDP. That amount is predicted to surpass \$10.5 trillion per year by 2025, ranking third globally behind only the US and China [1]. That statistic isn't taken from a dramatic report to make headlines. It symbolizes without any significant security built in, ransomware shutting down hospital operating rooms, and state-backed criminals stealing industrial control data.

The fundamental issue with IoT Security is an architectural one. The danger perimeter expands over heterogeneous, resources-constrained gear that was never intended for deep packet analysis as 5G develops and 6G approaches [2]. Conventional IDS makes

the assumptions is broken: centralizing telemetry produces a honeypot that smart attackers find tempting, and backhauling adds delay that turns “real-time” isolation into a marketing promise rather than a guarantee [3], [4] The stakes are real. An alarm and a fatality in a smart hospital are separated by a detection latency of 500 milliseconds. A poisoned detection model ensures that lights in an autonomous power grid remain on until they catastrophically go out. Engineers are currently forced to pick between data privacy and detection accuracy. The purpose of B-PriM is to eliminate such trade-off. The naïve trust assumption ingrained in the majority of collaborative IDS designs is what makes the current quo so vulnerable. The network as whole is oblivious to certain attack patterns when a single hacked device injects well-constructed gradient updates that gradually deteriorate a shared detection model over several rounds. This situation has been computational penalty that consensus places on systems with limited resources.

Four particular contributions have been made by this paper:

1. A lightweight edge-GRU that eliminates the GPU reliance that renders the majority of published real-time IDS designs unfeasible on actual IoT gateways by running inference in less than 10 ms on a commodity CPU.
2. An on-chain smart contract – enforced decentralized reputation score (DRS) system that finds and isolates Byzantine nodes without the need for a central arbiter.
3. local differential privacy layer that officially closes the membership-inference attack surface by providing verifiable (epsilon, delta)-DP guarantees on each gradient communication.
4. A thorough empirical analysis using eight experiments on two large-scale real IoT security datasets, supporting each claim with paired t-tests, one-way ANOVA, and Cohen’s d.




Fig. 1. Three- tier B-PriM architecture. EdgeGRU inference and LDP noise injection are locally executed by Edge Layer devices. Reputation – weighted federated aggregation is carried out by the orchestration Layer. On a PBFT blockchain , the integrity layer logs isolation instructions and DRS updates. The Edge- to – orchestration border is never crossed by raw traffic.

2. Related Work

The field has evolved from static signature libraries to adaptive machine learning models, and from isolated sensors to federated networks. Each step has created new capabilities and opened new attack surfaces that the literature has addressed inconsistently.

2.1 IoT IDS Federated Learning

Federated learning can train competitive detection models without sharing raw traffic, as Fed-IIoT [5] demonstrated persuasively. That was real development. Protection against an honest – but – curious aggregator was not addressed; in the event that the central server is compromised, individual traffic patterns can be reconstructed by inverting the gradient updates it has collected [8]. FedAvg was given to us as a workhorse protocol by McMahan et al. [10], and Yin et al. [18] thoroughly demonstrated how badly it breaks under Byzantine participants. FedAvg accuracy may approach random when 10% of nodes provide label – flipped gradients. The majority of IoT-IDS articles steer clear of that experiment.

2.2 IDS Integrity Via Blockchain

There is a lot of appeal to using a blockchain as an IDS audit trail. The tamper – evidence and immutability are real operational advantages. An example of this strategy is Block-IDS [6]. Latency is the enduring challenge. While proof – of – stake average about 12 seconds each finalized block, proof – of – work runs on timescales measured in minutes [13]. The sub-50 ms isolation requirement for real – time IoT traffic is two orders of magnitude away from these figures. By redefining the blockchain as an asynchronous reputation ledger and separating detection delay from consensus latency, B-PriM completely avoid this.

2.3 ML’S Differential Privacy

The strong mathematical underpinning of differential privacy was provided by Dwork and Roth [15]. These foundations were expanded to distributed machine learning by Bassily et al. [7]. The utility gap is a significant privacy budgets frequently reduces detection sensitivity gradient space instead of feature space. This offers the same formal assurance while maintaining the training signal’s statistical structure.

2.4 B-PriM ‘s Fitted

Table I compares five exemplary previous systems to seven practically essential capacity dimensions. There is currently no work that uses CPU- only hardware to deliver all seven at once on actual IoT traffic. B-PriM fills that exact gap.

Table 1.B-PriM Vs. Existing Literature – Capability Comparison

Features	Fed-IIoT [5]	Block-IDS [6]	Diff-priv [7]	B-PriM
Real-time Detection	Yes	No	Yes	Yes
Privacy (LDP)	No	No	Yes	Yes
Decentralised Agg.	Yes	Yes	No	Yes
Node Isolation	No	No	No	Yes
Byzantine Resilience	No	No	No	Yes
CPU-Only Deploy	Partial	No	No	Yes
Blockchain Audit	No	Yes	No	Yes

3. System Model and Problem Formulation

3.1 Architecture with Three Tiers

Three layers are used by B-PriM. IoT devices (sensors, actuators, embedded gateways) that each run a local Edge-GRU and an LDP module make up the Edge Layer, Fig. 2 Cloud Gateways that gather privacy – noised gradient updates and uphold the global model via reputation – weighted aggregation is housed in the orchestration Layer. A private PBFT-based blockchain that records isolation instructions and DRS updates is called the Integrity Layer. The crucial boundary conditions are that raw traffic never moves from Edge to orchestration. That interface is only used for (epsilon, delta)-private gradient modification.




Fig. 2. B-PriM hierarchical network architecture. Blockchain ledger nodes (top) are anchored by cloud gateways (middle), which are connected to IoT Edge nodes (bottom). Byzantine traffic , DRS transactions, and private gradient flows are all distinguished by color- coded arrows.

3.2 Adversary Capabilities of the

In our threat model, a Byzantine Edge Node is the enemy A. Fig. 3 links every kind of attack to its matching defense. Specifically, A can:

1. Traffic Injection: create malicious packets that suppress genuine anomaly signals or cause false positives.
2. Model Poisoning: deliver modified gradient updates that taint the global model and lead it to incorrectly identify targeted attack traffic as benign.
3. Collusion: To exceed per- node reputation threshold, coordinate poisoning across a fraction f of Byzantine nodes.
4. Membership Inference: try to reconstruct other nodes' private traffic patterns by inverting observed gradient updates.

The cloud Gateway is sincere but inquisitive; while it strictly follows the aggregation protocol, it might try to deduce personal information from received updates. It is unable to change, stop, or postpone those updates. In federated learning, such is the typical threat model [8],[19].




Fig 3. Attack-defense mapping and a threat model. The B-PriM framework is represented by the central box; four attack classes (red, left) and four corresponding B-Prim defenses (green, right). Each threat is mapped by arrows to its neutralizing mechanism.

3.3 Formal Optimization Goal

The local dataset at edge node is represented by D_i . The worldwide issues are:

$$\text{Min}_{\{\theta\}} \sum_{i=1}^N L(\theta; D_i) + \lambda R(\theta)$$

Subject to: (i) inference latency of no more than 50 ms per packet; (ii) privacy [identity revelation] of no more than $\exp(\epsilon)$; and (iii) reputation R_i of each active node i at or above the Reputation Floor, or isolate. R_i is the dynamic reputation score kept on the blockchain, while $R(\theta)$ is L2 regularization with weight λ .

4. Proposed Framework

B-PriM operates continuously, detecting locally, reporting privately, updating globally, and isolating reactively. The entire 12 step protocol, from raw packet to isolation decision, is shown in Fig. 4. The structure breaks down into four closely related parts.




Fig. 4. 12-step B-PriM protocol sequence. The cloud gateway is involved in stages 4-6, the blockchain and the isolation feedback path are involved in phases 7-12, and steps 1-3 are local to the edge node. One federation training round is represented by bracketed area.

4.1 The Edge-GRU Detector

Every device operates a lightweight GRU that outputs anomaly score s in the $[0,1]$ after processing a sliding window of $T=16$ consecutive packet feature vectors. Because combining the forget and input gates into a similar sequential modeling capacity, the GRU is favored over an LSTM. After INT8 quantization, Edge-GRU fits in 133.3 kb with 34,127 compact enough to operate on a Raspberry Pi with extra RAM.

4.2 Differential Local Privacy

Two operations are performed before any gradient leaves the device: (1) add Laplacian noise from $\text{Lap}(0, C/\epsilon)$; and (2) clip the gradient to ϵ — DP guarantee. The three processes are depicted in Fig.5: the sufficiently widened by calibrated Laplacian noise, to conceal the contribution of any one traffic record.




Fig. 5. Conduit for LDP gradient privatization. The initial sensitive gradient distribution is on the left. Centre : following L2 clipping at $C=1.0$. Right: demonstrating how the privacy – utility balance changes with the privacy budget following Laplacian at three epsilon levels.

4.3 Score for Decentralized Reputation

B-PriM substitutes a continuous, dynamic reputation per node for regular FL’s binary trust model. A node transmits a privacy – noised notice to the blockchain smart contract when its anomaly score surpasses the threshold τ . The score is then updated by the gateway using $\alpha=0.05$ and $R_{\text{new}}=R_{\text{old}}-(\alpha*s_{\text{reported}})$. Clean nodes recover at 0.02 per window, Reputation Floor (0.30) are instantly from false positives; node below the shown in Fig.6. Byzantine nodes degrade to isolation in about 18 rounds, while honest node hold above 0.085.




Fig. 6. Development of the DRS trust score over 60 federation rounds. Above 0.85, honest nodes (green bands) don’t change. At about round 18, byzantine nodes (red dashed lines) cross the isolation floor, causing automated traffic blockage.




Fig. 7. Blockchain DRS hash – link chain and block structure. Every block contains its predecessor's hash, timestamp, node ID, and privacy – noised score. When this score threshold is exceeded, the smart contract initiates the isolation directive.

4.4 Algorithm 1: B-PriM Isolation Logic

Algorithm 1: B-PriM Decentralised Detection and Isolation

Input: Packet stream P, model theta, privacy budget epsilon

Output: Node status S in {Active, Quarantined}

Initialise: $R_i = 1.0$ for all nodes i

For each incoming packet p in P:

$S = \text{GRU_theta}(p)$ // local anomaly inference

If $s > \text{tau_threshold}$:

$A = \{\text{ID}_i, s, \text{Timestamp}\}$

$g' = A + \text{Laplace}(0, c/\epsilon)$ //LDP privatization Broadcast g' to Blockchain

On Blockchain event received by Gateway:

$R_{\text{new}} = R_{\text{old}} - (\alpha * s_{\text{reported}})$

If $R_{\text{new}} < \text{Reputation_Floor}(0.30)$:

ISOLATE: drop all traffic from ID_i

Else: log for async global aggregation

Return S in {Active, Quarantined}

5. Security And Privacy Analysis

5.1 Guaranteed Differential privacy

Theorem 1 (privacy): For every communication round, the B-PriM gradient reporting method fulfills (epsilon, delta)- Differential privacy.

Proof sketch: Let k be the method for randomized reporting. If $\text{pr}[k(D) \text{ in } S] \leq \exp(\epsilon) * \text{Pr}[k(D') \text{ in } S]$ for every measurable s , then k fulfills epsilon -DP for any delta f is constrained by the clipping constant C , and the injected noise is $\text{Lap}(0, \text{delta } f/\epsilon)$. The ratio requirement is immediately satisfied by the output of $T * \epsilon$ (basic) or $0(\epsilon * \sqrt{T * \log(1/\delta)})$ under advanced composition [19].

Even a fully compromised gateway only learns what is revealed by the noisy output because the LDP process operates on- device- prior to any transmission. [QED].

The unanimous decision is contextualized in Fig.8 without the costly latency of proof-of-work or the security compromises of lighter alternatives, PBFT provides around 15 ms per block, which is quick enough to record isolated events in almost-real time.

All blockchain interactions are completely asynchronous, which is a crucial design feature. The block latency does not appear on the packet – processing critical path because the PBFT consensus path(15ms per block) and the Edge GRU inference path(3.2ms per packet) run on completely different threads and do not share any locks. The gateway implements a temporary soft -block rule in the worst situation. This easily meets the real- time limit of 50ms.

Using randomization {42,123,256,789,999} five independent runs were used to calculate the mean standard deviation of each result. Before testing , the following statistical assumptions were confirmed : Levene’s test verified variance homogeneity prior to ANOVA ($p=0.41$) ; paired t-tests were used when seeds were matched ; Shapiro-wilk normality was verified for all five-sample distributions($p>0.05$).




Fig. 8. Comparison of consensus mechanisms. PBFT(B-PriM) achieves about 15 ms per block. PoW averages about 600,000ms, while Pos averages about 12,000ms. An equivalent four -to – six order-of magnitude energy advantage for PBFT is displayed in the right panel.

5.2 Resilience of Byzantines

B-PriM aggregates as follow instead of using equal weighting: $\theta_{\text{global}} = \sum_i (R_i / \sum_j R_j) * \theta_i$. less than 2.1% of the total aggregation weight is made up of nodes with R_i values less than 0.5. In reality, the global choice boundary cannot be significantly shifted by their poisoned gradients. The approach allows for up to $f < N/2$ Byzantine participants nodes were isolated by the DRS in our tests in 3,2 consensus round (mean time: 8.7 s).

5.3 Complexity of computation:

On a 2.4GHz CPU with $T=16$ and $H=64$, edge-GRU inference takes $O(T*H^2)$ per packet, or about 3.2 ms. blockchain transactions add zero blocking time to the packet flow because they are asynchronous and $O(1)$ per event. For lengthy deployments, storage is limited by configurable pruning; the DRS ledger grows $O(N*E)$, where E is the number of flagged events per node.

6. Experimental Evaluation

6.1 Baselines, Dataset, and setup

Every experiment was conducted on an Intel core i7 with 16 GB RAM and no GPU, accurately simulating resources – constrained IoT gateway scenarios. The whole setup is listed in Table II. Five separate seed were used in each experiment; the mean and standard deviation are presented throughout.

Table II. Implementation Configuration

Parameter	Value/ Configuration
Framework	PyTorch2.1 + Ethereum Sidechain (PBFT consensus)
Detection model	Lightweight GRU (H=64, 1 layer, seq_len=16)
Datasets	Edge-IIoTset(2M+,61 feat,15types)+CICIoT2023(4.5M+,46 feat,33types)
Train/Test Split	70% /20% stratified
Optimiser	Adam (lr=0.001, weight decay=1e-5
Privacy budget	Epsilon in {0.1,0.5,1.0,2.0,5.0}; deita1e-5
Gradient Clip (C)	1.0
Hardware	Intel Core i7, 16GB RAM
Consensus	PBFT sidechain (~ 15 ms / block)
Rep.decay	Alpha=0.05 per event; floor=0.3 recovery=0.02
Repetitions/exp.	5 independent seeds (mean +/- std reported)

Two completing real -world benchmarks were employed. with more than two million labelled samples, Edges-IIoTset [5] captures 61 features across 15 attack types (DDoS, Ransomware, Backdoor, MITM, SQL Injection, and ten more). CICIoT2023 expends examination to 33 attack types and 46 attributes in seven categories. The experimental pipeline is depicted in fig.9 and simulation of genuine IoT data heterogeneity using Dirichlet partitioning is shown in Fig.10.




Fig. 9. Conduit for experimentation. Prior to federated training, raw data is cleaned, normalized, and Dirichlet- partitioned. For statistical analysis and figure creation, all outputs go into a structured result directory.




Fig. 10. Pipeline for experiments. Before federated training , raw data is cleaned, normalized, and Dirichlet-Partitioned. For statistical analysis and figure creation , all outputs are directed into an organized results directory.

The four baseline are summarized in Table III. In order to attribute performance variation to particular architectural decisions, each isolates a single design dimension.

Table III: Baseline Methods – Capability Summary

Method	Description	Privacy	Isolation
Centralized-DNN	MLP trained on pooled data	None	None
FedAvg [10]	Vanilla federated averaging	None	None
De-FedAvg [10]	FL + gradient differential privacy	LDP	None
BlockOnly [6]	FL + blockchain audit log only	None	Partial
B-PriM (Ours)	FL + LDP + DRS reputation	LDP	Full

6.2 E1: F1-Score and Detection Accuracy

B-PriM achieved mean accuracy of 96.2% $\pm$ 0.4 and F1 of 95.8% $\pm$ 0.5% at the operating privacy configuration $\epsilon=1.0$, $\delta=1e-5$ (gradient clip C=1.0) .

While keeping from LDP guarantees, B-PriM outperformed FedAvg by 8.2 percentage points on Edge-IIoTset, achieving mean accuracy of 96.2% $\pm$ 0.4% and F1 of 95.8% $\pm$ 0.5%. the benefit is statistically significant, according to paired t-test ($t(4) = 6.31$, $p=0.0032$, **; Cohen’s $d=2.14$, strong effect). The direction of the difference between B-PriM and each baseline is constant. In contrast to centralized-DNN, B-PriM forfeits 1.3 percentage points of F1 for complete privacy protection and Byzantine resistance. Two factors contribute to the 8.2pp difference in vanilla FedAvg: reputation- weighted aggregation filters an additional convergence benefit is obtained by DRS down-weighting of chronically high loss nodes($\sim 3.1pp$) and byzantine updates($\sim 5.1pp$).The 4.5pp advantage over DP-FedAvg at the same privacy budget results from gradient- space vs feature-space noise injection. The 8.7pp advantage over BlockOnly demonstrates that the accuracy benefit is driven by the DRS process itself rather than the existence of blockchain record keeping. As a result, Table III serves as a component ablation.

Figure12 shows the complete five -metric cross – method comparison, while Figure 11 track convergence.




Fig. 11. Convergence Curves : mean and ± 1 standard deviation across five seeds for each of the five approaches’ validation accuracy versus communication round. By round 10, B-PriM reaches its maximum accuracy which it sustains until round 25




Fig. 12. Accuracy, precision , recall F1, and AUC-ROC are the primary performance metrics compared. In every metric, B-PriM is ahead. One standard deviation over five seeds is displayed by error bars.

6.3 E2: Latency of Inference

Edge-GRU performed well under the 50 ms limit, averaging 3.2 ms $\pm$ 0.3 ms with a P95 of 3.8 ms across 1,000 consecutive single- packet queries on the CPU platform. On the same platform, LSTM alternatives needed 5.1 ms, which is 59% slower because of extra cell multiplications. The blockchain event handler contributed no blocking latency and operated asynchronously.

6.4 E3: Trade-Off Between privacy and Accuracy

B-PriM maintained 95.8% F1 at epsilon=1.0 compared to 91.3% for DP-FedAvg, a4.9% relative gain from gradient -space over feature noise injection, when sweeping epsilon across {0.1,0.5,1.0,2.0,5.0} with delta=1e-5. B-PriM remained over 89.2% F1 even at epsilon=0.1. As seen in Fig.13, the suggested operating point is epsilon=1.0.




Fig. 13. Trade-off between accuracy and privacy. At all epsilon values, B-PriM consistently performs better than DP-FedAvg, and the difference grows as privacy becomes more stringent. The suggested epsilon=1.0 operating point is indicated by the dotted line.

6.5 E4: Resilience to Byzantine Attack

FedAvg accuracy dropped below 76.3% at $f=20\%$ when the Byzantine fraction was increased from 0% to 30% utilizing label -flip poisoning. At $f=30\%$, B-PriM

maintained 91.0%, a 19.3% retention advantage. 94.3% of Byzantine nodes were isolated by the DRS in 3.2 consensus round (mean: 8.7 s). All approaches are overlaid in Fig.14, where the B-PriM -versus-FedAvg gap increases almost linearly with Byzantine fraction, in accordance with theory.




Fig. 14. Byzantine fortitude. Beyond 10%-20% adversarial nodes, FedAvg collapses. B-PriM functions and degrades gently. The crucial collapse zone for unprotected methods is indicated by the darkened area.

6.6 E5: Expandability

Accuracy remained above 94.9% throughout $N = \{10, 25, 50, 100, 200\}$ nodes. At $N=200$, the communication overhead reaches 4.2 MB every round, growing linearly with N . As seen on the dual axis in Fig.15, round time increase sub linearly as a result of concurrent client processing.




Fig. 15. Analysis of scalability. Across all node counts, the detection accuracy (blue circles, left axis) remains constant. At $N=200$ nodes, round time (red squares, right axis) increases essentially linearly to 210 s.

6.7 E6: Analysis of per- class confusion

The normalized confusion matrix for Edge -IIoTset is displayed in Fig.16. At $FPR=0.8\%$, DDoS achieves $TPR=97.1\%$. At $FPR=1.4\%$ MITM attained $TPR=94.6\%$. Backdoor was the hardest category, with a TPR of 81.3%. The purpose of low-rate, stealthy exfiltration is to imitate benign traffic.




Fig. 16. Normalized confusion matrix (5-run average, Edge-IIoTset). True Positive rates for each class are represented by diagonal entries. The most powerful are DDoS and DoS, Whereas the most difficult category to accurately identify is backdoors.

6.8 E7: IID and Non-IID Information or Data

B-PriM F1 decreased by 7.1% from 95.8% to 88.7% with high heterogeneity ($\alpha=0.1$), under the same circumstances, FedAvg drooped by 9.4%, which caused B-PriM to decline 2.3% less. All four setting are compared in Fig.17, where annotations highlight the F1 advantage’s consistency.




Fig. 17. Performance of IID against Non-IID (E7). In all four heterogeneity situations, B-PriM (dark blue) performs better than FedAvg (red). Annotations is percentage points demonstrate the steady benefit.

6.9 E8: Overhead in computation

Mean inference latency was 4.8 ms +/- 0.4 ms on an ARM Cortex -A53 mimicking an industrial gateway; peak RAM was 0.13 MB; model size after INT8 quantization was 33.3KB; B-PriM operates on actual edge hardware without a GPU or a sizable memory buffer.

7. Discussion

The findings clearly show that, at the privacy budgets in reality, privacy and detection quality are not inherently at odds. B-PriM runs within 1.2% of its non – private accuracy ceiling at $\epsilon=1.0$, a margin that most security teams would be noise overpowers the gradient signal, does the trade-off become uncomfortable. That area of the privacy dial is not required for the great majority of actual deployment..The

outcome of Byzantine resilience was more remarkable than anticipated. Up to $f=20\%$, the accuracy collapse is almost completely eliminated by the DRS process. The intuition is that a Byzantine node must either blend in with low scores (making its poisoned updates too weak to matter) or submit large anomaly scores (accelerating its own reputation depreciation). In any case, B-PriM prevails.

It is important to acknowledge five limitations:

- **Storage Growth:** The growth of the DRS ledger is $O(N \times E)$. 50 flagged per day at $N=200$ nodes equal about 40 MB per day, or 1.2GB over 30 days. There are three feasible pruning strategies: (i) sliding – window pruning (72-hour window); (ii) reputation – threshold pruning (archive nodes stable over 0.85 for 14+ days) ; (iii) off-chain archival (anchor only Merkle roots on -chain). In the current prototype , none were used.
- **Cold-Start Exposure:** New nodes that have no reputation history are introduced. Before the DRS has enough information to react, an adversary who is aware of the warm-up window-roughly three consensus rounds-can take advantage of this vulnerability.
- **LPWAN Bandwidth:** The per gradient payload is inflated by approximately 12% due to LDP noise. Insignificant over wi-fi or LTE, but maybe expensive on NB-IoT or LoRWAN networks.
- **Hardware Coverage:** only x86 and ARM Cortex-A53 were tested. The behavior of RISC -V or MIPS- based controllers, which are frequently used in industrial ICS, is still unknown.
- **Smart Contract Auditing:** The DRS contract is vulnerable to attack. Prior to any production deployment, formal verification with slither or Mythril is required.

8. Conclusion And Future Work

On commodity edge hardware, real-time IoT intrusion detection with autonomous isolation and explicit privacy assurances is possible. B-PriM enables provably private, Byzantine- resilient detection on a low – cost CPU by pushing detection to the device, limiting leakage reputation ledger. With inference well under 50ms and (epsilon, delta) -DP maintained throughout, eight runs on two real- world dataset verify statistically substantial improvements over all four baselines.

There are five worthwhile path to take next:

1. The DRS ledger's post -quantum signatures, which protect tamper- evidence from potential quantum adversaries.
2. The honest- but -curious aggregator assumption is eliminated with secure Multi-party Computation for gradient aggregation.
3. Cross-silo federated learning with cross-domain privacy budgets and synchronized DRS states across several cloud provider.

4. A methodical ablation study that eliminates B-priM components one at a time in order to measure the independent contribution of each component.
5. To bridge the gap between simulation and production, a live physical testbed in an actual industrial IoT environment with hardware heterogeneity and hostile tampering.

References

1. Cybersecurity Ventures:2025 official cybercrime Magazine (2025).
2. Kato,N, Mao, B.,Tang,F.,Kawamoto,Y.,Liu,: Ten challenges in advancing machine learning technologies toward 6G.IEEE Wireless Communication 27(3),96-103(2020).
3. Park,J.: Privacy -preserving edge computing: A survey. ACM Computing Surveys57(1)(2025).
4. European Union Agency for cybersecurity (ENISA): Data protection and privacy in emerging technologies. ENISA Technical Report, Publication Office of the EU (2023).
5. Ferrag, M.A.,Friha, O., Hamouda,D., Maglaras,L., Janicke,H.: Edge comprehensive realistic cyber security cyber security dataset of IoT and IIoT applications for centralized and federated learning. IEEE Access 10, 40281-40306 (2022).
6. Tanaka, S.: Blockchain-enable trust management for distributed IDS.IEEE Transactions on Information Forensics and Security 19(2025).
7. Bassily,R.: Practical differential Privacy for distributed ML. In: Proceedings of the IEEE Symposium on Security and Privacy,pp.88-105(2024).
8. Bnawitz,K., et al.: Practice secure aggregation for privacy-preserving machine learning.In: Proceedings of the ACM SIGSAC Conference on Computer and Communication Security (CCS),pp. 1175-1191 (2017).
9. Xu,X.: Real-time anomaly detection in cloud-edge environments. IEEE Transactions on Parallel and Distributed system 36 (2025).
10. McMahan, B., Moore, E., Ramage, D., Hampson, S., Aracas, B.A.y.: Communication- efficient learning of deep networks from decentralized data. In: Proceeding of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), vol.54,pp. 1273-1282(2017).
11. Papernot, N., Abadi, M., Erlingsson, U., Goodfellow, I., Talwar, K.: Semi-supervised knowledge transfer for deep learning from private training data. In: International Conference on Learning Representation(ICLR) (2017).
12. Zhang, H., et al.: Adversarial attack on IoT IDS: A Taxonomy. IEEE Internet of Things Journal 11(2024).
13. Liu, Y., et al.: Blockchain for 5G/6G networks (Springer) 30(2025).
14. Li, T., Sahu, A.K., Talwalkar, A., Smith, V.: Federated optimization in heterogeneous networks. In: Proceedings of Machine Learning Systems (MLSys) (2020).
15. Dwork, C., Roth, A.: The algorithmic foundation of differential privacy. Foundations and Trends in Theoretical Computer Science 9(3-4), 211-407 (2014).

16. Beltran, F., et al.: Resources- constrained ML on Edge Devices. IEEE Embedded Systems Letters 16 (2025).
17. Sun, G, et al.: Data poisoning attack on federated ML. IEEE internet of Things Journal 9 (2024).
18. Yin, D., Chen, y., Ramchandran, k., Bartlett, Byzantine -robust distributed learning : Towards optimal statistical rates. In : Proceedings of the 35th International Conference on Machine Learning (ICML),vol. 80,pp. 5650-5659 (2018).
19. Kairouz, P., McMahan, H.B, et al.: Advances and open problems in federated learning. Foundations and Trends in Machine Learning 14(1-2), 1-210 (2021).
20. Caldas, S., et al.: LEAF: A benchmark for federated settings. arXiv preprint arXiv:1812.01097 (2024).

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

