Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X Volume 4, Issue 3, 2026

PRIVACY-UTILITY TRADE-OFFS IN FEDERATED THREAT DETECTION FOR DISTRIBUTED IIOT AND SCADA SECURITY OPERATIONS

[1]Talha Bin Nusrat
[1]*M.Sc. Information Technology, Clark University, USA. Concentration in Cloud Computing Expertise: Cybersecurity, Security Operations, Network Security, Threat Intelligence, and Cloud Security*
[1]talhakhanghori1609@gmail.com

DOI: https://doi.org/10.5281/zenodo.21915463

Keywords:
Federated learning; industrial Internet of Things; SCADA; intrusion detection; differential privacy; secure aggregation; security operations; compliance

Article History
Received on: 15 Feb, 2026
Accepted on: 25 March, 2026
Published on 28 March, 2026

Corresponding Author:

Abstract
Industrial Internet of Things and supervisory control and data acquisition environments generate sensitive telemetry across sites that cannot always be centralised. This study evaluated federated learning as a privacy-preserving basis for collaborative threat detection and AI-assisted security operations. A reproducible Monte Carlo benchmark represented six industrial sites in Germany, the Netherlands, Finland, the United Kingdom and Canada. Six learning strategies were evaluated across 432 matched runs. Additional experiments included 360 differential-privacy runs and 1,440 stress-test runs. Outcomes covered macro-F1, receiver operating characteristic area under the curve, zero-day recall, false-positive rate, membership-inference exposure, latency, communication volume, energy demand and compliance readiness. The centralised model produced the highest macro-F1 of 0.958. FedProx reached 0.944. Secure aggregation reduced membership-inference AUC to 0.566. The differential privacy and secure aggregation configuration achieved macro-F1 of 0.933 and membership-inference AUC of 0.518. Its detection score did not differ from FedAvg after Holm adjustment. It also reached 95% compliance readiness in the study rubric. Repeated-measures analysis showed a strong strategy effect on macro-F1, $F(5, 355) = 7614.26$, $p < .001$, partial eta squared = .991. The results of the privacy sweep show that larger values of epsilon increase utility but also increase privacy risk. In terms of the poisoning attack, even the protected setting maintained a macro-F1 score of more than 0.92. Such results suggest that federated learning could be considered for industrial security in a distributed setting using privacy mechanisms such as secure aggregation, differential privacy and client authentication.

1 Introduction

Industrial connectivity has changed the boundary between information technology and operational technology. Sensors, PLCs, engineering workstations, historians and cloud services are now sharing information over wide-area connections. This helps increase efficiency and implement predictive maintenance. At the same time, it increases the attack surface. An infected endpoint can compromise process visibility or even availability. The problem is not limited to data loss. Physical operations and safety functions can be affected. Recent federated intrusion detection research indicates that machine learning can help detect attacks without transferring traffic to a centralized system [1]. Work on industrial control systems also shows that adversarial behaviour must be considered during model design rather than after deployment [2].

The current threat environment gives this problem practical urgency. The European Union Agency for Cybersecurity analysed 4,875 incidents from July 2024 to June 2025. Distributed denial-of-service attacks represented 77% of reported incidents. Ransomware remained the most impactful threat. Manufacturing represented 14.9% of ransomware claims in the 2025 landscape [3]. The prior report placed manufacturing at 17% of observed ransomware events [4]. Germany faces a similar operational risk. Bitkom reported that 87% of surveyed German companies were affected by theft, industrial espionage or sabotage during the preceding year. The estimated total damage reached EUR 289.2 billion in 2025. Cyberattacks accounted for EUR 202.4 billion of that total [5]. The Federal Office for Information Security also reported a tense situation. Its 2025 summary identified around 280,000 new malware variants per day and 119 newly disclosed vulnerabilities per day [6].

Centralised machine learning can combine data from several sites. Yet this design creates governance and security concerns. Raw packets may reveal production schedules, device identities, plant topology, maintenance patterns or employee activity. Cross-border transfer can create legal review. A central data lake also becomes a valuable target. Secure aggregation can lower exposure by protecting model updates during aggregation [7]. Federated learning is attractive because each site retains its records and shares model parameters. It can support data minimization. It can also reduce the need to transfer sensitive operational telemetry. Federation does not provide complete privacy by itself. Model updates can disclose properties of local training data. Malicious clients can poison the shared model. Non-identically distributed traffic can create unstable convergence. Slow links can delay security updates. Reviews of privacy-preserving federated intrusion detection emphasize the need to evaluate privacy, robustness and communication together [8]. Differential privacy can limit the influence of a single observation. Secure aggregation can hide individual updates. Robust aggregation can reduce poisoning risk. Each control has a cost. Strong privacy noise can reduce recall. Cryptographic protection can add latency. A model with high accuracy can remain unsuitable for industrial use.

The compliance context makes this trade-off important. Eurostat found that 92.76% of EU enterprises used at least one ICT security measure in 2024. Only 35.50% had documented security practices. Log retention was reported by 45.16%. Formal ICT risk assessment was used by 34.10% [9]. Federated threat detection may support data minimization and local control. It does not automatically meet the General Data Protection Regulation, the Network and Information Systems Directive 2 or the Cyber Resilience Act. Organisations still need lawful purpose, access control, retention rules, incident evidence and accountable human oversight. Privacy-preserving detection for cyber-physical systems should be evaluated as a technical and governance system [10]. Recent methods have explored transfer learning for heterogeneous Industry 4.0 data [11]. Blockchain-assisted federation has also been proposed for auditability and trust [12]. Differential privacy methods can assign different privacy levels to heterogeneous clients [13]. Bayesian privacy mechanisms can also reduce leakage risk [14]. Hierarchical designs can adapt privacy settings across tiers [15]. These studies show rapid progress. Yet many evaluations emphasize accuracy and omit security operations measures. Few combine false

alarms, zero-day recall, communication cost, membership-inference exposure and compliance readiness in one balanced benchmark.

This study addresses that gap through a reproducible Monte Carlo benchmark. It compares local-only training, centralised training, FedAvg, FedProx, secure FedProx and differential-private secure FedProx. The primary outcome is macro-F1. Secondary outcomes cover discrimination, zero-day recall, false alarms, privacy exposure, latency, communication, energy and compliance readiness. The study also tests non-IID traffic, client dropout, poisoning and concept drift. Its objective is not to claim field deployment. It is to identify a defensible privacy-utility operating point for future IIoT and SCADA pilots with special relevance to Germany and the wider European security environment.

2 Literature Review

Federated intruding detection is now at the system level of evaluation, not just a privacy argument. That change is all about secure aggregation. The main contributions of BalancedSecAgg are on the cost and scalability of secure update aggregation [16]. An immutable coordination layer is introduced through smart-contract-assisted aggregation, which can enhance update accountability [17]. These are the designs to handle the confidentiality during collaboration. They don't make the necessity to check clients and test tainted upgrades. Data heterogeneity, communication cost, privacy leakage and adversarial clients are all found to be related design constraints that are identified in a recent survey on federated intrusion detection [18].

The data, obtained from industrial network, are very non-IID. Repetitive control messages can be used within a manufacturing cell. Low event traffic at an energy site is possible. Wireless and mobile endpoints can be installed in a port. Local sample size is used to update standard FedAvg weights. Can be skewed towards large sites or often attacked classes. Ali [19] showed how federated learning with fog can be used for industrial threat detection. To address non-IID IoT intrusion data, Peng et al. [20] proposed knowledge distillation. Personalised hypernetworks are another approach, which involve customizing models for local distributions [21]. The scale industrial environments also require aggregation with the ability to handle device and traffic heterogeneity [22].

Protecting privacy isn't just about storing raw data on the local site. Gutti et al. [23] presented an intrusion detection problem that is distributed among the IoT nodes as a privacy preserving collaborative task. The idea of differential privacy is to introduce calibrated noise to gradients or model parameters. It minimizes the effect of each record. It can also cause a drop in minority attack recall when the privacy budget is too tight. Heterogeneous differential privacy has the ability to give different budgets to different clients [13]. Local risk/data sensitivity can be reflected in personalised privacy [15]. It is necessary to control the choice of epsilon. It may not be a hidden model hyperparameter.

A coordinator's update content is protected by secure aggregation. The research review indicates that the number of masking and secret-sharing, as well as homomorphic-encryption techniques is expanding [24]. Park et al. [25] implemented masking and homomorphic encryption. They give a direction that helps with confidentiality, at the cost of computational expense. Secure aggregation also needs poisoning resistance. A protected harmful update may be submitted by malicious client. SHIELD provides aggregation security in a hierarchical federation against poisoning [26]. This is a key aspect of industrial security operations. Integrity checks should not be hindered by confidentiality.

The other operating condition is "explainability. A security analyst should understand the reason for the alert raised by a model. Fatema et al. [27] combined federated learning with SHAP explanations. Kalakoti et. al. [28] investigated explainable federated AI for deep intrusion detection. Harshitha et al. [29] combined Explainable AI with hyperparameter optimisation. Sahoo et al. [30] concentrated on the veracity of explanations in edge-IoT intrusion detection. The findings indicate that the quality of explanations should be evaluated, not taken for granted, according to these studies. Explanation may be firm but incorrect. It may also show sensitive features if not shared with controls.

It is still challenging to achieve near real-time response. A Hybrid Federated (HF) design for Near Real-time (NRT) Intrusion Detection (ID) in IoT was studied by Rampone et al. [31]. An architecture can lessen the amount of raw data transfer. Still relies on local inference, edge capacity, communication windows. Secure aggregation adds the bytes and cryptographic operations. Gateways that share power or cooling with control equipment are of interest with regard to energy use. The realistic evaluation should include detection gain along with latency, communication and energy demand.

New priorities in the evolution of privacy and robustness are becoming increasingly apparent. SecuFL-IoT is a solution that combines privacy protection and adaptive anomaly detection in smart industrial networks [32]. The application of deep learning in the industrial control system also demonstrates the benefits of processing features that are familiar from that specific domain and multiple modalities [33]. The high benchmark accuracy may not be translated to a live plant. Scores can be distorted by dataset leakage, random record splitting and repeated attack patterns. Site-level partitioning is a more demanding test as it maintains local distributions separate.

There are four conclusions supported by the literature. First is the fact that when there are differences in sites, FedProx is better suited than FedAvg. First is that FedProx is more appropriate than FedAvg when sites are different. Second, there is a need for a utility analysis of differential privacy. Third, aggregation needs to be integrated with authentication and integrity checks. Fourth, Adoption of a SOC demands explanation of alerts and change records on which it can audit. The present benchmark combines these issues. It provides a detection quality, privacy leakage, robustness and operational burden as well as a compliance readiness mapping in the same study.

3 Methodology

This research used a reproducible Monte Carlo benchmark. It did not use observations collected from operating companies. The design represented six distributed industrial sites. Two sites were located in Germany. The remaining sites represented the Netherlands, Finland, the United Kingdom and Canada. Sectors included automotive manufacturing, energy distribution, water treatment, port logistics, automated warehousing and multi-utility services. Site profiles specified 87 to 148 IIoT or SCADA nodes. Traffic volumes ranged from 249,000 to 420,000 records. Attack prevalence ranged from 20.9% to 31.2%. Dirichlet alpha values of 0.20 to 0.40 generated non-IID client distributions. The design included common industrial protocols such as OPC UA, Modbus/TCP, PROFINET, DNP3, MQTT and IEC 61850.

Six learning strategies were compared. Local-only training represented isolated sites. Centralised training represented full raw-data pooling. FedAvg was the basic federated comparator. FedProx represented heterogeneity-aware optimisation. SecAgg-FedProx added protected aggregation. DP-SecAgg-FedProx added record-level differential privacy with a primary epsilon of 2.0. Each site completed 12 matched seeded runs per strategy. The main dataset contained 432 observations. A second experiment varied epsilon across 0.5, 1, 2, 4 and 8. It contained 360 observations. A third experiment tested baseline conditions, extreme non-IID traffic, 20% client dropout, 10% poisoning clients and concept drift. It contained 1,440 observations. The seed was fixed at 20260803.

The primary outcome was macro-F1 because attack classes were imbalanced. Secondary predictive outcomes were ROC-AUC, zero-day recall and false-positive rate. Privacy risk was measured by membership-inference AUC. A value near 0.50 indicates chance-level inference. Operational measures were local inference latency, model-update communication, energy consumption and rounds to convergence. Compliance readiness used ten controls. These covered raw-data minimization, update confidentiality, participant authentication, poisoning resilience, traceability, explainability, cross-border restraint, incident evidence retention, change control and privacy-budget governance. Each control was coded 0 for not met, 1 for partial and 2 for met. The resulting score is a technical readiness indicator. It is not a legal opinion or certification.

Matched site-seed blocks supported repeated-measures analysis. Means, standard deviations and 95% confidence intervals were calculated for each strategy. A repeated-measures analysis of variance tested the overall strategy effect on macro-F1. Partial eta squared measured effect size. The Friedman test and Kendall's W provided a distribution-free confirmation. The protected strategy was compared with each alternative through paired t tests and Wilcoxon signed-rank tests. Holm adjustment-controlled family-wise error. Cohen's dz quantified paired effect size. Spearman correlation assessed the association between epsilon and utility as well as privacy exposure. Stress-test results were summarized by strategy and scenario.

The analysis was completed with deterministic Python routines using NumPy, pandas and SciPy. A compatible SPSS syntax file was prepared for independent reruns. The spreadsheet keeps raw benchmark runs, stress tests, privacy sweeps and derived tables in separate sheets. This separation protects data lineage. All values in the article were produced from the supplied data files. No field effectiveness or organisational compliance is inferred.

Quality checks were applied before analysis. Metric values were bounded to their valid ranges. Every site-seed block contained all six strategies. Strategy labels and site identifiers were fixed before generation. Primary comparisons were specified before the privacy and stress analyses. The data did not tune values after hypothesis testing. Results were rounded only for presentation. The underlying files retain additional decimals. A two-sided alpha level of .05 was used. Effect sizes and confidence intervals were reported beside p values to limit significance-only interpretation. The proforma records scenario settings, privacy budget, aggregation mode and operational measures for future replication.

4 Results

The benchmarking approach generated 432 benchmarking runs, 360 privacy budget runs and 1,440 stress testing runs. In all these cases, all the six methods were assessed in each of the matched site-seed pairs. It is evident from the results that there was a trade-off between predictive performance and privacy exposure. Centralization led to better predictive performance. The protected federated learning led to lower privacy exposure and high readiness score.

4.1 Site Demographic and Traffic Profile

Table 1 describes the site demographic and traffic profile. The six sites represented 670 nodes and 1.92 million benchmark records. Germany contributed 785,000 records across manufacturing and energy settings. The German manufacturing site had the highest attack share at 31.2%. The United Kingdom logistics site had the lowest share at 20.9%. The Dirichlet alpha values created substantial distribution shift. DE-MFG had the strongest heterogeneity with alpha = 0.20. UK-LOG had the mildest heterogeneity with alpha = 0.40. This range tested whether aggregation remained stable when attack classes were concentrated at different sites.

Table 1: *Site Demographic and Benchmark Traffic Profile*

Site	Country	Sector	Nodes	Records	Attack %	Alpha	Protocols
DE-MFG	Germany	Automotive manufacturing	148	420,000	31.2	0.20	OPC UA; Modbus/TCP; PROFINET
DE-ENE	Germany	Energy distribution	126	365,000	27.8	0.25	IEC 60870-5-104; DNP3; MQTT
NL-WAT	Netherlands	Water treatment	94	292,000	22.6	0.35	Modbus/TCP; EtherNet/IP
FI-PRT	Finland	Port logistics	112	318,000	25.1	0.30	OPC UA; MQTT; CAN
UK-LOG	United Kingdom	Automated warehousing	103	276,000	20.9	0.40	PROFINET; MQTT; EtherCAT
CA-UTL	Canada	Multi-utility services	87	249,000	24.3	0.30	DNP3; Modbus/TCP; IEC 61850

Note. Alpha is the Dirichlet concentration parameter. Lower values indicate stronger non-IID partitioning.

4.2 Predictive Performance and Privacy Exposure

Table 2 compares predictive performance and privacy exposure. Centralised training produced the highest macro-F1 at 0.958 with a 95% CI from 0.956 to 0.959. Its ROC-AUC was 0.989. Zero-day recall reached 0.889. The false-positive rate was 0.024. This utility came with the highest membership-inference AUC at 0.843. The result indicates strong exposure under the benchmark attack.

FedProx was the strongest federated strategy for predictive utility. Its macro-F1 was 0.944 and its ROC-AUC was 0.983. Zero-day recall was 0.862. Secure aggregation caused a small decline. SecAgg-FedProx produced macro-F1 of 0.942 and membership-inference AUC of 0.566. DP-SecAgg-FedProx produced macro-F1 of 0.933 with a 95% CI from 0.931 to 0.934. Its membership-inference AUC was 0.518. This value was close to chance. The protected configuration reduced membership exposure by 0.325 AUC points compared with centralised training.

Local-only training was weakest. Its macro-F1 was 0.851. Its zero-day recall was 0.712 and its false-positive rate was 0.073. Local isolation protected raw traffic movement but did not provide strong privacy against model-level inference. Membership-inference AUC remained 0.708. The finding shows that data locality alone is not a complete privacy control.

The site-level pattern was consistent but not identical. DP-SecAgg-FedProx ranged from 0.925 macro-F1 at DE-MFG to 0.939 at UK-LOG. FedAvg ranged from 0.924 at DE-MFG to 0.940 at UK-LOG. FedProx ranged from 0.937 to 0.950 across the same sites. The lower German manufacturing values corresponded with the smallest alpha and the largest attack share. This profile is the most difficult in the benchmark. The protected strategy was still better than local only training at DE-MFG by 0.081 points. Its site-level membership AUC was between 0.516 and 0.521. The close range indicates steady protection of privacy across industries. The false positive rates of the protected strategy were between 0.034 and 0.038. The maximum rate was obtained at DE-MFG. This has implications on the workload of Security Operation Centre since a slight increase can lead to many false alerts. The protected strategy preserved its privacy benefit at every site. Its detection burden remained most visible in the strongest non-IID profile.

Table 2: *Detection Performance and Membership-Inference Exposure by Strategy*

Strategy	Macro-F1 mean +/- SD	95% CI	ROC-AUC	Zero-day recall	False-positive rate	Membership AUC
Local-only	0.851 +/- 0.007	0.849-0.853	0.924	0.712	0.073	0.708
Centralized	0.958 +/- 0.006	0.956-0.959	0.989	0.889	0.024	0.843
FedAvg	0.933 +/- 0.007	0.931-0.934	0.977	0.841	0.035	0.685
FedProx	0.944 +/- 0.006	0.943-0.946	0.983	0.862	0.030	0.656
SecAgg-FedProx	0.942 +/- 0.007	0.940-0.943	0.982	0.859	0.031	0.566
DP-SecAgg-FedProx	0.933 +/- 0.007	0.931-0.934	0.976	0.847	0.035	0.518

4.3 Omnibus Statistical Tests

Table 3 reports the omnibus tests. The repeated-measures analysis found a significant strategy effect on macro-F1, $F(5, 355) = 7614.26$, $p < .001$. Partial eta squared was .991. The magnitude reflects the large separation between local-only, centralised and federated strategies in a controlled simulation. The Friedman result confirmed the ordering, chi-square (5) = 334.25, $p < .001$. Kendall's W was .928. Strategy rankings were highly consistent across site-seed blocks.

The result should be interpreted with the design. Repeated seeds reduce random uncertainty. They do not represent independent plants. The effect size measures algorithm separation inside the benchmark. It is not a population estimate for all industrial networks.

The parametric and rank-based tests led to the same conclusion. This agreement reduces concern that the omnibus finding was produced only by normality assumptions. The centralised mean exceeded the protected mean by 0.0251. The protected mean exceeded the local-only mean by 0.0817. These gaps explain the very large omnibus effect. The narrow confidence intervals also reflect matched seeded runs. They should not be confused with uncertainty across the full industrial population. The corrected contrasts provide the more useful decision evidence. They show where the protected strategy lost utility and where it did not. The absence of a FedAvg difference was supported by both a Holm-adjusted t test and a Wilcoxon test. The two p values were .952 and .822. This convergence supports practical equivalence within the tested resolution. A formal non-inferiority trial would still need a prespecified margin and field observations.

Table 3: *Repeated-Measures and Distribution-Free Tests of Macro-F1*

Test	Statistic	df1	df2	p	Effect size	Interpretation
Repeated-measures ANOVA	F = 7614.26	5	355	< .001	Partial eta squared = 0.991	Very large benchmark strategy effect
Friedman test	Chi-square = 334.25	5	-	< .001	Kendall's W = 0.928	Highly consistent strategy ranking

4.4Corrected Pairwise Contrasts

Table 4 shows corrected paired contrasts for DP-SecAgg-FedProx. The protected configuration exceeded local-only training by 0.0817 macro-F1 points. The 95% CI ranged from 0.0803 to 0.0830. It was lower than centralised training by 0.0251 points. It was also lower than FedProx by 0.0116 points and SecAgg-FedProx by 0.0088 points. Each difference was significant after Holm adjustment. The comparison with FedAvg was different. The mean difference was 0.00004. The Holm-adjusted p value was .952. The protected model matched FedAvg utility within the resolution of this benchmark. It added strong privacy controls without a detectable macro-F1 penalty relative to basic federation.

Table 4: *Paired Contrasts for DP-SecAgg-FedProx*

Contrast	Mean difference	95% CI	t	Holm p	Wilcoxon p	Cohen's dz
DP-SecAgg-FedProx - Local-only	0.0817	0.0803 to 0.0830	119.37	< .001	< .001	14.07
DP-SecAgg-FedProx - Centralized	-0.0251	-0.0264 to -0.0237	-37.24	< .001	< .001	-4.39
DP-SecAgg-FedProx - FedAvg	0.0000	-0.0012 to 0.0013	0.06	.952	.822	0.01
DP-SecAgg-FedProx - FedProx	-0.0116	-0.0129 to -0.0102	-17.21	< .001	< .001	-2.03
DP-SecAgg-FedProx - SecAgg-FedProx	-0.0088	-0.0101 to -0.0076	-13.86	< .001	< .001	-1.63

4.4 Differential-Privacy Budget

Table 5 presents the privacy-budget sweep. An epsilon of 0.5 produced macro-F1 of 0.925 and membership-inference AUC of 0.514. Epsilon 1 improved macro-F1 to 0.929 and raised membership AUC to 0.528. The primary epsilon of 2 produced macro-F1 of 0.935 in the sweep and membership AUC of 0.548. Epsilon 8 achieved macro-F1 of 0.943 but membership AUC rose to 0.597. Epsilon had a positive association with macro-F1, Spearman rho = .739, $p < .001$. It had a stronger positive association with membership exposure, rho = .964, $p < .001$. The privacy gain was largest at the strict end of the budget. Utility gains became smaller at larger epsilon values.

The sweep contained 72 runs at each privacy budget. Standard deviations for macro-F1 remained between 0.006 and 0.007. Membership-AUC standard deviations remained between 0.006 and 0.009. The trade-off was not driven by one unstable budget. Moving from epsilon 0.5 to 2 added about 0.010 macro-F1 points. It also added about 0.034 membership-AUC points. Moving from epsilon 2 to 8 added only about 0.008 macro-F1 points while membership AUC increased by about 0.049. Latency declined from 15.72 ms to 14.79 ms across the sweep. The decrease reflected lower privacy-processing burden at larger budgets. Epsilon 2 was retained as the balanced setting because it remained below 0.55 membership AUC while recovering more utility than the two strictest budgets. This choice is a benchmark operating point. It is not a universal privacy threshold.

Table 5: *Privacy-Utility Sweep for DP-SecAgg-FedProx*

Epsilon	n	Macro-F1 mean +/- SD	Membership AUC mean +/- SD	Latency ms mean +/- SD	Interpretation
0.5	72	0.925 +/- 0.006	0.514 +/- 0.006	15.72 +/- 0.31	Strict privacy
1	72	0.929 +/- 0.007	0.528 +/- 0.008	15.54 +/- 0.28	Strict privacy
2	72	0.935 +/- 0.006	0.548 +/- 0.007	15.32 +/- 0.27	Balanced primary setting
4	72	0.940 +/- 0.006	0.575 +/- 0.009	15.11 +/- 0.25	Higher utility with higher exposure
8	72	0.943 +/- 0.007	0.597 +/- 0.008	14.79 +/- 0.28	Higher utility with higher exposure

4.5Robustness Stress Tests

Table 6 reports stress-test performance. Under extreme non-IID traffic, FedAvg fell to macro-F1 of 0.916. FedProx achieved 0.935. SecAgg-FedProx achieved 0.934. DP-SecAgg-FedProx retained 0.925. Under 20% client dropout, the protected configuration reached 0.927. Its membership-inference AUC remained 0.518.

Poisoning produced the most important security contrast. FedAvg fell to macro-F1 of 0.907 and membership-inference AUC rose to 0.714. DP-SecAgg-FedProx retained macro-F1 of 0.923 and membership-inference AUC of 0.535. SecAgg-FedProx achieved higher macro-F1 of 0.925 under the same condition. It had weaker privacy with membership AUC of 0.583. Concept drift lowered the protected model to 0.914. FedProx retained 0.927. These values show that privacy protection improved exposure but did not eliminate the need for drift detection and robust update screening.

Table 6: *Stress-Test Performance for Federated Strategies*

Scenario	Strategy	n	Macro-F1 mean +/- SD	Membership AUC	Communication MB
Baseline	FedAvg	48	0.934 +/- 0.007	0.686	407.3
Baseline	FedProx	48	0.944 +/- 0.008	0.657	434.3
Baseline	SecAgg-FedProx	48	0.942 +/- 0.008	0.566	508.3
Baseline	DP-SecAgg-FedProx	48	0.934 +/- 0.007	0.519	521.8
Extreme non-IID (alpha=0.10)	FedAvg	48	0.916 +/- 0.008	0.692	441.9
Extreme non-IID (alpha=0.10)	FedProx	48	0.935 +/- 0.007	0.663	465.1
Extreme non-IID (alpha=0.10)	SecAgg-FedProx	48	0.934 +/- 0.006	0.574	552.3
Extreme non-IID (alpha=0.10)	DP-SecAgg-FedProx	48	0.925 +/- 0.007	0.523	563.5
20% client dropout	FedAvg	48	0.920 +/- 0.007	0.688	352.3
20% client dropout	FedProx	48	0.937 +/- 0.007	0.661	372.4
20% client dropout	SecAgg-FedProx	48	0.934 +/- 0.008	0.570	434.3
20% client dropout	DP-SecAgg-FedProx	48	0.927 +/- 0.007	0.518	449.7
10% poisoning clients	FedAvg	48	0.907 +/- 0.008	0.714	424.1
10% poisoning clients	FedProx	48	0.917 +/- 0.008	0.685	445.9
10% poisoning clients	SecAgg-FedProx	48	0.925 +/- 0.008	0.583	528.9
10% poisoning clients	DP-SecAgg-FedProx	48	0.923 +/- 0.007	0.535	540.6
Concept drift	FedAvg	48	0.911 +/- 0.007	0.690	418.3
Concept drift	FedProx	48	0.927 +/- 0.007	0.664	438.6
Concept drift	SecAgg-FedProx	48	0.925 +/- 0.007	0.570	517.7

Scenario	Strategy	n	Macro-F1 mean +/- SD	Membership AUC	Communication MB
Concept drift	DP-SecAgg-FedProx	48	0.914 +/- 0.007	0.523	534.2

4.5Operational Efficiency

Table 7 presents operational cost. Local-only inference was fastest at 6.97 ms and required no collaborative communication. FedAvg required 12.42 ms and 429.23 MB of model-update traffic per benchmark cycle. FedProx required 12.97 ms and 452.34 MB. Secure aggregation increased latency to 15.14 ms and communication to 534.86 MB. Differential privacy and secure aggregation increased latency to 16.03 ms and communication to 551.46 MB. Its energy estimate was 11.59 Wh per training cycle. This was 30% above FedProx. It converged in 42.71 rounds compared with 38.86 rounds for FedProx. The privacy layer had a measurable but bounded operational cost.

Centralised training moved the largest volume. Its raw-data transfer estimate was 1,533.45 MB. This was almost three times the update traffic of the protected federated configuration. Centralised inference latency was lower than secure federation. Yet its data movement and membership exposure were much higher. The operational choice depends on whether privacy, network cost or maximum accuracy has priority.

Table 7: *Operational Burden and Convergence by Strategy*

Strategy	Latency ms mean +/- SD	Communication MB mean +/- SD	Energy Wh mean +/- SD	Rounds	Readiness score
Local-only	6.97 +/- 0.41	0.0 +/- 0.0	0.94 +/- 0.23	-	48.1
Centralized	13.68 +/- 0.50	1533.5 +/- 59.4	13.22 +/- 0.24	1.00	56.2
FedAvg	12.42 +/- 0.46	429.2 +/- 16.3	8.39 +/- 0.23	44.68	70.3
FedProx	12.97 +/- 0.49	452.3 +/- 13.1	8.92 +/- 0.23	38.86	75.1
SecAgg-FedProx	15.14 +/- 0.47	534.9 +/- 17.3	10.86 +/- 0.27	39.69	88.4
DP-SecAgg-FedProx	16.03 +/- 0.51	551.5 +/- 18.6	11.59 +/- 0.26	42.71	93.8

4.5 Compliance Readiness

Table 8 maps each strategy to ten technical controls. DP-SecAgg-FedProx fully met nine controls and partly met poisoning resilience. Its readiness score was 95%. SecAgg-FedProx reached 80%. FedProx reached 65% and FedAvg reached 60%. Centralised training reached 50%. Local-only training reached 45% because it lacked collaborative governance, protected updates and privacy-budget management.

The score indicates architecture elements that are present in the benchmark. Legal compliance is not demonstrated by the score. Any deployment would still need a DPIA, purpose documentation, processor responsibilities, incident reporting, assurance of suppliers, model monitoring and human accountability. The table is most suited to use as a pre-deployment gap analysis tool.

Table 8: *Technical Compliance-Readiness Mapping*

Strategy	Met	Partial	Not met	Readiness	Interpretation
Local-only	2	5	3	45%	Material control gaps
Centralized	4	2	4	50%	Material control gaps
FedAvg	5	2	3	60%	Material control gaps
FedProx	5	3	2	65%	Material control gaps
SecAgg-FedProx	7	2	1	80%	Strong technical base
DP-SecAgg-FedProx	9	1	0	95%	Pilot-ready with governance review

Note. Readiness is an internal technical rubric. It is not legal certification.

5 Discussion

The benchmark encourages studying federated learning as a viable intermediate solution between isolated detection and complete centralization. The most successful macro-F1 was Centralised Training. It also caused the maximum data movements and the maximum membership-inference exposure. FedProx was able to keep the most of the utility at heterogeneous sites. Less leakage with secure aggregation. The differential privacy achieves near chance membership-inference AUC. After multiplicity correction, the protected model obtained the same macro-F1 score as the FedAvg model. This combination is useful for security operations as it is a common baseline for FedAvg. The findings indicate that there need not be a great loss of utility for the adoption of greater privacy.

The utility pattern is consistent with recent research on personalised federation and heterogeneity-aware federation. Abhijit et al. [21] personalized hypernetworks to local distributions, intrusion models. Mao et al. [22] studied large-scale heterogeneous industrial environments. In the present study, FedProx improved the FedAvg by approximately 0.012 macro-F1 points. Under extreme non-IID conditions, the gap between the two increased. This aids proximal or personalised optimization for separate plants provided by customers.

The privacy result is also similar to adaptive secure federation. In 2026, Alqazzaz merged Adaptive Anomaly Detection and Privacy Controls in Smart Industrial Networks. The present benchmark is used to distinguish between those controls. Secure aggregation decreased membership AUC from 0.656 (FedProx) to 0.566. With differential privacy, it was further lowered to 0.518. Compared to FedProx, the cost was 0.0116 macro-F1 points. The epsilon analysis explains why one can't use one accuracy score. The macro-F1 score was improved by 0.0177 when epsilon was increased from 0.5 to 8. This also drove up the membership exposure by 0.0831 AUC points., a Security Operations Centre must see epsilon as a risk decision that needs to be approved and monitored.

Good robustness still is a concern. Khang et al. [34] designed an anti-poisoning defence for federated intrusion detection in non-IID data. The current poisoning situation lowered FedAvg to 0.907 macro-F1. The protected configuration resulted in a 923.0. The protected configuration kept 0.923. If the malicious update tries to hide each component, then secure aggregation on its own is not enough to identify a malicious update. Still need to be able to authenticate participators, limit the range of updates, assign some sort of anomaly score and rollback. This is in line with the federated anomaly detection trilemma, as described by Pinto et al. [35]. Privacy, usefulness and strength should be considered in tandem.

It's also crucial that the explanation is clear. Causal reasoning was introduced in the field of explainable federated intrusion detection by Asiri [36]. In an explainable privacy-preserving framework, Jacob and Sultana [37] reported. These methods can assist analysts to identify if process changes are the result of a cyberattack. They also can assist with alert triage. The benchmark readiness rubric was only awarded full credit if an architecture could maintain a human readable rationale. The fidelity of explanations should be tested with engineers in future field work.

The SCADA context requires safety-aware deployment. Etinge et al. [38] examined privacy-preserving federated analytics for smart-grid SCADA under coordinated attacks. Gulzar and Mustafa [33] showed the continuing role of deep anomaly detection in industrial control systems. The present latency results were between 12 and 16 ms for federated inference paths. These values are suitable for monitoring in many settings. They should not be used to justify direct closed-loop control. Detection models should remain outside safety instrumented functions unless separately certified.

The international comparison is important for Germany. ENISA recorded 4,875 incidents in its 2025 landscape and identified DDoS as 77% of reported incidents [3]. The protected configuration maintained macro-F1 above 0.92 under dropout, poisoning and extreme non-IID tests. This is relevant to distributed European infrastructure. It does not mean that 92% of real attacks would be detected. The metrics are results with controlled distributions.

There is also a compelling economic argument for collaboration in privacy-preserving ways in Germany. Bitkom estimated that EUR 289.2 billion is lost in the annual damage caused by theft, espionage and sabotage. It has estimated EUR 202.4 billion as costs due to cyberattacks [5]. According to BSI, they detected approximately 280,000 new malware variants every day and 119 new vulnerabilities every day [6]. These numbers facilitate a learning-by-shared approach between manufacturers, utilities and logistics operators. Raw telemetry pooling may not be preferred due to potential loss of operational and other information. Federation provides a path towards collective defensive knowledge without moving as much data. According to EU enterprise data, there is a governance gap. Almost 93% of enterprises used at least one security measure. The majority of respondents did not have documented practices (35.50%) or conduct an ICT risk assessment (34.10%) [9]. If ownership is not clear, a federated model can be another out-of-control situation. In this study, the 95% readiness score is based on documentation of privacy budgets, versions of models, incident evidence and change control. Scores may be lower in real organisations. The score should start a governance review rather than end zone.

To reduce operation costs on limited gateways, feature reduction can be used. Nguyen et al. [39] studied federated IoT intrusion detection with feature reduction. We used 551 MB of update traffic and 11.59 Wh per training cycle on the present protected model. These values are less than centralised data movement but still material. Tests should be conducted on compression, sparse updates and scheduled aggregation. FedSMOTE-DP also demonstrates the ability of privacy-aware learning to tackle class imbalance [40]. These procedures can be used to increase the recall of rare attacks while maintaining a reasonable value of epsilon.

The integrated decision view is the main contribution. Centralization was the only measure that was more strongly associated with detection accuracy. Strict differential privacy was preferred given privacy. Local inferences were the preferred operations. Protected federation was favored for compliance readiness. Selected operating point is dependent of threat level, safety criticality, available bandwidth and legal role. DP-SecAgg-FedProx is the most balanced starting point for Germany and the EU. Piloting should be done first with BSI aligned logging, NIS2 incident procedures and GDPR governance before wider roll-out.

6 Practical Implications and Policy Recommendations

ISOC should implement federation in stages. The first phase should be read-only monitoring and inference at the local site. Additional features for a second phase: authenticated model exchange, secure aggregation and signed version records. The differential privacy budgets should be set as risk parameters. Local explanation and model version should be kept in alerts. German operators should acknowledge the BSI guidelines and NIS2 procedures for evidence retention, escalation and supplier responsibilities. This protected configuration has been demonstrated to work well on a controlled pilot where it resulted in FedAvg utility and lower membership exposure. It cannot be used as a substitute for segmentation, patching, backups or safety controls. In a case of failure of federation, a rollback path or a local fallback model should be provided.

Procurement should also include model cards that can be reproduced and client software inventories and rotation of cryptographic keys. The false positive rate, drift, privacy spending and privacy participation failures should be displayed on operational dashboards. Security engineers should be involved in governance teams at regularly scheduled intervals to review these indicators.

7 Limitations and Areas for Future Research

Benchmark data was used in the study for IIoT or SCADA traffic. The reported confidence intervals are based on Monte Carlo variation. They do not cover all architectures of plants, attack tools and operator behaviour. Membership-inference AUC is one privacy test. Property inference was not performed and gradient inversion was not performed. The compliance score is a readiness rubric which is structured. It is not a legal article or a certificate. The energy and communication estimates are not platform specific and are only approximate. The benchmark was also based on

the assumption of stable labels. Unlabeled real zero-day events can last for extended periods of time. External validity entails the field test with approved telemetry, device-level measurements and independent red-team exercises.

Not all site profiles or industrial protocols can be captured by the six site profiles. Financial return or safety impact was not estimated. It also did not model dishonest coordinators, colluding clients or cryptographic key compromise. These gaps may change the preferred architecture.

8 Recommendations for Further Studies

Subsequent research needs to confirm this framework on German manufacturing, energy and water testbeds. Comparisons of trustworthy execution environments, homomorphic encryption and post-quantum secure aggregation must be considered. Another research area needs to consider the human element in AI-powered Security Operation Centers. Finally, drift, false alarm fatigue and NIS2 evidence quality need longitudinal research in real incidents.

9 Conclusion

Federated learning may enable threat detection in a privacy-friendly manner in distributed IIoT and SCADA environments. Its efficacy relies on the learning control measures. Simple federated learning leaves the raw data local. It cannot avoid update leakages, poisoning and lack of good governance. Secure aggregation is effective in protecting updates. Differential privacy minimizes record exposure. Proximal optimization provides better results for divergent site distributions.

Centralised learning had the highest macro-F1 of 0.958. FedProx had 0.944. The differential-private secure setup scored 0.933. It was not statistically different from FedAvg after the Holm correction. It had a membership inference AUC score of 0.518, which was near-chance and very low compared to the score for the centralised setup of 0.843. It also managed to maintain macro-F1 above 0.92 under the poisoned scenario. The privacy mechanism made the learning process costly in terms of latency, communication and energy. Such costs were relatively lower than the costs of raw-data transfer in centralised learning.

The compliance mapping found that the technical maturity level of DP-SecAgg-FedProx was 95%. This is due to the presence of privacy, traceability and governance elements. This is not an indication of compliance with the GDPR, NIS2, Cyber Resilience Act or the German regulation. Lawful purposes, supplier confidence, risk assessments, logging, human intervention and reporting incidents are still needed.

For Germany, this architecture presents a promising pilot direction. It allows manufacturers and CII owners to learn from attacks distributed across devices without sharing sensitive information through telemetry. The most robust deployment method would consist of multiple layers. These would include secure clients, secure aggregation, governed differential privacy, proper update verification, local explanation and monitoring of rollbacks. In future field tests, this architecture should be evaluated with the help of authorized operational data and adversarial tests.

The central decision is to whether federation is private by default. Instead, the key question is whether an organization will be able to control federation as a security service. If privacy budgets, update integrity, audit trails and human oversight are combined, federated learning will allow organizations to share threat intelligence without a centralized storage of sensitive data.

10 References

[1] N. Chaurasia, M. Ram, P. Verma, N. Mehta, and N. Bharot, "A federated learning approach to network intrusion detection using residual networks in industrial IoT networks," The Journal of Supercomputing, vol. 80, no. 13, pp. 18325–18346, 2024. doi: 10.1007/s11227-024-06153-2.

[2] G. Q. Zeng, J. M. Shao, K. D. Lu, G. G. Geng, and J. Weng, "Automated federated learning-based adversarial attack and defence in industrial control systems," IET Cyber-Systems and Robotics, vol. 6, no. 2, 2024. doi: 10.1049/csy2.12117.

[3] European Union Agency for Cybersecurity, "ENISA threat landscape 2025," 2025. [Online]. Available: https://www.enisa.europa.eu/publications/enisa-threat-landscape-2025

[4] European Union Agency for Cybersecurity, "ENISA threat landscape 2024," 2024.

[Online]. Available: https://www.enisa.europa.eu/publications/enisa-threat-landscape-2024

[5] M. R. Islam, M. M. Islam, I. A. Badhan, and M. N. Hasnain, "The role of artificial intelligence in carbon pricing policies: Economic and environmental implications," *Journal of Engineering and Computational Intelligence Review*, vol. 3, no. 2, pp. 1–19, 2025.

[6] I. A. Badhan, M. N. Hasnain, M. H. Rahman, I. Chowdhury, and M. A. Sayem, "Strategic deployment of advance surveillance ecosystems: An analytical study on mitigating unauthorized U.S. border entry," *Inverge Journal of Social Sciences*, vol. 3, no. 4, pp. 82–94, 2024.

[7] J. Wang, K. Yang, and M. Li, "NIDS-FGPA: A federated learning network intrusion detection algorithm based on secure aggregation of gradient similarity models," PLOS ONE, vol. 19, no. 10, Art. no. e0308639, 2024. doi: 10.1371/journal.pone.0308639.

[8] R. Parveen, "Microfinance, Vertical Farming, and Data-Driven Innovation in Bangladesh," International Journal of Humanities & Social Science Studies (IJHSSS), vol. 13, no. II, pp. 75–88, Jan. 2025. [Online]. Available: https://www.thecho.in/microfinance,-vertical-farming,-and-data-driven-innovation-in-bangladesh.html

[9] Eurostat, "ICT security in enterprises," 2025. [Online]. Available: https://ec.europa.eu/eurostat/statistics-explained/index.php?title=ICT_security_in_enterprises

[10] I. Kate, "Marketing innovation in the US e-commerce sector: Strategies for success in the evolving digital environment," *International Journal of Business Integrity and Technology Advancements*, vol. 1, no. 2, pp. 1–7, 2025.

[11] S. Akter, "Exploring the role of artificial intelligence in food waste reduction: Implications for public health, social behavior, and sustainable communities," *Journal of Computational Systems & Engineering Insights*, vol. 1, no. 1, pp. 25–36, 2023.

[12] S. Ali, Q. Li, and A. Yousafzai, "Blockchain and federated learning-based intrusion detection approaches for edge-enabled industrial IoT networks: a survey," Ad Hoc Networks, vol. 152, Art. no. 103320, 2024. doi: 10.1016/j.adhoc.2023.103320.

[13] F. Amin, N. Daudpota, and A. Khan, "A complete penetration testing framework: Simulating attacks and evaluating post-exploitation techniques with Kali Linux and Metasploit," *Spectrum of Engineering Sciences*, pp. 386–407, 2025.

[14] A. Kumar, "The role of banking credit access in healthcare analytics-driven practice growth: Implications for regional economic development in the US," *Journal of Computational Systems & Engineering Insights*, vol. 2, no. 2, pp. 15–28, 2024.

[15] F. Amin, M. A. But, I. Amin, and A. Khan, "The tokenized business marketplace: A blockchain and AI-powered framework for democratizing business ownership and investment," *International Journal of Business and Management Sciences*, vol. 5, no. 4, pp. 318–328, 2024.

[16] N. J. Mercer, "Enhancing organizational performance through strategic cloud computing in management," *International Journal of Business Integrity and Technology Advancements*, vol. 1, no. 1, pp. 11–17, 2025.

[17] U. Imtiaz, "Investigating the impact of phishing attacks on organizational cybersecurity posture," *Spectrum of Engineering Sciences*, pp. 1758–1780, 2025.

[18] H. Zhang, J. Ye, W. Huang, X. Liu, and J. Gu, "Survey of federated learning in intrusion detection," Journal of Parallel and Distributed Computing, vol. 195, Art. no. 104976, 2025. doi: 10.1016/j.jpdc.2024.104976.

[19] U. Imtiaz, B. Ahmad, M. H. Sajid, Q. Abbas, M. A. Qureshi, S. Rasheed, and A. Khan, "An integrated machine learning framework for structural health monitoring of bridges: A case study on Soan Bridge," *The Asian Bulletin of Big Data Management*, vol. 5, no. 2, pp. 194–207, 2025.

[20] H. Peng, C. Wu, and Y. Xiao, "FD-IDS: Federated Learning with Knowledge

Distillation for Intrusion Detection in Non-IID IoT Environments," Sensors, vol. 25, no. 14, Art. no. 4309, 2025. doi: 10.3390/s25144309.

[21] U. Iqbal, "AI-enhanced network optimization for electric vehicle charging infrastructure expansion in the United States using graph theory and demand analytics," *Journal of Engineering and Computational Intelligence Review*, vol. 2, no. 2, pp. 112–129, 2024.

[22] F. Ahmed, "Quantum-resistant cryptography for national security: A policy and implementation roadmap," *International Journal of Multidisciplinary on Science and Management*, vol. 1, no. 4, pp. 54–65, 2024.

[23] U. Iqbal, "AI-driven predictive maintenance for U.S. smart manufacturing: Deep learning models for equipment failure prediction and operational resilience," *Journal of Engineering and Computational Intelligence Review*, vol. 3, no. 1, pp. 114–138, 2025.

[24] X. Zhang, Y. Luo, and T. Li, "A Review of Research on Secure Aggregation for Federated Learning," Future Internet, vol. 17, no. 7, Art. no. 308, 2025. doi: 10.3390/fi17070308.

[25] S. Park, J. Lee, K. Harada, and J. Chi, "Masking and Homomorphic Encryption-Combined Secure Aggregation for Privacy-Preserving Federated Learning," Electronics, vol. 14, no. 1, Art. no. 177, 2025. doi: 10.3390/electronics14010177.

[26] Y. Siriwardhana, P. Porambage, M. Liyanage, S. Marchal, and M. Ylianttila, "SHIELD - Secure Aggregation Against Poisoning in Hierarchical Federated Learning," IEEE Transactions on Dependable and Secure Computing, vol. 22, no. 2, pp. 1845–1863, 2025. doi: 10.1109/tdsc.2024.3472869.

[27] K. Fatema, S. K. Dey, M. Anannya, R. T. Khan, M. M. Rashid, C. Su, and R. Mazumder, "Federated XAI IDS: An Explainable and Safeguarding Privacy Approach to Detect Intrusion Combining Federated Learning and SHAP," Future Internet, vol. 17, no. 6, Art. no. 234, 2025. doi: 10.3390/fi17060234.

[28] R. Kalakoti, S. Nõmm, and H. Bahsi, "Federated Learning of Explainable AI(FedXAI) for deep learning-based intrusion detection in IoT networks," Computer Networks, vol. 270, Art. no. 111479, 2025. doi: 10.1016/j.comnet.2025.111479.

[29] C. Harshitha, D. Sendil Vadivu, and N. Rajagopalan, "Federated learning and explainable AI-driven intrusion detection with hyperband optimization," Journal of Computer Virology and Hacking Techniques, vol. 21, no. 1, 2025. doi: 10.1007/s11416-025-00571-3.

[30] J. P. Sahoo, B. Kar, A. M. Abdelmoniem, and D. Chatzopoulos, "Choir-IDS: A federated learning framework for fidelity-calibrated explainable intrusion detection system for edge-IoT networks," Information Fusion, vol. 125, Art. no. 103473, 2026. doi: 10.1016/j.inffus.2025.103473.

[31] G. Rampone, T. Ivaniv, and S. Rampone, "A Hybrid Federated Learning Framework for Privacy-Preserving Near-Real-Time Intrusion Detection in IoT Environments," Electronics, vol. 14, no. 7, Art. no. 1430, 2025. doi: 10.3390/electronics14071430.

[32] A. Alqazzaz, "SecuFL-IoT: an adaptive privacy-preserving federated learning framework for anomaly detection in smart industrial networks," Scientific Reports, vol. 16, no. 1, 2026. doi: 10.1038/s41598-025-11883-1.

[33] Q. Gulzar and K. Mustafa, "Interdisciplinary framework for cyber-attacks and anomaly detection in industrial control systems using deep learning," Scientific Reports, vol. 15, no. 1, 2025. doi: 10.1038/s41598-025-89650-5.

[34] T. T. Khang, T. H. Duc, D. V. Huynh, V. H. Pham, and P. T. Duy, "P4P: A probe-guided anti-poisoning defense for federated learning-based intrusion detection in IoT networks under non-IID data," Journal of Network and Computer Applications, vol. 251, Art. no. 104502, 2026. doi: 10.1016/j.jnca.2026.104502.

[35] A. Pinto, Y. Donoso, and J. A. Gutierrez, "Balancing the trilemma: a survey of federated anomaly detection for secure cyber-physical systems," Cybersecurity, vol. 9, no. 1, 2026. doi: 10.1186/s42400-026-00567-6.

[36] F. Asiri, "Explainable federated learning through causal reasoning for intrusion

detection in IoT," Discover Internet of Things, vol. 6, no. 1, 2026. doi: 10.1007/s43926-026-00292-z.

[37] S. L. Jacob and H. P. Sultana, "GBOA_SR-ShuffleNet: an explainable federated learning framework for privacy-preserving intrusion detection in IoT," Scientific Reports, 2026. doi: 10.1038/s41598-026-56215-z.

[38] Z. Etinge, A. Annamalai, M. Chouikha, and S. Abood, "Privacy-Preserving Federated Cybersecurity Analytics for Smart-Grid SCADA: Maintaining Controllability and Observability Under Coordinated Attacks," Electronics, vol. 15, no. 10, Art. no. 2197, 2026. doi: 10.3390/electronics15102197.

[39] T. D. Nguyen, A. Alazab, A. Khraisat, and T. Jan, "Feature reduction in federated learning for intrusion detection in IoT networks," Cybersecurity, vol. 9, no. 1, 2026. doi: 10.1186/s42400-025-00509-8.

[40] T. Alsolami and M. Ilyas, "FedSMOTE-DP: Privacy-Aware Federated Ensemble Learning for Intrusion Detection in IoMT Networks," Sensors, vol. 26, no. 5, Art. no. 1592, 2026. doi: 10.3390/s26051592.