A Survey on DoS/DDoS Defense using Federated Learning in IoT Networks

2410216 Liang Zhenghao

In recent years, the Internet of Things (IoT) has expanded into a wide range of domains, including smart homes, healthcare and nursing care, agriculture, transportation, and industrial control (IIoT/OT), increasing its importance as a component of social infrastructure. Security issues in IoT networks are not limited to the compromise of information assets; they can also directly affect physical safety and business operations. In this context, DoS/DDoS remains a major threat, and recent industry reports repeatedly highlight both the growing scale and increasing frequency of attacks. Cloudflare's quarterly reports indicate not only an increase in observed large-scale DDoS incidents but also a shortening of peak durations, suggesting stronger requirements for rapid detection and response. Moreover, during Q2–Q3 of 2025, 8.3 million incidents were reported, including extremely large attacks reaching up to 29.7 Tbps and 14.1 Bpps. In addition, in OT/ICS domains, cyber incidents have been noted to lead to operational disruption such as business shutdowns and service outages, and cases have also been reported in which DoS attacks escalated to physical damage.

However, IoT networks operate in distributed settings where edge devices have limited CPU, memory, and power, and communications are often unstable (wireless links, bandwidth limits, variable delays, and loss). Multi-stakeholder operations further induce Non-IID data, making a single centralized model hard to apply uniformly. Under these constraints, DoS/DDoS can reduce availability via bandwidth exhaustion and latency spikes, disrupting monitoring/control. Thus, IoT/OT requires operational feasibility beyond accuracy—sustained detection/mitigation across devices, time bounds, and network conditions—because DDoS-driven congestion can break log collection, inference, and even model updates.

Against this backdrop, DoS/DDoS defense using federated learning (Federated Learning; FL) has attracted attention as a means to protect privacy, preserve data sovereignty, and mitigate data silos. FL enables collaborative learning without sharing raw data and thus helps protect information privacy; however, it depends on communication rounds and is sensitive to the status of participating devices and network conditions. In other words, when applying FL to IoT, the ability for training to "keep running" becomes a fundamental feasibility requirement, and communication degradation caused by DoS/DDoS can undermine FL's basic assumptions. This tension cannot be fully addressed by conventional IDS/classifier-centered evaluation frameworks alone, and a structured analysis is needed that includes behavior under degraded communication, the ability to sustain updates,

time requirements such as inference latency and end-to-end (E2E) response, and the feasibility of implementation on low-performance devices.

This survey aims to systematically organize recent research trends and practical operational challenges in studies that apply FL to DoS/DDoS detection and mitigation in IoT environments. We searched the literature using Google Scholar and ScholarGPT and selected 50 papers according to predefined criteria. The selected studies are concentrated in 2024–2025, confirming rapid growth of the field in recent years. We organized the literature using seven frameworks, including categories based on detector architectures (CNN/DNN/RNN/GNN/AE/Transformer) as well as system/network-mechanism-driven studies that emphasize operational feasibility—such as communication, participation, and trust—rather than detectors alone. We then extracted FL design aspects (e.g., aggregation, participation control, reliability and security) and operationally relevant metrics (inference latency, E2E response time, communication volume, evaluation under congestion/loss conditions, hardware-based measurement, and low-performance device evaluation), and analyzed reporting trends by publication year.

As a result, studies evaluated in computation-only settings (i.e., without confirmed implementations of update transmission/reception involving communication) were the most common(about 70%), indicating that many works do not experimentally examine how DoS/DDoS affects communication and latency. Dataset usage still heavily relies on public datasets from 2015–2019, with 68% of uses tied to datasets from 2019 or earlier. Among the 44 papers that explicitly mentioned IoT deployment, only eight reported real-device measurements (six on low-performance devices), while most relied on PC/server experiments. low-performance devices require model lightweighting, control of update frequency, and reductions in communication rounds due to limited computation and power, DoS/DDoS can impair training continuity through bandwidth exhaustion and increased latency; thus, beyond accuracy, design and evaluation must consider whether "training can continue running." We also found that some papers provide insufficient descriptions of measurement units and measurement points (device/gateway/server) for inference latency, E2E response, and communication volume, raising concerns about reproducibility. In addition, performance variability due to Non-IID data and environmental differences, zero-day attacks, the need to clarify re-training procedures under congestion, and the lack of cross-dataset evaluation were identified as cross-cutting challenges.

Based on these findings, this survey compiles reporting items into a checklist and organizes evaluation conditions and experimental environments in tabular form, with the expectation of providing practical guidance for research design aimed at real-world deployment in IoT environments.