

Blockchain for Federated Learning in the Internet of Things: Trustworthy Adaptation, Standards, and the Road Ahead

Farhana Javed*, Engin Zeydan*, Josep Mangués-Bafalluy*, Kapal Dev† and Luis Blanco[‡]

*Services as networkS (SaS), CTTC/CERCA, Castelldefels, Spain

†Department of Computer Science, Munster Technological University, Cork, Ireland

[‡]Space and Resilient Communications and Systems (SRCOM), CTTC/CERCA, Castelldefels, Spain
{farhana.javed, josep.mangués, ezeydan, lblanco}@cttc.es; kapal.dev@mtu.ie

Abstract—As edge computing gains prominence in Internet of Things (IoTs), smart cities, and autonomous systems, the demand for real-time machine intelligence with low latency and model reliability continues to grow. Federated Learning (FL) addresses these needs by enabling distributed model training without centralizing user data, yet it remains reliant on centralized servers and lacks built-in mechanisms for transparency and trust. Blockchain, a type of Distributed Ledger Technologies (DLTs) can fill this gap by introducing immutability, decentralized coordination, and verifiability into FL workflows. This article presents current standardization efforts from 3GPP, ETSI, ITU-T, IEEE, and O-RAN that steer the integration of FL and blockchain in IoT ecosystems. We then propose a blockchain-based FL framework that replaces the centralized aggregator, incorporates reputation monitoring of IoT devices, and minimizes overhead via selective on-chain storage of model updates. We validate our approach with IOTA Tangle, demonstrating stable throughput and block confirmations, even under increasing FL workloads. Finally, we discuss architectural considerations and future directions for embedding trustworthy and resource-efficient FL in emerging 6G networks and vertical IoT applications. Our results underscore the potential of DLT-enhanced FL to meet stringent trust and energy requirements of next-generation IoT deployments.

Index Terms—Federated learning, blockchain, Internet of Things (IoTs), standardization, trust, 6G networks.

I. INTRODUCTION

THE Third Generation Partnership Project (3GPP) has been at the forefront of the development of fifth generation (5G) and beyond 5G systems for various applications. In Release 16, 3GPP introduced the Network Data Analytics Function (NWDAF) in its Technical Specifications to enable data-driven insights and intelligent decisions in the core network. In Release 17, the analytical scope has been extended to meet the heterogeneous requirements of 5G services. Building on these improvements, Release 18 often referred to as 5G Advanced integrates Artificial Intelligence (AI) and Machine Learning (ML) across network operations to enable increasingly flexible and distributed provision of services [1].

This work was partially funded by Spanish MINECO grants TSI-063000-2021-54 and TSI-063000-2021-55 (6G-DAWN), Grant PID2021-126431OB-I00 funded by MCIN/AEI/10.13039/501100011033, by “ERDF A way of making Europe” (ANEMONE), Generalitat de Catalunya grant 2021 SGR 00770, and by UNITY-6G project, funded from the European Union’s Horizon Europe Smart Networks and Services Joint Undertaking (SNS JU) research and innovation programme under Grant Agreement No. 101192650.

These standards support the shift to edge computing in IoT, where applications like smart cities and healthcare require low latency, real-time responsiveness, and reliable models. Decentralized analytics at edge and fog nodes enable faster, context-aware processing while easing cloud load. Also, Federated Learning (FL) enables privacy-preserving intelligence by allowing devices to train local models and share only parameter updates. This approach suits IoT environments with diverse device capabilities, non-uniform data, and intermittent connectivity.

Despite its advantages, FL faces challenges in IoT [2]. Transmitting large model updates over constrained or intermittent links can increase overhead, and even sharing gradients may leak sensitive information. A centralized FL aggregator also introduces a potential bottleneck and single point of failure. For example, ETSI’s Zero-touch Network and Service Management (ZSM) group emphasizes the need for decentralized ML services with built-in *trustworthiness* and governance. Similarly, ETSI reports on collaborative AI scenarios stress that decentralized FL must include mechanisms to verify model integrity without centralizing data.

The integration of blockchain with FL and IoT can enhance trust and eliminate centralized aggregation servers. In IoT-driven FL, blockchain offers a tamper-evident, transparent infrastructure that reduces single points of failure. Contributions such as model parameters and device identifiers are securely recorded and verified, ensuring integrity and minimizing data leakage. 3GPP Release 18 extends the 5G System with features to monitor FL processes and evaluate participant performance. Meanwhile, ETSI explores DLTs to improve trust in FL ecosystems. Permissioned Distributed Ledger (PDL) technology supports secure, verifiable model exchanges, enabling reliable monitoring of FL participants.

While blockchain can replace the central aggregator in FL for IoT, it introduces challenges. IoT devices have limited energy and processing capacity, requiring lightweight integration. Not all data should be anchored on-chain, so gateways must filter and forward essential information. The blockchain must also manage resource constraints while taking on aggregation responsibilities. This shift necessitates on-chain verification and reputation mechanisms to ensure model integrity and trust.

Several works have explored blockchain and DLTs for decentralized FL. For example, [3] propose DLT-based ap-

TABLE I
SUMMARY OF CURRENT STANDARDIZATION EFFORTS FOR FEDERATED LEARNING IN IOT CONTEXT

Standardization Body	Working Group	Focus Area	Reference / Standard
3GPP	SA6 [*]	Adversarial FL threats, GDPR in IoT FL	TR 33.846
	RAN WG ⁺	FL architecture and resource allocation in RAN	TR 37.817 (Rel-17), NWDAF (Rel-18)
O-RAN Alliance	WG1, WG2 ⁺	FL in Near-RT RIC, AI/ML model exchange	O-RAN.WG1.AI/ML-v02.00
	SMO [†]	FL workflow orchestration across IoT domains	O-RAN SMO Framework
ETSI	ENI ⁺	Lightweight FL in constrained IoT environments	ETSI GS ENI 005
	TS [†]	Federated automation for cross-border IoT FL	ETSI TS 103 457
IEEE	P3652.1 ⁺	Heterogeneous IoT FL architecture	IEEE P3652.1
	P1934.1 [‡]	Industrial IoT	IEEE P1934.1
	1451-99 ⁺	Standardized APIs for FL-enabled IoT devices	IEEE 1451-99
IETF	DICE WG [*]	Lightweight security for constrained FL endpoints	DTLS profiles for IoT
IIC	Arch. Task Group [‡]	Reference architectures for FL in industrial IoT	IIC RA for Industrial Assets

Legend: * Security & Privacy (GDPR, adversarial FL, constrained environments), + Architectural Standards (FL system design, APIs, interoperability), † Orchestration & Automation (cross-domain FL, SMO, zero-touch workflows), ‡ Industry Solutions & Use Cases (industrial IoT, deployment reference architectures).

proaches to eliminate the need for a centralized server, while [4] describes an asynchronous FL framework that uses blockchain for flexible aggregation. In digital twin edge networks, [5] uses a permissioned blockchain to secure FL updates through smart contracts. Reputation mechanisms are also discussed or implemented in [6] and [7], which use voting-based methods to evaluate participant contributions. Further optimizations are proposed in [3], [8], [9], with a focus on scalability and IoT constraints. Related schemes in adjacent domains include BSIF for incentive-driven mobile crowdsensing [10] and BDSS for fine-grained medical data sharing [11]; both address secure one-shot data exchange but not iterative FL aggregation or on-chain reputation under IoT limits.

Although these works address partial aspects of the overall problem some remove the central server, others use reputation mechanisms, and still others optimize FL for the IoT no single approach covers all these aspects at once. Many omit the DLT-based reputation, retain elements of the centralized design, or overlook the strict resource constraints. Traditional FL still hinges on a trusted aggregator [2], DAG-based schemes such as DAGFL [4] decentralize aggregation but lack explicit trust scoring, and Ethereum-based solutions [12] introduce reputation yet incur gas fees unsuited to low-power nodes. Designing a framework that removes the central server, accounts for IoT constraints, incorporates reputation mechanisms, and keeps overhead low remains an open challenge in current research. The present work closes this gap by combining feeless IOTA transactions, an on-chain lightweight reputation manager, and gateway-assisted off-loading, thereby delivering server-less aggregation, verifiable trust, and energy awareness.

Therefore, the key contributions are:

- From a standardization perspective, we outline current SDO efforts shaping emerging standards and regulatory

frameworks for this convergence.

- We propose a framework to enable blockchain-based FL in IoT environments, highlighting a framework where blockchain serves as a replacement for aggregation servers, considering the challenges and limitations of IoTs and blockchain.
- We conclude with key findings and a 6G-oriented roadmap, highlighting trusted AI/ML integration and IoT vertical management based on the proposed architecture.

II. OVERVIEW OF STANDARDIZATION EFFORTS AND CURRENT CHALLENGES

A. Standardization Efforts for FL in IoT Context

The 3GPP mentions FL to enable distributed AI/ML in 5G/6G Radio Access Networks (RAN) and IoT systems. Its FL architecture defines roles such as Model Owner, Aggregation Server, and Edge Nodes, including User Equipment (UE), next-generation Node Bs (gNBs), and Distributed Units (DUs). The Service and System Aspects Working Group 6 (SA6) also highlights adversarial threats in FL and the need for General Data Protection Regulation (GDPR)-compliant processing in industrial IoT settings.

The Open Radio Access Network (O-RAN) Alliance complements 3GPP by embedding FL into open RAN architectures. O-RAN Working Groups 1 and 2 address FL interoperability through components supporting use cases and aligns with 3GPP's NWDAF for cross-domain FL workflows in IoT scenarios. Key challenges include adapting FL to O-RAN's service-based architecture (SBA) and achieving low-latency updates across DUs. Federated attestation using hardware root-of-trust is proposed to counter risks from rogue or biased nodes.

ETSI's Experiential Networked Intelligence (ENI) group addresses FL scalability and regulatory compliance in IoT,

proposing lightweight communication frameworks for dynamic edge environments. Additional ETSI efforts support cross-border FL and zero-touch workflows. The Institute of Electrical and Electronics Engineers (IEEE) contributes standards for heterogeneous FL architectures, industrial IoT, and interoperable APIs. The Internet Engineering Task Force (IETF), through its DTLS In Constrained Environments (DICE) group, defines lightweight encryption, while the Industrial Internet Consortium (IIC) offers FL-ready reference architectures for logistics and asset tracking. Table I summarizes the discussed standardization efforts.

B. Blockchain Standardization Efforts

A *blockchain* is a type of distributed ledger technology (DLT) where data is stored in cryptographically linked blocks. Each block references the previous one via a hash, making tampering easily detectable and preserving integrity.

Standards bodies like including the International Telecommunication Union - Telecommunication Standardization Sector (ITU-T), IEEE, ETSI, and China Communications Standards Association (CCSA) are developing specifications to support interoperability, security, and compliance in DLTs. Table II summarizes these efforts.

ETSI's Industry Specification Group on Permissioned Distributed Ledgers (ISG PDL) develops frameworks for secure, interoperable DLT in regulated domains. Its standards cover technology landscape, governance, use cases across sectors (e.g., supply chain, public services), and PoC guidelines. Recent work explores integrating PDL with the oneM2M IoT service layer and highlights AI/FL applications for decentralized, privacy-preserving model training in areas such as IoT security, fraud detection, and adaptive learning. Moreover, ITU-T coordinates blockchain standardization through Focus Groups (FGs) and Study Groups (SGs). FG DLT addresses terminology, architecture, and cross-sector use cases, while SG20 focuses on IoT and smart cities with recommendations for blockchain-based data management, secure processing, and next-generation networking.

IEEE defines blockchain standards for data management, IoT, identity, and cross-domain interoperability, including healthcare and energy. China's CCSA focuses on performance, Blockchain-as-a-Service (BaaS) scalability, and 5G/IoT integration. The CAMARA project standardizes APIs for multi-operator environments and explores blockchain-based identity and provenance. TM Forum develops DLT-based frameworks for billing, roaming, and identity in telecom networks.

C. Challenges

Consensus mechanisms in public blockchains, permissioned ledgers, and Layer-2 (L2) solutions pose challenges for FL in IoT due to the resource limits of edge devices. Although Proof of Stake (PoS) is more energy-efficient than Proof of Work (PoW), it still requires significant computation and incurs fees, especially during congestion. Frequent FL updates can overwhelm the network with validation demands, increasing latency. Full ledger replication also exceeds the storage capacity of typical IoT nodes.

While EVM-based blockchains and smart contracts offer flexible management, their computational overhead and complexity hinder adoption in IoT-based FL. Permissioned ledgers like Hyperledger Fabric improve throughput but rely on consensus protocols such as PBFT, which increase latency and bandwidth use at scale. The iterative nature of FL amplifies these limitations. Supporting FL in IoT requires lightweight blockchain designs that balance efficiency, selective on-chain storage, privacy, and decentralized verification.

III. BLOCKCHAIN-BASED FL IN IoT: SYSTEM ARCHITECTURE AND WORKFLOW

A. Integration of Blockchain in FL for IoT

The framework supports end-to-end FL across IoT devices using a DLT-based aggregator for tamper-proof, privacy-preserving model updates. It enables secure coordination of heterogeneous nodes for applications like reputation scoring, analytics, and business intelligence. Key components include:

IoT Orchestration and Context Manager: This component orchestrates interactions between IoT devices, FL processes, and blockchain validation. It includes three modules: IoT Data Analytics & Insights (IDAI), Trust & Reputation Management (TRM), and Federated AI & Optimization Engine (FAOE). IDAI evaluates device context (e.g., location, battery, speed) for FL participation. TRM assigns trust scores based on historical behavior and blockchain-verified compliance. FAOE manages training, device selection, and update coordination, adapting to network and resource conditions. This aligns with standardization efforts such as 3GPP's NWDAF and Policy Control Function (PCF).

Unlike centralized FL, which channels all updates through a cloud aggregator, and earlier blockchain-FL systems that rely on gas-priced smart contracts, our IOTA-based approach stores only hashed model metadata on a feeless ledger and off-loads computation from devices, thereby relieving battery-powered nodes of heavy cryptographic work.

Local Models and IoT Deployed Network: An IoT network comprises heterogeneous devices collecting domain-specific environmental and operational data for applications like smart cities, healthcare, industrial automation, and traffic management. For example, in traffic scenarios, connected vehicles, cameras, and sensors measure flow rates, speeds, and congestion levels. Each device trains a local model on its own non-independent and identically distributed data, reflecting unique conditions and usage patterns. Local training preserves data confidentiality by eliminating the need to transfer raw data. While this process typically occurs at the device level, similar concepts can apply within the network infrastructure for instance, a Network Function (NF), such as the Access and Mobility Management Function, could also train local models based on aggregated network data. At this stage, no aggregation occurs devices or NFs refine models independently based on local distributions. This approach aligns with 3GPP efforts in network data analytics and edge computing, enabling AI-driven optimizations and improved resource allocation through integration with components like NWDAF.

FL-IoT Decentralized Application (DApp): This layer operates below IoT orchestration and local model training, enabling

TABLE II
SUMMARY OF CURRENT STANDARDIZATION EFFORTS FOR BLOCKCHAIN

Standardization Body	Working Group	Focus Area	Reference / Standard
ETSI	ISG PDL*	PoC Framework	GS PDL 005
	ISG PDL*	Smart Contracts	GS PDL 011
	ISG PDL+	Landscape of Standards and Technologies	GS PDL 001
	ISG PDL+	Application Scenarios	GS PDL 003
	ISG PDL+	Inter-Ledger Interoperability	GS PDL 006
ITU-T	FG DLT+	DLT Terms and Definitions, DLT Overview, Concepts, Ecosystem, Standardization Landscape and Reference Architecture	FG DLT D1.1, D1.2, D1.3
	FG DLT+	DLT Use Cases	FG DLT D2.1
	FG DLT+	DLT Regulatory Framework	FG DLT D4.1
	FG DLT¶	Blockchain & DLT	REC-F.751.2
	SG20+	IoT	Y.dec-IoT-arch, REC-Y.4476
	SG20¶	NGN	REC-Y.2342
IEEE	P2144+	Data Management	Data Management Standards
	P2418+	IoT, CAV, Energy	IoT and Energy Standards
	P2958+	Identity & Access Management	Identity and Access Standards
	P3201-P3214+	Access Control, Interoperability	Access Control Standards
CCSA	TC1WG6*	Testing Methods	2017-0942T-YD
	TC1WG6†	Platform (BaaS)	2019-12527-YD
	TC5WG6†	Wireless Network Applications	2020B896
	TC10WG1‡	5G & Blockchain for IoT	2020B858

Legend: * PoC & Testing (PoCs, validation, testing), + Technical Standard (Protocols, security, interoperability), ¶ Reference Framework (Guidelines, architectural frameworks), † Industry Solutions (Application-driven solutions), ‡ Emerging Technologies (Study items for future).

secure data exchange and FL aggregation on a distributed ledger. Its core components the *DLT-Adapter*, *DLT-Reputation Mngr*, and *DLT-Aggregator* are managed on-chain by the *DLT-DApp Manager*. The *DLT-Adapter* acts as a gateway, filtering and forwarding essential data (e.g., hashed model weights) from authenticated devices to the blockchain, minimizing on-chain load. The *DLT-Reputation Mngr* scores devices and model updates based on reliability, guiding weighted contributions to improve global model quality. The *DLT-Aggregator* merges validated updates into the global model, stored on-chain for reference by devices or services. A permissioned IOTA Tangle limits participation reducing the likelihood of malicious actors. These trusted validators confirm transactions faster and help maintain Byzantine fault tolerance, meaning the system can still reach correct consensus even if some validators act arbitrarily or maliciously. Optional milestones special checkpoints issued by a coordinator lock in certain points of the Tangle, ensuring transactions are finalized quickly and making the overall process more reliable for sensitive or time-critical applications. Moreover, a permissioned IOTA Tangle supports FL-IoT particularly well: its feeless messages and negligible micro-PoW minimise energy draw on constrained devices, while milestone checkpoints provide deterministic finality even under bursty update traffic. Because IOTA stores only hashed model metadata, it avoids the gas-driven virtual-machine overhead typical of smart-contract platforms. In contrast, alternatives such as Hyperledger Fabric (endorsement-

heavy PBFT with TLS containers), Tendermint/Cosmos (PoS fees and signature traffic), or other private DAGs (token fees and gossip-intensive virtual voting) impose higher cost, lower throughput highlighting their limitations for resource-limited IoT devices. Compared with DAGFL [4], which decentralizes aggregation but treats all clients equally, our on-chain reputation module weights updates by historical honesty, strengthening robustness to Byzantine behaviour. Relative to Ethereum-based FL frameworks [12] the proposed Tangle workflow eliminates transaction fees and avoids virtual-machine overhead, resulting in faster confirmations and lower device workload.

B. FL Model Integration of IoT Devices with Blockchain

Our design leverages a permissioned variant of the IOTA Tangle, inherently reducing consensus overhead. However, direct interaction between IoT nodes and the ledger remains impractical, as repeated confirmations or computations strain energy-limited devices and disrupt critical operations.

To address this, we embed a *DLT-Adapter* within the *DLT-DApp Manager* (Figure 1), serving as a verified intermediary between IoT devices and the permissioned Tangle. This approach aligns with ETSI PDL recommendations (ETSI-PDL-028) advocating selective on-chain storage to prevent ledger bloat and reflects 3GPP guidelines on minimizing device-side resource consumption. The *DLT-Adapter* verifies device identity (by TRM module), filters redundant or invalid

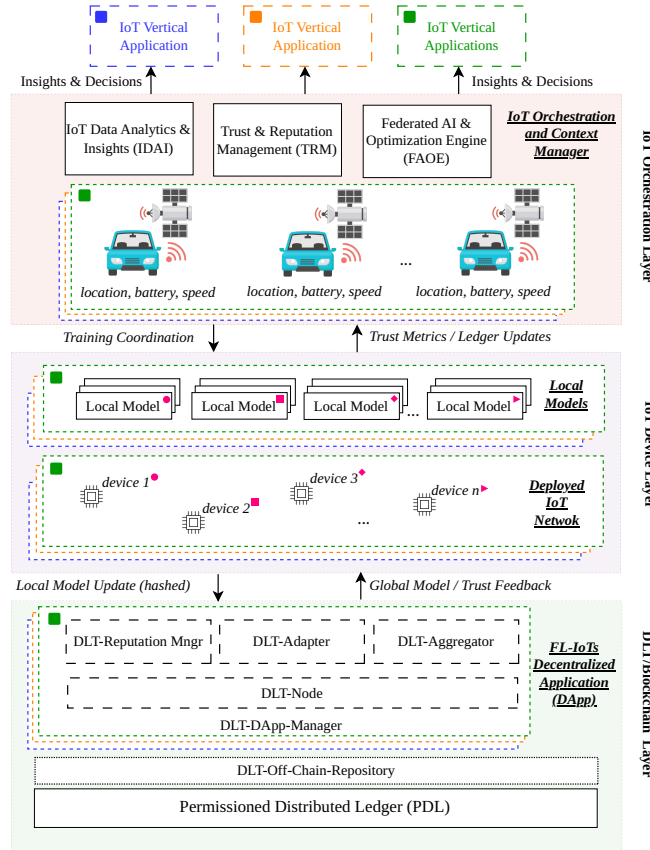


Fig. 1. High-level view of system architecture for DLT-enabled trustworthy and FL for IoT.

data, and coordinates further processing of valid local model updates.

In this architecture, IoT nodes send local model updates without interacting directly with the ledger. The DLT-Adapter forwards these to the DLT-Aggregator and anchors only a cryptographic hash on the Tangle. This approach preserves privacy, reduces on-chain storage, and ensures verifiable model exchange without burdening devices or exposing sensitive data.

C. Reputation Management

This architecture accounts for idle, unreliable, or malicious nodes. The *DLT-DApp Manager* includes the *DLT-Reputation Mngr* and *DLT-Aggregator*. Model updates are sent via the *DLT-Adapter* to the *DLT-Reputation Mngr*, which verifies contributions and assigns reputation scores based on quality. Invalid or outdated updates are penalized, while valid ones improve or maintain scores. Nodes below a trust threshold are down-weighted in future aggregations. Scores, updated by combining past reputation and current model accuracy, are stored on the ledger for accountability. The *DLT-Aggregator* computes the global model using weighted inputs, and anchors a hash of the result and references to verified updates on the ledger, supporting traceability in line with ETSI PDL guidelines.

D. Lightweight Transactions

In large-scale, resource-constrained IoT settings, storing full model updates on the ledger is impractical. Our lightweight framework addresses this by anchoring only essential metadata and hashes on the *IOTA Tangle*, with actual model data kept off-chain. The *DLT-Adapter* authenticates and validates submissions via low-overhead *MQTT*, allowing devices to focus solely on local training. Verified updates are passed to the *DLT-Aggregator*, while only hashes and minimal metadata are anchored using feeless *Zero-Value Transactions*. Remote *Proof-of-Work (PoW)* offloads computation to permitted *IOTA* nodes. Feeless zero-value messages, remote micro-PoW, and hash-only anchoring together keep ledger traffic minimal compared with endorsement-heavy Hyperledger Fabric or fee-based PoS sidechains, while still providing full auditability an essential benefit for constrained IoT deployments.

IV. EXPERIMENTAL EVALUATIONS

We evaluate the system using two metrics: average transactions per second (TPS) with standard deviation and variability, and block processing time defined as the delay between transaction submission and confirmation by a milestone in *IOTA*. Each transaction represents a model update from a simulated IoT client in an FL round. Experiments ran for 10, 30, and 50 rounds, repeated 10 times for consistency, on Ubuntu 22.04 LTS with an i7-1265U CPU (10 cores, up to

TABLE III
SIMULATION ENVIRONMENT SETUP AND SYSTEM THROUGHPUT DURING FL EXECUTION

Simulation Environment Setup			
Host OS	Linux (Ubuntu 22.04 LTS)		
CPU	12th Gen Intel(R) Core(TM) i7-1265U, 10 cores, max frequency 4.8 GHz		
Memory	31 GiB DDR4		
Docker Containers	Hornet Node 1: CPU 0.59%, 583.9 MiB Hornet Node 2: CPU 3.13%, 581.6 MiB INX Coordinator: 10.49 MiB Traefik (proxy): 21.27 MiB Monitoring and explorer services deployed		
Tangle Setup	Hornet v2.0.2 Base token: SandCoin (SAND) Milestone interval: default 10 s Block rate: ~0.2 blocks/s		
Payload Sizes	DLT-Adapter: 2–3 KB DLT-Aggregator: 2–3 KB DLT-Reputation Mngr: 1.5–2 KB		
MQTT Setup	Local broker (Mosquitto v2.0.11) on localhost:1883 Paho-MQTT Python client (v2.1.0) used for publishing		
Off-chain Storage	IPFS (v0.21.0)		
System Throughput and Variability During FL Execution			
FL Rounds	Average TX/sec	Std Dev	Variability (%)
10	1.82	0.56	30.8%
30	2.10	0.15	7.1%
50	2.12	0.10	4.90%

4.8 GHz) and 31 GiB RAM. A private, permissioned IOTA Tangle (Hornet v2.0.2) with two nodes handled milestones and validation. Node 1 received transactions via the DLT-Adapter; Node 2 managed gossip and confirmations. An INX coordinator handled milestones, and monitoring was included. All components ran in Docker (details in Table III). Transactions were sent using a Python client through a local Mosquitto MQTT broker (v2.0.11) and Paho-MQTT (v2.1.0). Payloads ranged from 1.5–3 KB, aligned with IOTA’s 32 KB limit, and weights were stored off-chain via IPFS (v0.21.0). The FL setup used public IoT data¹, partitioned across 20 simulated clients, each training a lightweight neural network (one hidden layer) for 20 local epochs. Updates were aggregated using FedAvg, with reputation scores based on validation accuracy.

A. Simulation Results

Table III shows system throughput and variability metrics for FL executions consisting of 10, 30 and 50 rounds. These results provide important insights into the scaling of transaction volume and node performance as FL update bursts increase, an important consideration for balancing low-latency responses and frequent model updates in IoT deployments. As shown in Table III, the average processed transactions per second (TX/sec) increase from 1.82 at 10 rounds to 2.10 and 2.12




Fig. 2. Distribution of block processing time across different FL rounds: 10, 30, and 50.

at 30 and 50 rounds, respectively. Parallel to this increase in throughput, the standard deviation and variability fall sharply from 30.8% at 10 rounds to 7.25% at 30 rounds and 4.9% at 50 rounds. This trend reflects an important observation: the system becomes more stable and predictable when it operates under higher, sustained load conditions. This stability stems from the fact that a well-configured IOTA node can accommodate predictable bursts without congestion or erratic queue buildup, confirming the resilience and scalability of the system for moderate- to large IoT FL workloads.

Figure 2 presents the distribution of block processing times for each FL scenario. Interestingly, we do not observe a significant upward shift in block processing times as the number of FL rounds increases. Instead, the block processing times remain consistent in all three scenarios, with median values clustering around 2 to 4 seconds and occasional outliers in the range of 10 to 16 seconds. These outliers correspond to the fixed milestone issuance intervals of nodes (approximately 10 seconds) and reflect transactions that just miss one milestone and wait for the next cycle. The fact that block processing times did not increase at higher loads indicates that the system has made good use of node processing capacity with no sustained backlog or congestion. This behavior underscores that block confirmation times are driven more by milestone issuance rates than by incremental increases in transaction volume as long as the system remains within operational limits.

The observed behaviour provides a realistic baseline for IOTA-based FL deployments in larger setups. However, factors such as network-propagation delay, bursty transaction spikes, and distributed milestone confirmations can still introduce fluctuations. Nonetheless, our results show that the node remained stable and responsive at moderate-to-high loads (up to 50 FL rounds with 20 devices each), confirming its suitability for FL scenarios. By contrast, Ethereum-based FL prototypes e.g., [12] and [6] inherit the chain’s block-confirmation interval and levy transaction fees on every model update. DAGFL [4] removes fees but, lacking on-chain reputation, must replay suspect updates under adversarial behaviour, thereby reducing effective throughput.

¹<https://iotanalytics.unsw.edu.au/iottraces.html>

V. DISCUSSION AND FUTURE DIRECTIONS

A. Discussion and Key Takeaways

Improving Blockchain Network Scalability: A key takeaway from our results is the need to optimize scalability when integrating FL with DLTs like IOTA. Our PoC treats each FL update as a separate transaction hash, which remains stable at moderate scale. As federations grow, *batch processing* can reduce ledger congestion and queuing delays, though optimal batch sizing is critical to limit latency. Larger deployments may also introduce overhead from gossip propagation and milestone syncing, pointing to the need for hierarchical clustering or multi-layer node architectures. While no severe congestion was observed in our tests, future systems must plan for potential bottlenecks under bursty loads through dynamic scheduling and resource provisioning.

Latency-Security Trade-offs: IOTA's milestone-based confirmation mechanism delivers strong consistency and data integrity for FL model updates but naturally introduces confirmation latency, particularly for transactions that narrowly miss milestone intervals. While this latency remained predictable in our controlled environment, larger and more dynamic deployments may benefit from selective mechanisms such as partial confirmations or fast-lane processing for time-sensitive IoT data. This highlights a trade-off between latency and integrity that becomes increasingly significant for real-time or mission-critical decision-making in large-scale FL deployments.

Energy Efficiency and Off-Chain Extensions: In resource-constrained IoT settings, energy efficiency and off-chain solutions are vital. Our setup simulated devices communicating through an MQTT broker, with payloads sized to reflect realistic FL model updates (200 KB to 2 MB). However, transmitting these updates over constrained wireless links (such as NB-IoT) would amplify energy costs and tail latencies. Solutions that rely on off-chain storage such as IPFS can reduce the complexity of on-chain data and the transaction load. As systems scale, future architectures may need to incorporate side-DAGs, sharding techniques, or advanced off-chain protocols to balance computational overhead, data availability, and ledger integrity while remaining energy-aware.

With a 10s milestone interval and a 32 kB payload ceiling in Hornet at most 16 model hashes can be confirmed per milestone, yielding ≈ 1.6 updates s^{-1} without batching or roughly 500 concurrent devices when those 16 hashes are aggregated into each message. Each FL round therefore uploads ≤ 0.5 MB. For example, to better reflect performance in realistic environments involving numerous heterogeneous IoT devices, we provide an upper-bound estimate of communication energy. Applying the 2 mWh MB^{-1} Cat-NB1 uplink figure reported by [13] gives an upper-bound radio cost of $\leq 1 \text{ mWh per round} < 0.7\%$ of a 150 mWh CR1225 coin cell and $< 0.15\%$ of a 700 mWh CR2032. Wi-Fi or Ethernet backhaul lowers this communication energy by roughly an order of magnitude. While this analysis offers a reproducible estimate relevant to constrained settings, we acknowledge that detailed energy consumption measurements across diverse hardware platforms remain an important direction for future work.

Threat Model and Mitigation Strategies: Threat models are essential in FL and IoT, where distributed, resource-constrained devices are vulnerable to attacks such as model poisoning, Sybil attacks, and data tampering. Combined attacks where adversaries inject poisoned models through multiple Sybil nodes pose a serious threat to global model integrity. Our system introduces the *DLT-Reputation Mngr*, which assigns dynamic scores based on contribution quality; the *DLT-Adapter*, which authenticates and filters inputs; and the *DLT-Aggregator*, which prioritizes high-reputation updates based on inputs from the other two components. This reputation-aware selection can help limit the influence of malicious or colluding nodes by down-weighting or excluding suspicious updates from aggregation. While these form a foundation for decentralized trustworthy architecture, future FL deployments will face increasingly complex threats. As systems scale, architectures must adopt hierarchical trust models and zones, energy-aware reputation decay, zero-trust techniques for unreliable participants, and dynamic anomaly detection. These features are critical for secure, low-latency FL in dense IoT environments.

Handling Convergence, Accuracy Trends, and Non-IID Challenges: FL in IoT environments faces several challenges, including non-IID data distributions, dynamic client participation, and the need for stable model convergence. Non-IID data where each device observes a distinct subset of the input space can cause model divergence, slower convergence, and degraded global accuracy due to conflicting or biased local updates. Our architecture addresses these challenges through three key mechanisms: (i) per-client validation accuracy is used to assess update quality; (ii) the *DLT-Reputation Mngr* assigns dynamic scores based on historical and current performance; and (iii) the *DLT-Aggregator* prioritizes updates from higher-reputation clients, thereby adjusting aggregation thresholds and reducing the influence of unreliable or skewed updates. As detailed in Section III-A and IV-A, our simulation employs IoT data partitioned across clients to reflect realistic non-IID distributions. Although model accuracy trends and convergence curves are not explicitly presented the architecture is designed to support stable learning via selective aggregation. Reputation-based filtering effectively down-weights untrustworthy updates, serving as a lightweight and scalable method to mitigate the effects of data heterogeneity. Future work will incorporate adaptive weighting schemes and decay models to further improve these aspects under extreme imbalance and participation variability.

B. Future Directions

Figure 3 outlines a 6G network architecture that integrates AI, FL, and DLT. The Function Layer (e.g., DLT-TF modules for Unified Data Management, AI/FL, and Sensing Functions) emphasizes native intelligence and distributed computing, supporting IoT services such as smart cities and Industry 4.0. The Trust Layer (DLT nodes and DApp-Manager) and the PDL Layer (integration of control/user/data planes with DLT and AI) emphasize decentralized security and transparency, which are essential for FL in IoT ecosystems. While this is in line

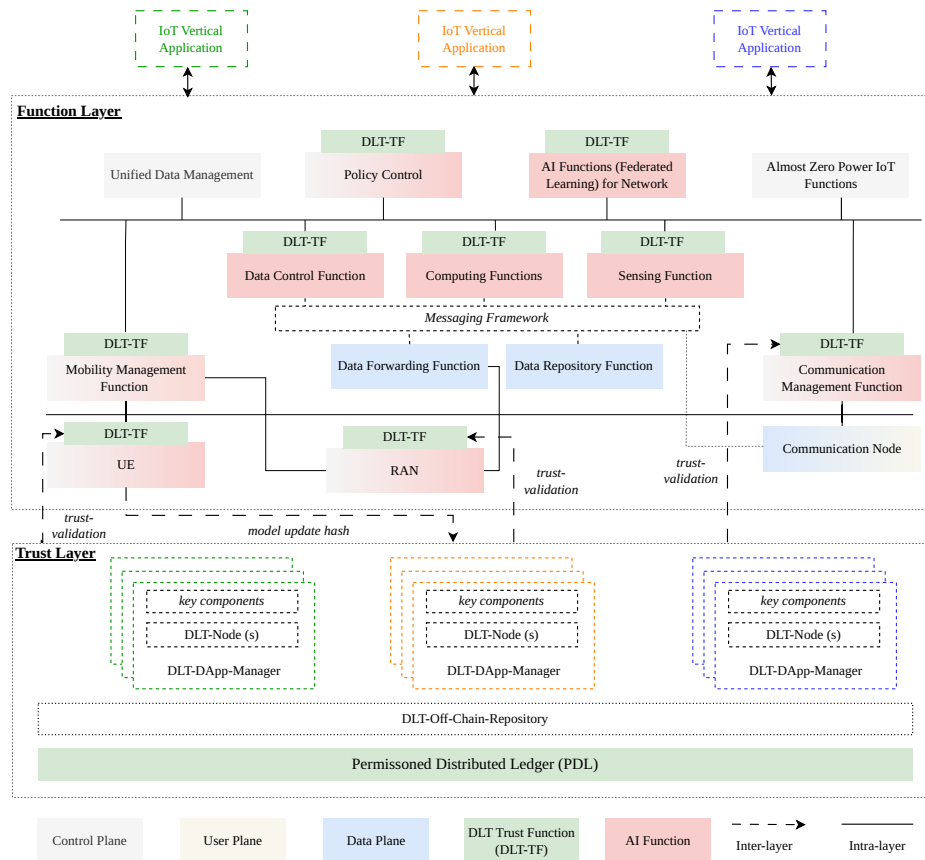


Fig. 3. 6G network architecture that integrates decentralized-ledger technology (DLT) and artificial intelligence (AI). Control-plane, user-plane, data-plane, and DLT-trust functions are color-coded. A solid line shows the inter layer data flow (i) path of a model-update hash from the UE, through the RAN, down to the Trust Layer for anchoring and reputation scoring. (ii) Dashed arrows show the resulting trust-validation messages returned to the UE, the RAN, and the data repository. The diagram highlights how decentralized trust and orchestration components support IoT vertical applications across all planes.

with the 6G vision of trustworthiness [14] as well as converged communications, sensing, and computing [15], several gaps in scalability, interoperability, and standardization remain.

Despite the promise of DLT-integrated FL for future 6G systems, several challenges remain before such architectures can be deployed at scale in real-world IoT environments. One critical limitation is the difficulty of scaling blockchain-based auditing mechanisms in highly resource-constrained edge devices. As quantified in Section V-A, sub-megabyte AZP-IoT nodes typically have only tens of kilobytes of headroom once trimmed implementations of SHA-256 and MQTT are linked. This makes secure element integration and flash-wear management ongoing concerns in evolving 3GPP Rel-18 IoT profiles. Furthermore, NB-IoT and LTE-M bearers particularly in deep-indoor or non-terrestrial 6G use cases can face multi-hour connectivity gaps. While our architecture supports local buffering of small update hashes (e.g., two 64-byte entries), long-period store-and-forward behavior and the use of alternate transport protocols like the Constrained Application Protocol (CoAP) or Quick UDP Internet Connections (QUIC) require further study.

At the scale of millions of devices, managing firmware and key updates without overwhelming the network will depend on automated lifecycle control. Standards such as

Lightweight Machine-to-Machine (LwM2M) 1.2 and Software Updates for Internet of Things (SUIT) with CBOR Object Signing and Encryption (COSE) offer promising paths for secure, PKI-free update handling. Their alignment with 3GPP NWDAF analytics and the ETSI-PDL Adapter framework further strengthens the case for integrating these mechanisms into the FL-DLT stack. A related barrier is the lack of unified standards for interoperability between DLT and FL systems. Without consistent APIs, identity management schemes, and data models, deployments will remain fragmented and vendor-specific. This issue is especially acute for AZP IoT devices, which require lightweight and energy-aware protocols to participate effectively in collaborative training while preserving battery life.

Addressing these issues requires a focused standardization effort across four pillars. First, lightweight consensus mechanisms (e.g., delegated proof-of-stake) and hierarchical architectures (e.g., parent-child chains) can improve scalability without compromising auditability. Second, unified interoperability protocols are essential for integrating DLT-Trust Functions (DLT-TFs) with 3GPP and O-RAN network functions. Third, energy-aware FL designs should support adaptive participation thresholds and model compression tailored to constrained endpoints. Finally, defining performance bench-

marks such as TPS and latency for FL-integrated DLTs can ensure alignment with standardization and promote practical applicability.

VI. CONCLUSION

This article summarizes the standardization efforts of 3GPP, ETSI, ITU-T, IEEE, and O-RAN that drive Federated Learning (FL) and blockchain convergence on the Internet of Things (IoT). These coordinated initiatives underscore how Distributed Ledger Technologies (DLTs) empower secure, decentralized intelligence in next-generation networks.

We also address the challenges inherent in blockchain-based FL for IoT and propose a framework that eliminates centralized aggregators by leveraging DLT-based components. Our system facilitates verifiable contributions, reputation-driven data integrity, and accommodates IoT resource constraints.

Finally, we explore future directions for embedding native trust within network architectures. Trust must be integral to FL, via decentralized model verification and alignment with evolving standards and regulations. We envision the synergy of blockchain, FL, and IoT fortified by interoperability and built-in trust as pivotal for scalable, secure future 6G and industrial applications.

REFERENCES

- [1] X. Lin, "Artificial intelligence in 3GPP 5G-Advanced: A survey," *arXiv preprint arXiv:2305.05092*, 2023. [Online]. Available: <https://arxiv.org/abs/2305.05092>, Accessed: Mar. 12, 2025.
- [2] M. Chen, Z. Yang, W. Saad, C. Yin, H. V. Poor, and S. Cui, "Federated learning-empowered mobile network management for 5G and beyond networks: From access to core," *IEEE Commun. Surv. Tutor.*, vol. 23, no. 4, pp. 2469–2501, 4th Quart., 2021. [Online]. Available: <https://ieeexplore.ieee.org/document/9449943/>, Accessed: Mar. 12, 2025.
- [3] R. Schmid *et al.*, "Tangle ledger for decentralized learning," in *Proc. IEEE IPDPS Workshops*, 2020.
- [4] M. Cao *et al.*, "Towards on-device federated learning: A DAG-based blockchain approach," *IEEE Trans. Neural Netw. Learn. Syst.*, 2021.
- [5] L. Jiang *et al.*, "Cooperative federated learning and model update verification in blockchain-empowered digital twin edge networks," *IEEE Internet Things J.*, 2022.
- [6] J. An *et al.*, "FREB: Participant selection in federated learning with reputation evaluation and blockchain," *IEEE Syst. J.*, 2024.
- [7] W. Sun *et al.*, "Blockchain-based federated learning: A survey and new perspectives," *Appl. Sci.*, vol. 12, no. XX, 2022.
- [8] S. Ren *et al.*, "A scalable blockchain-enabled federated learning architecture for edge computing," *PLoS One*, 2024.
- [9] C. Mazzocca *et al.*, "Enabling federated learning at the edge through the IOTA Tangle," *Future Gener. Comput. Syst.*, 2024.
- [10] Z. Wang *et al.*, "BSIF: A blockchain-enabled secure incentive framework for mobile crowdsensing," *IEEE Internet Things J.*, vol. 9, no. 1, pp. 486–499, Jan. 2022.
- [11] Y. Zhang *et al.*, "BDSS: Blockchain-based data sharing scheme for secure fine-grained access control of medical data," *IEEE Internet Things J.*, vol. 9, no. 10, pp. 7237–7250, May 2022.
- [12] F. Javed, J. Mangués-Bafalluy, E. Zeydan, and L. Blanco, "Trustworthy reputation for federated learning in O-RAN using blockchain and smart contracts," *IEEE Open Journal of the Communications Society*, 2025.
- [13] D. Yang, X. Zhang, X. Huang, L. Shen, J. Huang, X. Chang, and G. Xing, "Understanding power consumption of NB-IoT in the wild: tool and large-scale measurement," in *Proc. 26th Annu. Int. Conf. Mobile Computing and Networking (MobiCom '20)*, London, United Kingdom, Sep. 2020, Art. no. 55, 13 pp. [Online]. Available: <https://doi.org/10.1145/3372224.3419212>
- [14] NGMN Alliance, "6G Trustworthiness Considerations," Final Deliverable (approved), Version 1.0, Oct. 2023. Approved by NGMN Board, 14 September 2023. [Online]. Available: <https://www.ngmn.org> [Accessed: Mar. 20, 2025].
- [15] VIVO Communications Research Institute, "6G Network Architecture," White Paper, Version 1.0, 2023. [Online]. Available: <https://www.vivo.com> [Accessed: Mar. 20, 2025].