Elmi Qaynaqlar XXVIII Respublika Elmi Konfransı / 15 İyul 2026 / 46-51
Scientific Resources XXVIII Republican Scientific Conference / 15 July 2026 / 46-51
DOI: https://doi.org/10.36719/2663-4619/2/46-51

Architectural Design and Security Challenges of Decentralized Artificial Intelligence Systems

Turan Jafar-zada

Keywords: *decentralized AI, federated learning, blockchain, edge computing, security, privacy*
Açar sözlər: *mərkəzsiz süni intellekt, federativ öyrənmə, Blokçeyn, kənar hesablama, təhlükəsizlik, məxfilik*

Over the past decade, AI has shifted from centralized pipelines to distributed systems, driven by demands for data sovereignty, regulation, and resilience. Three pillars support this change: federated learning, which allows training without moving raw data; blockchain, which records updates immutably; and edge computing, which brings processing closer to where data is generated.

Together, these technologies form the backbone of decentralized AI and even hint at the future of the Internet, where intelligence and governance are shared rather than centralized. Yet their security challenges remain underexplored. Distributing computation across diverse devices creates vulnerabilities that traditional defenses cannot easily handle.

Azerbaijan University of Languages, PhD in Technical Sciences, Baku, Azerbaijan
E-mail: turanjafarzade1@gmail.com

Federated learning is especially useful in sensitive domains like healthcare and finance, where raw data cannot leave local servers. But the same mechanism that protects privacy also opens the door to malicious updates. Blockchain adds accountability by logging every interaction, while edge computing reduces latency by handling tasks locally. This stack is powerful, but its security implications demand deeper study.

A decentralized AI architecture is best understood as a multi-layer system in which federated learning, blockchain, and edge computing each assume a distinct functional role and in which the absence of any single layer diminishes the robustness of the whole. At its foundation, federated learning governs how intelligence is collectively constructed. Classical federated learning retains a central parameter server that aggregates participant updates; decentralized variants eliminate this bottleneck entirely by routing updates through peer-to-peer communication graphs. This structural change confers meaningful fault tolerance, since no single node failure can halt the training process and eliminates a natural target for attackers seeking to corrupt the global model at its source (Hallaji et al., 2024).

Positioned above the learning layer, blockchain fulfills the role of a shared coordination medium. Its append-only ledger records every consequential event in the system—model checkpoint submissions, aggregation outcomes, participant registrations—in a form that is publicly verifiable and resistant to retroactive modification. This property is essential for establishing a trustworthy audit trail in an environment where participants may have competing interests and where no central authority exists to adjudicate disputes. Smart contracts further extend blockchain's utility by encoding aggregation rules and incentive

mechanisms in self-executing code, reducing reliance on off-chain governance (Qammar et al., 2023).

Edge and fog computing nodes occupy the outermost layer, interfacing directly with data sources and serving as the sites of local training and real-time inference. Their physical proximity to sensors, mobile devices, and industrial equipment makes them well suited to latency-sensitive applications, but their resource constraints—limited processor cycles, restricted memory, and finite battery capacity—impose hard limits on the complexity of security measures that can be realistically be deployed. This trade-off between computational frugality and security robustness is one of the central design tensions in decentralized AI (Ahmed et al., 2025).

The three-layer architecture as a whole achieves a degree of resilience and scalability that centralized systems cannot easily match. Eliminating the central server removes the most obvious single point of failure; distributing storage across the blockchain prevents data monopolization; and processing data at the edge curtails the need to transmit sensitive information over wide-area networks. Nevertheless, the integration of these three components introduces a design space of considerable complexity, within which engineers must navigate competing objectives of performance, scalability, and security (Qu et al., 2022).

Security challenges in decentralized AI

Decentralized AI systems are appealing because they are open, distributed, and free of central authority—but these same qualities also make them vulnerable. Each layer carries its own risks, and together they amplify exposure. In federated learning, the main danger is model poisoning. A malicious participant can distort local updates so the global model learns incorrectly or

loses accuracy. Since honest and dishonest updates look alike, detection is hard, especially because the manipulation happens during training. Another risk, inference attacks, uses shared gradients to reconstruct sensitive data, undermining privacy.

In the blockchain layer, Sybil attacks exploit open participation: one actor creates many fake identities to gain unfair influence in consensus. Preventing this without closing the network is difficult. The 51% attack is another threat—if one party controls most of the network's resources, they can rewrite transaction history, erasing the supposed immutability of model updates. Even subtler tactics, like delaying messages or targeting validators, can quietly disrupt consensus.

At the edge computing level, the problem is physical. Devices are built for efficiency, not heavy cryptography, so strong protections are often impractical. If a device is compromised, it can inject false data or leak sensitive information. Because these devices are scattered and often unmonitored, spotting and isolating compromised nodes is a serious challenge.

Mitigating these risks requires action across protocols, hardware, and algorithms. In federated learning, robust aggregation methods like median, trimmed mean, or Krum help reduce the impact of malicious updates, especially when paired with anomaly detectors that flag suspicious gradients before aggregation.

For Sybil attacks, defenses rely on making identity creation costly or tying participation to resources or reputation. Proof-of-work, proof-of-stake, and reputation systems each balance openness and resistance differently, and the right choice depends on context—particularly when edge devices have limited capacity to stake.

Securing edge nodes means accepting hardware limits and using lightweight cryptography suited for IoT, or hardware modules and trusted execution environments where possible. Software-level safeguards, such as anomaly detection models that spot compromised behavior, add another layer of protection.

Looking ahead, the strongest path lies in hybrid architectures that integrate AI, distributed systems, and threat modeling as a single design challenge. Formal verification adapted to machine learning and game-theoretic models of adversaries could provide the foundation for certifying security at scale. Decentralized AI can deliver privacy, resilience, and scalability—but only if its security is treated as a core design priority.

References

1. Abreha, H. G., Hayajneh, M., & Serhani, M. A. (2022). Federated learning in edge computing: A survey. *Sensors*.
2. Ahmed, I., Syed, M. A., Maaruf, M., & Khalid, M. (2025). Distributed computing in multi-agent systems: A survey of decentralized machine learning approaches. *Computing*.
3. Alhawas, S., & Rassam, M. A. (2026). Federated learning in edge computing: Vulnerabilities, attacks and defenses—A survey. *Sensors*.
4. Hallaji, E., Razavi-Far, R., & Saif, M. (2024). Decentralized federated learning: A survey on security and privacy. *IEEE Transactions*.
5. Issa, W., Moustafa, N., Turnbull, B., & Sohrabi, N. (2023). Blockchain-based federated learning for securing IoT: A comprehensive survey. *ACM Computing Surveys*.

6. Nguyen, D. C., Ding, M., Pham, Q. V., et al. (2021). Federated learning meets blockchain in edge computing: Opportunities and challenges. *IEEE Internet of Things Journal, 8*(16), 12806–12825.
7. Qammar, A., Karim, A., Ning, H., & Ding, J. (2023). Securing federated learning with blockchain: A systematic literature review. *Artificial Intelligence Review*.
8. Qu, Y., Uddin, M. P., Gan, C., et al. (2022). Blockchain-enabled federated learning: A survey. *ACM Computing Surveys*.