ISSN: 2769-545X

Research Article

Advances in Machine Learning & Artificial Intelligence

Federated Learning Explained for Business Leaders: Strategic Foundations for Regulated Industries

Fabrizio Benelli[1], Redvin Marku[2] and Franco Maciariello[3,4*]

[1]*Zetta Software Tlc Shpk, Albania*

[2]*Department of Economics and Business, Tirana Business University College, Albania*

[3]*New Generations Sensors, Italy*

[4]*Marketing Area, Santa Maria la Fossa (CE) Italy*

[*]Corresponding Author
Franco Maciariello, New Generations Sensors, Italy.

Submitted: 2025, Nov 21; **Accepted**: 2025, Dec 22 **Published**: 2025, Dec 31

Citation: Benelli, F., Marku, R., Maciariello, F. (2025). Federated Learning Explained for Business Leaders: Strategic Foundations for Regulated Industries. *Adv Mach Lear Art Inte, 6*(4), 01-09.

Abstract

The increasing restrictions surrounding the use and movement of data across organizations, sectors and jurisdictions are fundamentally reshaping how advanced AI can be designed and deployed within regulated environments such as energy infrastructures, healthcare ecosystems and financial systems. For more than a decade, dominant AI architectures have relied on centralized data aggregation to feed large-scale machine learning algorithms, assuming that raw data could be transferred to secure environments and processed in bulk. This assumption is rapidly becoming untenable. The tightening of privacy rules, the reinforcement of national and sectorial security constraints, and the growing ethical awareness around human autonomy and data ownership are together limiting traditional data-driven strategies and demanding new approaches grounded in privacy-preserving AI, distributed AI and digital sovereignty principles. Federated Learning represents one of the most relevant architectural responses to these challenges, offering a scalable way to orchestrate model training across multiple organizations without forcing data to leave their original environment. This article offers business leaders an intuitive and strategically oriented interpretation of Federated Learning, avoiding mathematical explanations and instead focusing on managerial applicability, organizational implications and regulatory alignment. The goal is not to provide a technical deep dive into distributed algorithms, but rather to clarify how Federated Learning enables cross-organizational cooperation, respects privacy-by-design requirements, and supports new forms of Human-AI collaboration in sectors where data cannot be moved freely for legal, ethical or strategic reasons. The argument builds upon consolidated literature, industry whitepapers from leading players such as Google and Apple and relevant documents emerging from European policy initiatives and the evolving AI Act regulatory framework. Readers will find a structured introduction to the architectural logic and business rationale of Federated Learning; an intuitive conceptualization that highlights key differences between centralized and federated approaches; and an initial methodology enabling decision-makers to evaluate whether specific use cases benefit from this paradigm, especially when data sovereignty, confidentiality and regulatory compliance are critical. Although Federated Learning has been discussed extensively at technical level, its organizational, managerial and inter-institutional implications are less explored. This work addresses that gap by describing fundamental business trade-offs and illustrating why Federated Learning may soon become indispensable for large-scale AI innovation in healthcare, energy, finance and public services

Keywords: Federated Learning, Distributed AI, Privacy-Preserving AI, Digital Sovereignty, Human-AI, Explainable AI, Collaborative Intelligence

Acronyms

AI - Artificial Intelligence
DER - Distributed Energy Resources
DSO - Distribution System Operator
EU - European Union
FL - Federated Learning
GDPR - General Data Protection Regulation
H-AI - Human-AI
IoT - Internet of Things
TSO - Transmission System Operator
XAI - Explainable Artificial Intelligence

1. Introduction: Why Data No Longer Circulates Freely

Business leaders have long assumed that data, once generated, could be captured, stored, transferred and aggregated without significant frictions, provided that encryption, access control and security policies were in place. Such assumptions originated during the early digital transformation, when large platforms successfully consolidated data pipelines and built centralized AI models capable of extracting value from massive datasets. However, this traditional view is being challenged by multiple convergent forces that make data movement increasingly difficult and strategically undesirable [1].

Legislation is the first and perhaps most visible driver. European regulation has introduced strict requirements regarding personal data, sensitive healthcare information, critical infrastructure monitoring, and industrial knowledge considered as strategic national assets. New data governance frameworks are being discussed that extend sovereignty principles even to industrial telemetry and operational data coming from energy networks or financial systems. In parallel, the AI Act is defining restrictions on opacity and automated decision-making, making transparency and explainability structural requirements rather than optional technical features. Businesses cannot simply aggregate everything into centralized repositories, because doing so exposes organizations to substantial compliance, liability and reputational risks [2-4].

A second driver is geopolitical. Countries and regions increasingly perceive data as a strategic resource connected to security, autonomy and technological competitiveness. Decisions concerning whether data may leave national borders are no longer purely technical or legal, but geopolitical and industrial. National cybersecurity strategies often discourage the externalization of specific categories of data and encourage domestic control of digital and AI infrastructures. Under these conditions, cross-border data sharing becomes not only complex, but in some cases structurally impossible.

A third driver is ethical and societal [5]. The public expects more transparency, user control and fairness concerning the use of personal and behavioural information. The perception that organizations can collect and process information without limitations has generated distrust, amplified by the opacity of some AI architectures and the rise of black-box models that produce decisions difficult to scrutinize. Public pressure, media attention and ethical debates reinforce the need for privacy-preserving AI and for architectures that respect autonomy and human dignity.

At the same time, new collaborative business scenarios demand coordinated AI capabilities. Healthcare networks must monitor chronic diseases across distributed care providers; utilities must exchange operational insights to manage distributed energy resources; financial institutions must detect fraud across multiple actors without exposing sensitive customer information. In these contexts, the traditional requirement of aggregating data into centralized repositories becomes the main barrier preventing multi-party collaboration [6-12].

Federated Learning emerges as a response. Instead of moving the data to the algorithm, Federated Learning moves the algorithm to the data. Each participating organization trains the model locally, contributing only model updates rather than raw information. This apparently subtle inversion has profound architectural and governance implications: privacy-preserving AI becomes a structural characteristic, not a layer added afterwards, and collaboration becomes feasible even in highly regulated environments.

2. An Intuitive Explanation of Federated Learning

Federated Learning is frequently introduced in technical communities through mathematical formulations describing gradient sharing, distributed optimization and aggregation mechanisms. While these notions are essential for specialists, they do not necessarily clarify the strategic nature of the paradigm from a business perspective. The essential point is that Federated Learning decentralizes model training. Instead of collecting datasets from multiple sources into a single environment, Federated Learning orchestrates training processes across distributed nodes, each of which holds its own data locally and never exposes sensitive information to external parties [13,14].

Conceptually, one can imagine a central coordinating instance responsible for initializing the model and orchestrating training

rounds. Participating nodes—such as hospitals, energy operators or financial institutions—receive a model version, train it locally using their private data, and send back only the learned updates. The central coordinator aggregates those updates, producing a new, more accurate version of the global model. No raw data moves: only learning insights travel [15,16].

From a business standpoint, this means eliminating the need for sensitive information to cross institutional boundaries. A healthcare provider does not need to transfer clinical data to external servers; a utility operator does not need to share industrial telemetry outside its controlled perimeter; a bank does not need to expose customer behaviour to other actors. Yet, these organizations still contribute to a shared AI capability capable of capturing global patterns, improving predictive accuracy and enabling system-level intelligence. Federated Learning therefore introduces a new equilibrium between collaboration and confidentiality [17,18]. On one side, collaboration increases because multiple organizations contribute information that enriches the model. On the other side, confidentiality is preserved because data remains localized, under the governance of each data owner. Privacy-preserving AI, distributed AI and digital sovereignty principles become embedded in the architecture, rather than delegated to external processes or contractual agreements. Another essential component is Human-AI collaboration.

Federated Learning does not remove human oversight; on the contrary, it requires governance, coordination and alignment mechanisms that must be consistent with business strategy and regulatory obligations. Decisions about which models to federate, which partners to include, what data categories can be used, and under what conditions updates may be shared, are managerial decisions rather than purely technical ones. Such considerations point to the importance of organizational readiness, not only technological capability.

3. Business Methodology: Evaluating Federated Learning Suitability

Not all AI applications are appropriate for Federated Learning. In some scenarios, traditional centralized methods may remain more economical and operationally efficient. The question for business leaders is therefore not whether Federated Learning is revolutionary—which it is in many respects—but whether a specific situation demands such a paradigm. This section proposes a managerial methodology useful for evaluating applicability. The methodology is not a mathematical decision tree but a conceptual framework grounded on governance, regulatory constraints and risk considerations. The first dimension concerns data sensitivity. Federated Learning is most beneficial when data includes personal, clinical, financial or operational information considered critical under regulatory or cybersecurity perspectives. If data can be easily anonymized without losing its analytical value, or if it does not contain sensitive attributes, centralized architectures may remain sufficient. By contrast, if data is inherently sensitive or subject to stringent confidentiality rules, Federated Learning provides an immediate advantage by eliminating the need for data transfer.

The second dimension concerns regulatory boundaries. In sectors such as healthcare and finance, regulations define not only how data is processed but also where it may physically reside. Federated Learning aligns with the principle of data localization and makes compliance less burdensome by ensuring that information remains within authorized environments. Business leaders should therefore assess whether the regulatory landscape imposes constraints that make centralized aggregation problematic or legally risky.

The third dimension concerns collaboration requirements. Federated Learning is designed for scenarios where multiple organizations must contribute to a shared intelligence without merging their data. If a use case involves a single organization controlling all data sources, traditional centralized processing may be simpler. However, if value creation depends on multi-institutional insights—such as detecting pandemic trends across hospitals, optimizing energy distribution across operators, or identifying cross-border financial fraud—Federated Learning becomes essential.

The fourth-dimension concerns model performance versus privacy trade-offs. Federated Learning introduces technical complexities that may affect convergence speed, communication overhead and computational efficiency. Organizations must therefore evaluate whether the privacy benefits justify potential performance trade-offs. In regulated industries where compliance violations carry severe penalties, the answer is typically affirmative. In less regulated environments, the calculus may vary [19].

Table 1 provides a structured comparison between centralized and federated approaches across key dimensions relevant to business decision-makers. The comparison highlights fundamental differences in architecture, governance, risk profile and organizational implications. While centralized approaches offer simplicity and established operational patterns, Federated Learning introduces strategic advantages in contexts where privacy, sovereignty and regulatory compliance are non-negotiable requirements.

Dimension	Centralized Approach	Federated Learning
Data Movement	All data aggregated in central repository	Data remains distributed; only model updates travel
Privacy Risk	High exposure: single point of vulnerability	Reduced: data never leaves local environment
Regulatory Compliance	Complex: requires extensive data transfer agreements	Simplified: aligns with data localization principles
Governance	Centralized control; single entity responsible	Distributed governance; shared responsibility
Cost Structure	High infrastructure investment; lower coordination cost	Lower infrastructure needs; higher coordination overhead
Time to Value	Faster if data transfer is permitted	Slower initially; faster long-term in regulated contexts
Digital Sovereignty	Limited: data ownership transferred to central entity	Preserved: data ownership remains with originating party

Table 1: Comparative Analysis of Centralized vs. Federated Learning Approaches

The comparative analysis demonstrates that Federated Learning introduces fundamental architectural shifts with significant strategic implications. While centralized approaches remain operationally simpler in unregulated environments, Federated Learning offers decisive advantages when privacy preservation, regulatory compliance and digital sovereignty are essential. The trade-off between operational complexity and strategic benefit becomes favourable in sectors where data cannot legally or ethically be centralized, making Federated Learning not merely an option but increasingly a necessity.

4. Cross-Sector Use Cases and Strategic Applications

Federated Learning applicability extends across multiple regulated sectors where data sensitivity, compliance requirements and collaborative intelligence converge. This section examines representative use cases in healthcare, energy and finance, illustrating how Federated Learning enables innovation while preserving privacy, sovereignty and regulatory alignment. These examples are not exhaustive but demonstrate the breadth of strategic opportunities enabled by distributed AI architectures.

Table 2 presents concrete use cases organized by sector, highlighting specific applications, expected benefits and implementation considerations. The selection emphasizes scenarios where Federated Learning delivers measurable value beyond what centralized approaches could achieve given existing regulatory and ethical constraints. Each use case reflects real-world demands articulated by industry practitioners, regulatory authorities and research communities focused on responsible AI deployment.

Sector	Use Case	Expected Benefit	Implementation Note
Healthcare	Multi-hospital chronic disease monitoring	Improved diagnostic accuracy without patient data transfer	Requires GDPR-compliant governance and ethical review boards
Healthcare	Pharmaceutical research across clinical trial sites	Accelerated drug development with privacy preservation	Coordination with regulatory agencies and trial sponsors essential
Energy	Distributed energy resource optimization	Grid stability and renewable integration without exposing operational data	Coordination between TSOs, DSOs and independent prosumers required
Energy	Predictive maintenance across utility networks	Reduced downtime and cost through multi-operator learning	Requires standardized telemetry formats and secure communication channels
Finance	Cross-institution fraud detection	Enhanced detection patterns without exposing customer identities	Regulatory approval and inter-bank agreements necessary
Finance	Credit risk assessment across regional banks	Improved underwriting accuracy through shared learning	Compliance with banking secrecy and anti-money laundering regulations

Table 2: Representative Federated Learning Use Cases Across Regulated Sectors

The use cases presented demonstrate that Federated Learning applicability transcends individual sectors, addressing a common pattern: collaborative intelligence requirements constrained by regulatory and ethical boundaries. In each scenario, traditional centralized approaches would require extensive legal frameworks, risk management protocols and potentially compromise either innovation velocity or privacy guarantees. Federated Learning resolves this tension by embedding privacy preservation into the architectural design, enabling organizations to collaborate on AI development while maintaining data sovereignty and regulatory compliance.

5. Federated Learning Architecture: A Conceptual Overview

Understanding Federated Learning from an architectural perspective clarifies how the paradigm enables distributed collaboration while preserving data locality. This section presents a high-level conceptual diagram that illustrates the fundamental components and information flows characteristic of a multi-organization federated system. The diagram avoids technical implementation details, focusing instead on the logical organization and governance elements essential for business decision-makers.

Figure 1 depicts a generic Federated Learning architecture where multiple participating organizations maintain local data repositories, train models independently, and contribute learning updates to a central coordination layer responsible for aggregation and model distribution. The absence of raw data transfer is the defining characteristic that distinguishes this architecture from traditional centralized AI systems [20].

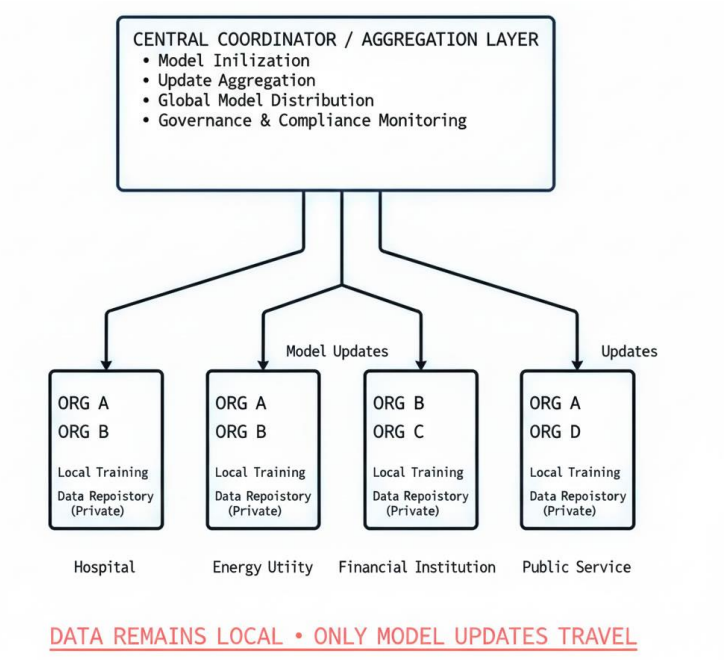

Figure 1: Conceptual Architecture of Multi-Organization Federated Learning System

The architectural diagram illustrates several critical design principles. First, data sovereignty is maintained through local storage and processing, ensuring that sensitive information never crosses organizational boundaries. Second, the central coordinator serves purely as an orchestration and aggregation layer, receiving only encrypted model parameters rather than raw datasets [21]. Third, governance and compliance monitoring occur at both local and central levels, enabling regulatory authorities to verify adherence to privacy and security requirements without compromising operational efficiency.

From a business perspective, this architecture enables organizations to participate in collaborative AI initiatives while retaining control over their most valuable asset: proprietary data. The model update mechanism ensures that learning insights are shared without exposing competitive intelligence, strategic information or sensitive customer details. This balance between collaboration and confidentiality represents the fundamental value proposition that makes Federated Learning strategically relevant for regulated industries.

6. Managerial and Societal Implications

Federated Learning introduces profound implications that extend beyond technical architecture, reshaping organizational strategies, governance models and societal relationships with AI systems. This section examines both managerial considerations relevant to business leaders and broader societal impacts that influence regulatory frameworks, public trust and ethical AI deployment. From a managerial perspective, Federated Learning demands fundamental reconsideration of how organizations create and capture value through data-driven innovation. Traditional AI strategies rely on competitive advantage derived from exclusive data ownership and centralized analytical capabilities. Federated Learning disrupts this logic by enabling collaborative intelligence where multiple organizations contribute to shared models without surrendering proprietary information. This shift requires new mental models regarding partnership structures, competitive positioning and strategic alliances.

Business leaders must also reconsider resource allocation and capability development. Implementing Federated Learning requires investment in local computational infrastructure, secure communication protocols and governance frameworks capable of ensuring compliance across distributed environments. While these investments may increase initial deployment costs, they reduce long-term exposure to regulatory penalties, data breach liabilities and reputational risks associated with centralized data aggregation. Moreover, Federated Learning enables participation in collaborative ecosystems that would otherwise be legally or ethically inaccessible, expanding innovation opportunities. Organizational readiness represents another critical managerial consideration. Successful Federated Learning deployment requires not only technical capabilities but also governance structures, inter-organizational coordination mechanisms and Human-AI collaboration frameworks. Decision-making authority must be clearly defined across participating entities, with explicit protocols governing model selection, update validation, quality assurance and regulatory compliance. These governance requirements transform Federated Learning from a purely technical initiative into a comprehensive organizational transformation [22,23].

At societal level, Federated Learning addresses growing concerns regarding privacy, autonomy and digital sovereignty. Public discourse increasingly questions whether AI innovation necessarily requires sacrificing individual privacy and organizational autonomy. Federated Learning demonstrates that collaborative intelligence and privacy preservation are not mutually exclusive, offering an architectural response that aligns innovation with ethical principles and regulatory requirements. This alignment strengthens public trust in AI systems, particularly in sensitive domains such as healthcare and financial services.

Federated Learning also influences the distribution of AI capabilities across society. Centralized AI architectures tend to concentrate power and capability within large technology platforms possessing resources necessary for massive data aggregation. Federated Learning democratizes AI by enabling smaller organizations, public institutions and regional healthcare networks to participate in advanced analytics without building centralized infrastructures [24-26]. This redistribution of capability supports more equitable innovation ecosystems and reduces dependency on dominant platforms.

Regulatory frameworks are adapting to recognize and encourage privacy-preserving AI architectures. The European AI Act explicitly emphasizes transparency, human oversight and privacy protection, creating structural incentives for adopting Federated Learning in regulated sectors. As regulatory pressure increases, organizations that proactively implement distributed AI architectures position themselves favourably for compliance, potentially gaining competitive advantages through earlier adaptation to evolving legal requirements. Consequently, Federated Learning transitions from optional technical consideration to strategic necessity for organizations operating under strict regulatory oversight [27].

7. Strategic Roadmap for Federated Learning Adoption

Implementing Federated Learning requires systematic planning, organizational alignment and phased execution. This section presents a strategic roadmap structured in five phases, designed to guide business leaders through the transition from centralized AI architectures to distributed, privacy-preserving systems. Each phase addresses distinct organizational, technical and governance challenges, enabling progressive capability development while maintaining operational continuity.

Phase 1: Strategic Assessment and Use Case Identification. Organizations must begin by evaluating their current AI landscape, identifying use cases where data sensitivity, regulatory constraints or collaborative requirements make Federated Learning strategically beneficial. This phase involves stakeholder consultation, regulatory analysis and feasibility assessment. Key activities include mapping data flows, identifying potential collaboration partners, evaluating regulatory implications and establishing initial business cases. Success criteria include documented use case prioritization, stakeholder alignment and executive sponsorship.

Phase 2: Governance Framework Development. Before technical implementation, organizations must establish governance structures defining roles, responsibilities and decision-making authority across participating entities. This phase addresses inter-organizational coordination, compliance monitoring and Human-AI oversight mechanisms. Key activities include drafting data-sharing agreements, defining model governance protocols, establishing quality assurance procedures and creating compliance reporting frameworks. Success criteria include approved governance documentation, legal review completion and partner agreement finalization.

Phase 3: Infrastructure and Capability Building. Organizations must invest in technical infrastructure supporting local model training, secure communication and coordination with external parties. This phase involves technology selection, infrastructure deployment and capability development. Key activities include deploying secure compute environments, implementing communication protocols, training technical teams and establishing monitoring systems. Success criteria include operational infrastructure, trained personnel and validated security controls.

Phase 4: Pilot Implementation and Validation. Organizations should begin with controlled pilot projects demonstrating feasibility, validating governance frameworks and identifying operational challenges. This phase enables learning and refinement before full-scale deployment. Key activities include pilot use case execution, performance measurement, governance validation and stakeholder feedback collection. Success criteria include demonstrated technical feasibility, validated governance effectiveness and documented lessons learned.

Phase 5: Scaling and Continuous Improvement. Following successful pilot validation, organizations can scale Federated Learning across additional use cases and expand collaboration networks. This phase emphasizes operational excellence, continuous governance refinement and ecosystem development. Key activities include expanding use case coverage, onboarding additional partners, optimizing performance and evolving governance frameworks. Success criteria include operational maturity, expanded ecosystem participation and demonstrated business value delivery.

This phased approach enables organizations to manage complexity, build capabilities progressively and demonstrate value incrementally. Executive leadership must maintain sustained commitment throughout the journey, recognizing that Federated Learning represents fundamental transformation rather than isolated technical project. Organizations that successfully navigate this roadmap position themselves as leaders in privacy-preserving AI, regulatory compliance and collaborative innovation.

8. Executive Takeaways / Consulting Pill

Federated Learning represents a strategic inflection point for organizations operating in regulated industries where data sensitivity, compliance requirements and collaborative intelligence converge. The following executive takeaways synthesize essential insights for business leaders considering Federated Learning adoption:

• Federated Learning enables collaborative AI development without compromising data confidentiality, regulatory compliance or digital sovereignty, making it essential for healthcare, energy and financial sectors where traditional centralized approaches face insurmountable legal and ethical barriers.

• Successful implementation requires investment in local computational infrastructure, secure communication protocols and governance frameworks that coordinate distributed decision-making while maintaining accountability and regulatory compliance.

• Organizations must evaluate Federated Learning suitability based on data sensitivity, regulatory boundaries, collaboration requirements and performance trade-offs, recognizing that not all AI applications benefit equally from distributed architectures.

• Inter-organizational agreements, clearly defined governance structures and Human-AI oversight mechanisms represent critical success factors that transform Federated Learning from technical capability into strategic business advantage.

• Regulatory frameworks including the EU AI Act and evolving data-sharing guidelines create structural incentives favouring privacy-preserving AI architectures, making early Federated Learning adoption strategically advantageous for compliance positioning.

• The competitive differentiator will not be technical implementation alone, but rather organizational capability to orchestrate multi-party collaboration, shape governance standards and maintain Human-AI alignment across complex regulatory environments.

• Organizations that proactively develop Federated Learning capabilities position themselves as innovation leaders in privacy-preserving AI, regulatory compliance and collaborative ecosystems, securing competitive advantages as regulatory pressure intensifies and societal expectations evolve toward responsible AI deployment.

9. Conclusions and Future Directions

Federated Learning is rapidly emerging as the reference paradigm for AI deployment across regulated environments where privacy preservation, digital sovereignty and regulatory compliance are non-negotiable requirements. By shifting architectural focus from centralized data aggregation to distributed collaboration, Federated Learning enables industries such as healthcare, energy and finance to benefit from shared intelligence without exposing sensitive information or violating ethical principles. This transformation requires new governance structures, Human-AI oversight mechanisms and inter-institutional agreements reflecting the complexity of distributed ecosystems.

The future trajectory of Federated Learning will be shaped by regulatory frameworks, collaborative innovation models and industry standards developed through real-world implementation. As legislative initiatives such as the AI Act continue emphasizing transparency, privacy and fairness, Federated Learning will become an essential pillar of compliant AI deployment. Moreover, as business leaders across regulated sectors increasingly adopt distributed AI strategies, Federated Learning is likely to evolve into foundational infrastructure for cross-sector innovation, public-private collaboration and digital sovereignty.

Over time, the combination of regulatory pressure, technological maturity and societal expectations will consolidate Federated Learning as more than an architectural option. It will function as a precondition for responsible and scalable AI, enabling organizations to reconcile innovation with compliance and values with profitability. While considerable managerial, technical and ethical challenges remain, the direction is clear: Federated Learning is not only technically feasible but strategically necessary for AI initiatives operating under complex regulatory and societal constraints.

Business leaders must recognize that Federated Learning represents fundamental transformation rather than incremental improvement. Organizations that successfully navigate this transition will position themselves as innovation leaders in privacy-preserving AI, regulatory compliance and collaborative ecosystems. The competitive advantages will flow not only from technical implementation but from organizational capabilities enabling multi-party coordination, governance excellence and sustained Human-AI collaboration aligned with evolving regulatory requirements and societal values.

References

1. Singla, A., Sukharevsky, A., Yee, L., Chui, M., & Hall, B. (2025). The state of AI. *How Organizations are Rewiring to Capture Value. Publisher: McKinsey*.
2. Union, E. (2021). Proposal for a regulation of the european parliament and of the council laying down harmonised rules on artificial intelligence (artificial intelligence act) and amending certain union legislative acts. *COM/2021/206final*, 1-107
3. OECD (Organisation for Economic Co-operation and Development). (2019). Artificial Intelligence in Society.
4. Benelli, F., Maciariello, F., & Salvadori, C. (2024). The influence of technologies on organizational culture in innovative SMEs.
5. European Data Protection Board. (2021). Guidelines on virtual voice assistants. EDPB Document Reference: Version 2.0.
6. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., ... & Cardoso, M. J. (2020). The future of digital health with federated learning. *NPJ digital medicine, 3*(1), 119.
7. Xu, J., Glicksberg, B. S., Su, C., Walker, P., Bian, J., & Wang, F. (2021). Federated learning for healthcare informatics.

Journal of healthcare informatics research, 5(1), 1-19.
8. ENTSO-E. (2020). Digital roadmap for European transmission system operators. European Network of Transmission System Operators for Electricity.
9. Benelli, F., Maciariello, F., Marku, R., & Stile, V. (2025). Towards an Energy Physical Internet: Open Business Models and Platforms for Electricity Distribution Enabled by IoT, Blockchain, and Conditional Payments. In *Conference Book of Abstracts of the 4th International Conference Creativity and Innovation in Digital Economy (CIDE 2025).* Universitatea Petrol-Gaze (UPG).
10. Maciariello, F., Benelli, F., Sangiuolo, G., Lorenzi, E., Caponio, C., & Salvadori, C. (2025, September). TrackOne: Smart Logistics for a Sustainable and Interoperable Agricultural Supply Chain in the Era of Digitization. In *2025 International Conference on Software, Telecommunications and Computer Networks (SoftCOM)* (pp. 1-7). IEEE.
11. World Economic Forum. (2020). Data free flow with trust: Paths towards free and trusted data flows. WEF White Paper.
12. Benelli, F., Këlliçi, E., Maciariello, F., & Stile, V. (2025). Artificial Intelligence for Decentralized Orchestration in the Physical Internet: Opportunities, Business Trade-offs, and Risks in Road Freight Logistics. In *Conference Book of Abstract of the 4th International Conference Creativity and Innovation in Digital Economy (CIDE 2025).*
13. McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017, April). Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics* (pp. 1273-1282). PMLR.
14. Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST),* 10(2), 1-19.
15. Google AI Blog. (2017). Federated learning: Collaborative machine learning without centralized training data. Retrieved from https://ai.googleblog.com/2017/04/federated-learning-collaborative.html
16. Hard, A., Rao, K., Mathews, R., Ramaswamy, S., Beaufays, F., Augenstein, S., ... & Ramage, D. (2018). Federated learning for mobile keyboard prediction. *arXiv preprint arXiv:1811.03604.*
17. Apple, D. (2017). Learning with privacy at scale. *Apple Machine Learning Journal, 1*(8), 71
18. Benelli, F., Caronna, M., Këlliçi, E., & Maciariello, F. (2025). Leveraging the urban physical internet for sustainable heritage management: Edge AI, federated learning, and digital twins. In *HERITAGE CAPITALISATION AND DEVELOPMENT-IDENTITY, INNOVATION, DIGITALISATION, ENVIRONMENT, AWARENESS AND SECURITY" HERITAGE–IIDEAS.*
19. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE signal processing magazine, 37*(3), 50-60.
20. Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., ... & Roselander, J. (2019). Towards federated learning at scale: System design. *Proceedings of machine learning and systems, 1*, 374-388.
21. IEEE Standards Association. (2020). IEEE P3652.1 - Guide for architectural framework and application of federated machine learning. IEEE Standards.
22. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... & Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and trends® in machine learning, 14*(1–2), 1-210.
23. Benelli, F., Maciariello, F., Salvadori, C., Kelliçi, E., & Stile, V. (2025). Human-AI Collaboration in SMEs: A Role-Sensitive Framework for Cognitive Enterprise Hubs. In *Proceedings of the 22nd Conference of the Italian Chapter of the Association for Information Systems (ITAIS 2025).*
24. Deloitte. (2020). Federated learning: A new era of collaborative artificial intelligence. Deloitte Insights.
25. Nedic, B. (2019). Gartner's top strategic technology trends. *Proceedings on Engineering Sciences.*
26. Benelli, F., Këlliçi, E., Maciariello, F., Salvadori, C., & Stile, V. (2025). Enhance Student Well being and Digital Literacy with Machine Learning and Spatial Analysis. In *The 2nd Workshop on Education for Artificial Intelligence (EDU4AI 2025).*
27. Caroline, B., Christian, B., Stephan, B., Luis, B., Giuseppe, D., Damiani, E., ... & Nguyen, D. C. (2020). Artificial intelligence cybersecurity challenges. *Threat Landscape for Artificial Intelligence*.