

A Primer on Federated Learning for Enterprise

Niklas Bitzer

Co-Founder and Chief Scientist, Datapool

February 2, 2026

Niklas Fischer

Co-Founder and CEO, Datapool

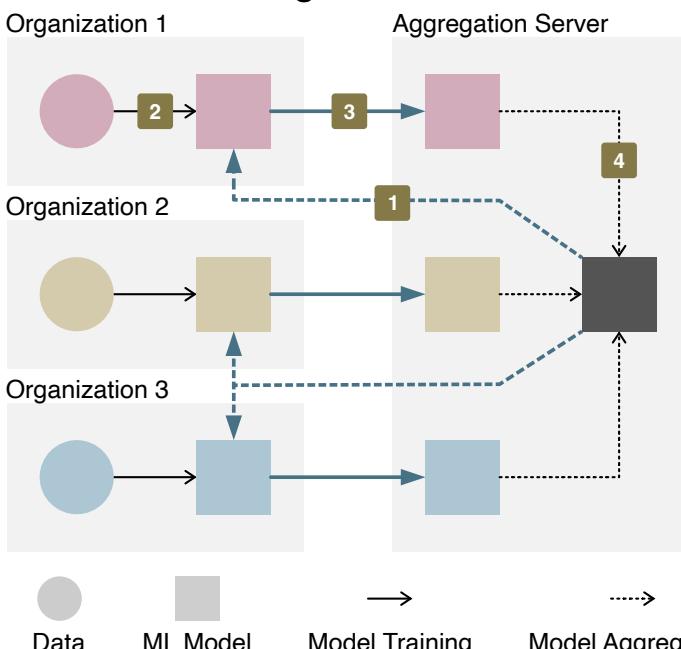
Traditional machine learning requires data centralization, creating compliance bottlenecks and security risks. Federated Learning offers a paradigm shift by moving the computation to the data rather than the data to the computation. Consequently, by aggregating learning across a consortium, businesses can achieve the data scale necessary to train enterprise-grade AI while maintaining data sovereignty.

Introduction

Artificial Intelligence (AI) and its subset, Machine Learning (ML), have evolved into essential tools for modern enterprises to maintain economic competitiveness. However, the substantial volume of diverse, high-quality (structured) training data required to build accurate models presents a critical barrier that often prevents widespread adoption. For medium-sized companies, this creates a fundamental dilemma: **The data needed to train competitive AI models is distributed across multiple organizations, yet sharing that data is neither legally permissible due to strict regulatory frameworks (e.g., GDPR) nor commercially desirable (e.g., risks of exposing trade secrets or lack of trust in external providers).**

Federated Learning (FL) offers a resolution to this conflict by taking a fundamentally different approach compared to centralized solutions. The paradigm has seen a sharp increase in production deployments worldwide over recent years. **It enables (non-competing) organizations to collaborate and train a shared, powerful model without their raw data ever leaving their on-premise servers.** This document introduces FL and provides an overview for organizations considering collaborative AI initiatives. The first part (*Technical*) explains the core mechanisms and training process, alongside privacy and security safeguards. The second part (*Business Case*) examines the strategic implications for SMEs and explores applications across industries.

Federated Learning



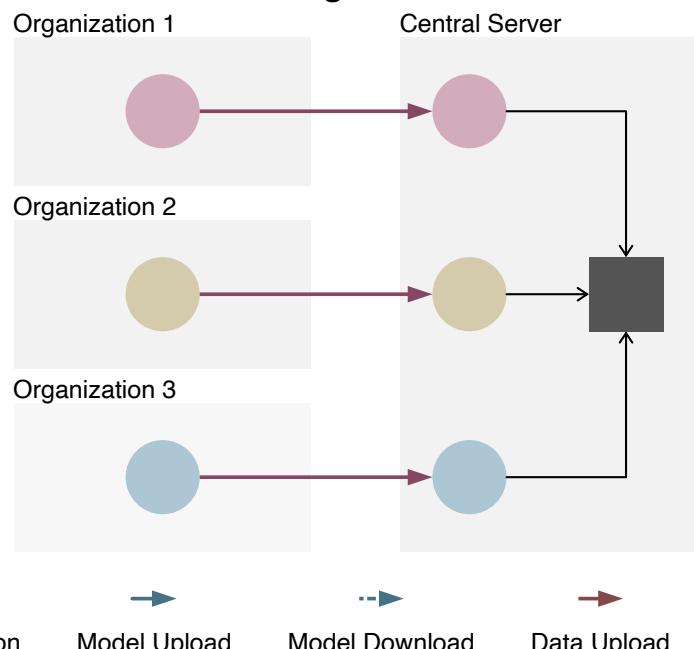
Federated Learning

FL builds on standard ML models to enable collaborative model training across organizations while mitigating data leakage risks. Instead of requiring data to be centralized for training, FL distributes the computation to where the data naturally resides, ensuring information remains within organizational boundaries [1]. While originally developed for mobile and edge devices (cross-device FL) [2], the architecture has evolved to support robust enterprise collaboration (cross-silo FL) [3, 4].

CORE MECHANISM

To understand the value of FL, one must first recognize a fundamental principle of AI: **Model performance scales with data volume.** If two companies in the same sector train separate models based solely on their own respective data, each individual model will be less accurate and robust compared to a single model trained on their combined data. Furthermore, isolated datasets often suffer from local bias; by accessing diverse data sources, the resulting model can achieve increased generalization. In conventional ML, organizations would be forced to pool their data on a central cloud server where a single model is trained on the combined dataset. FL essentially inverts this architecture by **moving the training to the data rather than the data to the training** [2]. Consequently, it enables companies to benefit from each other's data without ever exchanging the data itself.

Centralized Learning



In a business context, participants act as a stable consortium of organizations that share an incentive to train a superior model, yet cannot legally or commercially share their raw data directly [5, 3]. We focus specifically on **Horizontal Federated Learning**, where datasets overlap in features but differ in samples [1]. For Horizontal FL to work effectively, clients must have data with similar structures (same features/labels) but different entities. For example, two regional utility companies may serve completely different user groups (different samples), but their infrastructure and business logic are identical (same features). By aggregating the participants' updates, **FL allows the consortium to leverage a massive dataset to produce a global model that significantly outperforms what any single participant could achieve in isolation** [4, 6].

THE TRAINING & INFERENCE PROCESS

The FL training cycle is an iterative process consisting of rounds of local training and global aggregation, designed to improve the global model's accuracy while maintaining data residency. Understanding this workflow clarifies both the technical feasibility and the privacy guarantees of the approach. The standard workflow, often based on the FedAvg algorithm [2] or robust variants like FedProx for heterogeneous enterprise data [7], is visualized and compared to centralized learning in the graphic on the previous page. It proceeds as follows:

The Federated Learning Cycle

Initialization: Before the training loop starts, the central server initializes a global model with random or pre-trained parameters.

- 1 Distribution:** The current model parameters are distributed to selected authorized clients in the consortium.
- 2 Local Training:** Each client trains the model locally within their secure infrastructure using their private data.
- 3 Reporting:** The updated parameters, representing what was learned locally, are sent back to the central server.
- 4 Aggregation:** The server collects updates from participants and aggregates them (e.g., averaging) into a new, smarter global model.

Convergence: This cycle repeats until the global model reaches satisfactory performance, a predetermined number of rounds is completed or early stopping triggers.

After training concludes, the final model is deployed to all members for inference on their local data or for further company-specific fine-tuning, granting every participant the collective intelligence of the entire network. Furthermore, if new organizations join or new data is collected, the model can be updated continuously starting from the existing baseline. This allows for **continuous improvement and adaptation, resulting in a flexible solution for collective intelligence** [8].

This mechanism provides several fundamental advantages: (1) **Physical Data Isolation** ensures raw data never crosses organizational boundaries [2]; (2) **Reduced Communication Costs**, as model updates are orders of magnitude smaller than raw datasets [9]; (3) **Data Sovereignty** allows organizations to satisfy legal requirements [10]; (4) **Dynamic Scalability** supports continuous improvements and flexible consortium expansion [8]; and (5) **Risk Mitigation** through decentralization, lowering the impact of potential security breaches [11].

For further reading, see the following surveys: [3, 7, 1].

Privacy & Security

While the decentralized nature of FL inherently reduces the attack surface by keeping data on-premise, its viability in enterprise settings fundamentally depends on providing robust privacy and security guarantees—including protection against indirect information leakage. **Security in FL is not a single feature but a multi-layered architecture combining cryptographic protocols and privacy-preserving mechanisms** [3].

TECHNICAL SAFEGUARDS

In a consortium, the primary technical security objective is to ensure that neither the central aggregation server nor other participants can inspect or reconstruct proprietary data. While we have already highlighted the most fundamental safeguard, **Physical Data Isolation** (raw data never leaves each organization's infrastructure), further measures are required to ensure that the shared model updates for aggregation do not inadvertently reveal information about the underlying training data to the aggregation server. Several mechanisms have been developed to guarantee that raw data remains confidential while still enabling effective collaborative learning.

A sophisticated solution is provided by **Secure Aggregation (SecAgg)** protocols. Based on cryptographic techniques, SecAgg prevents the server from analyzing individual updates [12, 8]. It ensures the server can only decrypt the aggregated sum of all updates, but never the individual contributions of a single client. Each organization's update is encrypted before leaving the local infrastructure, remaining mathematically inaccessible even to the coordinating server. For applications requiring even stronger guarantees, FL can incorporate **Differential Privacy (DP)** [13, 14]. By injecting a statistically calculated amount of noise into the model updates, DP provides a formal mathematical guarantee that the presence or absence of any single data point cannot be reliably inferred from the trained model. While DP comes with a model-accuracy trade-off, it effectively neutralizes reconstruction attacks and ensures the final model is safe for deployment.

REGULATORY COMPLIANCE

Beyond technical security, legal compliance is a key aspect for organizations to consider and take into account when choosing a new technology. **FL aligns naturally with European data protection regulations**, making it particularly attractive for operations falling under the GDPR and related legislation [10]. Furthermore, the architecture offers a distinct advantage over centralized cloud AI by adhering to the principles of **Data Minimization, Purpose Limitation, and Privacy by Design**.

Data Localization & Sovereignty are inherently supported because only model parameters are exchanged and personal data never traverses the network or leaves the company's premises. The data controller retains full physical control at all times. This enables collaborative analytics while significantly reducing compliance friction surrounding complex legal constructs like data sharing agreements or cross-border transfers¹. Beyond GDPR, FL aligns with emerging frameworks like the EU Data Governance Act (DGA), which encourages data sharing for innovation while protecting individual and commercial interests, and the EU AI Act, which emphasizes transparency and data governance.

¹While FL provides strong technical and architectural privacy guarantees, organizations must still ensure that their specific implementation complies with applicable regulations. Data processing agreements, participant consent mechanisms (where required), and documentation of the FL process remain necessary components of a compliant deployment.

References

- [1] Qiang Yang et al. “Federated Machine Learning: Concept and Applications”. In: *ACM Trans. Intell. Syst. Technol.* 10.2 (Jan. 2019), 12:1–12:19. issn: 2157-6904. doi: 10.1145/3298981. url: <https://dl.acm.org/doi/10.1145/3298981>.
- [2] Brendan McMahan et al. “Communication-Efficient Learning of Deep Networks from Decentralized Data”. en. In: *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*. PMLR, Apr. 2017, pp. 1273–1282. url: <https://proceedings.mlr.press/v54/mcmahan17a.html>.
- [3] Peter Kairouz et al. “Advances and Open Problems in Federated Learning”. In: *Found. Trends Mach. Learn.* 14.1-2 (June 2021), pp. 1–210. issn: 1935-8237. doi: 10.1561/2200000083. url: <https://doi.org/10.1561/2200000083>.
- [4] Kristina Müller and Freimut Bodendorf. “Cross-silo federated learning in enterprise networks with cooperative and competing actors”. eng. In: *The Human Side of Service Engineering*. Vol. 108. 108. AHFE Open Acces, 2023. isbn: 978-1-958651-84-1. doi: 10.54941/ahfe1003126. url: https://openaccess.cms-conferences.org/publications/book/978-1-958651-84-1/article/978-1-958651-84-1_25.
- [5] Yufeng Zhan et al. “A Survey of Incentive Mechanism Design for Federated Learning”. In: *IEEE Transactions on Emerging Topics in Computing* 10.2 (Apr. 2022), pp. 1035–1044. issn: 2168-6750. doi: 10.1109/TETC.2021.3063517. url: <https://ieeexplore.ieee.org/abstract/document/9369019>.
- [6] Chao Huang et al. “Promoting Collaboration in Cross-Silo Federated Learning: Challenges and Opportunities”. In: *IEEE Communications Magazine* 62.4 (Apr. 2024), pp. 82–88. issn: 1558-1896. doi: 10.1109/MCOM.005.2300467. url: <https://ieeexplore.ieee.org/document/10373828>.
- [7] Tian Li et al. “Federated Learning: Challenges, Methods, and Future Directions”. In: *IEEE Signal Processing Magazine* 37.3 (May 2020), pp. 50–60. issn: 1558-0792. doi: 10.1109/MSP.2020.2975749. url: <https://ieeexplore.ieee.org/document/9084352>.
- [8] Keith Bonawitz et al. “Towards Federated Learning at Scale: System Design”. en. In: *Proceedings of Machine Learning and Systems* 1 (Apr. 2019), pp. 374–388. url: https://proceedings.mlsys.org/paper_files/paper/2019/hash/7b770da633baf74895be22a8807f1a8f-Abstract.html.
- [9] Jakub Konený et al. *Federated Learning: Strategies for Improving Communication Efficiency*. arXiv:1610.05492 [cs]. Oct. 2017. doi: 10.48550/arXiv.1610.05492. url: <http://arxiv.org/abs/1610.05492>.
- [10] Nguyen Truong et al. “Privacy preservation in federated learning: An insightful survey from the GDPR perspective”. In: *Computers & Security* 110 (Nov. 2021), p. 102402. issn: 0167-4048. doi: 10.1016/j.cose.2021.102402. url: <https://www.sciencedirect.com/science/article/pii/S0167404821002261>.
- [11] Viraaji Mothukuri et al. “A survey on security and privacy of federated learning”. In: *Future Generation Computer Systems* 115 (Feb. 2021), pp. 619–640. issn: 0167-739X. doi: 10.1016/j.future.2020.10.007. url: <https://www.sciencedirect.com/science/article/pii/S0167739X20329848>.
- [12] Keith Bonawitz et al. “Practical Secure Aggregation for Privacy-Preserving Machine Learning”. In: *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. CCS ’17. New York, NY, USA: Association for Computing Machinery, Oct. 2017, pp. 1175–1191. isbn: 978-1-4503-4946-8. doi: 10.1145/3133956.3133982. url: <https://dl.acm.org/doi/10.1145/3133956.3133982>.
- [13] Cynthia Dwork. “Differential Privacy”. en. In: *Automata, Languages and Programming*. Ed. by Michele Bugliesi et al. Berlin, Heidelberg: Springer, 2006, pp. 1–12. isbn: 978-3-540-35908-1. doi: 10.1007/11787006_1.
- [14] Martin Abadi et al. “Deep Learning with Differential Privacy”. In: *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. CCS ’16. New York, NY, USA: Association for Computing Machinery, Oct. 2016, pp. 308–318. isbn: 978-1-4503-4139-4. doi: 10.1145/2976749.2978318. url: <https://dl.acm.org/doi/10.1145/2976749.2978318>.