FEDERATED LEARNING | PUBLIC SECTOR | PRIVACY BY DESIGN | AI GOVERNANCE | GDPR | DIFFERENTIAL PRIVACY | DATA SILOS

Federated Learning: The Future of Privacy-Preserving Public-Sector AI

KATIE GRILLAERT · CEO GOOSE FL

Government agencies face a structural policy trap: citizens expect modern public services, from predictive health care to adaptive transit, yet strict privacy laws like the EU's General Data Protection Regulation (GDPR) prevent agencies from pooling the personal data needed to train AI. For years, public-sector AI adoption has stalled between legal compliance and isolated agency databases.

Federated Learning (FL) breaks this deadlock by inverting the standard machine learning workflow. Instead of collecting sensitive citizen records into central servers, FL sends the algorithm to the data, training models locally across hospital workstations, municipal sensors, and personal devices without ever transmitting raw, identifiable files[1][2]. Paired with advances in machine unlearning and 2026 hardware acceleration, FL gives public agencies a practical way to comply with privacy laws, collaborate across departments, and deploy effective public-sector AI, provided engineers and policymakers can navigate a few key trade-offs in bandwidth, security, and accuracy along the way.

A GDPR-compliant architecture for government agencies

The most immediate barrier to public-sector AI is legal rather than technical. While GDPR sets a global standard for data protection, its strict rules on data minimization and purpose limitation have slowed government initiatives that rely on centralized databases. Federated Learning satisfies these legal requirements by design: because personal data never leaves its local source, agencies share only encrypted model updates rather than underlying citizen records[1][2].

Rather than building centralized data lakes that create single points of failure, standards bodies point to decentralized architectures as a safer baseline for public institutions. Guidance from the National Institute of Standards and Technology (NIST), including its Privacy Framework and Privacy-Preserving Federated Learning (PPFL) testbed initiatives, highlights how pairing federated workflows with formal protections like differential privacy allows agencies to extract population-level insights while shielding individual records[3]. Across health care, education, and municipal infrastructure, departments can therefore train shared models without exposing sensitive files[1][2].

© 2026 Equitable Labs, dba Goose FL, A Delaware Public Benefit Corporation

Machine unlearning: operationalizing the “right to be forgotten”

Keeping data localized during training solves the first half of the compliance puzzle, but public-sector rules also govern what happens after a model is built. Under GDPR Article 17, citizens hold the “right to be forgotten,” obligating institutions to erase personal data upon request[4]. In traditional AI, removing a single person’s influence from a trained model typically requires wiping the system and retraining it from scratch, an impractical cost for government budgets[4].

Integrating machine unlearning into federated workflows closes this gap. Research published in the IEEE Transactions on Neural Networks and Learning Systems highlights how federated unlearning develops methods to surgically strip a specific participant’s data contributions from a global model without full retraining[4]. Although guaranteeing verifiable removal without degrading model accuracy remains an active engineering challenge, the approach gives public agencies a concrete mechanism for enforcing citizen data rights post-deployment[2][4].

Hardware acceleration removes the bottleneck

Even when legal and deletion requirements are met, a privacy-preserving architecture cannot scale if everyday devices cannot run it. Historically, the primary obstacle to deploying FL outside of research labs was computational overhead: running cryptographic protections like homomorphic encryption required far more processing power and battery life than consumer or municipal hardware could spare[5][6].

That constraint eased significantly in April 2026, when MIT researchers demonstrated a method that accelerates privacy-preserving AI training on everyday devices by roughly 81%[5]. Alongside hardware-level innovations like SmartNIC-based offloading for encryption workloads, these efficiency gains allow smartphones, laptops, and local municipal sensors to participate in model training without noticeable lag or battery drain[5][6].

Breaking down inter-agency data silos

Once individual devices and local servers can run encrypted training efficiently, FL can tackle one of government’s oldest operational problems: institutional data silos. Public health departments, transit authorities, and school districts routinely collect complementary information, yet privacy statutes and security risks prevent them from pooling records into a single warehouse[1][7].

Federated transfer learning bridges these boundaries by allowing insights learned in one domain to improve models in another without exchanging raw datasets[7][8]. For example, a municipal transit agency and a public health department can jointly model commuter patterns and air-quality health outcomes while keeping citizen records strictly partitioned[1][7]. The UK Government’s Responsible Technology Adoption Unit has specifically examined this capability, noting that privacy-preserving FL offers a practical framework for cross-agency public service delivery despite ongoing governance complexities[2].

© 2026 Equitable Labs, dba Goose FL, A Delaware Public Benefit Corporation

Technical hurdles to scale

While cross-agency collaboration demonstrates FL's operational potential, moving from isolated pilots to nationwide infrastructure requires overcoming two persistent engineering constraints:

- Infrastructure and data heterogeneity. Coordinating millions of devices with varying processing power, battery levels, and network reliability creates synchronization and bandwidth bottlenecks[8][9][10]. Furthermore, local datasets across different regions or demographics are rarely uniform (non-IID), which can slow down model convergence and skew accuracy[8][9][10].
- Security and privacy trade-offs. Stronger mathematical privacy guarantees, such as adding statistical noise via differential privacy, inevitably reduce a model's predictive accuracy[2][3]. Systems must also defend against model poisoning attacks, in which malicious actors inject corrupted updates to degrade or manipulate the shared model[3][11].

Moving from pilots to public infrastructure

As technical advancements continue to reduce the trade-offs, government adoption is underway. Public agencies and financial regulators are actively testing FL for joint fraud detection, risk modeling, and demand forecasting[9], supported by open-source validation environments like NIST's PPFL testbed and ongoing research from groups like Google Research[3][12].

For the public sector, the path forward is no longer blocked by a binary choice between operational capability and citizen privacy. By adopting federated architectures, funding research into verifiable unlearning, and establishing clear NIST-aligned security standards, agencies can deploy AI that improves public services while keeping citizen data firmly under local control.

© 2026 Equitable Labs, dba Goose FL, A Delaware Public Benefit Corporation

Sources

[1] Zhang, Q. (2026). "Federated learning for breaking data silos in smart governance: a privacy-preserving framework for cross-agency collaboration." *Future Technology*, 5(2), 179–188. https://fupubco.com/futech/article/view/770

[2] UK Government Responsible Technology Adoption Unit (RTAU). (2024, February 22). "Privacy-Preserving Federated Learning: Understanding the Costs and Benefits." *GOV.UK Blog*. https://rtau.blog.gov.uk/2024/02/22/privacy-preserving-federated-learning-understanding-the-costs-and-benefits/

[3] Near, J., & Darais, D. (2024, July 15). "Protecting Trained Models in Privacy-Preserving Federated Learning." National Institute of Standards and Technology (NIST) Cybersecurity Insights. https://www.nist.gov/blogs/cybersecurity-insights/protecting-trained-models-privacy-preserving-federated-learning

[4] Romandini, N., Mora, A., Mazzocca, C., Montanari, R., & Bellavista, P. (2024). "Federated Unlearning: A Survey on Methods, Design Guidelines, and Evaluation Metrics." *IEEE Transactions on Neural Networks and Learning Systems*, 36(7), 11697–11717. https://doi.org/10.1109/TNNLS.2024.3478334

[5] MIT News. (2026, April 29). "Enabling privacy-preserving AI training on everyday devices." Massachusetts Institute of Technology. https://news.mit.edu/2026/enabling-privacy-preserving-ai-training-everyday-devices-0429

[6] Choi, S., et al. (2024). "FedNIC: enhancing privacy-preserving federated learning via homomorphic encryption offload on SmartNIC." *Frontiers in Computer Science*, 6, 1465352. https://www.frontiersin.org/journals/computer-science/articles/10.3389/fcomp.2024.1465352/full

[7] Flock.io. (2025, September 1). "Federated Learning: the future of privacy-preserving public sector AI." *Flock.io Blog*. https://www.flock.io/blog/federated-learning-the-future-of-privacy-preserving-public-sector-ai

[8] Preprints.org. (2026, January). "Federated Learning: A Survey of Core Challenges, Current Methods, and Emerging Trends." *Preprints*, Manuscript 202601.0271. https://www.preprints.org/manuscript/202601.0271

[9] Baduwal, P., et al. (2026). "Federated Learning: A Survey of Core Challenges, Current Methods, and Emerging Trends." *MDPI Computers*, 15(3), 155. https://www.mdpi.com/2073-431X/15/3/155

[10] Yurdem, B., Kuzlu, M., Gullu, M. K., Catak, F. O., & Tabassum, M. (2024). "Federated learning: Overview, strategies, applications, tools and future directions." *Heliyon* / PubMed Central (PMC11466570). https://pmc.ncbi.nlm.nih.gov/articles/PMC11466570/

[11] Palo Alto Networks. (n.d.). "What Is Federated Learning? A Guide to Privacy-Preserving AI." *Cyberpedia*. https://www.paloaltonetworks.com/cyberpedia/what-is-federated-learning

[12] Bonawitz, K., Kairouz, P., McMahan, B., & Ramage, D. (2021). "Federated Learning and Privacy: Building privacy-preserving systems for decentralized data." *Google Research / ACM Queue*. https://research.google/pubs/federated-learning-and-privacy/

About. Goose FL is building AI-integrated infrastructure that equips institutions with key insights for innovative and inclusive outcomes. Discover how Goose FL can help your organization at www.goosefl.com.

Share Intelligence · Not Data

Working through public-sector AI under privacy law? Start here.

TALK TO

© 2026 Equitable Labs, dba Goose FL, A Delaware Public Benefit Corporation