WHITEPAPER

Privacy-Preserving Machine Learning in Global Healthcare AI

Breaking the Clinical Validation Bottleneck Without Breaking the Law

Prepared for the Northwestern University MS in Data Science and MS in Information Systems community

Primary audience: students, alumni, faculty, health-tech leaders, privacy professionals, cybersecurity leaders, and technical governance stakeholders

Rainier Anthony M. Milanes, JD | CIPM | CC

www.linkedin.com/in/milanesram

May 2026

Core thesis

Healthcare AI needs a validation architecture that lets models move, evidence accumulate, and patient data remain protected. PPML is that architecture when it is engineered with governance, patient agency, infrastructure equity, and measurable privacy controls.

Contents

How to read this whitepaper

The paper is written for technical leaders who must translate machine-learning design into business credibility, clinical evidence, and defensible privacy architecture. It avoids legal formalism and treats law, infrastructure, and model performance as one operating system.

Executive Summary

Healthcare AI has entered a credibility crisis. Public data, open-source tooling, and agile product methods allow startups to build predictive models faster than hospitals and regulators can validate them. The result is a shortage of clinically credible, regulatorily defensible, and cyber-resilient validation pathways.

The evidence gap is material. Chouffani El Fassi and colleagues reported in Nature Medicine that not all AI health tools with regulatory authorization are clinically validated and warned that devices lacking adequate clinical validation pose risks for patient care.[1] A public summary of that research reported that, among 521 FDA-authorized AI medical device authorizations reviewed, 226 or approximately 43 percent lacked published clinical validation data; only 22 were validated through randomized controlled trials.[2] A separate npj Digital Medicine scoping review of 692 FDA-approved AI/ML-enabled medical devices from 1995 to 2023 found reporting gaps in demographic data, socioeconomic data, age reporting, and performance-study details.[3]

The central conflict is that the models need hospital-grade validation, but the data needed for validation sits inside highly regulated institutional environments. United States hospitals must manage protected health information under HIPAA. European institutions must process health data under the GDPR, where health, genetic, and biometric data are special categories of personal data subject to heightened restrictions.[5] Cross-border transfer rules add another layer of friction.[6]

Privacy-Preserving Machine Learning (PPML) provides a practical path through the deadlock. Federated Learning sends computation to hospital-controlled environments. Differential Privacy mathematically limits what can be inferred about any individual from outputs, statistics, or model updates. Fully Homomorphic Encryption (FHE) enables selected computation over encrypted data or encrypted model assets. Used together, these controls allow startups and research hospitals to generate evidence without making raw clinical data transfer the default operating model.

Executive conclusion

PPML should not be treated as a niche privacy enhancement. It is strategic validation infrastructure for global healthcare AI. It turns privacy and cybersecurity from late-stage blockers into design conditions for clinical evidence, market access, and institutional trust.

Stakeholder Stress Test

The whitepaper integrates four industry stakeholder corrections that materially strengthen the architecture.

Stakeholder Lens	Core Objection	Whitepaper Response
Global healthcare expert	Federated Learning may exclude rural or lower-resourced hospitals.	Use edge-aware FL, containerized nodes, pruning, quantization, and site-readiness support.
Privacy advocate	Institutional computation agreements can erase the patient.	Treat consent, opt-in participation, and permitted-use logic as auditable governance controls.
CISO	FHE is expensive and may be too slow for hospital budgets.	Reserve FHE for high-sensitivity validation and roadmap hardware acceleration.
Data scientist	Non-IID hospital data can distort global models.	Use distance-aware non-IID assessment, client clustering, and localized personalization.

1. The Clinical Validation Crisis

Healthcare AI has a fundamental asymmetry. Model development is global, fast, and increasingly accessible. Clinical validation is local, slow, regulated, and institutionally constrained. A startup may train a model on public imaging datasets, open biomedical repositories, synthetic patient records, or academic challenge datasets. That model may achieve impressive benchmark performance. Yet public-data performance is not proof of clinical readiness.

Clinical validation requires the model to be tested against the messy conditions of real care delivery specificaly local patient populations, device variation, encounter documentation practices, institutional coding differences, language and demographic diversity, workflow interruptions, and distribution shifts that were not present in the public training corpus. Without that validation, a model can be technically impressive and clinically fragile.

The FDA maintains a public list of AI-enabled medical devices authorized for marketing in the United States, but that list is not a substitute for a complete clinical evidence package. FDA public summaries provide releasable information such as safety and effectiveness summaries, but they do not include most information submitted in an application and the list is not comprehensive.[4] Authorization, publication, and clinical validation are therefore related but not interchangeable.

The validation bottleneck becomes severe when startups seek to collaborate with major research hospitals in the United States and Europe. Hospitals have the clinical data needed to validate models, but the data cannot be casually exported. Startups have the models, but not the lawful and secure access path to test them across diverse populations. Regulators and clinicians need evidence, but the evidence pipeline is blocked by legal risk, security risk, institutional caution, and patient trust concerns.

[Exhibit A: The AI Validation Gap] - The chart visualizes the gap between FDA-authorized AI medical device authorizations, published clinical validation, and randomized controlled trial validation.

Exhibit A: The AI Clinical Validation Gap

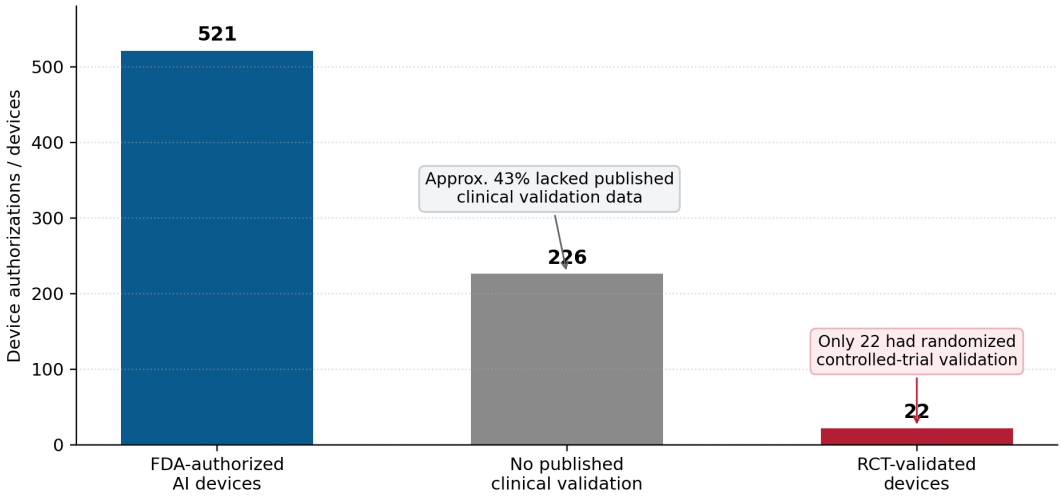

Key point: market authorization and public clinical validation are not equivalent.

Source basis: Chouffani El Fassi et al.; UNC Health Newsroom summary. Visualization prepared for this whitepaper.

The linkage-attack problem

Traditional anonymization fails because clinical data is dense. Even after direct identifiers are removed, combinations of age, diagnosis, medication history, encounter date, ZIP code, provider location, rare condition, imaging modality, and treatment sequence can point back to a person. Linkage attacks exploit this density by combining the target dataset with public records, breached datasets, social media, voter files, genealogy data, insurance information, or other quasi-identifiers.

This problem is worse in AI workflows because the attack surface is not limited to records. Model gradients, embeddings, confidence scores, model updates, validation statistics, and repeated aggregate queries may leak sensitive information. A membership inference attack may reveal that a specific person was part of a cancer cohort, a behavioral health dataset, or a rare-disease validation group.

2. Analysis

Business impact: validation is a market-access requirement

Healthcare AI vendors often speak about accuracy, automation, and scalability. Enterprise buyers ask a harder question, that is where is the clinical evidence? Hospitals, payers, regulators, professional societies, insurers, and boards need proof that the product works across realistic patient populations and clinical settings. Without that evidence, the vendor is asking the market to absorb unpriced risk.

Weak validation produces commercial friction. It slows procurement, triggers extended legal review, complicates cybersecurity due diligence, undermines reimbursement arguments, and increases product liability exposure. In healthcare, trust is part of the product. A company that cannot evidence safety, generalizability, privacy, and operational resilience will struggle to sell into serious clinical environments.

Legal and regulatory risk: HIPAA and GDPR are different constraint systems

HIPAA and GDPR are often casually grouped together as privacy laws. That shorthand is dangerous. HIPAA is sectoral and focuses on covered entities, business associates, and protected health information. HHS explains that HIPAA de-identification may be achieved through Expert Determination or Safe Harbor, and that information is not individually identifiable only when it does not identify an individual and there is no reasonable basis to believe it can be used to identify an individual.[7]

GDPR is broader, rights-based, and extraterritorial. It applies to organizations established in the EEA and can also apply to non-EEA organizations offering goods or services to individuals in the EEA or monitoring their behavior.[8] Article 9 prohibits processing special categories of personal data, including health data, unless a specific exception applies.[5] Article 32 requires security measures appropriate to risk, including pseudonymization, encryption, resilience, restoration capability, and regular testing of safeguards.[9]

Technical vulnerabilities: centralized PHI creates a compound attack surface

A centralized validation repository concentrates several risks at once like regulated clinical data, startup model assets, hospital participation metadata, validation logs, encryption keys, and performance outputs. A breach of that environment would expose the evidence-generation process itself.

ENISA warns that pseudonymization must be evaluated against adversarial models and techniques such as brute force attacks, dictionary search, and guesswork.[10] In clinical AI, pseudonymization

is control, and often an insufficient one, unless paired with computation governance, leakage limitation, encryption, and evidence-grade auditability.

3. Why Traditional Anonymization Fails

Traditional anonymization assumes that risk can be managed primarily by suppressing or generalizing fields before data moves. That assumption is increasingly weak in healthcare AI. Clinical data carries high-dimensional signals. Public data is abundant. Breached data is persistent. Models themselves can encode sensitive information. The risk is therefore is whether a person can be inferred from the combined behavior of data, model, output, and surrounding context.

This is where the legal and technical definitions of privacy collide. Under HIPAA, the question is whether there is a reasonable basis to believe the information can identify an individual. Under GDPR, pseudonymized data can remain personal data if it can be attributed to a natural person using additional information. In both regimes, the relevant question is practical identifiability, not whether a dataset appears clean to a developer.

The clinical validation crisis therefore requires a change in design philosophy. Instead of asking how to move clinical data safely, technical teams should ask whether clinical evidence can be generated without raw data movement. That shift is the foundation of PPML.

Design principle

The mature posture is not anonymize-and-export. It is govern-the-computation, minimize-the-output, prove-the-privacy-loss, and preserve-the-evidence.

4. The PPML Architecture

Privacy-Preserving Machine Learning is a layered control architecture that changes how data science work is performed across institutional boundaries. Federated Learning, Differential Privacy, and Fully Homomorphic Encryption each solve a different part of the validation crisis. Used together, they enable clinical evidence generation without defaulting to raw-data pooling.

The architecture is strongest when each technology is assigned a clear control function. Federated Learning is the collaboration layer. Differential Privacy is the leakage-control layer. FHE is the confidentiality layer. Governance, consent, auditability, and clinical oversight are the operating layer that makes the technical controls defensible.

4.1 Federated Learning: move the model, not the data

Federated Learning reverses the traditional data-sharing assumption. Instead of moving hospital datasets into a centralized repository, the model or validation workload moves to the hospital environment. The institution performs computation locally. Only model updates, aggregate metrics, encrypted outputs, or privacy-protected results return to the coordinator.

The foundational Federated Learning formulation by McMahan and coauthors described the problem clearly that data may be privacy-sensitive, large, or both, which can preclude logging it to a central data center for conventional training. Their alternative leaves training data distributed and learns a shared model by aggregating locally computed updates.[11]

In global healthcare validation, this matters because the hospital remains the custodian of clinical data. A German hospital does not need to export raw health data into a United States startup cloud. A United States academic medical center does not need to pool identifiable patient records with a European partner. Each institution can validate the model under its own governance, access controls, logging, and ethics framework.

[Exhibit B: Centralized Data Sharing vs. Federated Clinical Validation] - The diagram contrasts raw-data centralization with a federated architecture that sends computation to hospital nodes and returns only privacy-protected outputs.

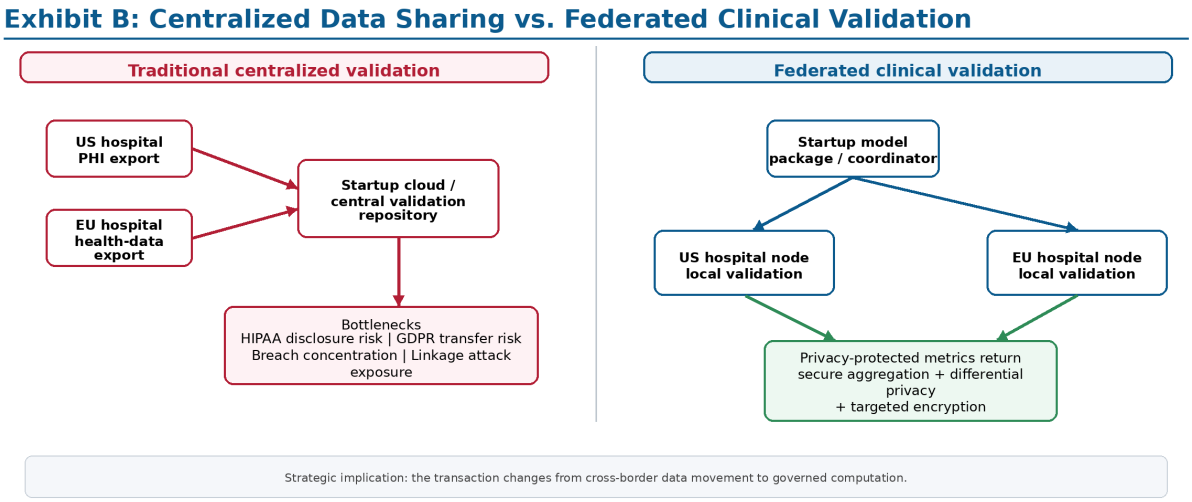

Federated Learning does not eliminate risk. Model updates can leak information if not protected. Malicious or compromised participants can poison updates. Non-independent and non-identically distributed hospital datasets can distort model performance. Governance must define participant eligibility, local compute requirements, model versioning, update integrity checks, secure aggregation, audit rights, and procedures for anomalous validation behavior.

The first operational correction is infrastructure equity. If FL nodes require massive GPU clusters, validation will be restricted to wealthy urban academic medical centers and will reproduce the same demographic blind spots that clinical validation is supposed to correct. The architecture should therefore include edge-aware deployment, lightweight model compression, pruning, quantization, sparse model execution, and containerized node packaging so lower-resourced hospitals can participate without massive capital expenditure.[14], [15]

The second correction is mathematical. Distributed clinical datasets are often non-IID. A pediatric hospital, a geriatric specialty center, an oncology institute, and a rural community hospital do not generate interchangeable data distributions. Basic FedAvg can stall, converge poorly, or produce a global model that underperforms for minority clinical contexts. The consortium should require distance-aware non-IID assessment, embedding-space divergence measurement, client contribution weighting, and clustering for localized personalization.[16]

4.2 Differential Privacy: make leakage measurable

Differential Privacy determines how much information a disclosed output can reveal about any individual in the dataset. NIST SP 800-226 describes differential privacy as a mathematical framework that quantifies privacy loss to entities when their data appears in a dataset and helps practitioners evaluate promises made by differentially private systems.[12]

In clinical validation, Differential Privacy can be applied to aggregate performance reports, subgroup metrics, model updates, query systems, or validation statistics. The objective is to prevent an attacker from reliably determining whether a specific patient was included in a hospital validation cohort. That protection matters because membership in a cohort can itself reveal sensitive information.

The design question is calibration. Differential Privacy is not achieved by adding arbitrary noise. The noise must be calibrated to query sensitivity and the privacy budget selected for the use case. Too little noise creates privacy theater. Too much noise destroys clinical utility. A serious deployment requires a documented privacy budget, query governance, cumulative-loss tracking, utility testing, stakeholder review, and clear explanation to legal, clinical, and technical reviewers.

4.3 Fully Homomorphic Encryption: compute while encrypted

Fully Homomorphic Encryption is the most technically demanding layer, but also the most strategically important for high-trust validation. NIST describes FHE as a main tool of privacy-enhancing cryptography and explains that it allows arbitrary functions to be computed over encrypted data without knowing the secret key.[13]

Hospitals do not want to expose patient data. Startups do not want to expose proprietary model weights, architecture, or feature logic. FHE creates a path where selected computations can be performed over encrypted inputs, with encrypted outputs decrypted only by authorized parties under an agreed protocol.

The CISO objection must be addressed directly. Software-only FHE is not yet a general-purpose answer for public hospital workloads. Research on programmable FHE acceleration describes software FHE as several orders of magnitude slower than cleartext computation, making indiscriminate deployment operationally unrealistic.[18] The strategic response is not to oversell FHE. Hospitals should reserve it for high-value targeted inference scoring, encrypted validation checkpoints, and workflows where both PHI exposure and model-IP exposure are unacceptable.

At the same time, vendors should roadmap heterogeneous compute and specialized hardware acceleration. Emerging FHE accelerators increasingly target the dominant bottlenecks in encrypted computation, including Number Theoretic Transform operations, modular arithmetic, and bootstrapping.[19] The practical position is disciplined adoption, that is to use FHE where the risk justifies the cost today, and design the architecture so hardware acceleration can be adopted as the market matures.

[Exhibit C: Legal and Technical Liability Reduction Matrix] - The matrix maps residual risk when moving from traditional pseudonymization toward Federated Learning, Differential Privacy, FHE-enabled validation, and layered PPML.

Exhibit C: Liability Reduction Matrix for PPML Controls

High residual risk Medium Low Lowest practical risk

Control pattern	Re-identification / linkage	Cross-border exposure	Model IP exposure	Operational cost
Traditional pseudonymization	High	High	Medium	Low
Federated learning	Medium	Low	Medium	Medium
Differential Privacy	Low to medium	Low	Medium	Medium
FHE-enabled validation	Low	Low	Low	High
Layered PPML architecture	Lowest practical risk	Lowest practical risk	Lowest practical risk	Managed through scoping

Reading the matrix: Federated Learning reduces transfer exposure; Differential Privacy reduces inference risk; FHE protects model and data confidentiality but remains the costliest control. The operating answer is layered PPML, not one privacy technology alone.

5. Implementation Blueprint

The implementation strategy should be organized around evidence generation. The objective is to produce a validation path that legal, clinical, cybersecurity, data science, patient-representation, and executive stakeholders can approve.

[Exhibit D: Implementation Blueprint] - The implementation path translates PPML into a repeatable validation operating model, from readiness through lifecycle assurance.

Exhibit D: Whitepaper Implementation Blueprint - From Model to Evidence

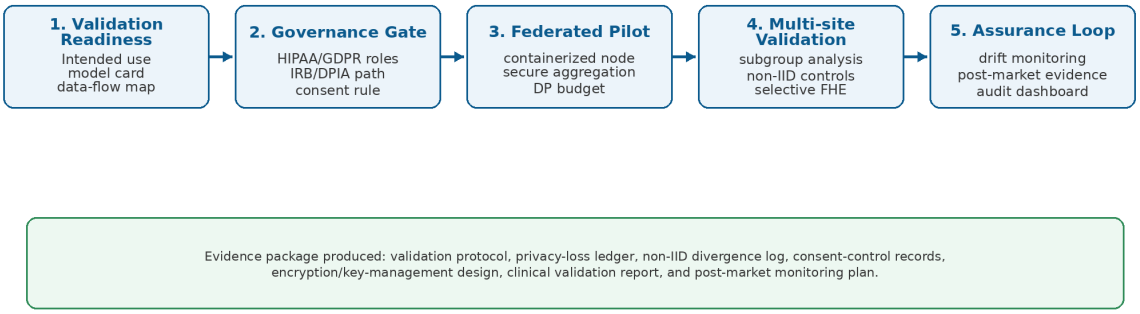

Phase	Timeframe	Primary Objective	Critical Activities	Exit Criteria
1	0-60 days	Validation readiness	Define intended use; identify target hospitals; complete privacy/security review; draft model card, data-flow map, and rural-node readiness profile.	Approved validation protocol, stakeholder map, and infrastructure inclusion baseline.
2	60-120 days	Legal and institutional alignment	Complete HIPAA/GDPR role analysis; DPIA/IRB pathway; SCC or transfer analysis if applicable; consortium agreement; dynamic-consent pathway where applicable.	Executed agreements, approved governance path, and consent-control decision record.
3	120-210 days	Technical pilot	Deploy federated node; test secure aggregation; configure differential privacy; run synthetic dry run; validate logging, rollback, containerized deployment, and non-IID divergence assessment.	Pilot report with security, privacy, infrastructure, and non-IID evidence.
4	210-360 days	Multi-site clinical validation	Run retrospective validation across hospitals; evaluate subgroup performance; apply distance-aware aggregation; use FHE for high-sensitivity scoring where warranted.	Clinical validation report, subgroup findings, FHE feasibility note, and limitations register.

Phase	Timeframe	Primary Objective	Critical Activities	Exit Criteria
5	Ongoing	Lifecycle assurance	Monitor drift; perform post-market surveillance; refresh privacy budget; update model cards; audit consent, consortium controls, and rural-node availability.	Repeatable evidence cycle and executive risk dashboard.

6. Governance and Evidence Model

A defensible PPML program requires more than technical integration. It requires a governance model that aligns legal authority, clinical oversight, security architecture, product accountability, patient agency, infrastructure equity, and evidence generation.

Layer	Primary Function	Key Controls	Evidence Produced
Governance	Define authority and accountability	Consortium charter; IRB/ethics alignment; DPIA review; model-use boundaries; escalation rights.	Signed protocol; accountability matrix; approved validation plan.
Legal and Regulatory	Map processing, transfer, and role allocation	HIPAA business associate analysis; GDPR controller/processor or joint-controller assessment; SCCs where required; Article 9 basis.	Legal basis memorandum; transfer impact assessment; contract exhibits.
Security Architecture	Protect local nodes, model movement, and update channels	Zero trust access; secure enclaves where appropriate; encryption; key management; secure aggregation; audit logs.	Security design review; control evidence; penetration test summary.
PPML Engineering	Limit exposure and inference risk	Federated orchestration; differential privacy budget; targeted FHE; update integrity; drift monitoring.	Privacy-loss ledger; model cards; validation metrics; cryptographic attestations.
Patient Consent and Equity	Preserve agency and inclusion	Dynamic consent logic; smart-contract records; rural-node readiness; model-compression baseline; non-IID assessment.	Consent audit trail; site-readiness score; equity participation report; non-IID divergence log.
Clinical Evidence	Generate defensible validation outcomes	Cohort definition; subgroup analysis; prospective validation where feasible; randomized evaluation for high-risk use cases.	Clinical validation report; subgroup performance; limitations and post-market monitoring plan.

Computation governance, not just data sharing

Traditional data-sharing agreements assume the key transaction is data movement. PPML allows the parties to negotiate a different transaction which is governed computation. The agreement should specify what code runs, where it runs, what data remain local, what outputs leave, how privacy loss is managed, who holds keys, how logs are retained, how consent or participation rules are enforced, and how adverse findings are escalated.

The privacy advocate correction is critical. Computation governance protects institutional data sovereignty, but it can still ignore the patient whose health data is being used to train or validate proprietary AI. Where consent or research participation is part of the lawful and ethical basis for processing, the consortium should treat consent as a dynamic governance control rather than a static signature. Smart contracts and consent ledgers can encode patient preferences, automate permitted-use checks, document consent state, and support revocation or preference changes without requiring raw-data exchange.[17]

Regulatory evidence package

A PPML-enabled validation project should produce an evidence package that can withstand regulatory, procurement, and clinical scrutiny. That package should include the validation protocol, technical architecture, privacy-risk analysis, cybersecurity control mapping, data-flow diagrams, model-card documentation, subgroup performance results, privacy-budget records, non-IID divergence logs, encryption and key-management design, consent-control decision records, and post-market monitoring plan.

For HIPAA, the package should address whether data are identifiable, de-identified, limited, or otherwise protected and how the Expert Determination or Safe Harbor analysis is supported. For GDPR, it should address Article 9 conditions, Article 32 security, international transfer posture, data minimization, purpose limitation, and DPIA outcomes where required. For FDA-facing strategy, it should distinguish retrospective validation, prospective validation, randomized controlled evaluation, and post-market surveillance.

7. Strategic Recommendations

Recommendation 1: Treat PPML as core infrastructure, not a research feature

Privacy-preserving architecture belongs in the product roadmap, validation strategy, security architecture, and enterprise sales motion from the beginning.

Recommendation 2: Design for validation before commercial scale

A model that cannot be validated across institutions is not ready for enterprise healthcare. Validation cohorts, subgroup reporting, privacy budgets, model update rules, non-IID controls, and audit evidence should be defined before broad deployment.

Recommendation 3: Use PPML to make equity measurable

Privacy should not become an excuse for avoiding demographic and subgroup analysis. The architecture should support measurement across age, sex, race, ethnicity, socioeconomic context, geography, and clinical subpopulations without unnecessary raw-record disclosure.

Recommendation 4: Separate legal transfer risk from computation risk

Organizations should stop assuming validation requires cross-border data transfer. Federated Learning and FHE allow parties to ask whether evidence can be generated without exporting raw data.

Recommendation 5: Build an audit trail for every material trust claim

Every claim should produce evidence: validation sites, cohorts, model versions, privacy budgets, encryption design, key custody, permitted outputs, and known limitations.

Recommendation 6: Prevent Federated Learning from becoming an academic-center privilege

Validation networks should require an infrastructure inclusion plan with lightweight FL nodes, containerized deployments, model compression, and support for lower-resourced hospitals.

Recommendation 7: Treat patient consent as a programmable governance control

When consent, opt-in research participation, or institutional policy governs data use, patient preferences should be encoded, auditable, updateable, and enforceable.

Recommendation 8: Be honest about FHE economics

FHE is appropriate for targeted, high-sensitivity validation workflows. It should not be sold as a universal replacement for normal compute.

Conclusion

The next era of healthcare AI will not be won by the fastest model builders alone. It will be won by professionals who can validate models in the real world without compromising privacy, security, intellectual property, patient agency, or public trust. The technical frontier is no longer separated from the legal frontier. The model, the contract, the data-flow diagram, the encryption design, the clinical validation protocol, and the board-risk narrative are now parts of the same system.

PPML matters because it makes responsible healthcare AI operationally possible. Federated Learning converts hospital data isolation into a controlled validation network. Differential Privacy turns privacy from vague assurance into measurable leakage control. FHE challenges the assumption that useful computation requires plaintext exposure. Edge-aware deployment, dynamic consent controls, selective FHE acceleration, and distance-aware non-IID reconciliation make the architecture credible beyond the first-generation PPML story.

Final position

The point is not to slow innovation rather it is to make innovation survivable. Healthcare AI needs innovators who can build systems that are clinically credible, regulatorily defensible, commercially viable, and secure by design.

References

[1] Chouffani El Fassi, S., Abdullah, A., Fang, Y., et al. (2024). Not all AI health tools with regulatory authorization are clinically validated. Nature Medicine, 30, 2718-2720. https://doi.org/10.1038/s41591-024-03203-3

[2] UNC Health Newsroom. (2024, August 26). Researchers highlight need for published validation data as artificial intelligence is thrust into patient care. https://news.unchealthcare.org/2024/08/researchers-highlight-need-for-published-validation-data-as-artificial-intelligence-is-thrust-into-patient-care/

[3] Muralidharan, V., Adewale, B. A., Huang, C. J., et al. (2024). A scoping review of reporting gaps in FDA-approved AI medical devices. npj Digital Medicine, 7, Article 273. https://www.nature.com/articles/s41746-024-01270-x

[4] U.S. Food and Drug Administration. Artificial Intelligence-Enabled Medical Devices. https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-enabled-medical-devices

[5] General Data Protection Regulation, Article 9, Processing of special categories of personal data. https://gdpr-info.eu/art-9-gdpr/

[6] European Commission. Standard Contractual Clauses for data transfers between EU and non-EU countries. https://commission.europa.eu/law/law-topic/data-protection/international-dimension-data-protection/standard-contractual-clauses-scc_en

[7] U.S. Department of Health and Human Services. Guidance regarding methods for de-identification of protected health information in accordance with the HIPAA Privacy Rule. https://www.hhs.gov/hipaa/for-professionals/special-topics/de-identification/index.html

[8] European Data Protection Board. Data protection basics. https://www.edpb.europa.eu/sme-data-protection-guide/data-protection-basics_en

[9] General Data Protection Regulation, Article 32, Security of processing. https://gdpr-info.eu/art-32-gdpr/

[10] European Union Agency for Cybersecurity. (2019). Pseudonymisation techniques and best practices. https://www.enisa.europa.eu/publications/pseudonymisation-techniques-and-best-practices

[11] McMahan, B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. y. (2017). Communication-Efficient Learning of Deep Networks from Decentralized Data. PMLR 54:1273-1282. https://proceedings.mlr.press/v54/mcmahan17a.html

[12] Near, J., Darais, D., Lefkovitz, N., & Howarth, G. (2025). Guidelines for Evaluating Differential Privacy Guarantees. NIST Special Publication 800-226. https://doi.org/10.6028/NIST.SP.800-226

[13] National Institute of Standards and Technology. Privacy-Enhancing Cryptography: Fully-Homomorphic Encryption. https://csrc.nist.gov/projects/pec/fhe

[14] Jiang, Y., Wang, S., Valls, V., Ko, B. J., Lee, W. H., Leung, K. K., & Tassiulas, L. (2023). Model pruning enables efficient federated learning on edge devices. IEEE Transactions on Neural Networks and Learning Systems. https://pubmed.ncbi.nlm.nih.gov/35468066/

[15] Brecko, A., Kajati, E., & Papcun, P. (2022). Federated Learning for Edge Computing: A Survey. Applied Sciences, 12(18), 9124. https://www.mdpi.com/2076-3417/12/18/9124

[16] Alekseenko, J., Karargyris, A., & Padoy, N. (2024). Distance-Aware Non-IID Federated Learning for Generalization and Personalization in Medical Imaging Segmentation. PMLR, 250:33-47. https://proceedings.mlr.press/v250/alekseenko24a.html

[17] Kostick-Quenet, K. M., et al. (2025). Patient-centric federated learning: automating meaningful consent to health data sharing with smart contracts. Journal of Law and the Biosciences. https://academic.oup.com/jlb/article/12/1/lsaf003/8123125

[18] Feldmann, A., Samardzic, N., Krastev, A., et al. (2021). F1: A Fast and Programmable Accelerator for Fully Homomorphic Encryption. MICRO-54. https://dl.acm.org/doi/10.1145/3466752.3480070

[19] Meng, X., Jiang, Z., & Lyu, Y. (2024). HF-NTT: Hazard-Free Dataflow Accelerator for Number Theoretic Transform. arXiv. https://arxiv.org/abs/2410.04805

Appendix: Citation-to-Claim Map

Claim Area	Source	Use in Whitepaper
Clinical validation deficit	[1], [2], [3]	Supports the central claim that authorization and transparent clinical validation are not equivalent.
FDA list limitations	[4]	Supports the discussion that public FDA summaries are useful but incomplete evidence packages.
HIPAA de-identification	[7]	Supports the analysis of Expert Determination, Safe Harbor, and reasonable-basis standards.
GDPR health data and security	[5], [8], [9]	Supports special category data treatment, extraterritorial reach, and risk-appropriate controls.
International transfer mechanisms	[6]	Supports cross-border analysis and the limits of contract-only approaches.
Pseudonymization attack surface	[10]	Supports linkage-attack and re-identification discussion.
Federated Learning	[11]	Supports model-to-data architecture and distributed learning rationale.
Differential Privacy	[12]	Supports privacy-loss and implementation-governance discussion.
FHE and acceleration	[13], [18], [19]	Supports encrypted-computation discussion and hardware-aware scoping.
Rural participation	[14], [15]	Supports edge-aware FL, pruning, quantization, and resource-conscious deployment.
Dynamic consent	[17]	Supports smart-contract and consent-ledger governance for patient-centric FL.
Non-IID data reconciliation	[16]	Supports distance-aware assessment and client clustering for distributed clinical heterogeneity.