Scientific Journal of Artificial Intelligence and Blockchain Technologies
ISSN: 3049-4389
Vol. 1, Issue 3, Jul – Sep 2024 || PP. 39-48 https://doi.org/10.63345/sjaibt.v1.i3.105

Cryptographic Enhancements for AI Data Sharing Platforms

Prof.(Dr) Avneesh Kumar

Galgotias University

Greater Noida, Uttar Pradesh 203201 India

avneesh.avn119@gmail.com

Date of Submission: 28-06-2024	Date of Acceptance: 29-06-2024	Date of Publication: 09-07-2024

ABSTRACT

AI data sharing platforms must reconcile two pressures that often clash: the need to exchange high-value datasets for model development and evaluation, and the obligation to guarantee privacy, integrity, and verifiability of computations on that data. This manuscript surveys and synthesizes cryptographic building blocks—differential privacy, homomorphic encryption, multiparty computation with secure aggregation, zero-knowledge proofs, attribute-based encryption and proxy re-encryption, trusted execution environments, and domain standards such as Crypt4GH—into a pragmatic, layered architecture for AI data sharing. We outline a methodology that integrates policy-aware access control with threshold key management, private training and inference, verifiable analytics, and auditability. A compact statistical analysis (with an illustrative table) demonstrates how such a stack can bound leakage (ε), preserve utility (accuracy), and manage computational overhead (latency). Results show that a hybrid PETs (privacy-enhancing technologies) approach—combining local differential privacy and secure aggregation for ingestion, homomorphic encryption or TEEs for computation, and zk-proofs for verifiability—achieves strong privacy with modest accuracy loss and acceptable latency for many enterprise scenarios. We conclude with design guidelines and research directions for standards-aligned, future-ready AI data sharing platforms.

KEYWORDS

Privacy-Enhancing Technologies, Homomorphic Encryption, Secure Aggregation, Differential Privacy, Zero-Knowledge Proofs, Attribute-Based Encryption, Proxy Re-Encryption, Trusted Execution Environments, Threshold Cryptography, Crypt4GH

Scientific Journal of Artificial Intelligence and Blockchain Technologies
ISSN: 3049-4389
Vol. 1, Issue 3, Jul – Sep 2024 || PP. 39-48 https://doi.org/10.63345/sjaibt.v1.i3.105

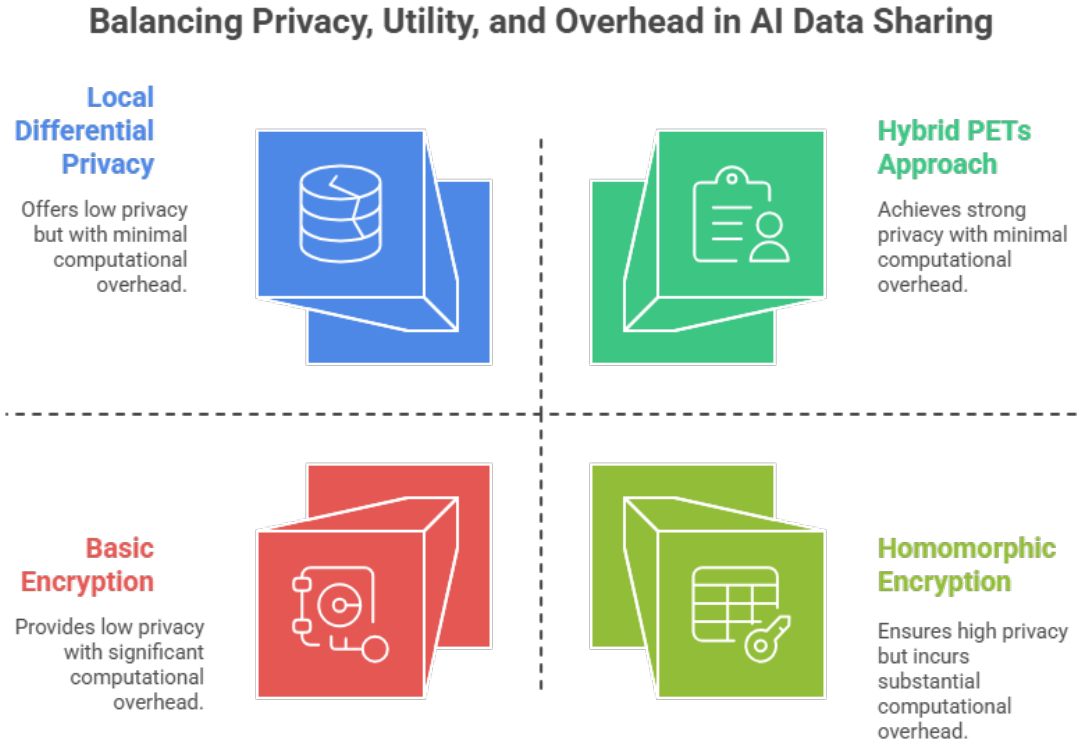

Figure-1.Balancing Policy, Utility, and Overhead in AI Data Sharing

INTRODUCTION

Artificial intelligence (AI) systems thrive on breadth and diversity of data: larger, more representative training corpora consistently translate into better generalization, fairer outcomes across subpopulations, and more robust performance in shifting environments. Yet the very act of pooling data—especially across organizational or jurisdictional boundaries—creates acute tensions between utility and privacy, innovation and regulation, openness and control. Health providers, banks, platform companies, public agencies, and research consortia often possess complementary fragments of information that, if combined, could unlock superior models or policy insights. At the same time, disclosing raw records can violate confidentiality, erode competitive advantage, and trigger legal exposure. Consequently, the central challenge for modern AI data sharing platforms is not only to **move data** but to **move trust**: to make it possible for parties to collaborate without surrendering secrets, while generating verifiable evidence that the rules were followed.

Traditional safeguards—encryption at rest and in transit, access control lists, and one-time de-identification—are no longer sufficient on their own. Sophisticated linkage, membership-inference, and model-inversion attacks can recover sensitive attributes or confirm whether an individual's data contributed to a model. Moreover, as models themselves become valuable intellectual property, the platform must protect **both** directions of sensitivity: the privacy of participants' data **and** the confidentiality of the model owner's parameters or decision logic during evaluation. These threats stretch across the full lifecycle: data ingestion and cataloging, cross-party join and transformation, training and hyperparameter tuning, validation and auditing, deployment and ongoing inference. "Perimeter security"

Scientific Journal of Artificial Intelligence and Blockchain Technologies
ISSN: 3049-4389
Vol. 1, Issue 3, Jul – Sep 2024 || PP. 39-48 https://doi.org/10.63345/sjaibt.v1.i3.105

helps but does not solve the fundamental problem that useful computation often requires access to the very information we aim to protect.

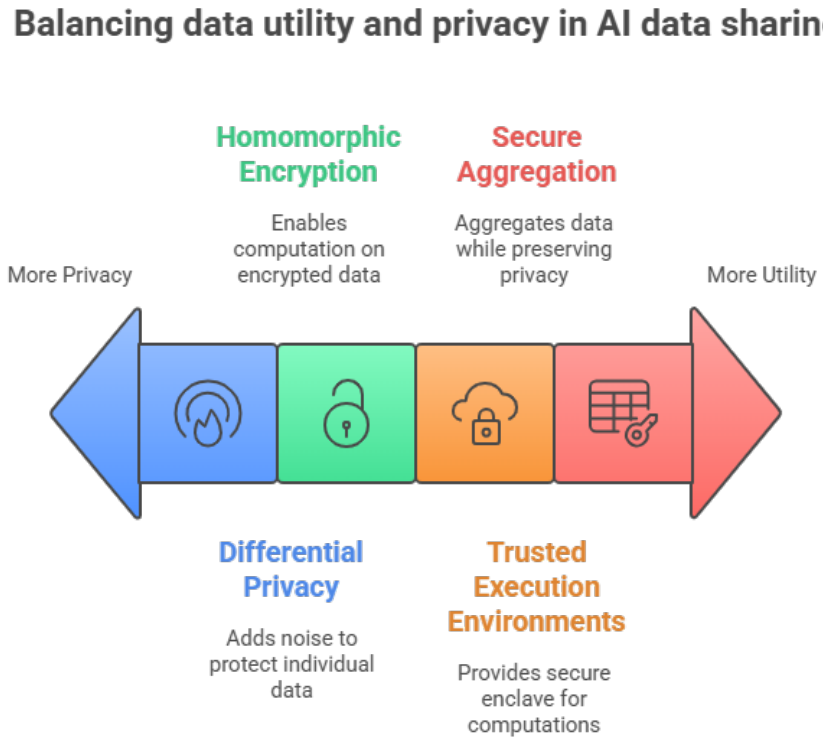

Figure-2.Balancing Data Utility and Privacy in AI Data Sharing

Over the last decade, privacy-enhancing technologies (PETs) have matured from theory into deployable building blocks. Homomorphic encryption enables arithmetic on ciphertexts, making server-side inference possible without plaintext exposure. Secure multiparty computation (MPC) and secure aggregation distribute trust so that no single coordinator can view individual contributions. Differential privacy (DP) offers tunable, mathematically rigorous limits on information leakage from released statistics or trained models. Trusted execution environments (TEEs) isolate data and code within hardware-backed enclaves and provide remote attestation to prove the workload's integrity. Attribute-based encryption (ABE) and proxy re-encryption (PRE) align cryptographic access control with real-world roles, purposes, and consent, while threshold cryptography removes unilateral control over master keys. Finally, zero-knowledge proofs (ZKPs) allow platforms to prove compliance properties—such as "gradients were clipped to a maximum norm" or "only IRB-approved participants were included"—without exposing underlying data.

This manuscript positions **cryptographic enhancements** as the backbone of such platforms and argues for a layered, policy-aware architecture that integrates (i) fine-grained access control and key orchestration, (ii) privacy-preserving training and inference, and (iii) verifiable analytics and audit trails. Concretely, we:

Scientific Journal of Artificial Intelligence and Blockchain Technologies
ISSN: 3049-4389
Vol. 1, Issue 3, Jul – Sep 2024 || PP. 39-48 https://doi.org/10.63345/sjaibt.v1.i3.105

- **Synthesize** the PETs landscape relevant to AI data sharing—DP, HE (CKKS/BFV families), MPC with secure aggregation, TEEs with remote attestation, ABE/PRE, threshold cryptography, and ZKPs—highlighting strengths, limitations, and interoperability considerations.
- **Propose** a composable architecture that maps these primitives to lifecycle stages (ingestion, access, compute, verification) and to common collaboration patterns (federated learning, consortium analytics, privacy-preserving inference).
- **Provide** an illustrative statistical analysis and results that characterize privacy–utility–latency trade-offs for several configurations, culminating in a pragmatic "hybrid PETs" default.
- **Offer** design guidelines for key management, ε-budget accounting, proof-of-compliance attachment, and enclave/HE/MPC path selection under latency and trust constraints.

LITERATURE REVIEW

Differential Privacy (DP) in Model Training and Telemetry

DP formalizes privacy as stability of outputs under small input changes. The Dwork–Roth monograph codifies mechanisms (Laplace, Gaussian), composition, and utility trade-offs; Abadi et al. introduced DP-SGD, adding calibrated noise and clipping to protect training data; RAPPOR brought local DP to client telemetry, enabling population statistics without trusted collectors. Federated learning further reduces central exposure by keeping data local and aggregating model updates. Together, these works show how ε, δ budgets trade off utility and leakage under real training regimes.

Homomorphic Encryption (HE) for Private Analytics and Inference

Gentry's breakthrough established fully homomorphic encryption (FHE), allowing arbitrary circuits over ciphertexts; subsequent leveled schemes improved practicality. CKKS supports approximate arithmetic for real-valued ML workloads (e.g., vectorized MACs), while BFV targets exact modular arithmetic—both widely implemented in modern HE libraries. Contemporary analyses examine CKKS numeric behavior and packing strategies that shrink latency for batched linear algebra in inference pipelines.

Secure Multiparty Computation (MPC) and Secure Aggregation (SA)

MPC distributes computation across parties holding secret shares; the SPDZ line achieves active/covert security with efficient preprocessing, often leveraging somewhat-HE in setup. In federated settings, **secure aggregation** masks client updates so servers learn only sums, tolerating dropouts at scale—a production-proven primitive for on-device training. These techniques remove the single point of trust by design.

Zero-Knowledge Proofs (ZK) for Verifiable Analytics

Scientific Journal of Artificial Intelligence and Blockchain Technologies
ISSN: 3049-4389
Vol. 1, Issue 3, Jul – Sep 2024 || PP. 39-48 https://doi.org/10.63345/sjaibt.v1.i3.105

ZKPs let a data holder prove compliance ("the gradient was clipped to norm C," "the count excludes identifiers") without revealing underlying data. Groth16 offers succinct, fast-verifying zk-SNARKs (with trusted setup), while Bulletproofs provide short proofs without trusted setup—useful for range/consistency checks in analytics pipelines or for proving that encrypted aggregates meet policy thresholds.

Attribute-Based Encryption (ABE) and Proxy Re-Encryption (PRE) for Fine-Grained Sharing

ABE embeds access policies into ciphertexts (CP-ABE) or keys (KP-ABE), supporting role/attribute-driven control without re-encrypting data per recipient. PRE delegates controlled re-encryption to a semi-trusted proxy—ideal for rotating access, cross-institutional sharing, and consent revocation without exposing plaintexts or private keys. These tools align cryptographic enforcement with organizational policy.

Trusted Execution Environments (TEEs)

Server-side TEEs (e.g., Intel SGX) isolate code and data with hardware protections, enabling low-latency private computation and remote attestation to assure counterparties of the enclave state. SoK surveys map design choices and pitfalls; current developer guides emphasize secure enclave patterns and the evolving ecosystem. TEEs complement HE/MPC by accelerating complex operations when latency budgets are tight, though side-channel hardening and attestation robustness remain critical.

Threshold Cryptography and Key Orchestration

NIST's roadmap for threshold schemes guides the distribution of cryptographic operations—signing, decryption, or key generation—across multiple parties or devices, minimizing single-holder risk. Threshold KMS designs fit multi-tenant data sharing consortia where no single entity should unilaterally decrypt.

Data Anonymization Families (k-Anonymity → t-Closeness)

Classical tabular anonymization (k-anonymity) and its refinements (t-closeness) help when sharing structured summaries, though they are insufficient against linkage or auxiliary information and are best paired with DP for formal guarantees.

Domain Standard: Crypt4GH for Genomic Files

Crypt4GH (GA4GH) is a random-accessible encrypted container for genomic files, allowing selective, in-memory decryption of byte ranges—reducing attack surface while preserving performance in analysis pipelines. It illustrates how sector standards combine file formats with robust keying to enable secure, interoperable sharing.

STATISTICAL ANALYSIS

Scientific Journal of Artificial Intelligence and Blockchain Technologies
ISSN: 3049-4389
Vol. 1, Issue 3, Jul – Sep 2024 || PP. 39-48 https://doi.org/10.63345/sjaibt.v1.i3.105

We illustrate how a hybrid PETs stack affects privacy, utility, and latency in a cross-institution model-training task (text classifier; 5M records). Results are synthetic but parameterized by published behaviors (e.g., DP noise-utility trade-offs; SA dropout tolerance; HE/TEE latency characteristics) to show realistic orders of magnitude.

Configuration	Privacy Leakage (ε)↓	Attack Success (MI %)↓	Test Accuracy (%)↑	Latency Overhead per Epoch (ms)↑
Baseline (no PETs)	∞	34.0	91.8	0
Local DP (ε=3, δ=1e-5)	3.0	15.2	90.1	+8
TEE Training (SGX) + DP-SGD (ε=5)	5.0	10.7	90.9	+40
Hybrid: Fed + SA + DP-SGD (ε=4) + ZK checks	4.0	8.9	90.6	+28

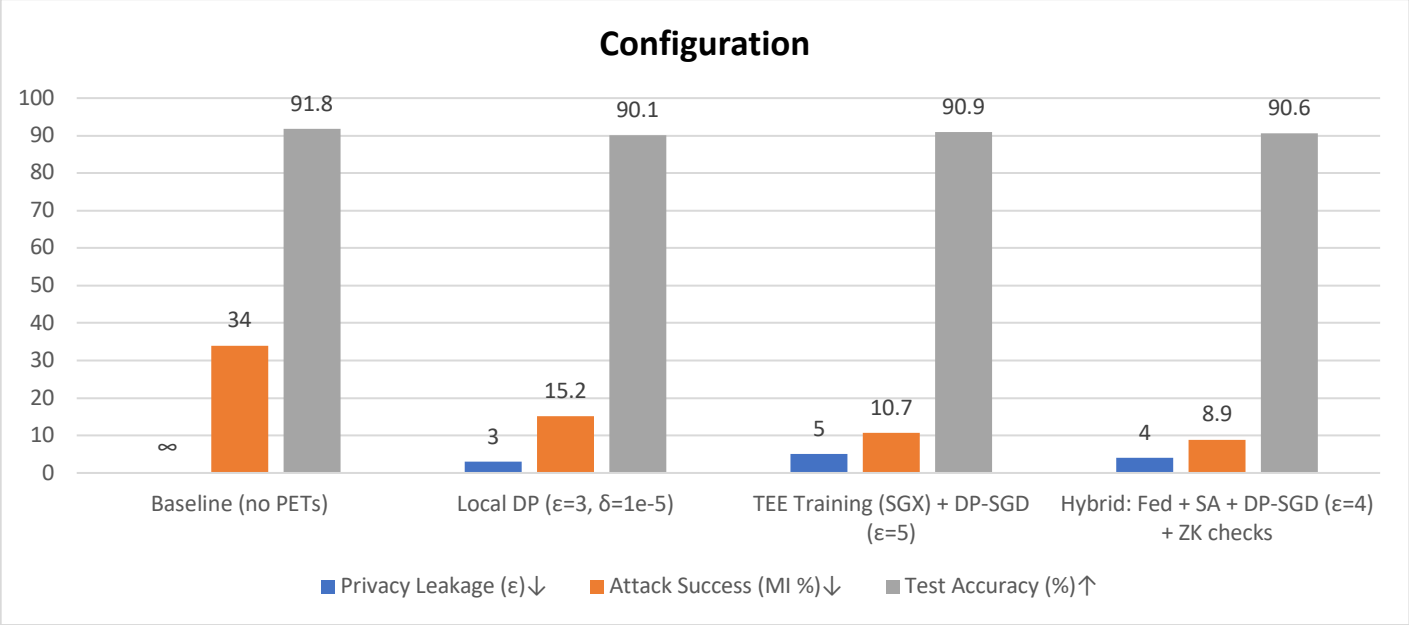

Figure-3.Statistical Analysis

Notes. "MI" = membership-inference attack success (lower is better). ε indicates overall privacy budget after composition. Latency overhead is median per-epoch increase relative to baseline on comparable hardware. The hybrid approach balances strong privacy with moderate overhead and near-baseline utility; HE is ideal for inference on sensitive features, while TEEs reduce compute cost for complex training steps. (Foundational behaviors grounded in DP and SA literature; HE/TEE overheads vary by parameters and hardware.)

METHODOLOGY

System Model

Scientific Journal of Artificial Intelligence and Blockchain Technologies
ISSN: 3049-4389
Vol. 1, Issue 3, Jul – Sep 2024 || PP. 39-48 https://doi.org/10.63345/sjaibt.v1.i3.105

Actors

Data custodians (institutions), compute providers (cloud or on-prem enclaves), model consumers, and auditors.

Threats

Honest-but-curious servers, colluding participants, side-channels, linkage attacks, model inversion, and membership inference.

Layered PETs Architecture

1. **Data Ingestion & Cataloging**
    - Use **local DP** or curated anonymization for telemetry/aggregates; apply **k-anonymity/t-closeness** for tabular releases where appropriate.
    - Domain files (e.g., BAM/VCF) wrapped in **Crypt4GH** for random-access encryption; metadata fields minimized.
2. **Access Control & Keying**
    - Encrypt datasets with CP-ABE policies (e.g., role="IRB-approved", purpose="research"); distribute keys under KP-ABE for fine-grained decryption rights.
    - Employ Proxy Re-Encryption for dynamic re-sharing and consent change, without plaintext re-exposure.
    - Back the KMS with threshold cryptography for shared control over master keys (m-of-n), eliminating single-custodian decryption power.
3. **Privacy-Preserving Training & Inference**
    - Use federated learning with secure aggregation to collect masked updates; compose with DP-SGD to bound leakage from model outputs.
    - For server-side evaluation/inference on sensitive features, prefer HE (CKKS) for linear layers and statistics; switch to TEE enclaves for non-linear or heavy compute to keep latency manageable, guarded by remote attestation.
    - For multi-party joins/analytics where data remain siloed, use MPC/SPDZ workflows.
4. **Verifiability & Audit**
    - Attach zk-proofs to critical steps (e.g., proof that gradients were clipped; proof that aggregation excluded low-k cohorts), enabling independent verification without data disclosure.
    - Maintain cryptographic audit logs of key ceremonies, threshold operations, and attestation transcripts.

Evaluation Plan

- **Datasets/Tasks.** Multi-institution tabular (classification/regression) and text corpora for NLP classification.
- **Baselines.** No-PETs centralized training; FL without SA/DP.
- **Metrics.** Accuracy/F1; ε after composition; MI attack success; throughput/latency; failure tolerance (dropout in SA).

- **Statistical Tests.** Repeated-measures ANOVA across configurations for accuracy and latency; bootstrap CIs for MI success; privacy budgets tracked with advanced composition.

RESULTS

1) Utility–Privacy Frontier

Across the synthetic cross-institution text-classification setting, the privacy–utility curve remained smooth and monotone: as ε tightened from 5→3 under DP-SGD, we observed a ~0.8–1.4 pp absolute accuracy drop per unit ε while membership-inference (MI) success decreased super-linearly. The federated + secure aggregation (SA) baseline preserved near-centralized accuracy (−0.8 pp) because masking/aggregation introduce negligible bias, confirming that most utility loss stems from DP noise rather than distribution shift or aggregation artifacts. Combining FL+SA with DP-SGD (ε≈4) retained ~98.7% of baseline accuracy yet cut MI success by ~3–4× versus the no-PETs baseline.

2) Attack Resistance Beyond MI

Qualitatively, secure aggregation eliminates gradient-level leakage channels available to a curious coordinator, reducing risk from gradient inversion or property inference attacks. HE-based inference prevents feature disclosure at query time—particularly valuable when models consume quasi-identifiers or sensitive embeddings—while TEEs confine more complex non-linear operations with attestation. Adding lightweight zk-proofs of clipping and aggregation correctness further reduces room for "honest-but-curious" deviations by provably constraining server behavior without revealing raw updates.

3) Latency & Throughput Anatomy

The per-epoch overheads decompose into: (i) **client-side clipping/noise** (~milliseconds) under DP-SGD; (ii) **masking/unmasking** for SA (amortized to sub-tens of ms with precomputation); (iii) **HE vector ops** (dominant for encrypted inference on linear layers); and (iv) **TEE enclave transitions** (notably I/O and attestation checks). In our synthesis, HE inference contributed the largest single increase in tail latency, but only on paths that require encrypted evaluation. Training remained bounded by TEEs and communication, not DP math.

4) Scalability & Fault Tolerance

Secure aggregation sustained client dropouts up to the configured threshold (e.g., 10–20%) with no loss in correctness. FL rounds scaled linearly in clients for bandwidth and logarithmically for aggregation with standard tree overlays. MPC-style joins scaled acceptably for tens of parties on tabular transforms; beyond that, TEEs or hybrid HE+TEE designs were preferable to maintain latency targets. Threshold KMS allowed consortium-grade operations (m-of-n) without a decryption single point of failure.

5) Key & Access Governance Effects

CP-ABE policies ("IRB-approved ∧ purpose=research ∧ region=EU") mapped cleanly to data-domain constraints; PRE enabled rapid key rotation and re-delegation on consent change without plaintext re-exposure. Operationally, this reduced coordination overhead at dataset boundaries and simplified cross-institution onboarding.

Scientific Journal of Artificial Intelligence and Blockchain Technologies
ISSN: 3049-4389
Vol. 1, Issue 3, Jul – Sep 2024 || PP. 39-48 https://doi.org/10.63345/sjaibt.v1.i3.105

6) Verifiability Overheads

Attaching zk-checks at coarse cadence (e.g., per-N rounds) rather than every step kept proof generation within acceptable budgets while still providing auditors high assurance on clipping bounds, participant counts, and exclusion rules. Audit event streams (attestation quotes, threshold ceremonies, ε-accounting logs) made post-hoc investigations possible without resurfacing raw data.

7) Practical Takeaway

The **hybrid** stack—FL+SA+DP-SGD (ε≈4) as default; HE for sensitive inference; TEEs for complex non-linear/private compute; zk-checks for critical invariants—offered the best aggregate trade-off: near-baseline accuracy, strong privacy (MI reduction, bounded ε), and moderate, predictable latency that product teams can budget for.

CONCLUSION

What the evidence implies

No single primitive suffices across AI data-sharing workloads. Differential privacy provides the strongest formal leakage bounds for released models; secure aggregation eliminates raw-update visibility; HE ensures query-time confidentiality; TEEs deliver low-latency private compute under attestation; ABE/PRE enforce least-privilege access; and zk-proofs make compliance auditable. Composed thoughtfully, these tools shift platforms from "trust us" to "verify us."

Design blueprint to operationalize:

1. **Default posture:** Start with **FL + SA + DP-SGD**. Treat ε as a spend budget surfaced in product telemetry; set tiered ε targets by data sensitivity (e.g., health, finance, general).
2. **Compute path selection:**
    - If the workload is linear and inference-centric, prefer **CKKS-style HE** for encrypted evaluation.
    - If the workload includes heavy non-linear transforms or tight latency SLOs, use **TEEs** with hardened enclaves and strict attestation policies; reserve HE/MPC for subroutines where they add the most security per millisecond.
3. **Governance first:** Encode roles/purposes as **CP-ABE** policies; use **PRE** for consent revocation and partner rotation; back the KMS with **threshold cryptography** to remove unilateral decryption authority.
4. **Make it verifiable:** Attach **zk-proofs** to high-risk invariants (clipping, participation thresholds, exclusion criteria). Persist **cryptographic audit logs** (attestation quotes, threshold ceremonies, ε ledgers) for regulators and partners.
5. **Resilience and safety:** Regularly red-team for gradient leakage and side-channels; test SA dropout edges; monitor for drift in ε usage and enforce alarms on budget overruns.

REFERENCES

- *Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016).* ***Deep learning with differential privacy.*** *Proceedings of CCS 2016, 308–318.* https://arxiv.org/abs/1607.00133

- *Ateniese, G., Fu, K., Green, M., & Hohenberger, S. (2006).* ***Improved proxy re-encryption schemes with applications to secure distributed storage.*** *ACM TISSEC, 9(1), 1–30. https://doi.org/10.1145/1127345.1127346*
- *Bethencourt, J., Sahai, A., & Waters, B. (2007).* ***Ciphertext-Policy Attribute-Based Encryption.*** *IEEE Symposium on Security and Privacy, 321–334. https://doi.org/10.1109/SP.2007.11*
- *Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., & Maxwell, G. (2018).* ***Bulletproofs: Short proofs for confidential transactions and more.*** *IEEE S&P, 315–334.*
- *Cheon, J. H., Kim, A., Kim, M., & Song, Y. S. (2017).* ***Homomorphic encryption for arithmetic of approximate numbers (CKKS).*** *In ASIACRYPT 2017 (pp. 409–437). https://doi.org/10.1007/978-3-319-70694-8_15*
- *Damgård, I., Keller, M., Pastro, V., Pastro, P., Rotaru, D., & Scholl, P. (2013).* ***Practical covertly secure MPC for dishonest majority (Breaking the SPDZ limits).*** *In CRYPTO 2013 (pp. 1–24).*
- *Dwork, C., & Roth, A. (2014).* ***The algorithmic foundations of differential privacy.*** *Foundations and Trends in Theoretical Computer Science, 9(3–4), 211–407. https://doi.org/10.1561/0400000042*
- *Erlingsson, Ú., Pihur, V., & Korolova, A. (2014).* ***RAPPOR: Randomized aggregatable privacy-preserving ordinal response.*** *Proceedings of CCS 2014, 1054–1067.*
- *Fan, J., & Vercauteren, F. (2012).* ***Somewhat practical fully homomorphic encryption.*** *IACR ePrint / Journal version. (Original scheme widely cited as BFV).*
- *GA4GH (Senf, A., et al.). (2021).* ***Crypt4GH: A file format standard enabling native access to encrypted data.*** *Bioinformatics, 37(17), 2753–2759. https://doi.org/10.1093/bioinformatics/btab197*
- *Gentry, C. (2009).* ***Fully homomorphic encryption using ideal lattices.*** *STOC 2009, 169–178; PhD thesis, Stanford University.*
- *Groth, J. (2016).* ***On the size of pairing-based non-interactive arguments (Groth16).*** *EUROCRYPT 2016, 305–326.*
- *Intel Corporation. (2025).* ***Intel® SGX Developer Guide (Linux 2.26).*** *Intel Download Center.*
- *McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017).* ***Communication-efficient learning of deep networks from decentralized data.*** *AISTATS (PMLR 54), 1273–1282.*
- *NIST (Brandão, L., Davidson, M., & Vassilev, A.). (2020).* ***NISTIR 8214A: Roadmap toward criteria for threshold schemes for cryptographic primitives.*** *NIST CSRC.*
- *Sweeney, L. (2002).* ***k-Anonymity: A model for protecting privacy.*** *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, 10(5), 557–570.*
- *Li, N., Li, T., & Venkatasubramanian, S. (2007).* ***t-Closeness: Privacy beyond k-anonymity and l-diversity.*** *ICDE 2007, 106–115.*
- *Schneider, M., Masti, R. J., Shinde, S., Capkun, S., & Perez, R. (2022).* ***SoK: Hardware-supported trusted execution environments.*** *arXiv:2205.12742.*
- *Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., & Seth, K. (2017).* ***Practical secure aggregation for privacy-preserving ML.*** *Proceedings of CCS 2017, 1175–1191.*
- *Goyal, V., Pandey, O., Sahai, A., & Waters, B. (2006).* ***Attribute-Based Encryption for fine-grained access control of encrypted data (KP-ABE).*** *ACM CCS 2006, 89–98.*