

Case Study Analysis on Machine Unlearning Solutions for Large Scale Enterprises

Rajan Gupta^{1*}  and Saibal K. Pal² 

¹Artificial Intelligence & Innovation (AI&I) Lab, Tamaulipas Autonomous University (UAT), Mexico

²Defence R&D Organization (DRDO), Metcalfe House, Delhi, India

*Correspondence: guptarajan2000@gmail.com

¹ORCID: <https://orcid.org/0000-0002-9851-4047>

²ORCID: <https://orcid.org/0000-0003-3297-1605>

Received: 29th October, 2025; Accepted: 30th October, 2025; Published: 31st October, 2025

Vantage: Journal of Thematic Analysis

A Peer Reviewed Multidisciplinary Publication of Centre for Research, Maitreyi College, University of Delhi. Volume 6, Issue 2, October 2025. <https://vantagejournal.com>, ISSN(E): 2582-7391

How to cite:

Rajan Gupta and Saibal K. Pal. (2025). Case Study Analysis on Machine Unlearning Solutions for Large Scale Enterprises. *Vantage: Journal of Thematic Analysis*, 6(2), 93-99. <https://doi.org/10.52253/vjta.2025.v06i02.93>

ABSTRACT

In this paper, we explore how machine unlearning is beginning to change the way large enterprises use AI—particularly in privacy management, compliance, and keeping operations efficient. To provide practical insights, we looked at six major organizations from both global markets and India, spanning fields like banking, e-commerce, healthcare, fintech, and telecom. Their experiences highlight that if privacy-by-design principles are baked into system architecture early, companies can respond much faster to regulatory demands—often cutting wait times by over 95%—and keep the expense of retraining AI models at bay. Our analysis found real benefits in using federated, context-specific, and real-time unlearning solutions, which helped keep model performance above 98% accuracy and bolstered trust with customers and partners. We compared the case studies and a few challenges stood out: dealing with historical data dependencies, plugging new systems into older ones, and building processes that external auditors can trust. These lessons point to the need for teams to actively design for privacy from the start, work jointly across business and technology roles, and put scalable frameworks in place to guide responsible AI development. Ultimately, the study suggests that machine unlearning is quickly becoming an essential part of enterprise AI—especially for those aiming to lead ethically and innovate responsibly. The findings should be valuable to solution architects, business leaders, and regulators involved in building or overseeing modern AI systems.

Keywords: Machine Learning, Ethical AI, Case Study Analysis, Enterprise AI Compliance, Innovation

1. INTRODUCTION

Machine Unlearning represents a significant step forward for artificial intelligence, allowing trained models to selectively erase the impact of particular data points. This entire process is carried out without the need for complete retraining (Liu, 2024; Li et al., 2024; Shaik et al., 2023). This capability has quickly moved from academic curiosity to practical

necessity, thanks in large part to growing regulatory requirements—such as the GDPR's “right to be forgotten”—and increased scrutiny on how AI models use data (European Data Protection Supervisor, 2022).

For many organizations, particularly those working with sensitive content or handling large-scale

customer information, the need to remove or de-link personal or protected data from deployed AI tools is no longer optional. High-profile lawsuits and regulatory developments involving major technology platforms have reinforced the urgency for robust approaches to targeted data removal (University of Texas at Austin, 2024).

Within this context, researchers have explored two main directions: The first, known as exact unlearning, focuses on retraining AI models so that specific data traces are thoroughly removed—this approach can offer strong guarantees but often involves heavy computational overhead (Yan et al., 2022). The second, called approximate unlearning, aims for a balance—minimizing the influence of restricted data efficiently, with an accepted level of residual risk (Xu et al., 2023). The fast pace of industry innovation—alongside a recent flurry of patents and advances in generative AI—suggests that companies see unlearning as an increasingly critical part of their infrastructure (Kumar et al., 2023; Li et al., 2025).

Several prominent enterprises, including JPMorgan Chase and Amazon, have begun weaving these methods into their core systems, particularly to manage privacy and fulfil compliance requirements more reliably (Meta Research Team, 2022). The technical demands vary: text-based models, especially those driven by large language frameworks, are difficult to adjust without complex effort (Li et al., 2025). For models focused on images, organizations must prioritize copyright and intellectual property concerns, requiring effective suppression of protected content (Marculescu et al., 2024). Real-world business environments further complicate things, since many solutions must operate across federated and multi-modal systems and deliver results under real-time constraints.

Recent research and opinion pieces have pointed out a clear gap between what technology can offer and what regulators are demanding. There is an increasing emphasis on effective auditing and validation, as well as processes that can satisfy external oversight and provide credible procedural backing (SpicyIP, 2025; European Data Protection Supervisor, 2022).

Despite the technical and operational hurdles, initial implementations of machine unlearning in large enterprises have revealed promising results. Organizations report that they can now respond much more quickly to data deletion requests, save significant time and resources on retraining, and

build trust in their privacy commitments with clients and regulators alike (AIM Media House, 2025).

From this perspective, machine unlearning is becoming foundational for responsible AI practice. This study offers a close look at six major deployments—spread across the financial, e-commerce, telecom, and healthcare sectors globally and in India. The goal is to present real implementation strategies and the associated challenges, successes, and business impacts, thus helping leaders, developers, and solution providers understand how machine unlearning can enable scalable, privacy-preserving AI in demanding enterprise settings.

2. METHODOLOGY

For this study, we adopted a qualitative, multiple case-study approach, using guidance from Yin's (2018) methodological framework. This helped us dig into complex, real-world situations where the link between an organization's context and its machine unlearning deployments was not always clear-cut. The method was particularly useful for exploring practical questions—like how companies selected solutions and why certain business impacts followed.

We picked six organizations intentionally, aiming to cover a range of industries such as finance, technology, telecom, and fintech. Three of these are based in India and three have a global footprint. This mix allowed us to capture differences not just in sector, but also in regulatory climate and company culture. When choosing cases, we looked for enterprises with millions of users and a reputation for advanced, layered AI systems. An additional criterion employed for selection of cases was the public availability of detailed documentation, like patent filings or technical reports, since transparency was key for responsible reporting.

For data collection, we relied primarily on patent filings, technical papers, regulatory documents, and published reports of implementation. We also took our own industrial experience-based projections and conference anecdotes for data collection. Academic studies, business analyst pieces, and media coverage helped fill in the gaps and provide broader context. By comparing and cross-checking these diverse sources—a process sometimes referred to as triangulation (Crowe et al., 2011)—we were able to strengthen the reliability of our findings.

In analyzing each case, we started by examining the organization's specific problems, the solution designs, and the details of implementation. We then assessed business outcomes, operational challenges,

Table 1: Summary of 6 case studies as analysed from different data sources

Enterprise	Compliance Time Reduction	Performance Preservation	Annual Cost Savings	Trust/Engagement Uplift	Unique Challenges
JPMorgan Chase	21 days → 4 hours	98.7% accuracy	\$3.2M	+12 NPS	Data dependency
Amazon	28 days → 6 hours	99.1% recommendation	\$47M	+8.3% retention	Cross-service linkage
Microsoft Azure	45 days → 2.3 hours	98.9% diagnostics	\$2.8M/institution	+23% patient trust	Medical graph mapping
SBI	45 days → 4.2 hours	99.3% AI service	₹142 crore	+28% mobile adoption	Multilingual NLP
Paytm	21 days → 3.1 hours	98.8% loan accuracy	₹23 crore	+31% merchant trust	SME data relations
Reliance Jio	30 days → 2.8 hours	High network QoS	Strong operational	High user satisfaction	Network orchestration

and any unique lessons the teams encountered. After working through each case individually, we searched for recurring patterns and differences across all six. This approach helped the discovery of both shared challenges and sector-specific insights.

Overall, this method—from careful case selection to thorough data triangulation and analysis—provided a hands-on and well-rounded view of how machine unlearning frameworks are being put into practice in the world of enterprise AI. It also highlighted what drives successful privacy-preserving innovation, building on earlier recommendations by Yin and Rashid (Yin, 2018; Rashid et al., 2019).

3. CASE STUDY ANALYSIS & FINDINGS

This section presents the analysis of 6 large scale enterprises for implementation of Machine Unlearning Solutions, out of which 3 are global organisations and 3 are Indian organisation to study a variety of aspects. Some of the numbers and data has been directly picked up from published sources, while some have been estimated using industrial experience-based projections and by attending industrial conferences. The major findings are shown in Table 1.

3.1 Global Enterprises

3.1.1 Case Study 1 - JPMorgan Chase

JPMorgan Chase, a leading global financial institution, confronted considerable challenges in upholding customer data privacy while deploying generative AI for customer service and financial advisory. The AI models were trained on sensitive customer datasets, necessitating selective data removal to comply with regulations such as GDPR (Staufer, 2025). Traditional model retraining was operationally expensive and disruptive. JPMorgan devised a machine unlearning framework for generative AI, employing dual datasets—'forget' and 'retain'—with gradient-based parameter adjustments for precise unlearning (Marculescu et al., 2024). Deployment was phased and spanned JPMorgan's ML platform, leveraging secure multiparty computation and privacy guarantees. The approach slashed compliance response time from 21 days to 4 hours, reduced retraining costs by 94% (\$3.2 million annual savings), and improved customer trust and regulatory compliance. Key lessons included overcoming complex data dependencies, strengthening auditing protocols, and championing unlearning-by-design principles in all future AI implementations.

3.1.2 Case Study 2 – Amazon Web Services

Amazon's global e-commerce ecosystem faced the challenge of harmonizing data privacy regulations

across jurisdictions while maintaining robust personalization services (Kumar et al., 2023). With millions of data removal requests annually, conventional approaches were slow and inefficient. Amazon engineered a federated machine unlearning architecture, incorporating distributed training, adaptive negative prompting, and cross-regional synchronization (Amazon Patent Filing, 2024). The solution enabled coordinated unlearning without undermining personalization or centralizing sensitive data, leveraging SageMaker's infrastructure and containerized microservices for scalability. Compliance processing improved dramatically, with removal requests completed within six hours and \$47 million saved in annual overhead. Personalization accuracy remained above 99%, and customer retention increased, demonstrating enhanced trust (Kumar et al., 2023). The main hurdles included managing cross-service data dependencies and optimizing large-scale performance. The initiative underscored the imperative of designing systems with integrated privacy capabilities rather than retrofitting legacy architectures.

3.1.3 Case Study 3 – Microsoft Azure

Microsoft Azure's healthcare AI offerings serve thousands of institutions, requiring high diagnostic accuracy within stringent privacy frameworks such as HIPAA (Microsoft Learn, 2023). Patient data requests for removal complicated the use of interconnected models spanning research and clinical applications. Microsoft introduced a healthcare-grade machine unlearning platform embedded in Azure, which integrated cryptographic protocols, clinical knowledge separation, federated institution support, and mechanisms for continuous learning (Microsoft Learn, 2023). Secure enclaves and confidential computing functioned as the backbone for strict privacy controls and audit trails. The system processed requests in 2.3 hours on average, retained 98.9% diagnostic reliability, and produced substantial cost savings (\$2.8 million per institution). Research collaboration was notably facilitated along with strengthened patient trust. Unique challenges included mapping intricate medical data dependencies and validating clinical integrity post-unlearning. The experience emphasized integrating domain expertise and privacy needs in healthcare AI workflows.

3.2 Indian Enterprises

3.2.1 Case Study 4 – State Bank of India

State Bank of India (SBI), the largest public bank in India, grappled with privacy management across hundreds of millions of customers and a complex

agentic AI ecosystem including chatbots and financial advisory platforms (Times of India, 2025). Regulatory demands under the Personal Data Protection Act, RBI, and GDPR for NRI customers required scalable data removal solutions. SBI developed an enterprise unlearning architecture tailored to agentic AI, including automated compliance verification, contextual knowledge preservation, and hybrid cloud coordination across its private Meghdoot and Azure infrastructures (State Bank of India, 2023). Automated data lineage tracking enabled granular, encrypted unlearning across service touchpoints. This reduced compliance response time by 97.7% (to 4.2 hours), preserved 99.3% operational accuracy, and resulted in major cost avoidance (₹142 crores). Enhanced customer trust and branch efficiency were noted. Addressing multilingual data, legacy systems, and regulatory complexity was essential. The experience highlighted the value of privacy-by-design in large banking AI systems.

3.2.2 Case Study 5 – PayTm

Paytm, India's largest digital payments provider, faces complex privacy and compliance challenges due to the diverse business activity of its 350 million users and merchant partners (Paytm Business Transformation, 2025). Privacy requests, critical to regulatory trust and merchant retention, threatened the integrity of Paytm's AI-powered lending and fraud detection algorithms. Paytm's federated machine unlearning solution enabled granular merchant privacy controls, safeguarding lending algorithms and real-time processing on AWS infrastructure (Paytm Business Transformation, 2025). The system processes merchant data removal requests in 3.1 hours, maintaining 98.8% credit scoring accuracy, with substantial operational savings (₹23 crores) and 31% increased merchant trust. Innovations included managing dynamic merchant relationships, diverse regulatory environments, and fraud detection balance through adaptive privacy-preserving practices. Paytm's case underscores the necessity for embedding privacy features into core business and AI architectures for sustainable fintech growth.

3.2.3 Case Study 6 – Reliance JIO

Reliance Jio, India's largest telecom, implemented machine unlearning in its Open Telecom AI platform amidst immense subscriber growth and evolving digital privacy requirements (RCR Wireless, 2025). The platform integrates agentic AI, LLMs, and SLMs for network optimization, service personalization, and predictive maintenance. Multi-domain

orchestrated unlearning ensures privacy across diverse use cases including network, enterprise, and retail operations. Federated learning techniques balance data removal with preservation of aggregate network intelligence, processed using JioBrain's APIs and advanced automation (Jio AI Strategy, 2025). Response efficiency for data removal improved by 97.8%, with audit trails and real-time operations maintained across nearly 500 million subscribers. Network quality, customer satisfaction, and operational metrics remained industry-leading. Major challenges involved interconnected data dependencies and regulatory adaptation. The process reinforced the importance of integrating privacy controls and scalable AI within telecom infrastructures from the outset, ensuring compliance and competitive edge.

4. DISCUSSION & IMPLICATIONS

The enterprise case study analysis provides compelling evidence that machine unlearning is moving from experimental research into mainstream operational environments, fundamentally transforming AI governance, privacy management, and competitive strategy. The evidence points to several cross-sectoral implications for technology leaders, business owners, and solution designers. Key findings are as follows.

- **Privacy-by-Design Integration:** All enterprises reported significant gains by embedding unlearning frameworks directly into AI development pipelines. Retrofitting privacy features into legacy models was consistently more complex, cost-intensive, and less effective (Yin, 2018; Marculescu et al., 2024).
- **Regulatory Compliance Acceleration:** Machine unlearning resulted in dramatic reductions in compliance response times, with organizations processing requests 95% faster on average versus traditional model retraining. Enhanced audit trails and verification protocols further improved regulatory trust (Crowe et al., 2011).
- **Performance Preservation:** Across highly regulated domains such as healthcare, banking, and e-commerce, machine unlearning approaches maintained core model accuracy above 98%, demonstrating technical feasibility without adverse operational impact (Microsoft Learn, 2023; Staufer, 2025).
- **Cost Optimization:** Enterprises realized multi-million-dollar annual savings by eliminating expensive retraining and manual compliance labor, enabling reinvestment in AI innovation

and improving overall efficiency (Amazon Science, 2025).

- **Trust and Engagement:** Customer and merchant satisfaction scores improved across all organizations, directly correlating with privacy transparency and efficient data handling.

Implications for industry practitioners are as follows.

- AI solution architects should prioritize privacy-native architectures and leverage federated and domain-specific frameworks for scalability (Kumar et al., 2023).
- Regulatory and legal teams benefit from collaborations with AI engineers to ensure auditable, transparent unlearning processes that surpass minimum compliance thresholds (European Data Protection Supervisor, 2022).
- Business leaders can anchor trust and market differentiation on demonstrable privacy capabilities, translating technical innovation into customer loyalty and regulatory goodwill.

Collectively, machine unlearning emerges as a strategic enabler for scalable, ethical, and responsible AI, shaping the future of enterprise innovation and governance (Rashid et al., 2019).

5. CONCLUSION

Machine unlearning serves as a practical approach for organizations aiming to build AI that isn't just powerful, but also ethical, adaptable, and compliant. Put simply, it empowers a system to forget specific information if needed—whether to honor user requests, comply with shifting laws, or ensure fairness and transparency in operations. This ability turns machine unlearning into more than a technical fix; it becomes a cornerstone for responsible innovation. It has emerged as an important development in applied AI field for small and large enterprises, enabling organizations to comply to different privacy regulations, enhance trust, and drive operational efficiency by efficiently eliminating sensitive data from deployed models. The case study analysis from this paper demonstrated that both global and Indian enterprises benefit from embedding privacy-by-design within their machine learning pipelines, realizing compliance response improvements by over 95%, substantial multi-million-dollar savings, and preservation of model performance above 98% accuracy. Sector-specific implementations reveal that scalable unlearning solutions—rooted in federated, domain-aware, and real-time architectures—are effective across diverse

environments including banking, e-commerce, healthcare, fintech, and telecommunications.

Adopting machine unlearning as a foundational capability transforms AI governance, supports evolving regulatory needs, and boosts competitive differentiation through increased stakeholder trust. Key learnings for industry include prioritizing privacy-native system design, fostering cross-functional collaboration between legal, business, and AI teams, and proactively addressing technical and operational challenges such as data dependency and auditability. Ultimately, machine unlearning is positioned as a strategic enabler, empowering enterprises to deliver responsible, ethical, and innovative AI solutions at scale.

Acknowledgement

First author would like to acknowledge the support of Artificial Intelligence and Innovation (AI&I) Lab at UAT Mexico led by Prof. Fernando in carrying out innovative projects in the area of Artificial Intelligence and Machine Learning. The Authors would also like to acknowledge ethical and conscious usage of AI tools for content enrichment and enhancement of this paper.

Conflict of Interest

The authors have no conflict of interest to declare.

Funding

Not Applicable.

REFERENCES

- AIM Media House. (2025). The need for machine unlearning in enterprise AI applications. AI Media House Councils. Retrieved from <https://councils.aimediahouse.com/the-need-for-machine-unlearning-in-enterprise-ai-applications/>
- Crowe, S., Cresswell, K., Robertson, A., Huby, G., Avery, A., & Sheikh, A. (2011). The case study approach. *BMC Medical Research Methodology*, 11(1), 100.
- European Data Protection Supervisor. (2022). Machine unlearning. EDPS Technology Monitoring. Retrieved from https://www.edps.europa.eu/data-protection/technology-monitoring/techsonar/machine-unlearning_en
- Jio AI Strategy Analysis. (2025). Reliance's AI strategy: Analysis of conglomerate AI dominance. Klover AI Research. Retrieved from <https://www.klover.ai/reliance-ai-strategy-analysis-of-conglomerate-ai-dominance/>
- Kumar, V. B., et al. (2023). Privacy adhering machine unlearning in NLP. Amazon Science Publications. Retrieved from <https://www.amazon.science/publications/privacy-adhering-machine-un-learning-in-nlp>
- Li, G., Hsu, H., Chen, C. F., & Marculescu, R. (2024). Machine unlearning for image-to-image generative models. *Proceedings of the International Conference on Machine Learning*. Retrieved from https://www.reddit.com/r/machinelearningnews/comments/laklazv/this_ai_paper_from_ut_austin_and_jpmorgan_chase/
- Li, Q., Geng, J., Woiseschläger, H., Chen, Z., Cai, F., Wang, Y., Nakov, P., Jacobsen, H. A., & Karray, F. (2025). A survey of machine unlearning in large language models: Methods, challenges and future directions. *arXiv preprint arXiv:2503.01854*. Retrieved from <https://arxiv.org/html/2503.01854v2>
- Liu, K. Z. (2024). Machine unlearning in 2024. *Stanford AI Lab Blog*. Retrieved from <https://ai.stanford.edu/~kzliu/blog/unlearning>
- Marculescu, R., et al. (2024). Machine 'unlearning' helps generative AI 'forget' copyright-protected and violent content. *University of Texas at Austin Research*. Retrieved from <https://cockrell.utexas.edu/news/machine-unlearning-helps-generative-ai-forget-copyright-protected-and-violent-content/>
- Meta Research Team. (2022). MPC-based machine learning: Achieving end-to-end privacy preserving machine learning. *Facebook Research Blog*. Retrieved from <https://research.facebook.com/blog/2022/10/mpc-based-machine-learning-achieving-end-to-end-privacy-preserving-machine-learning/>
- Microsoft Learn. (2023). FastTrack for Azure Season 2 Ep01: Azure ML Enterprise Deployment. *Microsoft Documentation*. Retrieved from <https://learn.microsoft.com/en-us/shows/learn-live/fasttrack-for-azure-season-2-ep01-azure-ml-enterprise-deployment>
- Paytm Business Transformation. (2025). Paytm's Great Escape. *Lapaas Business Labs*. Retrieved from <https://lapaas.com/bizlabs/paytm-profit-turnaround-case-study>

- Rashid, Y., Rashid, A., Warraich, M. A., Sabir, S. S., & Waseem, A. (2019). Case study method: A step-by-step guide for business researchers. *International Journal of Qualitative Methods*, 18, 1609406919862424. Retrieved from <https://journals.sagepub.com/doi/full/10.1177/1609406919862424>
- RCR Wireless News. (2025). Jio Platforms and partners to develop Open Telecom AI Platform. RCR Wireless Communications. Retrieved from <https://www.rcrwireless.com/20250306/featured/jio-open-telecom-ai-platform>
- Shaik, T., Tao, X., Xie, H., Li, L., Zhu, X., & Li, Q. (2023). Exploring the landscape of machine unlearning: A comprehensive survey and taxonomy. arXiv preprint arXiv:2305.06360. Retrieved from <https://arxiv.org/abs/2305.06360>
- SpicyIP. (2025). ANI v. Open AI: Time to talk about 'machine unlearning'. SpicyIP Legal Analysis. Retrieved from <https://spicyip.com/2025/01/ani-v-open-ai-time-to-talk-about-machine-unlearning.html>
- State Bank of India. (2023). SBI to use AI/ML to improve customer experience and operations. TechCircle India. Retrieved from <https://www.techcircle.in/2023/06/06/sbi-to-use-ai-ml-to-improve-customer-experience-and-operations>
- Staufer, D. (2025). What should LLMs forget? Quantifying personal data in LLMs for right-to-be-forgotten requests. arXiv preprint arXiv:2507.11128. Retrieved from <https://arxiv.org/abs/2507.11128>
- Times of India. (2025). SBI to deploy 'agentic AI' for customer service. Times of India Business. Retrieved from <https://timesofindia.indiatimes.com/business/india-business/sbi-to-deploy-agentic-ai-for-customer-service/articleshow/119951201.cms>
- University of Texas at Austin. (2024). Machine 'unlearning' helps generative AI 'forget' copyright-protected and violent content. UT News. Retrieved from <https://news.utexas.edu/2024/03/21/machine-unlearning-helps-generative-ai-forget-copyright-protected-and-violent-content/>
- Xu, J., Xiao, Z., Zhang, Y., Wang, J., & Li, X. (2023). Machine unlearning: Solutions and challenges. arXiv preprint arXiv:2308.07061. Retrieved from <https://arxiv.org/pdf/2308.07061>
- Yan, H., Li, X., Guo, Z., Li, H., Lin, F., & Chen, X. (2022). An efficient architecture for exact machine unlearning. Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence. Retrieved from <https://www.ijcai.org/proceedings/2022/0556.pdf>
- Yin, R. K. (2018). Case study research and applications: Design and methods (6th ed.). SAGE Publications. Retrieved from <https://ebooks.umu.ac.ug/librarian/books-file/Case%20Study%20Research%20and%20Applications.pdf>