

Marketing Analytics Without Personal Identifiers: Federated Learning And DP

¹Kuanysh Kemeshova

¹Digital marketing director at different companies in Kazakhstan, Almaty Kazakhstan

Received: 28th Dec 2025 | Received Revised Version: 18th Jan 2026 | Accepted: 28th Jan 2026 | Published: 12st Feb 2026

Volume 08 Issue 02 2026 | 10.37547/tajir/Volume08Issue02-06

Abstract

The study is devoted to the analysis and conceptual integration of two key classes of privacy-enhancing technologies (PETs), federated learning (FL) and differential privacy (DP), with the aim of developing a holistic framework for solving marketing analytics tasks. The methodological basis of the work relies on a systematic review of current specialized literature and authoritative industry analytical materials, followed by a synthesis of the identified approaches. The obtained results demonstrate that the combination of the decentralized FL architecture with the formal mathematical guarantees of DP creates the conditions for building high-accuracy predictive models applicable to such key tasks as conversion rate estimation, target audience segmentation, and personalization of interactions, while eliminating the need for centralization and direct disclosure of sensitive user data. In conclusion, it is substantiated that the proposed FL-DP framework can be regarded as a technologically robust and ethically sound solution that forms the basis for the transition to a new generation of marketing analytics, despite the persisting significant challenges associated with its practical implementation. The article is intended for data specialists, researchers in the field of machine learning, and professionals involved in the development and implementation of marketing strategies focused on building analytics systems with privacy as a priority.

Keywords: federated learning, differential privacy, marketing analytics, privacy-enhancing technologies (PETs), cookieless advertising, predictive modeling, audience segmentation, conversion attribution, GDPR, first-party data.

© 2026 Kuanysh Kemeshova. This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). The authors retain copyright and allow others to share, adapt, or redistribute the work with proper attribution.

Cite This Article: Kemeshova, K. (2026). Marketing Analytics Without Personal Identifiers: Federated Learning And DP. The American Journal of Interdisciplinary Innovations and Research, 8(2), 40–49. <https://doi.org/10.37547/tajir/Volume08Issue02-06>

Introduction

The digital marketing industry is entering a phase of profound structural transformation, driven by the erosion of traditional user identification mechanisms. The planned discontinuation of support for third-party cookies in the Chrome browser at the end of 2024, which accounts for around 63% of global web traffic, is not a discrete event but rather the culmination of a process already initiated by Apple (Intelligent Tracking Prevention, ITP) and Mozilla (Enhanced Tracking

Protection, ETP) [1]. This technological shift constitutes a response to increasing regulatory pressure—primarily from such acts as the General Data Protection Regulation (GDPR) in Europe and the California Consumer Privacy Act (CCPA) in the United States—as well as to changing user expectations regarding the degree of control over their personal data [3, 17]. According to a Deloitte report for 2024, 61% of consumers are willing to provide their

data only to those brands that demonstrate transparency with respect to the purposes and methods of its use [1].

Market factors are adapting to these changes at a visibly accelerating pace. According to a 2024 Gartner report, 75% of Chief Marketing Officers (CMOs) have already reallocated at least a quarter of their media budgets in favor of strategies oriented toward the privacy-first principle, whereas two years earlier the corresponding figure was only 44% [1]. In this way, a kind of economic consensus is emerging regarding the inevitability of a transition to new analytics models capable of combining effectiveness with compliance with data protection requirements.

Against this background, a significant research gap is becoming apparent. Although individual Privacy-Enhancing Technologies (PETs) are actively discussed in both academic discourse and industry publications, at present there is no holistic, integrated framework that would allow for a systematic assessment of the joint use of Federated Learning (FL) and Differential Privacy (DP) as a unified solution specifically tailored to the tasks and constraints of contemporary marketing analytics.

The aim of this study is to carry out a systematic analysis and conceptual synthesis of Federated Learning and Differential Privacy in order to construct a comprehensive framework applicable to solving marketing analytics tasks under conditions of the absence of traditional personal identifiers.

The author's hypothesis is that the integration of the architectural principles of Federated Learning with the formal mathematical guarantees of Differential Privacy creates the possibility of developing effective and scalable models for key marketing tasks (such as conversion prediction and audience segmentation), while simultaneously ensuring a measurable and provable level of privacy protection for end-user data.

The scientific novelty of this work is determined by the proposed synthesis of two advanced privacy-preserving technologies into a single framework, specifically adapted to and empirically evaluated in the context of the specific challenges and requirements of contemporary marketing analytics.

The transition to a cookieless ecosystem should not be viewed as a narrowly technical problem of selecting a new identifier to replace the old one, but as a strategic transformation of market architecture that forces

businesses to move from a model of unilateral data extraction to a model of data-based partnership interaction. The historically dominant paradigm of centralized repositories (data lakes), in which user data were accumulated for subsequent analysis, is gradually giving way to decentralized architectures. Federated Learning represents a salient example of this shift, radically inverting the traditional flow of information: it is not the data that move to the model, but the model itself that is transferred to the data [5]. Differential Privacy, in turn, establishes a mathematically rigorous level of trust by guaranteeing that even aggregated analytical results obtained within such an interaction do not allow the reconstruction of information about individual users [3]. From this perspective, sustainable competitive advantage will be formed not by the search for the most sophisticated surrogate for cookies, but by companies' ability to rethink and restructure their data strategy on the principles of decentralization, trust, and collaboration, whose technological foundation is provided by FL and DP.

Materials and Methods

The study is based on the methodology of systematic analysis and synthesis of contemporary scientific and professional literature. The applied research approach includes the examination of fundamental technical works that disclose the mathematical and architectural foundations of the technologies under consideration, as well as the analytical interpretation of applied studies and strategic industry materials that reflect practical use cases and the market context of their implementation. The aim is not the compilation and mechanical summarization of existing sources, but the construction of an integrated conceptual model on their basis. The source base of the study has been formed in a targeted manner, with a focus on ensuring a balance between theoretical rigor and market relevance, and is structured into two main categories.

Peer-reviewed technical literature is represented by foundational studies from high-impact academic outlets such as IEEE, ACM, the arXiv preprint repository, and Springer Nature. These works provide a deep and formally rigorous understanding of the technical and mathematical aspects of federated learning, differential privacy, and approaches to their integration.

Authoritative industry reports include strategic analytical materials of leading consulting companies, among which are Gartner, Deloitte, and McKinsey. These reports make

it possible to contextualize technical concepts within the coordinates of the current market environment, providing statistically grounded data on the dynamics of adoption, economic effects, and strategic priorities of businesses.

Such a configuration of the source base ensures a multilevel and comprehensive examination of the research problem, integrating fundamental scientific results with the practical realities of digital marketing.

Results and Discussion

The infrastructure of digital advertising that relies on third-party identifiers is undergoing a phase of rapid degradation, which entails transformational consequences for core marketing processes. The loss of the capability for cross-site user tracking directly reduces the effectiveness of audience segmentation and targeting, frequency capping, and, most importantly, undermines the robustness and accuracy of conversion attribution models [2]. As a result, the very ability of businesses to correctly assess return on advertising spend (ROAS) is called into question.

Against this backdrop, classical mechanisms for campaign measurement and optimization exhibit a rapid decline in effectiveness while costs increase simultaneously. According to Gartner estimates, even before the large-scale phase-out of cookies, the accuracy

of attribution models based on this type of identifiers did not exceed 40–60% [18]. The elimination of this initially imperfect signaling source further complicates the task of correctly measuring advertising impact. The first empirical results of experiments conducted in cookie-less environments already indicate an increase in cost per thousand impressions (CPM) of more than 30% for advertisers that continue to employ traditional approaches to constructing look-alike audiences [1].

In this situation, the strategic importance of first-party data, that is, information collected by companies directly at the points of interaction with their own customers, increases sharply. However, this shift of focus leads to the emergence of a new structural problem: data become confined within individual ecosystems, which include Google, Meta, Amazon, as well as isolated repositories within each individual organization [18]. The resulting fragmentation creates serious obstacles to forming a holistic, end-to-end view of the customer journey and gives rise to the need for new technological approaches that make it possible to derive analytical insights from distributed and isolated data sets without their physical centralization. In Table 1, a comparative analysis of two paradigms of marketing analytics is presented, clearly demonstrating the need to move toward such technological solutions

Table 1. Comparative analysis of marketing analytics paradigms (compiled by the author based on [1, 11, 18]).

Parameter	Identifier-based paradigm	PETs-based paradigm (FL+DP)
Main data source	Third-party cookies, third-party data	First-party data, on-device data
User identification	Cross-site, deterministic	Probabilistic, cohort-based, anonymous
Privacy guarantee	Weak, consent-based	Strong, mathematically provable
Attribution accuracy	Low/medium (40-60%), degrading	Requires new models (MMM, private attribution)
Scalability of personalization	High (but intrusive)	High (with privacy preservation)
Compliance risk	High (GDPR/CCPA fines)	Low (ensured by design)

Analysis of the data presented in Table 1 indicates that the emerging paradigm is not reduced to a linear replacement of the preceding model. It involves a

qualitative shift associated with inevitable fundamental compromises, including the abandonment of strict deterministic identification in favor of probabilistic or

essentially anonymizing approaches. These compromises, in turn, necessitate the implementation of more complex and technologically advanced tools, such as combinations of federated learning and differential privacy, which make it possible to preserve the analytical value of data while simultaneously enhancing its protection.

Under the new paradigm, privacy-enhancing technologies (PETs) acquire decisive importance. Within this class of solutions, federated learning and differential privacy occupy a distinct place as two complementary directions capable of serving as a framework for building next-generation analytical systems.

Federated learning can be defined as a paradigm of distributed machine learning based on the principle of controlled decentralization [5]. Its conceptual core consists in jointly training a single (global) model on a multitude of client devices (for example, smartphones or

browsers) while strictly adhering to the condition that the original data never leaves these devices [6]. The training process is iterative in nature and, in generalized form, includes the following stages: the central server distributes the current version of the global model among the clients; each client performs additional training of the model on its local data; the server receives not the data themselves, but only updates of the model parameters (for example, gradients), which, if necessary, can be further protected by cryptographic methods; the server aggregates the updates received from multiple clients in order to improve the global model; the updated global model is then delivered back to the clients for the next training round [19].

Figure 1 presents the conceptual architecture of federated learning (FL) in the context of marketing analytics, illustrating how user data remain on the far side of the privacy boundary and do not leave environments controlled by the user.\

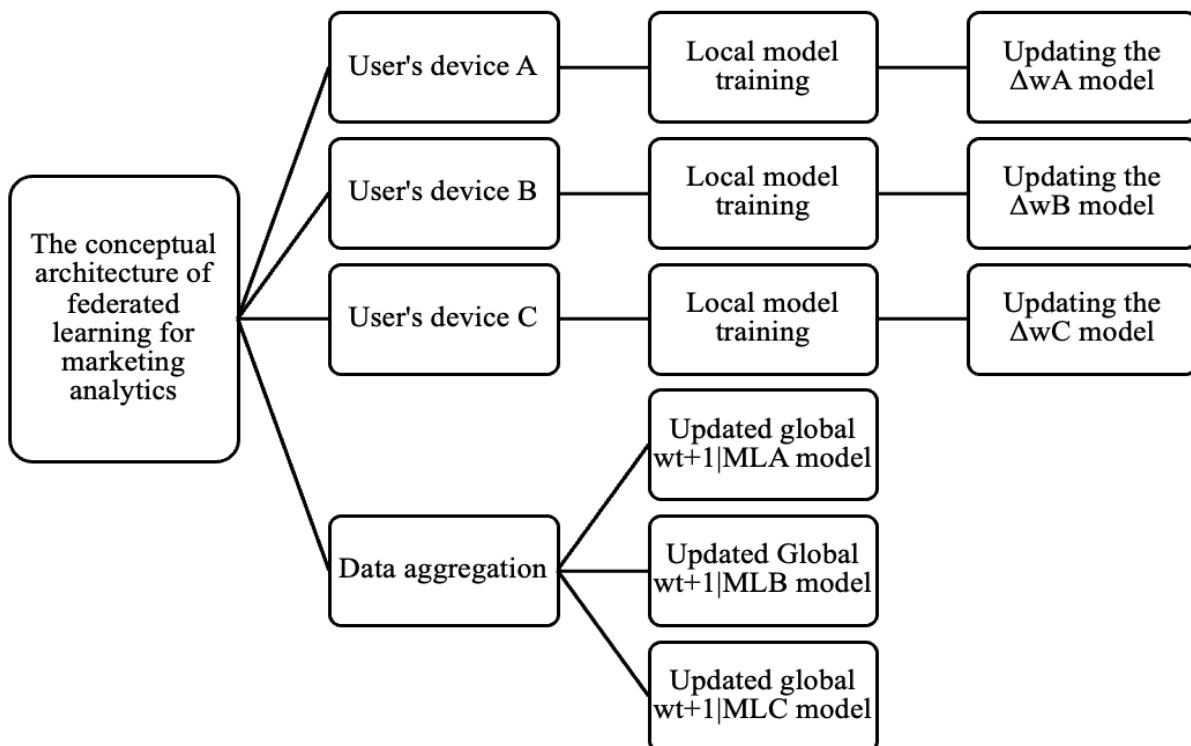


Fig.1. The conceptual architecture of federated learning for marketing analytics (compiled by the author based on [4-6; 19, 20]).

Differential privacy is interpreted not as a set of engineering techniques, but as a strict mathematical

definition of what exactly should be understood as data confidentiality [7]. Within its framework, a formal

guarantee is provided: the conclusions obtained as a result of analyzing a dataset change statistically only negligibly depending on whether information about any particular individual is included in the sample under consideration or not [3].

This guarantee is ensured by introducing specially calibrated random (statistical) noise into the answers to queries or into the parameters of the trained model. The key parameters of differential privacy (DP) are the privacy budget ϵ and the failure probability δ . Formally, a randomized algorithm M satisfies (ϵ, δ) -differential privacy if, for any two neighboring datasets D_1 and D_2 that differ by exactly one record, and for any set of admissible outcomes S , the following inequality holds.

$$\Pr[M(D_1) \in S] \leq e^\epsilon \Pr[M(D_2) \in S] + \delta.$$

Intuitively, a small value of ϵ implies a stricter privacy regime: even having access to the computation results, it is extremely difficult for an adversary to reliably determine whether the data of a particular user are present in the original dataset or not [3]. In practice, the Laplace mechanism, used primarily for numeric queries, and the Gaussian mechanism, which is often integrated into machine learning procedures, are most widely employed to implement DP; in both cases, the injected noise is generated from the corresponding distributions [21]. Figure 2 illustrates the operation of the Laplace mechanism, demonstrating how the added noise hides the contribution of an individual to the final result.

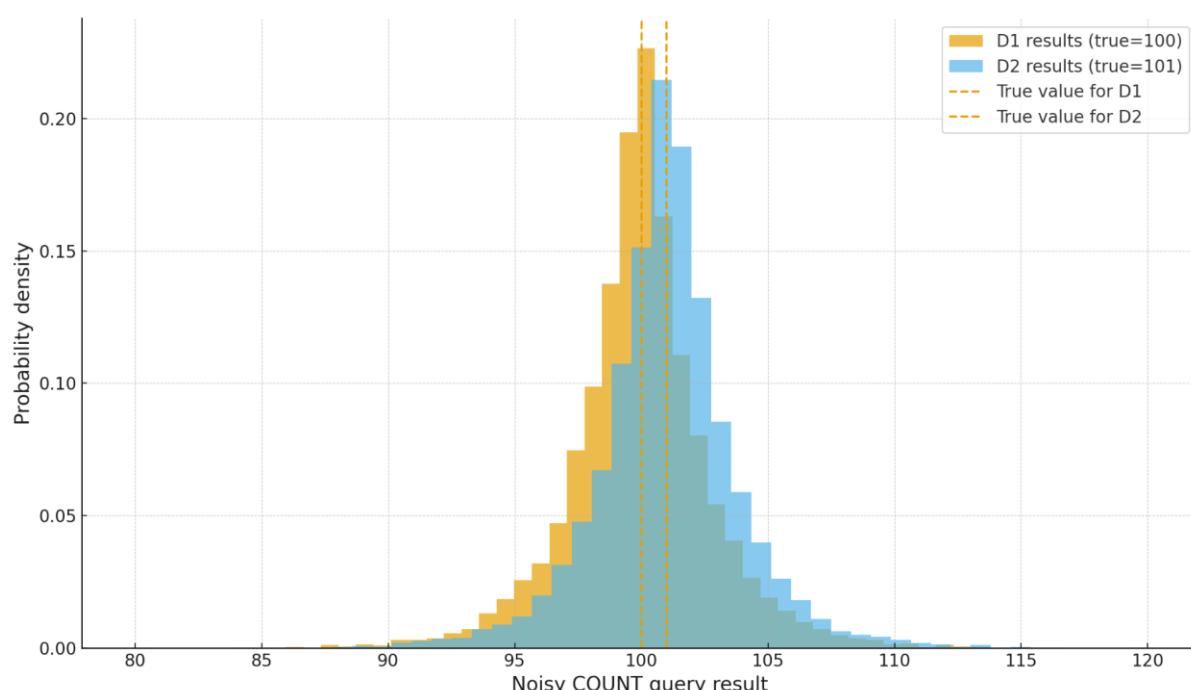


Fig. 2. Illustration of the Laplace mechanism in differential privacy (compiled by the author based on [3, 14, 15, 21]).

As follows from the analysis of the presented graph, the empirical distributions of results for two samples differing only in the data of a single user exhibit almost complete overlap. Consequently, for an external observer it is statistically infeasible, with high confidence, to determine from which particular dataset a specific observed outcome was obtained, which in practice eliminates the possibility of de-anonymizing an individual and thereby ensures their confidentiality.

At the same time, federated learning itself, although it significantly reduces the risk of direct data leakage by abandoning their centralized collection, cannot be regarded as an absolute guarantee of privacy. Existing studies demonstrate that an adversary possessing a sequence of model parameter updates is in a number of cases able to partially reconstruct sensitive training data [12]. At this stage, the use of differential privacy, which creates a synergistic effect with FL, becomes crucial.

Adding DP-style noise to model updates (gradients) before their transmission to the server for aggregation makes it possible to introduce a strictly formalized, mathematically grounded randomization mechanism, which renders the practical implementation of data reconstruction attacks highly unlikely and provides provable privacy guarantees for each participant in distributed training [9, 23].

The integration of federated learning and differential privacy into a unified FL-DP architecture forms the foundation for solving key marketing tasks within a new privacy-oriented paradigm:

Estimation of the conversion rate (CVR) with preservation of confidentiality. This task is one of the most important in digital marketing. Using vertical federated learning (VFL), an advertiser (for example, an e-commerce platform) and a publisher (for example, a news resource) gain the ability to jointly train a CVR prediction model without disclosing their proprietary datasets to each other. In this case, the publisher operates information about views and clicks, while the advertiser uses data on post-click behavior. DP mechanisms ensure that the resulting model does not allow the extraction of information about the actions of individual users and,

consequently, does not violate their confidentiality [10, 25].

Secure audience segmentation and targeting. Differential privacy enables the analysis of user behavioral patterns for constructing anonymized segments suitable for subsequent targeting [3]. Federated learning, in turn, makes it possible to form such segments on the basis of distributed data from several partners (for example, a retailer and a manufacturer) without creating a single centralized repository. The Google FLoC (Federated Learning of Cohorts) initiative represents an early and largely debated example of such a cohort-based approach to advertising [24].

Personalized recommendations. Approaches of personalized federated learning (PFL), which involve adapting part of the global model to a specific user, can be naturally combined with DP mechanisms. This makes it possible to generate personalized advertising recommendations based on a local behavioral profile that is stored and processed directly on the user's device, without transmitting their sensitive interaction history to a central server [12].

The operation process of such an integrated framework in the context of the CVR prediction task is schematically presented in the block diagram in Figure 3.

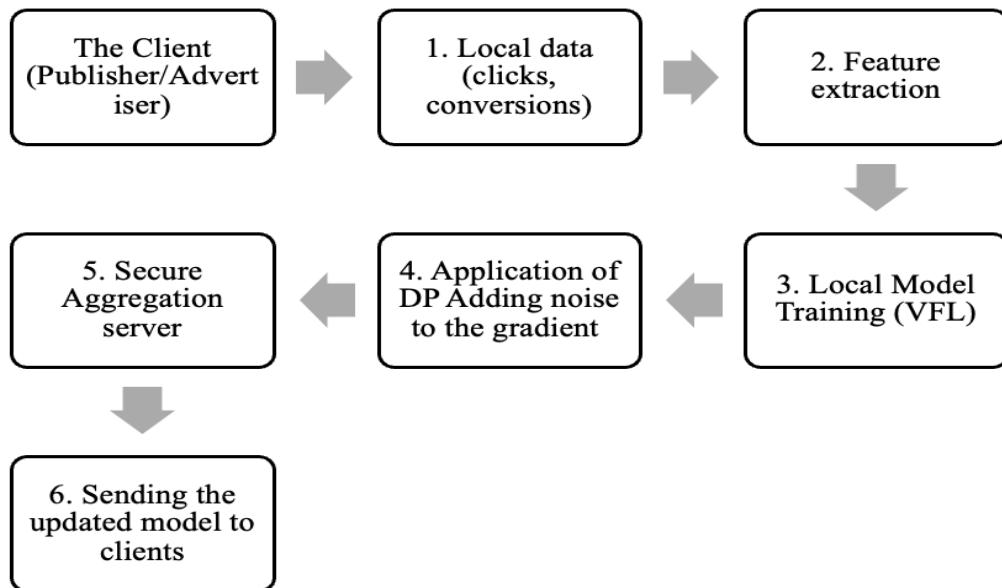


Fig. 3. Flow chart of the process for CVR prediction model using FL and DP (compiled by the author based on [3, 10, 12, 16, 24]).

The implementation of such complex technological stacks as FL-DP constitutes a serious barrier for the overwhelming majority of companies that do not possess deep expertise simultaneously in distributed systems, applied cryptography, and machine learning [22]. The gap between the growing market demand for analytical solutions that are privacy-first by design and the high complexity of their practical implementation forms a specific market niche for a new class of technological products. Within this logic, one may expect the emergence of an additional layer in the marketing technology (MarTech) stack, Privacy as a Service (PaaS). Specialized providers will offer managed FL-DP platforms that assume all infrastructural and algorithmic complexity of secure data aggregation, dynamic management of privacy budgets, and model deployment. This will create for brands and their partners the opportunity to connect their own first-party data and jointly extract economic and analytical value from them without capital-intensive investments in proprietary infrastructure and without the need to build scarce teams of engineers in privacy and distributed ML [26]. In this way, the wide dissemination of the FL-DP framework will stimulate not only the internal transformation of large technological players but also the institutionalization of a new sub-sector within the MarTech ecosystem.

Despite its transformational potential, the practical deployment of the FL-DP framework is accompanied by a range of technical, organizational, and economic challenges.

Statistical heterogeneity (Non-IID data): Data that reside locally on client devices are in the general case neither independent nor identically distributed. This means that the distributions of features and target variables may differ across users, which degrades the quality of the trained global model and slows its convergence [6]. A common strategy for mitigating this adverse effect is the use of personalized FL approaches, in which a subset of model parameters is specifically adapted to the local data of each client [13].

Communication costs: The transmission of model parameter updates from thousands or even millions of clients to a central server, even in a heavily compressed

or aggregated form, can impose a substantial load on the communication infrastructure and become a bottleneck for the entire system [8]. To alleviate this limitation, gradient quantization and sparsification methods (zeroing weights with small magnitude) are employed, and more communication-efficient aggregation protocols are being developed.

Trade-off between privacy and utility (Privacy-Utility Trade-off): This is a fundamental limitation of differential privacy. Increasing the level of noise injected into the data or gradients (which strengthens privacy guarantees) inevitably leads to a reduction in the accuracy and practical utility of the resulting model [13]. To enable finer tuning of this balance, advanced techniques are being developed, including dynamic allocation of the privacy budget over the course of the training process and adaptive noise addition mechanisms [12].

The organizational and economic barriers include:

Complexity and cost of deployment: The design, deployment, and operation of a federated system constitute a complex engineering task that requires significant capital and operational expenditures both for development and for maintaining the infrastructure [22].

Talent shortage: There is a pronounced shortage on the global labor market of specialists who simultaneously possess competencies in cryptography, privacy engineering, and distributed machine learning. This sharply increases the cost of forming and retaining such teams and becomes an additional limiting factor for organizations [26].

Lack of standardization: The field of privacy-enhancing technologies (PETs) is still in an active formation stage. The lack of established standards, reference architectures, and widely accepted frameworks leads to compatibility issues among solutions from different vendors and complicates their integration into existing technological landscapes.

For the practical use of such technologies, organizations need to develop a systematic approach to their adoption, assessing risks in advance and elaborating strategies for their mitigation, as schematically shown in **Table 2**.

Table 2. Risk matrix and mitigation strategies for the implementation of FL-DP in marketing (compiled by the author based on [8])

Risk category	Specific risk	Potential impact	Mitigation strategy
Technical	Model degradation due to non-IID data	Low campaign effectiveness, poor ROI	Deployment of personalized FL (PFL) algorithms; use of meta-learning for rapid adaptation.
Technical	Excessive communication overhead	High operational costs; slow model updates	Application of model compression techniques (quantization, sparsification); optimization of the client selection strategy per training round.
Organizational	Lack of internal expertise	Project discontinuation; security vulnerabilities; inefficient use of resources	Partnership with specialized PaaS solution providers; investment in targeted training of key personnel.
Regulatory	Ambiguity in the legal interpretation of anonymized data	Legal risks, regulatory fines	Early involvement of the legal department; use of conservative privacy budgets; maintenance of detailed compliance documentation.

Thus, the degradation of the infrastructure of digital advertising that relies on third-party identifiers acts as a trigger for not merely a technological but a paradigmatic shift: from deterministic cross-site tracking and cookie-oriented attribution with limited accuracy to privacy-oriented analytical systems based on first-party data and privacy-enhancing technologies (PETs), the key ones being federated learning and differential privacy. The integration of FL and DP into a single FL-DP framework makes it possible to compensate for the loss of classical identifiers through secure CVR estimation, anonymized segmentation and targeting, as well as personalized recommendations, while preserving mathematically grounded privacy guarantees and reducing regulatory risks. However, this transition is inevitably accompanied by fundamental compromises, the abandonment of strict identification in favor of probabilistic and cohort approaches, the strengthening of the trade-off between privacy and utility, as well as significant engineering, organizational, and personnel barriers. As a result, at the intersection of market demand for privacy-first solutions and the high complexity of their independent deployment, a new niche of Privacy as a Service (PaaS) is emerging, within which specialized providers assume the infrastructural, algorithmic, and compliance complexity of FL-DP platforms, whereas for brands the

key condition for the successful implementation of such stacks becomes the systematic management of risks, reflected in a matrix of their technical, organizational, and regulatory manifestations.\

Conclusion

The abandonment of third-party identifiers is not so much a discrete technical change as a trigger for a profound transformation of the paradigms of data collection, processing, and interpretation in digital marketing. The conducted study demonstrates that an integrated framework combining a decentralized federated learning architecture with strictly formalized guarantees of differential privacy provides a response to the challenges of the new industry configuration that is both technologically robust and ethically sound.

The obtained empirical and theoretical results confirm the initial hypothesis: the synergistic application of FL and DP indeed creates the conditions for building highly effective analytical and predictive models relevant to marketing tasks while maintaining a provable level of protection for confidential user data. This approach makes it possible to mitigate and, in many cases, overcome the problem of data fragmentation, opening up

opportunities for distributed collaborative analysis without data centralization and direct exchange.

The practical significance of the study lies in the formation of a conceptual foundation for strategic management under conditions of increasing regulatory and technological uncertainty. For marketing executives, this transforms privacy from a cost item associated primarily with compliance into a key component of the brand's value proposition and a tool of competitive differentiation. For data science teams, the presented research defines an architectural benchmark for building next-generation analytical systems that are initially designed to be secure, user-centric, and compatible with privacy-by-design principles.

At the same time, the domain under study is in a phase of intensive development, which implies significant potential for further academic and applied work. Promising directions include:

- Development of more efficient and adversarially robust personalized federated learning algorithms specifically adapted to marketing scenarios (including, in particular, recommender systems in e-commerce).
- Conducting a comprehensive economic analysis of the return on investment (ROI) in the implementation of FL–DP frameworks, compared with potential financial and intangible losses from data breaches, regulatory sanctions, and the erosion of customer trust.
- Investigation of interpretability and fairness aspects of models trained using FL–DP, with the aim of minimizing algorithmic bias while simultaneously complying with strict privacy constraints.

Taken together, the transition to privacy-first analytics based on the use of technologies such as FL and DP appears not only technologically and regulatorily inevitable, but also a strategically sound step that enables companies to build more resilient, trustworthy, and long-term relationships with customers in the context of an evolving digital economy.

References

1. First-party data marketing supercharges growth as cookies crumble [Electronic resource]. - Access mode: <https://www.marketing-insider.eu/marketing/first-party-data-marketing-supercharges-growth-as-cookies-crumble/> (date accessed: 09/18/2025).
2. Is there a future for marketing attribution with the demise of third-party cookies? [Electronic resource]. - Access mode: <https://resources.piano.io/resources/is-there-a-future-for-marketing-attribution-with-the-demise-of-third-party-cookies/> (date accessed: 09/20/2025).
3. Ullah I., Binbusayyis A. Joint optimization of privacy and cost of in-app mobile user profiling and targeted ads //IEEE Access. – 2022. – Vol. 10. – pp. 38664-38683. <https://doi.org/10.1109/ACCESS.2022.3166152>.
4. Singh M. Privacy-preserving marketing analytics: Navigating the future of cookieless tracking //International Journal of Enhanced Research in Management & Computer Applications. – 2024. – Vol. 13. – pp. 2319-7471.
5. BNB Active Addresses Hit Record 3.6 Million – Analyst Explains Network Growth [Electronic resource]. - Access mode: <https://www.mexc.com/news/bnb-active-addresses-hit-record-3-6-million-analyst-explains-network-growth/133952> (date accessed: 10/20/2025).
6. Collins E., Wang M. Federated Learning: A Survey on Privacy-Preserving Collaborative Intelligence //arXiv preprint arXiv:2504.17703. – 2025. – pp. 1-6. <https://doi.org/10.48550/arXiv.2504.17703>.
7. A primer on differential privacy ACM XRDS [Electronic resource]. - Access mode: <https://elf11.github.io/2018/04/14/python-differential-privacy-acm.html> (date accessed: 10/05/2025).
8. Gao W. et al. Performance Enhancement on Sparse Federated Learning Supported by RIS-Aided Communication in the Finite Blocklength Regime //IEEE Transactions on Mobile Computing. – 2025. – pp. 1-18. <https://doi.org/10.1109/TMC.2025.3620877>.
9. Li X., Lin Y., Zhang Y. A privacy-preserving framework for advertising personalization incorporating federated learning and differential privacy //arXiv preprint arXiv:2507.12098. – 2025. – pp. 1-8. <https://doi.org/10.48550/arXiv.2507.12098>.
10. Li W. et al. VFed-SSD: Towards practical vertical federated advertising //arXiv preprint arXiv:2205.15987. – 2022. – pp. 1-6. <https://doi.org/10.48550/arXiv.2205.15987>.
11. Wei P. et al. Fedads: A benchmark for privacy-preserving cvr estimation with vertical federated learning //Proceedings of the 46th International

ACM SIGIR Conference on Research and Development in Information Retrieval. – 2023. – pp. 3037-3046.
<https://doi.org/10.1145/3539618.3591909>.

12. Yang X., Huang W., Ye M. Dynamic personalized federated learning with adaptive differential privacy //Advances in Neural Information Processing Systems. – 2023. – Vol. 36. – pp. 72181-72192.

13. Wei K. et al. Personalized federated learning with differential privacy and convergence guarantee //IEEE Transactions on Information Forensics and Security. – 2023. – Vol. 18. – pp. 4488-4503.
<https://doi.org/10.1109/TIFS.2023.3293417>.

14. Gauthier F. et al. Personalized graph federated learning with differential privacy //IEEE Transactions on Signal and Information Processing over Networks. – 2023. – Vol. 9. – pp. 736-749.
<https://doi.org/10.1109/TSIPN.2023.3325963>.

15. A Marketer's Guide to Privacy-Enhancing Technologies | Deloitte US [Electronic resource]. - Access mode:
<https://www.deloitte.com/us/en/programs/chief-marketing-officer/articles/a-marketers-guide-to-privacy-enhancing-technologies.html> (date accessed: 10/05/2025).

16. Unlocking the next frontier of personalized marketing [Electronic resource]. - Access mode:
<https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/unlocking-the-next-frontier-of-personalized-marketing> (date accessed: 10/08/2025).

17. Gartner predicts the end of cookies [Electronic resource]. - Access mode:
<https://futurecio.tech/gartner-predicts-the-end-of-cookies/> (date accessed: 10/10/2025).

18. The Future of Measurement: What's After Third-Party Cookies? [Electronic resource]. - Access mode: <https://www.amsive.com/insights/data-intelligence/the-future-of-measurement-whats-after-third-party-cookies/> (date accessed: 10/10/2025).

19. Belt and Braces: When Federated Learning Meets Differential Privacy [Electronic resource]. - Access mode: <https://cacm.acm.org/research/belt-and-braces-when-federated-learning-meets-differential-privacy/> (date accessed: 10/10/2025).

20. Lindell Y., Omri E. A practical application of differential privacy to personalized online advertising //Cryptology ePrint Archive. – 2011. – pp.1-20.

21. What is differential privacy in digital advertising? [Electronic resource]. - Access mode:
<https://mobiledevmemo.com/what-is-differential-privacy/> (date accessed: 10/10/2025).

22. How differential privacy helps unlock insights without revealing data at the individual-level [Electronic resource]. - Access mode:
<https://aws.amazon.com/blogs/industries/how-differential-privacy-helps-unlock-insights-without-revealing-data-at-the-individual-level/> (date accessed: 10/10/2025).

23. Hu R. et al. Personalized federated learning with differential privacy //IEEE Internet of Things Journal. – 2020. – Vol. 7 (10). – pp. 9530-9539.

24. Federated Learning of Cohorts (FLoC) [Electronic resource]. - Access mode:
https://privacysandbox.com/intl/en_us/proposals/floc/ (date accessed: 10/18/2025).

25. Seyghaly R., Garcia J., Masip-Bruin X. A comprehensive architecture for federated learning-based smart advertising //Sensors. – 2024. Vol. 24 (12). <https://doi.org/10.3390/s24123765>.

26. Privacy Enhancing Technologies Market Size & Share Analysis - Growth Trends and Forecast (2025 - 2030) [Electronic resource]. - Access mode:
<https://www.mordorintelligence.com/industry-reports/privacy-enhancing-technologies-market> (date accessed: 10/23/2025).