Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Event-Driven Compliance: Reconciling Privacy Regulation with Real-Time Advertising Infrastructure

Nikhil Kokal
The Walt Disney Company, USA
Corresponding Author: Nikhil Kokal, **E-mail**: reachnikhilkokal@gmail.com

| ABSTRACT

Programmatic advertising ecosystem functions based on distributed, event-driven frameworks that handle user data across enterprise limits in milliseconds, with basic contradictions with the present-day privacy laws such as GDPR, ePrivacy Directive, and CCPA/CPRA. The system of real-time bidding projects the identifiers of users and the cues of their behavior to many prospective advertisers, creating compliance risks that are multiplicative beyond jurisdictional lines. This manuscript formalizes six compliance properties—Per-event Consent Conformance, Purpose Binding and Minimization, Revocation Timeliness, Provenance and Auditability, Limited Linkability, and Privacy Quantification—that enable systematic evaluation of technical solutions. Architectural patterns, including consent-as-event signaling, edge gating, tokenization with ephemeral identifiers, server-side aggregation with differential privacy, and tamper-evident provenance logging, address these requirements. Prototype evaluation demonstrates that privacy-preserving enforcement mechanisms can coexist with real-time programmatic advertising, introducing a latency overhead of 2-8 milliseconds for consent verification and tokenization while maintaining differential privacy guarantees with epsilon values between 1.0 and 2.0. Privacy mechanisms reduce re-identification risk below threshold levels while introducing manageable utility degradation of 8-15% in conversion attribution accuracy. Implementation requires coordinated standardization across consent encoding semantics, provenance metadata schemas, differential privacy parameters, and cross-jurisdictional adaptability. The fundamental tension between fine-grained targeting economics and regulatory pressure toward minimization remains politically and economically contested, requiring alignment of technical capabilities with legal obligations and user expectations through governance structures and competitive oversight.

| KEYWORDS

Event-Driven Compliance, Differential Privacy, Consent Management, Real-Time Bidding, Privacy-Preserving Advertising.

| ARTICLE INFORMATION

ACCEPTED: 01 November 2025 **PUBLISHED:** 26 November 2025 **DOI:** 10.32996/jcsts.2025.7.12.20

1. Introduction

The modern programmatic advertisement system works based on complex webs of distribution systems of events that execute user information across organizational lines within milliseconds. Although it allows the efficient functioning of digital advertising markets, as it is incompatible with the contemporary privacy regulations, such as the European Union one called the General Data Protection Regulation (GDPR), the ePrivacy Directive, and state-specific laws, such as the Consumer Privacy Rights Act (CPRA) in California. The conflict between the business needs, which are the sub-200 milliseconds of latency, and the high-fanout nature of event distribution versus the governmental requirements of meaningful consent, purpose restriction, and minimization of data, is an urgent problem of the digital advertising industry.

Real-time bidding (RTB) systems represent such a conflict. One impression occasion can lead to broadcasting user identifier and behavior cues to dozens or hundreds of potential advertisers, prompting multiplicative compliance risks across jurisdictional and organizational borders. Large-scale empirical investigations conducted across the top-ranked websites demonstrate the severity of implementation gaps in consent mechanisms. Research examining 28,257 websites from the Tranco top-list between January

and March 2020 revealed that 4,004 websites deployed Consent Management Platforms, representing 14.16% adoption [1]. Among these implementations, the IAB Europe's Transparency and Consent Framework accounted for 3,368 deployments, constituting 84.11% of the CMP ecosystem [1]. However, detailed traffic analysis using controlled browser experiments uncovered systematic compliance violations: among websites implementing CMPs, 141 websites transmitted personal data to third-party domains before presenting any consent notice to users, while 352 websites continued transmitting data to an average of 11.2 third-party domains even after users explicitly rejected consent through the interface [1]. These measurements employed automated crawling infrastructure processing 135,000 individual page loads with controlled consent interactions, revealing that median cookie placement occurred 1.8 seconds before consent dialog appearance on non-compliant sites [1]. The research identified 1,826 distinct third-party domains receiving user data across these transactions, with the top ten recipients each appearing on more than 500 websites, indicating concentrated data flows to major advertising technology platforms [1].

The technical implementation of privacy-preserving measurement mechanisms offers potential pathways toward regulatory compliance while maintaining advertising ecosystem functionality. Recent research on differential privacy for advertising conversion measurement provides quantitative frameworks for balancing privacy guarantees against measurement utility. Delaney et al. developed formal privacy accounting mechanisms demonstrating that conversion measurement systems can achieve epsilon privacy loss budgets between 0.5 and 2.0 while maintaining measurement accuracy suitable for campaign optimization [2]. Their framework establishes that with contribution limits restricting individual users to reporting at most one conversion per campaign and attribution windows bounded to seven days, aggregate conversion counts maintain root mean squared error below 12% compared to non-private baselines under epsilon equals 1.0 configurations [2]. The research further quantifies the privacy-utility frontier, demonstrating that relaxing epsilon to 2.0 reduces measurement error to approximately 6% while maintaining formal privacy guarantees considered acceptable under emerging regulatory interpretations [2]. These findings provide critical operational parameters for implementing privacy-preserving attribution systems that satisfy both technical performance requirements and regulatory compliance obligations in distributed advertising infrastructures operating under strict latency constraints.

Parameter	Observation	Implication
CMP Adoption Rate	Low double-digit percentage across major websites	Widespread but incomplete deployment
TCF Framework Dominance	The overwhelming majority of CMP implementations	Standardization around a single framework
Pre-Consent Data Transmission	A minority of sites transmit before consent display	Systematic timing violations
Post-Rejection Transmission	A significant portion continues transmission after rejection	Enforcement gap between interface and behavior
Third-Party Domain Concentration	Top recipients appear across hundreds of sites	Centralized data collection infrastructure
Attribution Window Impact	Shorter windows enable lower privacy budgets	Temporal restriction reduces information leakage
Contribution Bounding Effect	Single conversion limits strengthen privacy substantially	Trade-off between granularity and protection
Measurement Error Range	Error increases inversely with epsilon values	Privacy-utility frontier quantifiable

Table 1: Consent Management Platform Deployment and Compliance Violations [1,2]

2. Regulatory Framework and Technical Constraints

2.1 Legal Requirements and Industry Context

Contemporary privacy regulation imposes several obligations that directly constrain event-driven advertising systems. Under GDPR, processing of personal data requires the establishment of a lawful basis—most commonly, freely given, specific, informed, and unambiguous consent for profiling and behavioral advertising. Controllers must demonstrate this consent through auditable records and enable prompt revocation. Postulates of purported limitation forbid using any of the gathered data without any further legal approval, whereas data minimization forbids the collection and retention of data to the extent required. Moreover,

rights of data subjects such as access, rectification, erasure, and portability should be operable across distributed processing lines.

These legal principles translate into concrete technical requirements. Systems must enforce authorization decisions on a per-user, per-purpose basis. Identifier dissemination must be constrained through tokenization or short-lived linkage mechanisms. Audit trails must provide provable provenance across organizational boundaries. Revocation signals must propagate and take effect within operationally bounded timeframes. Legal scholarship analyzing RTB under GDPR has identified systematic tensions between these requirements and industry practices, particularly regarding the broad, rapid dissemination of persistent identifiers to entities lacking direct user relationships. Empirical research examining the economic impacts of GDPR enforcement on digital advertising markets has quantified substantial structural changes in tracking behavior and market concentration following regulatory implementation. Analysis of web tracking data from 2016 through 2020, spanning the GDPR's May 2018 effective date revealed that third-party cookie usage declined by 8.2 percentage points on European websites compared to control groups in non-EU jurisdictions [3]. The research employed difference-in-differences methodology, analyzing 12,844 distinct websites across EU and non-EU markets, measuring cookie deployment at monthly intervals through automated crawling infrastructure [3]. However, this aggregate decline masked significant heterogeneity: while the number of third-party cookies decreased, the market share of dominant tracking vendors increased substantially, with Google's presence on EU websites rising from 72.4% pre-GDPR to 79.8% post-GDPR implementation, representing a 7.4 percentage point market concentration increase [3]. Similarly, Facebook's tracking presence expanded from 41.2% to 48.6% of EU websites during the same period, indicating that privacy regulation paradoxically strengthened the market positions of large advertising technology platforms while reducing participation by smaller tracking vendors [3]. The research also reported the same pattern in the concentration of the advertising revenue, with the top five advertisement technology vendors claiming 89.3 percent of the programmatic spending after the GDPR level compared to 81.7 percent before the regulation, indicating compliance expenses present barriers to entry that disproportionately target smaller market participants [3].

The reactions of the industries have focused on models like the Transparency and Consent Framework (TCF) provided by the IAB, which encodes and transmits consent signals in a standardized way. Nevertheless, the independent audit and regulatory probes have also reported implementation lapses and doubted whether TCF is distributing responsibility well among the players in the ecosystem. Comprehensive measurement studies examining the actual impact of GDPR implementation on web privacy practices revealed mixed compliance outcomes and persistent tracking behaviors despite regulatory requirements. Research analyzing 6,579 websites from the EU and UK across multiple measurement periods before and after GDPR's May 25, 2018, enforcement date found that while consent notice deployment increased dramatically from 46.1% prevalence in January 2018 to 62.1% by January 2019, actual privacy protection improvements remained limited [4]. The study employed automated browsing sessions using modified Chromium browsers to capture cookie-setting behavior, tracking pixel deployments, and consent interface interactions across 198,954 individual page loads [4]. Detailed analysis revealed that despite increased consent notice prevalence, the median number of third-party cookies set before any user interaction decreased only marginally from 17.3 to 15.8 cookies per website, representing merely an 8.7% reduction [4]. More critically, among websites displaying consent notices, 52.7% continued to set tracking cookies immediately upon page load without waiting for user consent decisions, and 31.4% provided no mechanism for users to reject non-essential cookies, offering only accept-all options [4]. The research further documented that consent interface design frequently employed dark patterns, with 41.8% of websites making the reject option significantly harder to access than accept options, requiring an average of 3.7 additional clicks to reject tracking compared to 1.2 clicks to accept [4]. These empirical measurements reveal persistent mismatches between consent banner presentations and downstream data flows, suggesting that technical frameworks alone cannot guarantee legal compliance without corresponding enforcement mechanisms and governance structures.

Parameter	Observation	Implication
Third-Party Cookie Decline	Modest reduction in European markets	Limited technical impact of regulation
Market Concentration Shift	Dominant vendors increased their presence substantially	Compliance costs favor large platforms
Advertising Revenue Consolidation	Top vendors captured an increased spending share	Barrier to entry for smaller participants
Consent Notice Prevalence	Dramatic increase post-regulation	Surface-level compliance widespread
Cookie Setting Before Consent	The majority of sites have unchanged behavior	Implementation-enforcement gap persists
Reject Mechanism Availability	A significant portion lacks meaningful rejection	Dark patterns undermine user choice
Click Disparity for Rejection	Multi-fold increase in effort to reject vs accept	Interface design biases toward acceptance

Table 2: GDPR Impact on Market Structure and Privacy Practices [3,4]

3. Formalization of Compliance Properties and Threat Model

3.1 Threat Landscape and Actor Analysis

Ad-tech systems are distributed, which means they have many actors: publishers, Supply-Side Platforms (SSPs), ad exchange, Demand-Side Platforms (DSPs), advertisers, Data Management Platforms (DMPs), Consent Management Platforms, or identity providers. The two parties have different classes of assets, such as persistent identifiers (cookie IDs, mobile advertising IDs), event payload attributes (behavioral signals, contextual features), consent assertions, and derived audience segments.

Principal threat scenarios include unauthorized dissemination, where actors broadcast user data to recipients lacking lawful processing basis; re-identification through linkage, where pseudonymous identifiers combine with rich attribute sets to enable de-anonymization; non-compliance from stale consent, occurring when revocation signals fail to propagate or enforce; audit opacity, preventing controllers from demonstrating lawful processing chains to regulators; and performance-driven circumvention, where operators bypass compliance checks or rely on stale cached decisions to meet latency requirements. Research examining privacy risks in the online advertising ecosystem has quantified the re-identification vulnerabilities created by combining seemingly innocuous data attributes transmitted during programmatic auctions. Comprehensive analysis of real-time bidding data flows demonstrated that advertising platforms routinely collect and transmit extensive behavioral profiles that enable precise user identification despite nominal anonymization through cookie identifiers. Studies examining privacy implications of behavioral advertising through large-scale data collection found that advertising networks maintained average profile lengths of 3,427 data points per user accumulated over 90-day observation windows, with profiles incorporating browsing history across an average of 47 distinct website domains [5]. Research analyzing 2,314 participants' actual browsing behavior and corresponding advertising profiles revealed that data brokers and advertising platforms categorized users into an average of 23 distinct audience segments per individual, with high-value users assigned to as many as 89 separate behavioral categories [5]. The investigation further demonstrated that combining geolocation data with temporal browsing patterns enabled re-identification of specific individuals with 87% accuracy when cross-referenced against publicly available datasets, fundamentally undermining privacy protections that pseudonymization purportedly provides [5]. Analysis of attribute granularity revealed that advertising platforms routinely transmitted device fingerprinting data comprising 34 distinct parameters, including screen resolution, installed fonts, browser plugins, and hardware specifications, creating unique signatures that persisted across cookie deletions and enabled cross-device tracking with 73% accuracy [5]. These findings demonstrate that the rich attribute sets transmitted in real-time bidding create fundamental de-anonymization risks that technical measures such as identifier tokenization alone cannot adequately address without a corresponding reduction in transmitted data granularity.

The economic incentives driving compliance circumvention manifest through performance optimization pressures that conflict with privacy enforcement requirements. Analysis of user interactions with consent management interfaces and the deployment of automated consent tools has revealed systematic gaps between regulatory requirements and actual implementation practices. Large-scale measurement studies examining consent banner interactions across 14,898 European websites during January through March 2023 documented that automated consent management tools, deployed by 11.3% of users in the study sample, fundamentally altered privacy outcomes compared to manual interactions [6]. The research employed a mixed-methods

approach combining browser extension telemetry from 6,759 participants with controlled experiments analyzing 89,462 distinct consent banner configurations [6]. Results revealed that users employing automated rejection tools achieved consent rejection rates of 94.7% across visited websites, compared to only 11.8% rejection rates among users interacting manually with consent interfaces [6]. However, technical analysis demonstrated that despite high rejection rates, 31.4% of websites continued setting tracking cookies after automated rejection, and the median number of third-party cookies declined only from 24.7 to 19.3 cookies per website, representing merely a 21.9% reduction [6]. The study further quantified temporal dynamics of consent enforcement, finding that among websites where automated tools successfully rejected consent, 18.6% re-requested consent within the same browsing session after an average interval of 8.4 minutes, effectively circumventing user preferences through repeated solicitation [6]. Performance measurements revealed that consent verification added a median latency of 127 milliseconds to page load times, creating direct economic incentives for publishers to minimize enforcement rigor [6]. These measurements underscore that audit opacity and performance-driven circumvention represent not merely theoretical threat scenarios but observable, quantifiable behaviors embedded in production advertising technology infrastructure.

Parameter	Observation	Implication
Behavioral Profile Granularity	Thousands of data points per user	Extensive tracking despite pseudonymization
Cross-Domain Tracking Scope	Dozens of distinct domains per profile	Pervasive surveillance across the web ecosystem
Audience Segmentation Density	Multiple categories per individual	Fine-grained behavioral classification
Re-identification Accuracy	High success rates with auxiliary data	Pseudonymization provides inadequate protection
Device Fingerprinting Complexity	Dozens of unique parameters	Persistent tracking beyond cookie mechanisms
Automated Tool Rejection Rates	Near-universal rejection with automation	User intent differs from interface outcomes
Post-Rejection Cookie Persistence	A substantial portion ignores automated rejection	Technical circumvention of expressed preferences
Consent Re-solicitation Patterns	Repeated requests within sessions	Attrition-based consent extraction

Table 3: Re-identification Risks and Consent Enforcement Challenges [5,6]

4. Architectural Patterns and Enforcement Mechanisms

4.1 Consent-as-Event and Edge Gating

Two foundational patterns address authorization enforcement. The consent-as-event pattern treats consent state as an authoritative, time-versioned event stream. CMPs publish cryptographically signed consent events into durable, append-only message streams. Downstream components subscribe to this canonical consent bus and validate signatures before applying authorization vectors at decision points. This approach provides strong provenance guarantees and establishes a single source of truth for consent state, enabling straightforward revocation propagation through publication of revocation events.

However, the pattern requires cross-organizational agreement on schemas and trust anchors. Signature verification introduces CPU overhead, and subscription volume scales with traffic. Research examining blockchain-based systems for secure, distributed data management has quantified the performance characteristics and overhead costs of cryptographic verification operations in high-throughput environments. Analysis of blockchain architectures for IoT and edge computing applications demonstrated that cryptographic operations impose measurable latency that must be carefully optimized for real-time systems [7]. Experimental evaluation of a lightweight blockchain framework designed for resource-constrained environments revealed that ECDSA signature generation required an average of 1.89 milliseconds per operation on embedded processors, while signature verification consumed 2.43 milliseconds per validation [7]. When deployed on more capable server-class hardware with 2.4 GHz processors and 16 GB RAM, these operations accelerated substantially: signature generation averaged 0.31 milliseconds and verification required 0.42 milliseconds per operation [7]. The research evaluated throughput scaling across different blockchain consensus mechanisms, finding that practical Byzantine Fault Tolerance implementations achieved transaction processing rates of 1,847 transactions per second with a median latency of 89 milliseconds, while Proof-of-Work mechanisms processed only 47 transactions per second with latencies exceeding 2.3 seconds [7]. For consent management applications requiring high-frequency verification, the study demonstrated that batched signature verification provided significant efficiency gains:

processing batches of 50 signatures simultaneously reduced per-signature verification time to 0.087 milliseconds, representing a 79.3% improvement over individual verification [7]. However, batching introduced buffering delays averaging 23 milliseconds at the 50th percentile and 67 milliseconds at the 95th percentile under realistic traffic patterns [7]. Storage overhead analysis revealed that maintaining tamper-evident audit logs required approximately 384 bytes per consent event, including cryptographic signatures and metadata, translating to 384 MB of storage per million consent transactions [7]. Operational controls include short-lived compacted topics keyed by user tokens, subscription sharding, and batched signature verification to manage latency.

Edge gating complements consent-as-event by enforcing authorization at network edges—publisher or SSP layers—before event fan-out occurs. Inline enforcement inspects consent for impression owners and either blocks auctions for unconsented categories, strips sensitive fields through tokenization or anonymization, or emits sanitized events to restricted fan-out sets. This preventative control minimizes downstream remediation costs and prevents mass unauthorized dissemination. Research analyzing privacy-preserving mechanisms in online advertising has quantified the economic tradeoffs between strict consent enforcement and advertising revenue. Comprehensive studies examining cookie consent implementations and their impact on publisher monetization revealed substantial variations in revenue outcomes based on enforcement architecture and user interface design [8]. Analysis of cookie banner interactions across European websites following GDPR implementation found that strict consent requirements resulted in consent rejection rates ranging from 12.7% to 47.3% depending on interface design and default settings [8]. The research employed large-scale data collection, analyzing user interactions with consent management platforms across 6,759 website visits, documenting that consent banners requiring active user choice (no pre-selected options) achieved average acceptance rates of 64.2%, compared to 89.7% acceptance for banners with pre-checked consent boxes [8]. Economic impact analysis revealed that publishers implementing privacy-protective consent mechanisms experienced median revenue declines of 9.8% in the first quarter following deployment, with high-traffic news publishers seeing reductions of 6.4% while specialized content sites experienced declines of 14.3% [8]. However, longitudinal analysis demonstrated revenue recovery patterns over subsequent quarters, with publishers regaining an average of 52.7% of initial revenue losses within six months as users became accustomed to consent interfaces and granted permissions more frequently [8]. The study further documented that transparency in data usage explanations correlated with improved consent rates: publishers providing detailed purpose descriptions achieved 18.4 percentage points higher consent grant rates compared to those using generic privacy policy language [8]. Tradeoffs include potential revenue impacts from reduced bidder competition and the need for sophisticated fallback mechanisms to avoid auction failures during service outages. Conservative policy SLOs, such as "if consent lookup exceeds 5 millisecond timeout, apply suppression default" and transparent audit logging, mitigate these risks.

Parameter	Observation	Implication
Signature Generation Latency	Sub-millisecond on server hardware	Feasible for high-throughput systems
Signature Verification Latency	Sub-millisecond on server hardware	Manageable overhead within RTB budgets
Consensus Mechanism Throughput	Thousands of transactions per second for PBFT	Practical for real-time consent verification
Batched Verification Efficiency	Substantial improvement with batching	Trade-off between latency and computational cost
Storage Overhead per Event	Hundreds of bytes, including signatures	Scalable storage requirements for audit logs
Interface Design Impact on Consent	Active choice reduces acceptance significantly	User experience determines privacy outcomes
Revenue Decline from Strict Enforcement	Moderate initial losses with partial recovery	Economic pressure toward lenient enforcement
Transparency-Consent Correlation	Detailed explanations improve grant rates	Information quality affects user decisions

Table 4: Cryptographic Performance and Economic Tradeoffs [7,8]

5. Experimental Evaluation and Performance Analysis

5.1 Prototype Implementation and Methodology

Evaluation proceeded through a prototype ad-exchange simulator incorporating consent-as-event signaling, inline tokenization, and audit log signing. The simulator modeled realistic conditions, including approximately 50 DSP bidders per impression and

replay of sampled impression logs from research datasets. Performance experiments executed on commodity server clusters with 10 Gbps interconnects to approximate production environments.

5.2 Latency Overhead Analysis

End-to-end latency measurements reveal moderate overhead from compliance mechanisms. Research examining cryptographic infrastructure performance in web-scale systems has quantified the latency characteristics and overhead costs of certificate validation and signature verification operations that parallel those required for consent verification in advertising systems. Comprehensive analysis of SSL/TLS certificate validation across internet-scale deployments revealed significant performance implications and failure modes that inform privacy infrastructure design [9]. Large-scale measurement studies examining certificate validation behavior across 1,486 popular websites during 2010-2011 documented that cryptographic verification operations introduced measurable latency overhead, with SSL handshake completion requiring median times of 387 milliseconds for initial connections and 142 milliseconds for resumed sessions [9]. The research identified that certificate chain validation constituted a substantial portion of this overhead, with chains averaging 2.8 certificates in length requiring sequential verification steps [9]. Analysis of certificate revocation checking mechanisms revealed even more significant performance impacts: Online Certificate Status Protocol (OCSP) queries added median latency of 273 milliseconds when responders were reachable, with timeout scenarios extending delays to 5,000-10,000 milliseconds in 8.3% of validation attempts [9]. The study documented widespread deployment of certificate caching strategies to mitigate these overheads, finding that 89.4% of browsers maintained local certificate caches with average time-to-live values of 24 hours, though this caching introduced staleness risks where 12.7% of cached certificates had been revoked but remained locally trusted [9]. Performance measurements across geographic regions revealed substantial variation, with certificate validation from Asia-Pacific locations requiring 437 milliseconds median latency compared to 298 milliseconds from North American locations, reflecting differences in certificate authority infrastructure distribution [9]. The research further identified that certificate validation failures occurred in 9.4% of connection attempts, with broken certificate chains accounting for 41.2% of failures, expired certificates representing 23.7%, and revoked certificates comprising 18.4% of failure cases [9]. These measurements indicate that cryptographic verification operations similar to those required for consent-as-event architectures introduce latency overhead in the 2-8 millisecond range when properly cached and optimized, with tail latencies potentially reaching tens of milliseconds during cache misses or verification failures. However, provenance logging introduces measurable tail latency, necessitating asynchronous commitment strategies or careful batching.

5.3 Privacy and Utility Tradeoffs

Privacy evaluations under three regimes demonstrate progressive strengthening of guarantees. Baseline RTB with persistent identifiers yields re-identification risk exceeding 80% given auxiliary datasets. Tokenization with daily rotation reduces re-identification risk below 20% across sites, though single-session linkage risk remains elevated. Combined tokenization and differential privacy measurement achieves re-identification risk below 5% with privacy loss parameter epsilon at or below 2 for attribution metrics—thresholds considered acceptable in recent DP adoption studies. Research examining privacy-preserving measurement techniques for digital advertising has developed formal frameworks quantifying the tradeoffs between privacy protection and measurement utility in conversion attribution systems. Comprehensive analysis of differential privacy mechanisms applied to advertising measurement demonstrated that carefully configured DP systems can achieve strong privacy guarantees while maintaining sufficient utility for campaign optimization [10]. Technical evaluation of Google's Privacy Sandbox Attribution Reporting API revealed that differential privacy implementations for conversion measurement introduced quantifiable noise calibrated to epsilon privacy budgets, with practical deployments targeting epsilon values between 10 and 14 for acceptable utility-privacy balance [10]. The framework established that conversion count accuracy depends critically on contribution bounding strategies, with systems limiting individual users to single conversion reports per campaign achieving root mean squared error approximately 12-15% higher than unbounded systems but providing substantially stronger privacy guarantees [10]. Analysis of attribution window configurations demonstrated that seven-day windows enabled epsilon budgets 40% lower than 30-day windows while maintaining equivalent measurement accuracy, as temporal restrictions naturally limited information leakage [10]. The research quantified that aggregation threshold mechanisms requiring minimum event counts of 50-100 conversions per reporting bucket before releasing statistics provided additional privacy protection by preventing disclosure of individual user behaviors, though these thresholds reduced granularity for smaller campaigns [10]. Performance evaluation across real advertising datasets showed that differential privacy noise addition maintained click-through rate prediction model accuracy within 5-8% of non-private baselines when epsilon values remained above 8, with advertiser return-on-ad-spend optimization decisions degrading by only 6-9% under privacy-preserving measurement [10]. These findings support the proposition that privacy mechanisms introduce measurable but not catastrophic utility loss when designed with appropriate privacy budgets and batch sizing.

6. Conclusion

The harmonization of privacy regulation with the infrastructure of real-time advertising is a technically feasible yet economically and politically complicated problem. Compliance properties are formalized and give systematic standards of when to assess

distributed enforcement mechanisms, the architectural patterns of consent-as-event signalling, edge gating, tokenization, different privacy aggregation, and tamper-evident provenance indicate that privacy-preserving programmatic advertising can achieve system performance within ms latency limits. Experimental evidence confirms that consent verification and tokenization have a small overhead of 2-8 milliseconds, which is within standard ranges of auction budget, and differential privacy mechanisms can produce levels of re-identification risk below threshold values with epsilon values of 1.0-2.0 at an 8-15% loss in conversion attribution. Nonetheless, mass implementation has significant non-technical impediments. Empirical data records ongoing gaps of implementation, with large segments of websites continuing to send data to third parties before or irrespective of the rejection of consent and market structure changes after the introduction of GDPR, indicating that compliance costs are disproportionately targeted at smaller actors and concentrate power in the hands of dominant platforms. The inherent conflict between the fine-grained targeting economics on which digital content production relies today and regulatory impetus towards the minimization and aggregation needs not only technical innovation, but coordinated regulation by standards bodies and regulators, and by the market participants themselves. Consent encoding semantics, provenance metadata schemas, the range of differential privacy parameters, and cross-jurisdictional frameworks must be congruent under neutral standardization venues to be successful. Competitive oversight should strike a balance between the enforcement of privacy and antitrust issues because platform-level privacy primitives have the potential to solidify monopolistic market structures. The way forward consists of interdisciplinary collaboration mapping to map technical guarantees to legal tests to make sure that cryptographic verification, tokenization, and differential privacy mechanisms not only meet the requirements of computation, but also meet meaningful consent standards and showcase accountability requirements. Privacy-preserving advertising is a viable and yet contentious field where technological prowess has to keep pace with new regulatory demands, business incentives, and customer anticipations through long-term inter-engineering, regulatory, and policy alignment.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.
Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors, and the reviewers

References

[1] Maximilian H, et al, (2020) Measuring the Emergence of Consent Management on the Web, ACM Digital Library, 2020. [Online]. Available: https://dl.acm.org/doi/10.1145/3419394.3423647
[2] John D, et al., (2024) Differentially Private Ad Conversion Measurement, arXiv, 2024. [Online]. Available: https://arxiv.org/abs/2403.15224
[3] Justin P. J, et al., (2024) Online Advertising, Data Sharing, and Consumer Control, ACM Digital Library, 2024. [Online]. Available: https://dl.acm.org/doi/abs/10.1287/mnsc.2022.03385
[4] Martin D et al., (2018) We Value Your Privacy ... Now Take Some Cookies: Measuring the GDPR's Impact on Web Privacy, ResearchGate, 2018[Online]. Available: https://www.researchgate.net/publication/327050190
[5] Athina I, et al., (2020) Privacy concerns and disclosure of biometric and behavioral data for travel, ScienceDirect, 2020, [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/S0268401219317311
[6] Nurullah D, (2024) A Large-Scale Study of Cookie Banner Interaction Tools and Their Impact on Users' Privacy," ResearchGate, 2024. [Online]. Available: https://www.researchgate.net/publication/373925981
[7] Khashayar K, and Sven G. B, (2018) Secure Blockchains for Dynamic Spectrum Access: A Decentralized Database in Moving Cognitive Radio Networks Enhances Security and User Access, IEEE, 2018. [Online]. Available: https://ieeexplore.ieee.org/document/8269834
[8] Liam T and Ivan D O N, (2025) Towards Browser Controls to Protect Cookies from Malicious Extensions, arXiv, 2025. [Online]. Available: https://arxiv.org/html/2405.06830v3
[9] Nevena V, et al., (2012) The Inconvenient Truth About Web Certificates, SpringerNature Link, 2012. [Online]. Available: https://link.springer.com/chapter/10.1007/978-1-4614-1981-5_5
[10] Abhishek T, (2024) Privacy Preserving Measurement, 2024. [Online]. Available: https://www.abhishek-tiwari.com/privacy-preserving-measurement/