COLOMBO SCIENTIFIC PUBLISHING

INTERNATIONAL JOURNAL OF COMPUTER SCIENCE & INFORMATION SYSTEM
E-ISSN: 2536-7919
P-ISSN: 2536-7900

PAGE NO: 24-36

Privacy by Design as a Constitutive Principle in the Architecture of Family Digital Platforms: Models of Data Minimization, Access Differentiation, and Fiduciary Responsibility

Mykola Nesvietaiev
Business owner, CreationJoy Art LLC Brooklyn New York

ARTICLE INFO

Article history:

Submission: May 21, 2025
Accepted: June 19, 2025
Published: July 31, 2025
VOLUME: Vol.10 Issue 07 2025

Keywords:
Privacy by Design, data minimization, access differentiation, fiduciary responsibility, digital platforms for children, ABAC, information fiduciaries, family cybersecurity.

ABSTRACT

This study is aimed at the theoretical conceptualization and engineering-methodological substantiation of Privacy by Design as a foundational architectural imperative in the development of digital platforms oriented toward children and family-centered usage scenarios. Against the backdrop of an accelerating escalation of cyber risks in both educational environments and the domestic sphere, where in the second quarter of 2025 the average number of attacks per organization reached 4,388 per week, the limitations of reactive and predominantly perimeter-based approaches to data protection become increasingly evident. The analysis focuses on current models of data minimization and localization implemented through federated learning and on-device computing, as well as on mechanisms of attribute-based access control (ABAC) modified to account for family hierarchy and the distribution of roles within the household. In addition, the article examines the concept of the fiduciary responsibility of technology companies as a normative and ethical superstructure that demands not mere formal compliance, but the prioritization of the child's interests within a logic of loyalty and the prevention of conflicts of interest. As a final result, a multi-level architectural model, the Family-Centric Privacy Framework (FCPF), is proposed, integrating technical privacy guarantees with ethical and legal obligations embedded throughout the product life cycle. The theoretical conclusions and architectural solutions are grounded in an analysis of recent shifts in international regulation (COPPA 2.0, GDPR, DSA) and in the findings of empirical studies published in recent years.

INTRODUCTION

In the middle of the third decade of the twenty-first century, the digital environment has been transformed into a stable space for the socialization, learning, and leisure of minors, where the personalization of educational trajectories and entertainment practices is ensured by high-intensity processing of user data. Behind the functional advantages of this approach lies a qualitatively new scale and depth of observation: family digital platforms that unite an ecosystem of devices ranging from "smart" toys to complex learning management systems (LMSs) operate on the most sensitive categories of information, including biometric identifiers, behavioral patterns, parameters of psycho-emotional condition, and real-time geolocation data [1]. Under these conditions, the principle of Privacy by Design (PbD), proposed by Ann Cavoukian and embedded in the normative logic of such instruments as the General Data Protection Regulation (GDPR), acquires the character of a mandatory architectural requirement: confidentiality must function not as an optional setting, but as an immanent property of the system, sustained throughout the entire life cycle-from conceptual design to operation and decommissioning [2, 3].

The relevance of this study is determined by the crisis configuration of information security in the child-centered digital segment, where data breaches have assumed a systemic character and lead to disproportionately high economic and social consequences. According to IBM's 2025 Cost of a Data Breach

Report, the average breach cost in the United States reached $10.22 million, while the global average stood at $4.44 million; these figures underscore the economic stakes of inadequate data governance in child-focused digital ecosystems [5]. The situation is made still more complex by the development of generative artificial intelligence (GenAI), which has expanded the spectrum of attack vectors and intensified the risks of unauthorized profiling: even where personally identifiable information (PII) has been removed, anonymity can no longer be presumed, since contemporary algorithmic methods demonstrate the capacity to re-identify up to 89% of data subjects [6].

The situation is further aggravated by the phenomenon of a "fiduciary vacuum," that is, the gap between the technologically expanded possibilities of algorithmic influence on a child's behavior and the insufficiently defined legal responsibility of platform developers and operators [7, 8]. Within the family context, traditional consent models reveal their limited applicability: they do not ensure effective control and, not infrequently, rely on mechanisms of behavioral exploitation that mobilize cognitive biases through "dark patterns," thereby undermining decisional autonomy and diminishing the quality of parental oversight [9]. The consequence is the need for an architectural reconsideration of family platforms, one in which the protection of the child's interests is fixed not as a declarative principle, but as a constitutive element of design that determines the structure of data, the logic of access, and the permissible modes of analytics.

The aim of this work is to develop and substantiate architectural models that implement the principles of data minimization and fiduciary responsibility in the interests of the family, shifting the emphasis from the formal acquisition of consent toward demonstrable technical constraints on the collection and use of information. A key direction of analysis is the transition from static role-based access management to dynamic, context-sensitive models capable of accounting for family hierarchy, situational conditions, and the risk profile of operations. In parallel, the study examines the implementation of privacy-enhancing technologies (PETs) designed to preserve service functionality while radically reducing the volume of collected data and minimizing the probability of secondary uses of information beyond declared purposes.

The scientific novelty of the study lies in the development of an integrated architectural model of a family digital platform in which privacy-by-design principles are implemented through a combination of local data processing, context-sensitive ABAC access, and a fiduciary model of platform responsibility.

The study's **hypothesis** is based on the assumption that if privacy-by-design is built into the architecture of a family digital platform through data minimization, contextual access control, and management accountability mechanisms, the risk of unauthorized processing and secondary use of children's data is reduced more effectively than with models based primarily on formal consent and ex post compliance.

Limitations of the study. The study has several limitations, as it is predominantly conceptual and theoretical in nature and relies on a comparative analysis of regulatory sources, architectural models, and secondary empirical data, without conducting independent testing of the proposed model on real family-oriented digital platforms. In addition, the rapid transformation of the technological environment and of regulatory approaches in the field of children's digital safety means that particular architectural solutions and the criteria used to evaluate their effectiveness require further empirical validation across specific platforms, jurisdictions, and user scenarios.

Significance of the approach for the United States. For the United States, the proposed approach has particular practical significance because the American model of regulating children's data has historically been built around sector-specific and procedural mechanisms, above all parental consent as established in COPPA, whereas the contemporary digital environment demonstrates the insufficiency of consent alone as an instrument of genuine child protection. In the context of the rapid expansion of educational platforms, family applications, connected toys, voice assistants, and AI-driven services, the architectural embedding of privacy by design makes it possible to shift the focus from formal compliance to the prevention of excessive data collection, the restriction of secondary uses of information, and the technical reduction of risks associated with the exploitation of children's vulnerability.

For the United States, this conclusion is especially important for another reason as well: the American legal system is increasingly moving toward stronger design-based obligations, in which the object of evaluation is no longer limited to the text of a privacy policy but extends to the logic of the product itself, its interface, and its algorithmic influence on the minor user. In this context, models of data minimization, context-sensitive access control, and fiduciary responsibility may be regarded as a promising foundation for the development of child-centered governance in the United States, compatible both with the federal requirements of COPPA and with the growing wave of state-level regulation in the areas of age-appropriate design, online safety, and the accountability of digital platforms.

Materials and Methods

The methodological design of the study was constructed on the basis of a systems paradigm for secure software design and a comparative legal analysis of regulatory regimes that establish mandatory requirements for the processing of children's data. The empirical and normative foundation was formed from a body of current materials and comprises several mutually complementary layers.

Within the framework of the legal and regulatory analysis, the study examined the April 2025 amendments to the Children's Online Privacy Protection Act (COPPA), the provisions of the EU Digital Services Act (DSA), and the British Age-Appropriate Design Code (AADC) [2]. The central analytical node was the concept of the Best Interests of the Child, interpreted not as a merely declarative principle, but as an operationalizable technical standard to be implemented at the level of architectural constraints, default settings, and data life-cycle procedures.

The architectural dimension of the research was realized through an audit of microservice-based solutions and decentralized processing practices aimed at reducing data concentration and shrinking the attack surface. Two metrics were used for the quantitative assessment of minimization efficiency. The Privacy Risk Expansion Factor (PREF) was employed as an indicator of data multiplication and duplication within the system, making it possible to identify points of unnecessary copying and secondary accumulation. The Privacy Exposure Index (PEI) was used as an index of the probabilistic disclosure of sensitive information, reflecting the aggregate risk of exposure under specified data flows and access regimes [13].

A separate methodological block consisted of modeling access control systems, within which a comparison was carried out between RBAC (Role-Based Access Control) and ABAC (Attribute-Based Access Control) mechanisms as applied to family hierarchies, role heterogeneity, and the contextual variability of permissions. The comparison criteria included flexibility and scalability, as well as resilience to the phenomenon of "role explosion," which emerges when complex family scenarios are described exclusively through static roles [14].

The statistical and empirical segment of the study was based on data from the Verizon DBIR 2025, IC3, and NCMEC reports, which were used to identify recurring patterns of attacks on children's and family data and to refine the typology of threats in educational and domestic digital environments [16, 17]. In addition, the findings of user-centered studies (PoPETs 2025) were analyzed in order to capture the attitudes of developers and parents toward the fiduciary model of privacy and the acceptable boundaries of algorithmic intervention in matters affecting minors [9].

The formalization of the proposed architectural solutions was carried out through the use of block-diagram notations and tables comparing technical parameters, which ensured the reproducibility of the conclusions and the comparability of alternatives. The evaluation of the developed models proceeded through a three-layer lens: the operational layer (access control and the governability of permissions), the tactical layer (technical data protection and the minimization of exposure), and the strategic layer (risk management, compliance, and the fiduciary duties of platform operators) [18].

Results and Discussion

The principle of data minimization presupposes that the collection and subsequent processing of information are restricted to the volume that is adequate, relevant, and necessary for a predetermined and

lawful purpose. In family digital platforms, this requirement frequently comes into tension with the demand for deep service personalization and with parental control functions, since both objectives stimulate the expansion of profiling practices directed at minors and the accumulation of behavioral data sets. In order to avoid replacing the original purpose of processing with the convenience of analytics, it is advisable to embed within the system an architectural pattern of a "privacy cascade from edge to cloud," under which the primary processing and inference of sensitive features are performed on the child's end device, while only derivative values that are strictly necessary for the functioning of the selected feature are transmitted to the external infrastructure [4].

For the practical implementation of minimization, privacy-enhancing technologies play a significant role. Among the most sought-after approaches is federated learning, which makes it possible to train models, including those used in adaptive educational solutions, without transferring the underlying data from the device: only model update parameters are sent to the infrastructure, which reduces the risk of disclosing substantive information about the child in the event of channel or storage compromise. Differential privacy is based on the introduction of mathematically grounded noise and ensures that the characteristics of a particular child cannot be isolated, while still preserving an acceptable degree of accuracy in aggregated indicators for legitimate monitoring and reporting purposes [19]. Zero-knowledge proofs are applicable as a cryptographic instrument for confirming age or another legally significant attribute without transmitting identity documents, which makes it possible to reduce the processing of identifying data to a level proportionate to the purpose [19].

The legal soundness of the "privacy cascade from edge to cloud" is strengthened by the fact that it transforms minimization from a declaration into a verifiable design standard: sensitive data, by default, remain within the device boundary, while the cloud component receives only a strictly limited set of processing results. Such an approach reduces the probability of "secondary use" of information, in which data initially collected for one purpose begin, almost imperceptibly, to be used for other tasks, such as marketing profiling, behavioral segmentation, or the construction of long-term trajectories of interests. In the family environment, this is especially critical, since a minor does not possess autonomy in expressing informed consent comparable to that of an adult; accordingly, priority should be given to minimization as a means of preventing excessive intrusion into private life, rather than to the subsequent "correction" of the consequences of a breach or an improperly configured access regime.

The technical implementation of minimization in family systems requires not only the selection of appropriate technologies, but also the institutionalization of restrictions within the data life cycle: the establishment of short retention periods, the prohibition of indefinite archiving of activity logs, strict separation of access roles, mandatory logging of operations, and cryptographic protection of both the data and the keys. In local processing scenarios, trusted hardware environments, on-device encryption, and deterministic rules for cloud transmission that exclude the "background" sending of telemetry acquire particular importance. As a result, minimization ceases to depend on the good faith of the developer and is transformed into a form of architectural coercion: even where errors arise in the application logic, the system structurally limits the exposure of personal data [11].

For the evaluation of an architectural solution for data minimization, the application of formalized metrics is well founded, since such metrics make it possible to compare different design variants and to identify latent redundancy. One such metric proposed here is the Privacy Exposure Index (PEI), which may take into account the number of data categories, their sensitivity, retention duration, the number of transmission environments, the number of access subjects, and also the probability of re-identification under aggregation. In practical terms, PEI should be linked to threat modeling and to data protection impact assessment: this makes it possible not merely to count "how much data has been collected," but to analyze what legally significant risk is generated by a specific architecture and which elements of the system increase exposure without adding functional value [13]. Minimization is achieved by reducing each of these parameters through automated deletion policies. A comparative analysis of data minimization methods for children's platforms is presented in Table 1.

Table 1. Comparative analysis of data minimization methods for children's platforms (compiled by the author based on [6, 7, 9, 19, 20, 21]).

Technology	Protection Mechanism	Impact on Functionality	Privacy Level	Applicability (2025)
Local processing (Edge computing)	Data do not leave the device	High (requires client-side capacity)	Maximum	Recommended for biometrics and voice
Anonymization (K-anonymity)	Removal of direct identifiers	Moderate (risk of re-identification up to 89%)	Low	Considered insufficient without PETs
ZKP (Zero-Knowledge)	Mathematical verification of a fact	Minimal	High	Standard for age verification
Synthetic data	Creation of copies based on statistical models	No effect on risk	Absolute	For AI model testing

The role-based access control (RBAC) model, which relies on predefined roles such as parent, child, and teacher, frequently leads to excessive permission grants in family digital systems and is poorly adapted to account for changing circumstances. The need to automatically restrict a child's access to entertainment content at certain times, for example after 9:00 p.m., within a pure RBAC framework in practice compels designers either to construct temporal roles or to complicate the permissions matrix, thereby increasing the risk of error and reducing the verifiability of compliance with the principle of least privilege. For this reason, a transition to attribute-based access control (ABAC) is well justified, since under this approach the decision to permit or deny an operation is derived from a combination of attributes pertaining to subjects, objects, and the conditions under which the request is made [14, 15].

In the attribute-based model, access is determined not by the "status of a role," but by a set of legally significant and technically verifiable attributes. Subject attributes include the age of the minor, the level of digital literacy, and also parameters characterizing the current state, provided that these are generated locally and do not become an independent corpus of sensitive data. Object attributes include the age labeling or rating of material, as well as the type of information being processed, for example educational data as compared with communications data. Environment attributes capture the time of day, the place of use, such as home or an educational institution, and the characteristics of the communication channel, whether a domestic wireless network or a public fifth-generation network. The action attribute reflects the substance of the operation itself: reading, making modifications, or completing an in-service purchase [22, 26].

From the standpoint of recognized approaches to information security, attribute-based access control is described as a methodology of logical access control in which authorization is determined by evaluating the attributes of the subject, the object, the requested action, and, where necessary, environmental conditions in relation to the rules and relationships established in the security policy. This approach has been formalized in the guidance materials of the U.S. National Institute of Standards and Technology and is applied as a basis for constructing more precise and context-dependent admission regimes.

The practical implementability of ABAC in applied systems is usually ensured through a separation between the function of decision-making and the function of enforcement: the mechanism makes an access decision

on the basis of policy and attributes, while the application component merely ensures the compulsory execution of that result. For the description and evaluation of such decisions, a standard architecture and policy language are widely used, within which an access request is interpreted through attributes and matched against "permit/deny" rules, including the handling of conflicts between rules and the prioritization of policies [27, 32].

In the family context, the attribute-based model requires heightened legal caution in the selection of the attributes themselves: access control must not be transformed into an instrument of expanded surveillance. Attributes must be proportionate to the purpose and, wherever possible, derived from less sensitive indicators, for example an "age category" rather than an exact date of birth; for dynamic parameters, local computation and short-term use without accumulation are preferable. Additional importance attaches to rules for resolving collisions between the powers of legal representatives and the autonomy of the minor within permissible limits: the policy must ensure predictability, meaning that it is clear why access is restricted, and verifiability, meaning that the rule that was triggered is recorded, because it is precisely these properties that make it possible to assess the lawfulness of restrictions in a qualified manner and to reduce the risk of abuse in parental control settings. The matrix of access attributes for a family digital ecosystem is demonstrated in Table 2.

Table 2. Matrix of access attributes for a family digital ecosystem (compiled by the author based on [10, 12, 14, 23, 24, 25]).

Attribute	Type	Description and Role in PbD	Data Source
Age_Range	Subject	Category under AADC (0–5, 6–9, 10–12, 13–15, 16–17)	Verified Age Service
Consent_Status	Subject	Presence of verified parental consent for a specific type of processing	Consent Manager
Location_Context	Environment	"Safe zone" (home/school) vs. "Unknown zone"	Edge Location
Safety_Rating	Object	Content risk level (automatic AI-based assessment)	Content Classifier
Device_Trust_Level	Environment	Presence of installed updates and MFA	Security Agent

The use of ABAC creates a foundation for the algorithmic operationalization of the principle of the Best Interests of the Child, translating it from the realm of declarations into the realm of executable access policies. Where the system detects indicators of an impulsive interaction pattern, or where it records the use of an application during nighttime hours, predefined restrictions are activated automatically: access to features with heightened addictive potential, including infinite feeds and push-notification mechanisms, is disabled or blocked. Such a regime of contextual enforcement is consistent with the requirements of the Maryland Kids Code of 2024 and is implemented as an embedded rule of interface behavior management at the level of attributes and execution-environment conditions [10].

Figure 1 demonstrates the logic of PDP (Policy Decision Point) decision-making in the ABAC model.

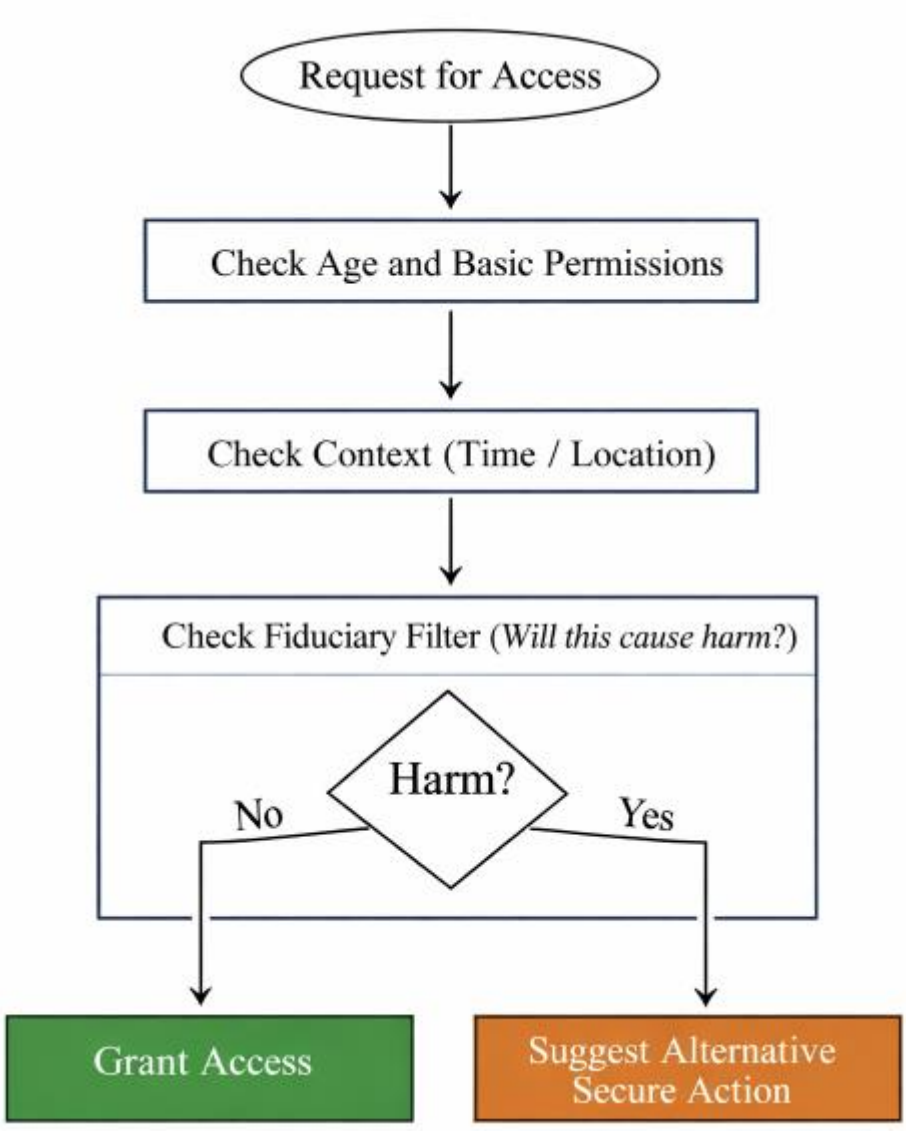

Figure 1. PDP (Policy Decision Point) decision-making logic in the ABAC model (compiled by the author based on [35, 36, 40]).

The concept of the "information fiduciary" proceeds from the premise that digital platforms occupy a position of structural asymmetry in relation to users: access to behavioral data, the ability to interpret such data, and the subsequent capacity to shape or influence choice together create a position comparable, in terms of the burden of trust involved, to professional relationships in medicine or law, where regimes of confidentiality and heightened standards of good faith are well established [28]. In the domain of children's digital products, this approach acquires a quasi-normative form through the Duty of Loyalty, understood as an obligation of loyalty that precludes deriving benefits from a child's data to the detriment of that child's well-being, even where consent has been formally obtained [9]. In this way, fiduciary responsibility is framed not as a declaration of corporate ethics, but as a design imperative that affects permissible processing purposes, data architecture, and the boundaries of functionality built around engagement.

The urgency of the problem is driven by the phenomenon of a "fiduciary vacuum," in which the technological capacity for algorithmic modeling and behavioral steering advances more rapidly than the consolidation of obligations of accountability and verifiability of comparable force. At the level of public perception, this appears as a persistent deficit of trust: studies indicate that in 2025, 77% of consumers did not trust companies to use AI responsibly [30]. As an architectural response to this gap, a layer of Fiduciary Oversight is formed, intended to institutionalize the verifiability of decisions and reduce the risk of concealed profiling. Within this layer, first, there is the mechanism of Algorithmic Mirroring, which presupposes providing

adolescents with a clear representation of the way in which an algorithm reconstructs their interests and what profiles it forms; such visualization increases transparency and functions as an instrument for developing a critical attitude toward recommendations and "personalization" [31]. Second, the Data Trusts model is employed, based on the participation of independent intermediaries (trusted stewards) who manage children's data in the interests of the family and are capable of negotiating with platforms regarding terms of access, thereby moving an individually vulnerable subject into a regime of collectively secured protection and contractual rebalancing [29]. The metrics of platform fiduciary responsibility are described in detail in Table 3.

Table 3. Platform fiduciary responsibility metrics (compiled by the author based on [9, 10, 12, 23, 33, 38]).

Indicator	Measurement Method	Target Value	Architectural Implication
Dark Pattern Index	Automated UI/UX audit for manipulative design elements	< 0.01 (zero tolerance)	Elimination of autoplay and infinite scrolling
Ad-to-Content Ratio	Ratio of educational content to commercial content	Age-dependent (0% for children under 13)	Architectural prohibition on embedding advertising SDKs in children's modules
Transparency Score	Comprehensibility of privacy notices for the child (readability score)	Correspondence to the reading level of the age group	Visual interfaces instead of text-only policies
Retention Precision	Accuracy of data deletion once the purpose has been fulfilled or consent has been withdrawn	100% within 48 hours	Automated data deletion triggers

Fiduciary responsibility in corporate governance presupposes that cyber risks fall within the sphere of the board of directors' direct oversight mandate rather than being delegated exclusively to IT divisions or compliance functions. Within this logic, data security acquires the status of a governance duty requiring regular monitoring, the setting of measurable objectives, verification of the adequacy of resourcing, and the formalization of incident escalation procedures. Judicial decisions in Delaware in 2025 reinforced this position, indicating that the absence of effective oversight over data protection may be qualified as an improper discharge of the fiduciary duties of corporate leaders [34].

The matrix for assessing the fiduciary risk of AI models in children's products is presented in Figure 2.

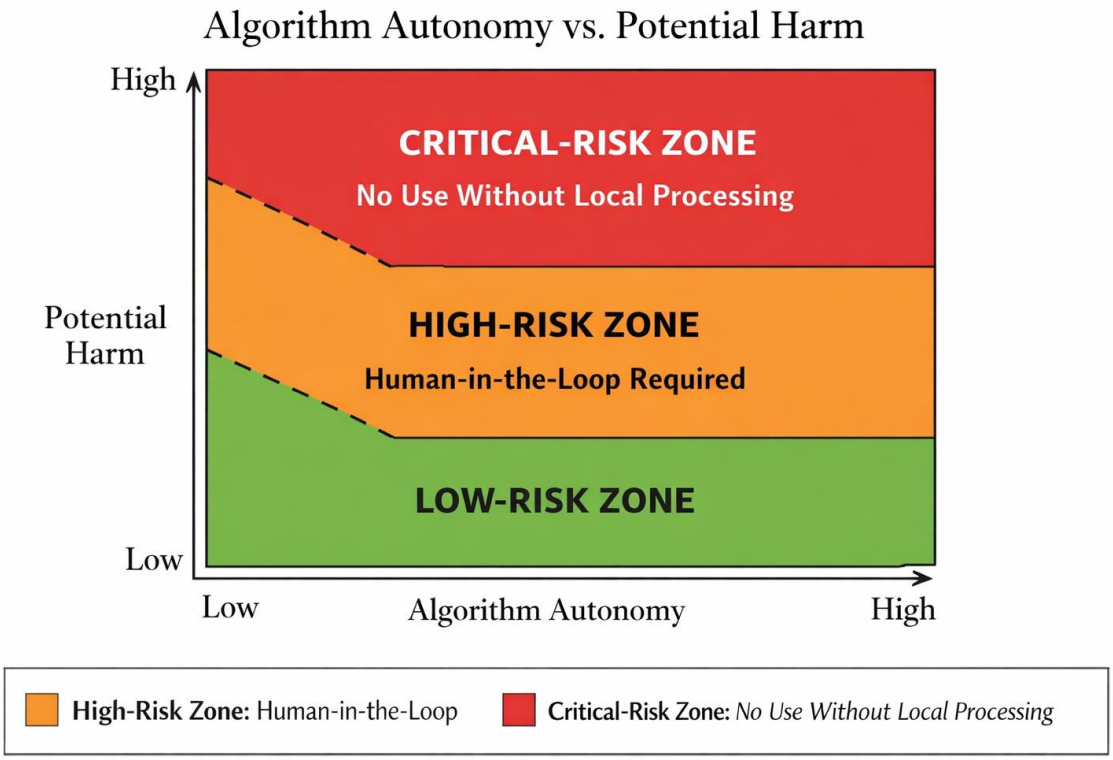

Figure 2. Matrix for assessing the fiduciary risk of AI models in children's products (compiled by the author based on [5, 34]).

The statistics for 2025 reveal troubling trends. In July 2025, educational institutions were recording an average of 4,210 attacks per week [5]. The principal vectors remain the following: Steganography in GenAI: the use of AI to generate malicious content disguised as children's images or videos [16].
Supply Chain Attacks: attacks delivered through third-party SDKs and cloud services, which accounted for 30% of all incidents [5].

The architectural implementation of Privacy by Design presupposes the adoption of a Zero Trust Architecture (ZTA), within which trust is not treated as a given for any subject or component, regardless of whether the entity in question is a parent, a child, or an internal service. In such a model, every data transaction is subject to a regime of Continuous Verification, which shifts control from a one-time act of authentication to an ongoing process of confirming the legitimacy of access [39]. One consequence is a reduction in the probability of cascading breaches, in which the compromise of a single account automatically exposes the resources of the entire family group, because the absence of "trusted zones" limits the chainwise propagation of privileges.

The practical applicability of this logic can be traced through the example of current solutions in the Family Link and Apple Screen Time class, considered through the lens of PbD-oriented approaches as of 2025. Screen-time control in such systems may be interpreted not as an instrument of disciplinary restriction, but as a mechanism for preventing risks associated with health and development: recent studies have documented a direct relationship between excessive screen time and delays in speech development [41]. Within a PbD architecture, restrictions of this kind acquire the status of a protective measure embedded in the system as an element of risk management rather than as an arbitrary setting imposed from outside.

The transparency of monitoring requires separate attention. Within the PbD framework, a shift from covert forms of observation to overt ones becomes justified, meaning arrangements in which the minor is informed of exactly what data are available to the parent and to what extent. Such a solution reduces the likelihood of an erosion of trust within the family and aligns with the ethical orientation of "Augmented Human Development," under which digital control should not become an implicit instrument of pressure, but must

remain within the boundaries of supportive development [1].

Biometric data are subject to the strictest requirements. In accordance with the 2025 COPPA updates, the collection of voiceprints or facial images for authentication purposes requires separate informed consent and strict local storage [2]. As a result, architectures developed in 2025 tend toward the use of biometric templates that do not permit reconstruction of the original image, which reduces the value of a stolen artifact and lowers the risk of secondary data use outside the context of the primary function [2].

The integrated combination of data minimization, ABAC, and fiduciary responsibility forms a mutually reinforcing configuration: reducing the volume of information collected and retained diminishes the "attack surface," attribute-based access control ensures precision and contextual sensitivity of permissions in real time, and the fiduciary framework limits processing purposes by subordinating them to the interests of the child as the priority criterion of permissibility. At the same time, the implementation of such a model is associated with a number of systemic obstacles. First, the effect of an "explosion of complexity" becomes apparent: ABAC requires substantial investment in policy formalization, ensuring the quality of attributes, and maintaining their integrity; errors in attributes, for example an incorrectly recorded age, can lead either to critical denials of access or to leaks caused by improper authorization of operations [14]. Second, latency increases: the verification of multiple ABAC conditions and the execution of ZKP protocols on mobile devices may degrade performance and reduce the quality of the user experience (UX), thereby creating pressure in favor of simplifying protective layers. Third, regulatory fragmentation constitutes a significant barrier: even where general principles coincide, differences in requirements across jurisdictions, including divergences between the Maryland Kids Code and the European DSA, compel global platforms to construct hybrid compliance architectures, increasing development costs and the risk of policy inconsistency [10].

Despite the limitations listed above, market dynamics indicate the presence of long-term benefits for organizations that invest systematically in PbD: a substantial share of the professional community points to the practical advantages of privacy protection efforts, including increased user loyalty and reduced cyber-risk insurance premiums [37].

Conclusion

As a final conclusion, it must be stated that Privacy by Design in the segment of family platforms has lost the status of an ethical aspiration and acquired the features of a mandatory architectural standard. Under the conditions of 2025, when children's digital products process information of critical sensitivity, an acceptable level of protection can be achieved only through a multi-layered model combining preventive constraints, demonstrable technical guarantees, and risk governability throughout the entire life cycle.

The proposed architectural construct, grounded in data minimization through local processing and privacy-enhancing technologies (PETs), dynamic attribute-based access control (ABAC), and the qualification of the platform as an information fiduciary, demonstrates the potential to overcome the vulnerabilities inherent in traditional protection schemes. Its system-forming element is the software-based automation of the principle of the Best Interests of the Child: design must exclude the possibility of harm through data as a technically realizable scenario, or else transform such harm into an economically irrational strategy for the provider through embedded constraints, auditability, and the inevitability of accountability.

Prospective directions for further research should reasonably be connected with the standardization of formal languages for describing security and privacy policies in family systems, ensuring unambiguous interpretation and the portability of requirements across platforms, as well as with the development of lightweight PETs suitable for operation on low-power devices, including the IoT environment. Under such conditions, the technological sovereignty of the family and the protection of the digital future of minors require not only formal compliance with legal norms, but also a deeper engineering reflection on how digital environments of development are constructed and which system properties make harm impossible by design.

References

1. Iqbal, M. Z., Xu, X., Nallur, V., Scanlon, M., & Campbell, A. G. (2023). Security, ethics and privacy issues in the remote extended reality for education. In Mixed reality for education (pp. 355-380). Singapore: Springer Nature Singapore.

2. European Data Protection Board. (2025). Statement 1/2025 on age assurance. Retrieved from: https://www.edpb.europa.eu/our-work-tools/our-documents/statements/statement-12025-age-assurance_en (date accessed: February 18, 2025).

3. Le Métayer, D. (2013). Privacy by design: A formal framework for the analysis of architectural choices. Proceedings of the 2013 ACM Conference on Computer and Communications Security. https://doi.org/10.1145/2435349.2435361

4. Bi, T., Yu, G., & Wang, Q. (2023). Privacy in Foundation Models: A Conceptual Framework for System Design. arXiv preprint arXiv:2311.06998.

5. Verizon. (2025). 2025 Data Breach Investigations Report. Retrieved from: https://www.verizon.com/business/resources/reports/dbir/ (date accessed: April 24, 2025).

6. Mireshghallah, N. (2025). Privacy and security challenges in machine learning systems [Conference presentation]. Retrieved from: https://mireshghallah.github.io/talks/camlis_2025.pdf (date accessed: June 3, 2025).

7. Gartner. (2025). Gartner identifies the top cybersecurity trends for 2025. Retrieved from: https://www.gartner.com/en/newsroom/press-releases/2025-03-03-gartner-identifiesthe-top-cybersecurity-trends-for-2025 (date accessed: March 19, 2025).

8. Benthall, S., & Shekman, D. (2023). Designing fiduciary artificial intelligence. In Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization. https://doi.org/10.1145/3617694.3623230

9. Harkous, H., et al. (2025). Evaluating a data fiduciary standard for privacy: Developer and end-user perspectives. Proceedings on Privacy Enhancing Technologies, 2025(4). https://doi.org/10.56553/popets-2025-0114

10. European Commission. (2025). Commission publishes draft guidelines on protection of minors online under the Digital Services Act. Retrieved from: https://digital-strategy.ec.europa.eu/en/news/commission-publishes-draft-guidelines-protection-minors-online-under-digital-services-act (date accessed: May 20, 2025).

11. European Data Protection Board. (2025). EDPB comments on the draft guidelines on protection of minors online under the Digital Services Act (DSA). Retrieved from: https://www.edpb.europa.eu/system/files/2025-06/edpb_comments_europeancommission_article_28_dsa_en.pdf (date accessed: June 18, 2025).

12. Federal Trade Commission. (2025). FTC finalizes changes to children's privacy rule limiting companies' ability to monetize kids' data. Retrieved from: https://www.ftc.gov/news-events/news/press-releases/2025/01/ftc-finalizes-changes-childrens-privacy-rule-limiting-companies-ability-monetize-kids-data (date accessed: January 22, 2025).

13. Chereja, I., Erdei, R., Delinschi, D., Pasca, E., Avram, A., & Matei, O. (2025). Privacy-conducive data ecosystem architecture: By-design vulnerability assessment using privacy risk expansion factor and privacy exposure index. Sensors, 25(11), 3554. https://doi.org/10.3390/s25113554

14. National Institute of Standards and Technology. (2025). Attribute Based Access Control (ABAC). Retrieved from: https://csrc.nist.gov/projects/attribute-based-access-control (date accessed: May 15,

2025).

15. Sandhu, R. (2025). Role-based access control. In Encyclopedia of Cryptography, Security and Privacy. https://doi.org/10.1007/978-3-030-71522-9_829

16. National Center for Missing & Exploited Children. (2025). NCMEC releases new data: 2024 in numbers. Retrieved from: https://www.missingkids.org/blog/2025/ncmec-releases-new-data-2024-in-numbers (date accessed: May 12, 2025).

17. Federal Bureau of Investigation. (2024). 2024 IC3 annual report. Retrieved from: https://www.ic3.gov/AnnualReport/Reports/2024_IC3Report.pdf (date accessed: January 29, 2025).

18. Alabdulatif, A. (2025). Blockchain-based privacy-preserving authentication and access control model for e-health users. Information, 16(3), 219. https://doi.org/10.3390/info16030219.

19. Kurian, N. (2025). 'No, Alexa, no!': Designing child-safe AI and protecting children from inappropriate AI interactions. Learning, Media and Technology. https://doi.org/10.1080/17439884.2024.2367052

20. Colnago, J., et al. (2020). Operationalizing the legal principle of data minimization for personalization. Proceedings on Privacy Enhancing Technologies, 2020(4), 6–25. https://doi.org/10.2478/popets-2020-0050

21. OECD. (2025). How's life for children in the digital age? Retrieved from: https://www.oecd.org/en/publications/how-s-life-for-children-in-the-digital-age_0854b900-en.html (date accessed: May 27, 2025).

22. National Institute of Standards and Technology. (2025). Role-based access control (RBAC). Retrieved from: https://csrc.nist.gov/glossary/term/role_based_access_control (date accessed: June 6, 2025).

23. Information Commissioner's Office. (2025). Age appropriate design: A code of practice for online services. Retrieved from: https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/childrens-information/childrens-code-guidance-and-resources/age-appropriate-design-a-code-of-practice-for-online-services/ (date accessed: February 27, 2025).

24. Federal Trade Commission. (n.d.). Protecting kids online. Retrieved from: https://consumer.ftc.gov/identity-theft-and-online-security/protecting-kids-online (date accessed: March 28, 2025).

25. Federal Trade Commission. (2025). How to use parental controls to keep your kid safer online. Retrieved from: https://consumer.ftc.gov/consumer-alerts/2025/04/how-use-parental-controls-keep-your-kid-safer-online (date accessed: April 30, 2025).

26. National Institute of Standards and Technology. (2025). Role-Based Access Control (RBAC) project. Retrieved from: https://csrc.nist.gov/projects/role-based-access-control (date accessed: June 11, 2025).

27. Federal Trade Commission. (n.d.). Children's privacy. Retrieved from: https://www.ftc.gov/business-guidance/privacy-security/childrens-privacy (date accessed: June 17, 2025).

28. Balkin, J. M. (2016). A duty of loyalty for privacy law. Boston University Law Review, 99(3), 1183–1227. https://doi.org/10.2139/ssrn.2790379

29. Delacroix, S., & Lawrence, N. (2020). Trust law, fiduciaries, and data trusts. Data Economy Lab Report. https://doi.org/10.2139/ssrn.3531568

30. Meacham, D., Gianni, R., Brüggen, E., Werf, M., & Post, T. (2025). AI-based financial advice: An ethical discourse on AI-based financial advice and ethical reflection framework. Journal of Public Policy &

Marketing, 44(3), 436–456. https://doi.org/10.1177/07439156241302279

31. Nassif, S. A., & Ben Moussa, M. (2024). Algorithm literacy among youth: Understanding and navigating social media algorithms. The Egyptian Journal of Media Research, 89, 33-72.

32. Delacroix, S., & Lawrence, N. (2019). Bottom-up data trusts: Disturbing the "one size fits all" approach to data governance. International Data Privacy Law, 9(4), 236–252. https://doi.org/10.1093/idpl/ipz014

33. Ministry of Electronics and Information Technology. (2025). Draft Digital Personal Data Protection Rules, 2025. Retrieved from: https://www.meity.gov.in/content/draft-digital-personal-data-protection-rules2025 (date accessed: February 11, 2025).

34. U.S. Securities and Exchange Commission. (2025). Cybersecurity risk management, strategy, governance, and incident disclosure. Retrieved from: https://www.sec.gov/rules-regulations/2023/07/s7-09-22 (date accessed: May 26, 2025).

35. Russell Reynolds Associates. (2025). Global corporate governance trends for 2025. Retrieved from: https://www.russellreynolds.com/en/insights/reports-surveys/global-corporate-governance-trends/2025 (date accessed: February 21, 2025).

36. IBM. (2025). Cost of a Data Breach Report 2025. Retrieved from: https://www.ibm.com/reports/data-breach (date accessed: June 20, 2025).

37. U.S. Department of Education. (n.d.). K-12 cybersecurity. Retrieved from: https://www.ed.gov/teaching-and-administration/safe-learning-environments/school-safety-and-security/k-12-cybersecurity (date accessed: June 23, 2025).

38. U.S. Department of Health and Human Services, Office for Civil Rights. (n.d.). Breach portal: Notice to the Secretary of HHS involving unsecured protected health information. Retrieved from: https://ocrportal.hhs.gov/ocr/breach/breach_report.jsf (date accessed: June 24, 2025).

39. U.S. Department of Education. (n.d.). Data breach. Retrieved from: https://studentprivacy.ed.gov/topic/data-breach (date accessed: June 25, 2025).

40. National Institute of Standards and Technology. (2025). Role-based access control (RBAC). Retrieved from: https://csrc.nist.gov/glossary/term/role_based_access_control (date accessed: June 28, 2025).

41. Li, X., Keown-Stoneman, C. D., Omand, J. A., et al. (2025). Screen time and standardized academic achievement tests in elementary school. JAMA Network Open, 8(10), e2537092. https://doi.org/10.1001/jamanetworkopen.2025.37092