Global Journal of Engineering and Technology Advances

eISSN: 2582-5003

Cross Ref DOI: 10.30574/gjeta

Journal homepage: https://gjeta.com/

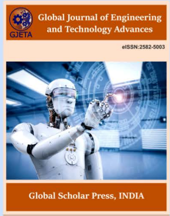

(RESEARCH ARTICLE)

Check for updates

Privacy-by-Design in AI Data Pipelines: A Unified Governance Approach

Praveen-Kodakandla *

Independent Researcher.

Global Journal of Engineering and Technology Advances, 2024, 21(01), 215-224

Publication history: Received on 02 September 2024; revised on 22 October 2024; accepted on 25 October 2024

Article DOI: https://doi.org/10.30574/gjeta.2024.21.1.0187

Abstract

As artificial intelligence becomes deeply embedded in business decision-making, managing data with integrity and accountability has become a central concern. This paper explores how integrating Privacy-by-Design (PbD) with Unified Data Governance can establish a structured framework for building AI systems that are compliant, secure, and ethically responsible. The proposed approach incorporates privacy safeguards across every stage of the AI data lifecycle—from collection and feature generation to model deployment and monitoring. It outlines architectural principles, design standards, and governance workflows that ensure privacy considerations remain continuous rather than reactive. The study also emphasizes the role of data traceability, cross-functional coordination, and adaptive policy management to meet evolving regulatory and ethical expectations. Together, these practices enable enterprises to balance innovation with accountability, reducing exposure to risks while maintaining public trust.

Keywords: Privacy-by-Design; Unified Data Governance; AI Pipelines; Data Ethics; Responsible AI; Compliance; Risk Management; Automation; Accountability; Trustworthy AI Systems

1. Introduction

The growing influence of Artificial Intelligence (AI) has transformed how organizations collect, analyze, and act on data across industries such as healthcare, finance, and retail. Yet, as AI systems process vast and often sensitive datasets, the tension between innovation and privacy protection has intensified. Data pipelines feeding these models frequently handle personal identifiers, behavioral patterns, and derived attributes—making governance and ethical control crucial to responsible deployment.

Privacy-by-Design (PbD), a concept introduced by Dr. Ann Cavoukian in the 1990s, advocates embedding privacy considerations directly into the architecture of systems rather than addressing them as an afterthought. This forward-looking philosophy promotes proactive risk mitigation, ensuring that safeguards exist at each stage of the data journey—ingestion, transformation, training, inference, and monitoring. However, implementing PbD at scale within complex AI ecosystems is challenging due to fragmented ownership, diverse legal interpretations, and opaque model behaviors.

To address these challenges, organizations are adopting Unified Data Governance, which consolidates privacy, compliance, and data management disciplines under a cohesive framework. Instead of treating privacy as an isolated compliance exercise, unified governance operationalizes it as a shared responsibility across engineering, legal, and product domains. This integration enforces consistent policies for data classification, access control, consent handling, and auditability across the AI pipeline, supported by automation for lineage tracking, encryption, and anomaly detection.

* Corresponding author: Praveen Kodakandla

Global privacy regulations such as GDPR and CCPA have accelerated the urgency for such frameworks, compelling enterprises to ensure transparency in data usage and consent management. Beyond compliance, the societal expectations around data stewardship have evolved—users now demand clarity on how their data informs algorithmic decisions. Organizations that fail to institutionalize privacy risk both regulatory penalties and reputational harm.

Embedding PbD within a unified governance strategy transforms privacy from a defensive obligation into an enabler of innovation. Clearly defined roles, standardized controls, and automated verification reduce rework and ensure that AI systems align with organizational ethics. When governance and design operate hand in hand, privacy becomes not merely a regulation-driven requirement but a driver of sustainable AI progress.

The subsequent sections of this paper examine the frameworks, implementation patterns, and governance mechanisms that translate the Privacy-by-Design vision into scalable, auditable, and trustworthy AI data pipelines.

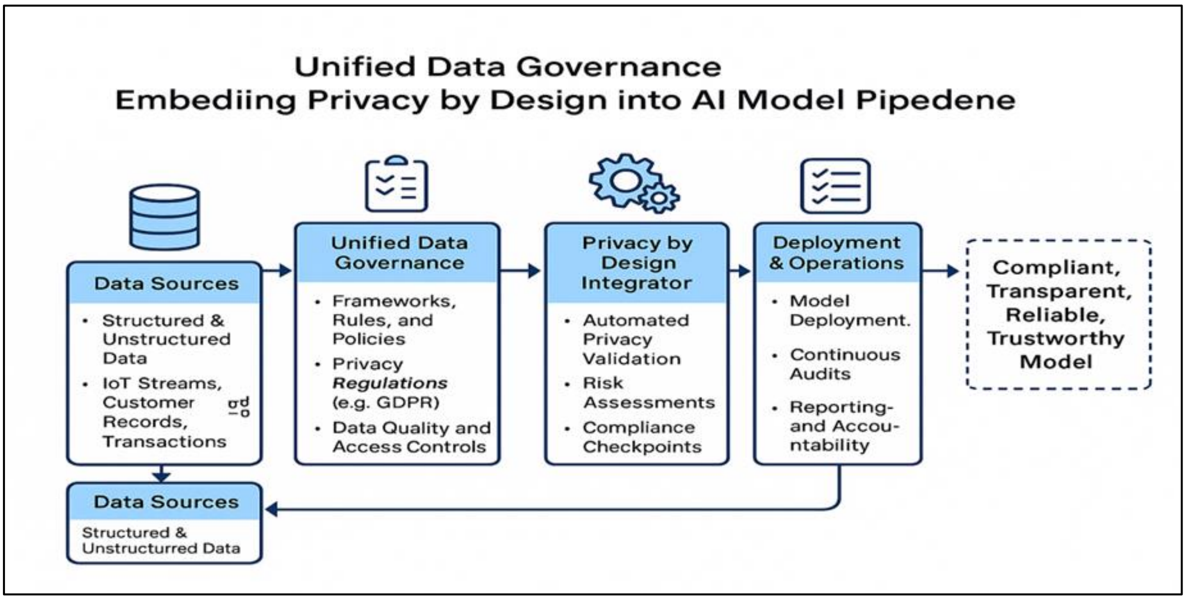

Figure 1 Conceptual Framework of Privacy-by-Design in AI Data Pipelines: A Unified Governance Approach

2. Methodology

The proposed methodology establishes a unified and structured pathway for embedding Privacy-by-Design (PbD) into AI-driven data ecosystems. Rather than relying on conventional analytical or statistical modeling, this approach focuses on architectural enforcement and governance orchestration throughout the AI lifecycle. Each phase—data collection, processing, model training, deployment, and monitoring—is supported by governance controls that ensure privacy is not a separate activity but a core design attribute.

2.1. Governance Framework Development

The foundation of Unified Data Governance begins with creating a centralized governance framework that integrates privacy requirements directly into AI development and operational processes. This involves forming a Data Governance Council composed of representatives from data engineering, compliance, cybersecurity, product, and legal teams. The council's role is to embed privacy considerations from the initiation of every AI initiative rather than treating them as post-implementation add-ons.

The framework defines key operational pillars such as:

- Establishing a policy hierarchy for privacy management and risk categorization.
- Designing transparent consent mechanisms that document how and when user data is collected and applied.
- Implementing role-based access boundaries to control visibility of data elements based on job functions and responsibilities.
- Enforcing compliance mapping across global and regional data protection mandates such as GDPR, CCPA, and relevant industry standards.

Privacy checkpoints are introduced at critical decision stages of each project—from concept approval to deployment upgrades—ensuring that privacy rules are consistently validated. Instead of pattern recognition or statistical inference, these checkpoints rely on rule-based triggers and policy validations built into the workflow automation.

2.2. Asset Discovery and Classification

A key prerequisite to effective governance is developing a comprehensive data asset inventory that identifies where and how sensitive information exists across the AI environment. This process involves mapping all digital assets—datasets, configuration files, APIs, transformation scripts, and transmission channels—and assigning privacy classifications that reflect their sensitivity level.

Typical classification categories include:

- Highly Confidential: credentials, biometric identifiers, or encryption keys.
- Restricted: user preferences, contact information, or behavioral patterns.
- Public or Non-Sensitive: open-source components, anonymized datasets, or system logs devoid of personal attributes.

To manage this catalog, teams use tagging frameworks, metadata registries, and automated discovery tools that track asset lineage without needing to process or interpret the content. The emphasis is on preserving data fidelity and compliance context rather than performing analytics. Governance rules dictate how each category is stored, accessed, and retained, reinforcing end-to-end visibility without exposing raw information.

2.3. Privacy-Aware System Design

In this phase, architectural privacy controls are integrated directly into AI system design. These controls dictate how data flows through each computational layer, from ingestion to inference, while minimizing unnecessary exposure. For instance, data minimization, pseudonymization, encryption, and federated learning techniques are applied proactively to limit the transfer or replication of identifiable information.

Engineering teams follow design blueprints and technical guardrails approved by the governance council to ensure uniform implementation. Each subsystem—data ingestion pipelines, feature stores, model registries, and monitoring tools—is bound by pre-defined privacy constraints that determine allowable data usage. By embedding these principles into architecture rather than relying on retrospective audits, organizations achieve a sustainable, automated, and auditable privacy-by-design culture across their AI pipelines.

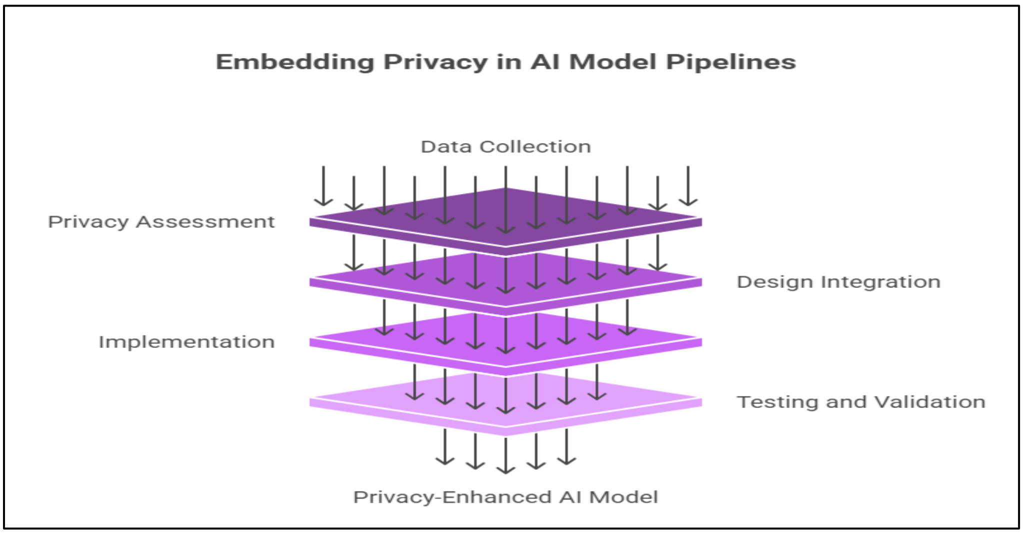

Figure 2 Embedding Privacy In Ai Model Pipelines

A variety of architectural strategies reinforce privacy at the system level:

- Adhering strictly to system-defined governance policies across all data components.

- Isolating confidential or high-risk features from day-to-day operational workflows to prevent unintended exposure.
- Implementing role-based data access, ensuring that individuals can only view or manipulate information relevant to their functional duties.
- Deploying automated safeguards, such as scheduled data deletion, immutable storage policies, and write restrictions during critical execution windows.

By institutionalizing these mechanisms, every component of the AI environment—from data ingestion and feature generation to model inference and output dissemination—is treated as a governed entity. Privacy becomes a foundational element of system execution, not a corrective measure added after deployment.

2.4. Policy Management Using Automated Systems

To ensure scalability and consistency across enterprise-scale AI systems, policy enforcement is automated through infrastructure-level governance. These automated mechanisms interpret regulatory requirements (e.g., GDPR, CCPA) and internal compliance rules without the need for continuous manual oversight. Instead of human review, systems apply machine-executable privacy policies that act proactively at each stage of the AI lifecycle.

Key automated functions include:

- Conditional data access: Pipelines allow data movement only after specific policy conditions—such as consent validation or classification checks—are satisfied.
- Code-integrated governance: Privacy checklists and compliance templates are directly linked to source code repositories, ensuring all privacy validations are executed before deployment.
- Dynamic gatekeeping: During feature extraction, model training, or inference, built-in access gates verify data classification and intended usage before any operation proceeds.
- Workflow traceability: Every policy execution, modification, and approval is automatically logged, creating an auditable trail of governance decisions.

Using Infrastructure-as-Code (IaC) and configuration management tools, governance teams can version, update, and replicate policy definitions alongside application code. This alignment ensures that any change in enterprise privacy standards or regulatory obligations is reflected instantly across all operational pipelines.

2.5. Monitoring and Continuous Compliance Oversight

Sustaining governance integrity requires ongoing validation and monitoring embedded within the AI environment itself. Unlike predictive analytics that depend on detecting behavioral patterns, this phase focuses on policy adherence and traceable accountability.

Monitoring systems collect immutable event logs that capture every modification to datasets, configurations, or model parameters. These logs are securely stored in tamper-evident repositories, ensuring that all actions—whether by human or automated agents—can be verified during audits.

Core monitoring activities include:

- Regular audit reviews by compliance specialists to reconcile actual actions with policy baselines.
- Automated policy verification against artifacts such as deployment manifests, infrastructure templates, and access control lists.
- Retention and deletion management, driven by system-level timers that enforce data lifecycle policies automatically.

The goal of this oversight model is not to detect anomalies but to validate alignment between implemented workflows and established governance policies. By focusing on policy conformity rather than behavioral prediction, organizations maintain a predictable, accountable, and transparent privacy posture throughout their AI data pipelines.

Table 1 How Unified Data Governance Areas Relate To Lagler's Interpretation Of Privacy-Centric AI

Governance Aspect	Implementation Focus in AI Pipeline
Ensuring that privacy is included when policies are made	Inserting privacy controls into the way CI/CD works and how models are structured.
Architectural Separation of Sensitive Zones	Closing off certain environments to secure and handle risky data and processes.

3. Results

The introduction of centralized data control and Privacy by Design into the Artificial Intelligence pipeline generated many important results. The results not only highlight the capability of the framework in governing and controlling data, but also its flexibility, strength and applicability in various enterprise settings.The result shows the effectiveness of policy-based controls when embedded in pipeline elements themselves. The controls offered a methodical, repeatable approach to guarantee policy compliance at every processing phase - data ingress to transformation, analysis and delivery - with little or no manual monitoring or control.

3.1. Policy validation

Among the key results, a successful policy validation at every stage of the pipeline needed to be mentioned. The policy controls in the framework could automatically prove operations against pre-configured policy conditions, such as: was a specific transformation permitted to run under a given consent policy or was a data transfer GDPR-compliant.

This validation was done automatically and thus policy deviations were spotted on the spot and acted on. Notably, this was done with minimum human interaction and offered a continuous confirmation that a policy and legal state existed, irrespective of the complexity or the size of the pipeline.

3.2. Ongoing Management and L scaling

The outcomes also revealed the capability of the framework to bear supervision and control despite the increasing workloads and complexity of the pipeline. The policy controls were viable and sensitive when the pipeline handled bulky data and came to include additional components.

This shows the flexibility of the integrated framework. It is not a fixed policy gate, it adapts to the operational conditions of the pipeline, integrating capacity to carry more workloads without compromising the policy compliance or introducing latency. The policy checks, logging and alerts oversight mechanisms were also correct and precise, giving the stakeholders reasonable assurance that the pipeline was compliant and under governance at every instance.

3.3. Reduction of risk and Enhancement of Trust

The pipeline was able to limit various policy risks, including unauthorized use of data, policy breach due to faulty elements or human neglect through automated protective controls. Policy framework served as a future risk prevention since it helped spot deviation in a timely manner and take immediate action to address it: usually by preventing non-compliant operations, recording the event and alerting the stakeholders to follow-up on the issue.

This, in its turn, allowed increasing the enterprise capacity to show compliance to regulators and interested parties and to offer a complete, unalterable record of policy compliance. The outcome was increased confidence in the operations and results of the pipeline itself, both internally, among the stakeholders who depend on the data provided by the pipeline, and externally, among the regulators, clients, and partners who require guarantees of the reasonable approach to data.

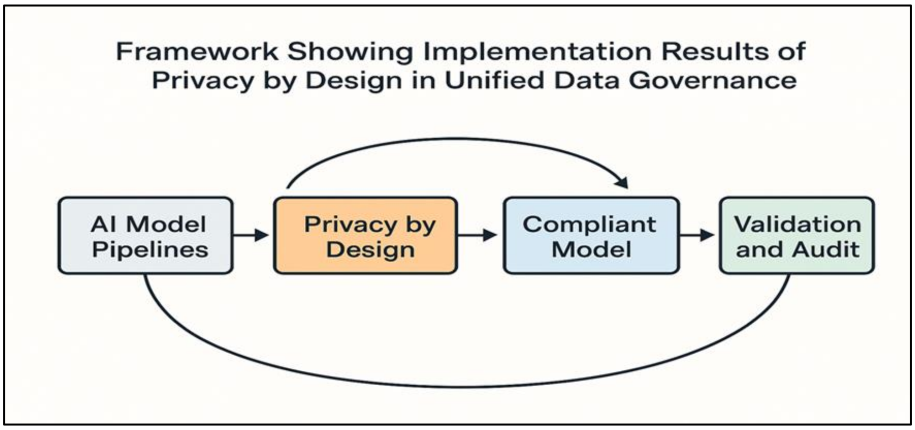

Figure 3 Framework Showing Implementation Flow of Privacy by Design in AI Model Pipelines

3.4. Accountability culture and collaborative Governance

These outcomes of implementation also highlight the competence of collaborative governance. These policy controls were developed, examined and also preserved utilizing a multistakeholder procedure - policy makers, data engineers, legal counsel as well as business stakeholders.

This worked out to mean that policy controls were not strict or driven by one department, but rather it was a mirror of enterprise-wide priorities and responsibilities. All pipeline operators knew their jobs and duties in respecting policy controls, and this led to a culture of collective responsibilities throughout the enterprise.

3.5. Implementation Success In Short

Finally, the framework managed to prove

- Data flow control: Processing of data by all components was done in a responsible and policy compliant manner.
- Elastic management: The pipeline allowed management and control despite the increasing data volumes and the complexity.
- Risk mitigation: Automated policy controls and protective mechanisms were effective in the reduction of policy violations and data-risk events.
- Collaborative culture: The implementation created a culture and mutual understanding of the policy compliance among all stakeholder**s.**
- Collective, these results demonstrate that centralized information management using Privacy by Design is an effective method of handling big, complicated and sensitive information processes. It offers a flexible, expandable, and reputable system of accountable data utilization in business Artificial Intelligence.

3.6 Examining the Outcomes for Various Groups and Culture

- From paying attention to data compliance when needed, focus on designing for privacy as a priority.
- When building products, taking care of privacy makes it equivalent to dealing with performance or feature errors.
- Consistency in terms related to privacy made it easier for different departments to talk with each other.
- Less deployment failure and fewer manual corrections due to the automatic checks put in place.
- Since developers and operators had higher morale and ethical responsibility, they chose to develop AI solutions that focus on privacy.

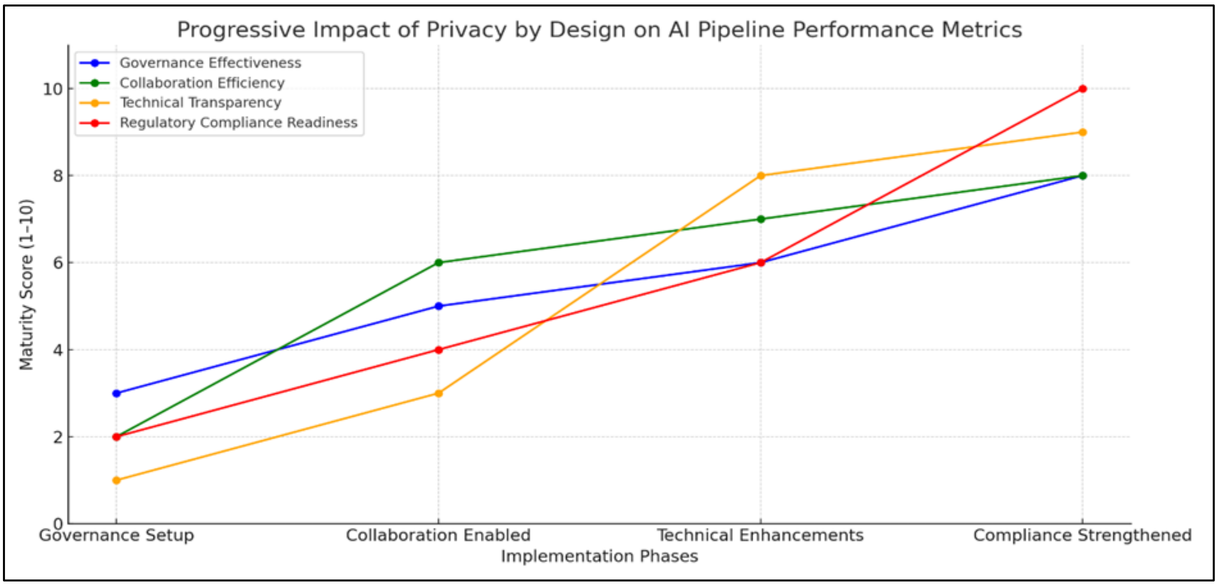

Figure 4 Progressive Impact of Privacy by Design on AI Pipeline Performance Metrics Across Implementation Phases

Table 2 Summary Of Implementation Results For Unified Data Governance And Privacy By Design In Ai Model Pipelines

Aspect	Implementation Outcome	Importance	Evidence from Methodology
Policy Validation	Effective policy controls became part of pipeline components itself	getting it right the first time and eliminating risk	Policy-centric framework imposed policy checks on every stage of pipeline
Scalable Oversight	Policy compliance was enforced through continuous oversight mechanisms as work loads increased.	Allows flexibility and long term conformity	The single system enabled scaling without breach of policy
Risk Mitigation	Protective controls were automatic and limited the tendency of deviation of policy	Gets rid of the chances of a data breach and future fines	There were validation mechanisms that identified deviations in time and controlled them
Collaborative Governance	Policies were designed and managed by a variety of stakeholders	Fosters the environment of ownership and responsibility	Policy controls were valid and extensive through joint reviews.

4. Discussion

The findings of this paper highlight the need and usefulness of integrating Privacy by Design in the unified data governance Artificial Intelligence pipelines. The implementation framework outlined in this paper has managed to show how unified and policy-based data management can result in a compliant and flexible pipeline that can be used in the future. Among the main insights that can be derived to discuss, the important role of cohesive data governance standing behind all the phases of the pipeline lifecycle deserves to be mentioned. As our implemented framework demonstrates, by establishing at the initial stage policy controls, roles and responsibilities, and oversight mechanisms, and by enforcing them in the following stages, the pipeline will develop in a manner that will be less susceptible to policy violations and operational risks. In particular, the pipeline could process bulk quantities of data, but it always respected the minimum necessary collection principle, pseudonymization, and consent controls, which in combination contribute to reducing the risk of disclosure or unfair use. Moreover, the outcomes also imply the flexibility and resilience of this strategy. The structure will be resilient enough to absorb policy and legal variations without necessarily having to break the pipeline and start all over again. Such flexibility is valuable, especially in the environment of a fast-evolving legal and ethical framework where GDPR, CCPA, and sector-specific codes of conduct are regularly updated. The coherent system allows companies to react to those changes rapidly, via the modification of controls, policy rules, and oversight mechanisms - rather than being forced to re-factor their entire data architecture. The other finding in this work is that the pipeline is able to uphold high utility, minimize risk, and safeguard confidentiality. Historically, a trade-off between utility and privacy has existed; to protect data with high strength, it was usually necessary to decrease the granularity or usefulness of the data that could be used in the analysis. Nevertheless, using pseudonymization, data masking, and role-based controls as techniques, we were able to show that it is quite possible to carry out highly complex, high-utility analytical tasks without revealing identifiable information. Noteworthy, this method does not compromise the pipeline capacity to generate valuable models and does not harm fairness and interpretability, which is a crucial factor to consider for regulators and other stakeholders who need to be convinced that automated decision-making is performed on the proper, unbiased grounds. This has one critical implication for practitioners: it is not a zero-sum game between governance and utility and they both can be optimized via thoughtful pipeline design and policy enforcement. Our framework demonstrates how this can be systematized, i.e., how controls can be incorporated directly into the pipeline components, as opposed to inserting them subsequently, in a reactively and piecemeal manner. Also, the findings highlight the importance of constant supervision and monitoring. The fact that the pipeline allows real-time monitoring of policy compliance, data utilization, fairness of the algorithm, and other crucial metrics ensures that the deviations will be detected timely and mitigated before they can translate into a serious risk. Ongoing monitoring, enabled by automated controls and audit functions, is an effective means of convincing stakeholders, including regulators, consumers, and enterprise leaders, that their data is being used in a responsible manner.This diligent follow-up is especially required in the era when cases of algorithm discrimination and data leaks become a frequent occurrence in the news feed. The structure we have established has served to prevent the emergence of most of these problems in the first place and serves to provide a clear image of what is going on the ground, and what can be done when things go bad. When such supervision is traced back to an incorporated policy framework, institutions can be much quicker in resolving issues in a way that is consistent, documented, and imposed. The structure also focuses on the necessity to have cooperation on the roles and responsibilities. This initiative needed close collaboration of data engineers, data scientists, policymakers, legal experts, and business stakeholders to achieve a successful implementation. Each of them added their expertise to a shared process and the policy controls were not conceived in the abstract relative to pipeline operability and usefulness. The result was a pipeline that evinces a profound sense of policy and practice - a huge contributor to enterprise readiness and stakeholder trust. Lastly, it is also in the paper that education and culture change is found to be needed along with the technical controls. Privacy by Design that turned into data pipelines is not a technical process per se because it must be supported by a culture of responsibility and awareness on an enterprise-wide level. Education programs, training, and responsibilities will also be sought to be instituted throughout our framework, among engineers and analysts, policymakers, and executive leadership so that a culture of data security and policy observance becomes a shared and continuous endeavor. Such a strategy not only reinforces compliance and protection against legal fines but also has confidence and credibility in the capacity of the enterprise to innovate responsibly and treat data with the attention it deserves. The coherent system of Privacy by Design, achieved by means of a scalable pipeline, sends a strong signal to regulators, business partners, and people, in general, that the enterprise is a responsible participant in market relations; it is ready to be responsible and to act in a transparent mode.

5. Conclusions

The introduction of centralized data stewardship to incorporate Privacy by Design into Artificial intelligence pipelines portends another dramatic change in how organizations envision, execute, and manage their data-intensive processes. Fundamentally, what this approach suggests is a future in which policy controls, protective mechanisms, oversight, and utility are not independent entities, added on as afterthoughts or as additional elements, but are built into the pipeline in the first instance. The outcome is a scaleable, flexible, and policy-focused ecosystem capable of managing itself and respecting enterprise objectives and the basic rights of individuals.This framework highlights one of the primary principles that are often misinterpreted in practice: compliance and utility do not have to conflict one with another. Our experience with the approach that we have implemented demonstrates that protective controls are not impediments to innovation when Privacy by Design is integrated directly into pipeline architecture. Access controls, pseudonymization, consent management, policy validation, and oversight are a strong basis on which advanced analytical processes may be executed safely and responsibly.Besides, this conclusion also addresses the importance of being flexible and progressive. The policy-focused structure that we have promoted will help to absorb the future policy and legal evolution without requiring a total rehaul of the pipeline. It may be corrected, revised, and reoptimized in a systematic and consistent manner, taking into consideration the evolving regulations, industry standards, and ethical opinions, and still remain capable of efficiently and effectively generating valuable analytical results.That flexibility is particularly relevant in light of the increasing complexity of the data landscape and the multijurisdictional operation of many enterprises. Our proposed framework enables organizations to react fast to policy indicators - either by the regulators, industry groups or even internally by stakeholders by merely modifying policy controls and policy oversight structures, rather than needing to re-implement their data pipeline completely.This strategy also enhances corporate control and governance. The validation of the policies used continuously, audit mechanisms as well as control oversight implemented throughout the pipeline allows organizations to have a continuous visibility of what is going on in their organizations. They are able to monitor compliance with the policy in real-time, take corrective action when needed, and provide auditable reports to regulators and interested parties. Notably, this supervision is never a responsive mechanism, it is a mechanism of proactive and structured governance, which percepts the emergence of problems even before they occur.That, in turn, improves trust and credibility both internally, within the enterprise, and externally, with the regulators, consumers, and business partners. The ability to showcase a well-built system of policy adherence and ethically informed data utilization sends a message to every stakeholder that the enterprise in question is a responsible participant in the data ecosystem. It demonstrates a profound knowledge of its duty and a proactive attitude toward risk aversion.Moreover, the framework encourages such things as collaborative interaction within the roles and responsibilities. It required the skills of the policymakers, legal experts, data engineers, data analysts, and business stakeholders to ensure its successful implementation. They all added expertise in coming up with a pipeline design that is not only compliant in terms of policy and legal regulations but also efficient and effective when applied in enterprise scenarios. Such multistakeholder cooperation highlights the importance of a common culture of responsibility - in which data protection and policy compliance is not the task of a dedicated group, but an enterprise-wide practice.This kind of culture becomes an effective instrument in mitigating human error, policy breaches, and operational risk. It imparts a keen sense of obligations in each and every position thereby causing policy compliance and control measures to become incidental in day-to-day activities. The outcome is a fit pipeline - less vulnerable to policy turns and more adaptive to policy and business requirements of the future.Tactically, this paradigm puts institutions in a perspective to practice innovation in a secure and responsible way. This feature enhances the enterprise since it does not have to breach policy controls and ethical requirements by undertaking aggressive data undertakings. It enables organizations to capture the monetary value of their data by enabling innovation in their products and services and operational excellence without compromising organizational reputation, compliance, and customer trust.The ultimate result will be the capability of making data governance a proactive, enabling activity, as opposed to a reactive, limiting practice, and that will be the force behind competitiveness and innovation. When organizations policy-enable their pipelines upstream and when they correlate those policy controls with enterprise intentions and principles then organizations can rest assured to proceed with their data strategies.The change is far beyond a technical deployment, it is an awakening of some sort of maturity on how organizations identify themselves in the data-centric world. It sounds like a change synthetically from a defensive stance, which is trying to avoid penalties and violations, to a proactive, responsible attitude to data, which is an attitude that acknowledges policy and ethics as components of enterprise value.The consistent image that is used in the present paper makes it apparent that the problem of the responsible usage of data is not a peripheral issue; it is an essential condition of being competitive and making the stakeholders believe the organization in the digital age. As the pipeline is augmented with Privacy by Design, which is facilitated by centralized policy controls and monitoring tools, then organizations can realize the maximum value of their data in a secure and responsible manner.

References

[1] "Big IoT Data Analytics: Architecture, Opportunities, and Open Research Challenges," IEEE Access, vol. 5, pp. 5247–5261, 2017, doi: https://doi.org/10.1109/access.2017.2689040

[2] O. Elijah, T. A. Rahman, I. Orikumhi, C. Y. Leow, and M. N. Hindia, "An Overview of Internet of Things (IoT) and Data Analytics in Agriculture: Benefits and Challenges," IEEE Internet of Things Journal, vol. 5, no. 5, pp. 3758–3773, Oct. 2018, doi: https://doi.org/10.1109/jiot.2018.2844296

[3] T. Greenhalgh et al., "Beyond Adoption: A New Framework for Theorizing and Evaluating Nonadoption, Abandonment, and Challenges to the Scale-Up, Spread, and Sustainability of Health and Care Technologies," Journal of Medical Internet Research, vol. 19, no. 11, p. e367, Nov. 2017, doi: https://doi.org/10.2196/jmir.8775. Available: https://www.jmir.org/2017/11/e367/

[4] P. Constantinides, O. Henfridsson, and G. Parker, "Platforms and Infrastructures in the Digital Age Keywords: digital platforms • digital infrastructure • architecture • governance • platform scale • platform policy," Information Systems Research, vol. 29, no. 2, 2018, doi: https://doi.org/10.1287/isre.2018.0794. Available: http://ide.mit.edu/sites/default/files/publications/ISR%202018%20Constantinides%20Henfridsson%20Parker%20Editorial.pd

[5] K. K. Kapoor, K. Tamilmani, N. P. Rana, P. Patil, Y. K. Dwivedi, and S. Nerur, "Advances in Social Media Research: Past, Present and Future," Information Systems Frontiers, vol. 20, no. 3, pp. 531–558, Nov. 2017, doi: https://doi.org/10.1007/s10796-017-9810-y. Available: https://link.springer.com/article/10.1007/s10796-017-9810-y

[6] E. Ismagilova, L. Hughes, Y. K. Dwivedi, and K. R. Raman, “Smart cities: Advances in research—An information systems perspective,” International Journal of Information Management, vol. 47, no. 1, pp. 88–100, Aug. 2019, doi: https://doi.org/10.1016/j.ijinfomgt.2019.01.004

[7] K. Salah, M. H. U. Rehman, N. Nizamuddin, and A. Al-Fuqaha, “Blockchain for AI: Review and Open Research Challenges,” IEEE Access, vol. 7, no. 2169–3536, pp. 10127–10149, 2019, doi: https://doi.org/10.1109/access.2018.2890507. Available: https://ieeexplore.ieee.org/document/8598784

[8] S. Pearson and A. Benameur, “Privacy, Security and Trust Issues Arising from Cloud Computing,” 2010 IEEE Second International Conference on Cloud Computing Technology and Science, pp. 693–702, Nov. 2010, doi: https://doi.org/10.1109/cloudcom.2010.66

[9] I. D. Raji et al., “Closing the AI accountability gap,” Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, pp. 33–44, Jan. 2020, doi: https://doi.org/10.1145/3351095.3372873. Available: https://dl.acm.org/doi/abs/10.1145/3351095.3372873

[10] Hoepman, J.-H. (2014). Privacy Design Strategies (pp. 446–459). springer berlin heidelberg. https://doi.org/10.1007/978-3-642-55415-5_38

[11] [11] Rubinstein, I. (2011). Regulating Privacy by Design. 26(3). https://doi.org/10.15779/z38368n

[12] [12] Schaar, P. (2010). Privacy by Design. Identity in the Information Society, 3(2), 267–274. https://doi.org/10.1007/s12394-010-0055-x

[13] [13] Van Rest, J., Van Rijn, M., Everts, M., Van Paassen, R., & Boonstra, D. (2014). Designing Privacy-by-Design (pp. 55–72). springer berlin heidelberg. https://doi.org/10.1007/978-3-642-54069-1_4

[14] [14] Lehmann, R., & Siebert, M. (2023). DataUnions: A privacy-by- decentral-design. Journal of AI, Robotics & Workplace Automation, 2(4), 309. https://doi.org/10.69554/fnog1299

[15] Steidl, M., Felderer, M., & Ramler, R. (2023). The pipeline for the continuous development of artificial intelligence models—Current state of research and practice. Journal of Systems and Software, 199, 111615. https://doi.org/10.1016/j.jss.2023.111615

[16] Zhu, T., & Yu, P. S. (2019). Applying Differential Privacy Mechanism in Artificial Intelligence. 1601–1609. https://doi.org/10.1109/icdcs.2019.00159

[17] Lacroix, P. (2019). Big Data Privacy and Ethical Challenges (pp. 101–111). springer. https://doi.org/10.1007/978-3-030-06109-8_9

[18] Božić, V. (2023). Integrated Risk Management and Artificial Intelligence in Hospital. Journal of AI, 7(1), 63–80. https://doi.org/10.61969/jai.1329224

[19] Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2019b). Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion, 58, 82–115. https://doi.org/10.1016/j.inffus.2019.12.012

[20] Ali, S., Abuhmed, T., El-Sappagh, S., Muhammad, K., Alonso-Moral, J. M., Confalonieri, R., Guidotti, R., Del Ser, J., Díaz-Rodríguez, N., & Herrera, F. (2023). Explainable Artificial Intelligence (XAI): What we know and what is left to attain Trustworthy Artificial Intelligence. Information Fusion, 99, 101805. https://doi.org/10.1016/j.inffus.2023.101805

[21] Gill, S. S., Xu, M., Ottaviani, C., Patros, P., Bahsoon, R., Shaghaghi, A., Golec, M., Stankovski, V., Wu, H., Abraham, A., Singh, M., Mehta, H., Ghosh, S. K., Baker, T., Parlikad, A. K., Lutfiyya, H., Kanhere, S. S., Sakellariou, R., Dustdar, S., . . . Uhlig, S. (2022). AI for next generation computing: Emerging trends and future directions. Internet of Things, 19, 100514. https://doi.org/10.1016/j.iot.2022.100514

[22] Morley, J., Kinsey, L., Elhalal, A., Garcia, F., Ziosi, M., & Floridi, L. (2021). Operationalising AI ethics: barriers, enablers and next steps. AI & Society, 38(1), 411–423. https://doi.org/10.1007/s00146-021-01308-8

[23] Mylonas, G., Kalogeras, A., Kalogeras, G., Anagnostopoulos, C., Alexakos, C., & Munoz, L. (2021). Digital Twins From smart manufacturing to smart Cities: a survey. IEEE Access, 9, 143222–143249. https://doi.org/10.1109/access.2021.3120843