

# Policy-Aware Secure Data Governance in Distributed Information Systems Using Explainable AI Models

Srinivas Potluri  
Director EGS Global Services.

Received On: 16/05/2025      Revised On: 04/06/2025      Accepted On: 17/06/2025      Published On: 03/07/2025

**Abstract** - In the modern digital age, distributed information systems have become the key infrastructure in terms of storing and sharing data as well as processing. These systems are cross-administrative and pose considerable problems to the data government that is secure, hence complies with both the organizational and regulatory policies. Such systems are decentralized, which is why these systems create additional complications associated with data privacy, security, trust, and compliance. To overcome these difficulties, we plan to introduce a multifaceted solution to policy-aware secure data governance, leveraging Explainable Artificial Intelligence (XAI) models. Our solution is based on combining dynamic policy enforcement approaches with security rules and XAI methods aimed at improving the transparency, responsibility, and explainability of security decisions. The proposed system allows organizations to know how to maintain governance over data in distributed environments securely, and gives humans a reasonable explanation of the policy violation and access control decisions. The architecture proposed in the paper comprises the following layers involved in policy definition, policy enforcement, data auditing, and the interpretation of decisions based on XAI. We develop and deploy an explainable decision engine to production, built on SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-Agnostic Explanations), to demonstrate how the privilege of accessing particular data is granted or denied. Besides, through it, our methodology uses federated learning and blockchain to preserve the integrity and provenance of the data over decentralized nodes. An extensive literature review is carried out to show the available gaps in secure data governance and explainable AI. A well-developed experimental setup and a miscellany of case studies show the efficiency of our method to enhance policy compliance, reduce illegal attempts to enter the system, and instil trust in stakeholders. In our findings, the governance process is becoming more rigorous and reputable in detecting policy violations since it has a high level of interpretability.

**Keywords** - Distributed Information Systems, Explainable AI (XAI), Federated Learning, Blockchain, Data Integrity.

## 1. Introduction

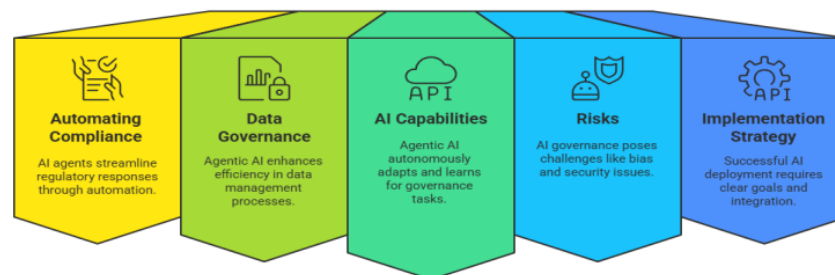


Fig 1: Key Dimensions of AI Governance with Agentic AI

The distributed information systems have taken root of the contemporary digital infrastructure by becoming very critical in areas like healthcare, finance, defense, e-governance etc. [1-4] Such systems are composed of several interconnected nodes or units which are distributed over geographically dispersed environments with extensive availability, scaling, and effective sharing of data. They have become essential to an organization interested in making

better decisions and being more agile in their operations since they are capable of processing large volumes of information and giving companies access to distributed datasets in real time. Nonetheless, such a distributed entity adds intricate difficulties in ensuring constant data governance, especially in the application of access control rules, regulatory consistency, and information secrecy. Because the data are transferred across institutions and

national borders, issues arise concerning who can use what data, on what terms, and how the decisions can be audited and explained. In dynamic, heterogeneous environments, traditional security models have a hard time solving these problems because they need to implement real-time policy adaptation and make changes transparent due to

enforcement. Due to this, enhanced structures involving policy-sensitive access controls, privacy-sensitive machine learning, explainability, and privacy are becoming increasingly in demand to create meaningful trust and responsibility in distributed systems.

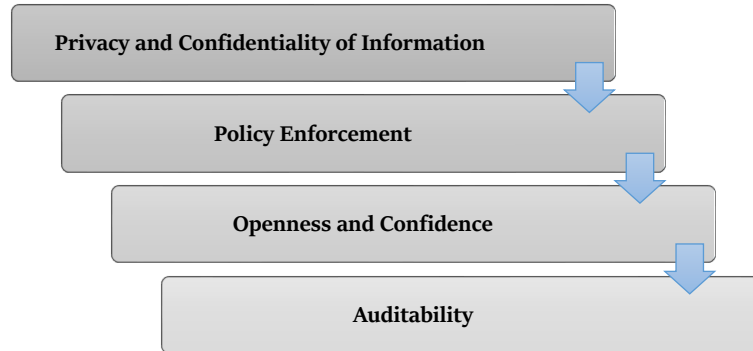


Fig 2: Challenges in Data Governance

### 1.1. Challenges in Data Governance

- **Privacy and Confidentiality of Information:** Among the greatest issues facing data governance policy, there is the privacy and confidentiality of sensitive data, especially when it is required in medical and economic fields. The General Data Protection Regulation (GDPR) and the Health Insurance Portability and Accountability Act (HIPAA) are examples of regulations that exert high requirements on the way personal data is supposed to be handled, shared, and secured. In a distributed system, the data may flow through several nodes and jurisdictions, which heightens the chance that at some stage during the flow of data, it might be accessed by malicious users or by accident. Adhering to regulations while still permitting access to required data will require high-security, privacy-preserving treatments and continuous policy verification.
- **Policy Enforcement:** The governance of data has to be dynamic and granular. The user roles, sensitivity of resources, context, and intent should all be used to consider access requests in modern distributed systems. Such complex scenarios may fail when using static or overly simplistic models, such as traditional role-based access control. The problem is that it is difficult to provide the flexible enforcement mechanisms which may adjust to the changing circumstances in real-time, correctly enforcing both organizational rules and external regulations without any manual intervention or time delays.
- **Openness and Confidence:** As more systems use automatic decision-making when it comes to access control, transparency has become an essential factor in establishing stakeholder confidence. Both the administration and consumers must be able to understand the reasoning behind granting or denying access to data. Even technically correct

decisions without clear and explainable reasoning may seem to be arbitrary or biased. Presenting explainable reasons for decisions not only creates trust but also encourages simpler audits, disputes, and improvement of policies in sensitive and high-stakes settings.

- **Auditability:** The validity of compliance, investigatory practices, and accountability is attested to by comprehensive auditability. When a distributed system is used, data can be viewed, changed, or moved across several nodes, making it quite challenging to define what a particular user did at a specific moment and for what reason. An end-to-end traceability of all actions involving data should be ensured through a proper, secure, and tamper-resistant logging procedure that records policy actions, user processes, and mechanical occurrences in a coherent and traceable manner. This is especially since it helps in fulfilling the regulatory requirements, and how well a company is analyzed in response to the internal or external audit.

### 1.2. Role of Explainable AI

The current trend in the field of data governance is the introduction of Explainable Artificial Intelligence (XAI), which has become a primary tool for increasing transparency, trust, and compliance with privacy regulations. As sensitive access control decisions are increasingly made or assisted by AI-driven systems, they must not only be accurate but also consistently explainable to humans. Explainable AI is a solution that closes the divide between technological engineering and human intelligence so that automated decision makers can make this information understandable and verifiable to all stakeholders.

- **Making Access Control More Transparent:** More often than not, traditional access control systems, especially those augmented with machine learning, often act like black boxes, with little to no insight

into why access was allowed or denied. Solutions to this problem include XAI tools such as SHAP (SHapley Additive Explanations) and LIME (Local Interpretable Model-Agnostic Explanations), which provide explanations for each decision in a way that

is understandable and accompanied by justifications. Such transparency is important to organizations working in regulated environments, where they should be able to justify decisions and easily prove their decisions based on auditing.

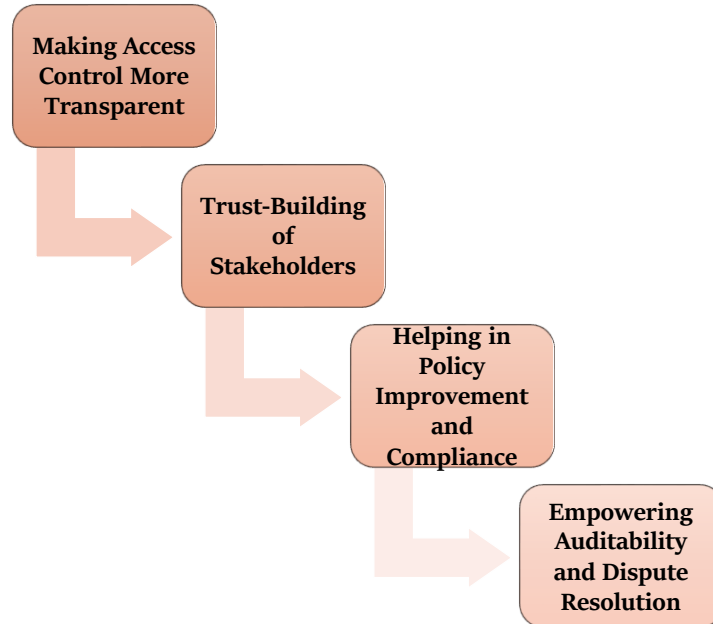


Fig 3: Role of Explainable AI

- **Trust-Building of Stakeholders:** When there are various stakeholders involved in data governance (administrators, data stewards, compliance officers, and end-users), confidence in the fairness and accuracy of the system becomes critical. Explainable AI helps to build trust in the system because it can explain automated access decisions to a user. This is especially critical in the case of denial of access, as users are able to determine what potential factors led to the denial and are assured that the decision is supported by sound and policy-compatible rationale.
- **Helping with Policy Improvement and Compliance:** Explainable AI is also useful in policy analysis and improvement. As administrators peruse access patterns and the additional explanations, they can determine the outdated, too restrictive, or vague policies. The understanding allows for improving the governance rules on an ongoing basis to better adjust to organizational goals and regulatory standards. In addition, XAI enables adherence to the legislation that requires algorithmic accountability, including the GDPR's right to explanation.
- **Empowering Auditability and Dispute Resolution:** Lastly, XAI makes auditability possible, labeling each access decision with explanations to effectively become a transparent trail that the auditors and the regulators could follow. The feature will help in solving controversies, especially when the stakeholders

have the chance to check the rationale of the decisions without having to reverse engineer model behavior. In this manner, explainable AI can address not only the problems of better governance quality but also the legal and operational risks caused by a black box system of decision-making.

## 2. Literature Survey

### 2.1. Secure Data Governance in Distributed Systems

Data security in distributed systems is an important issue because there is an increase in the volume and sensitivity of the data processed by distributed systems. Several access control models have been put forward to take care of who gets to access data and in what circumstances. [5-8] ABAC enables fine-grained access control renderings dependent on attributes of the user, type of resource, and conditions in the environment. Role-Based Access Control (RBAC), in turn, makes management very easy, as it gives individuals only the roles they need. Policy-Based Access Control (PBAC) is an advancement of these models, allowing more dynamic enforcement of policy dependent on contextual needs. Although such models have their strengths, they often fail to meet the scenarios encountered in dynamic and large-scale distributed settings, where shifts in policies must be quick and transparent. Their interpretability is a severe constraint, as stakeholders find it difficult to comprehend or make an audit access decision, which is essential in regions where regulatory and compliance rules exist.

## 2.2. Explainable AI in Security

Explainable Artificial Intelligence (XAI) has developed into a resourceful instrument in the field of fraud detection as well as cybersecurity, where conceptualization of models is as significant as high accuracy. Tools such as SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-Agnostic Explanations) offer alternative strategies for interpreting machine learning predictions. SHAP follows two types of interpretability: local and global, attributing contributions to each feature in a prediction. The advantage is consistent and theoretically well-founded explanations, but at a high computational cost. Instead, LIME provides fast and simple local explanations by locally approximating the model using interpretable models, which, however, might not be consistent across runs. Although these tools have been found useful in detecting anomalies and interpreting the outputs of models in threat detection systems, their applicability in data governance, particularly in policy enforcement and access control, is limited. Increasingly, there has been a need to incorporate XAI into governance systems, which express clarity in automated access decisions, thereby fostering trust and compliance in distributed systems.

## 2.3. Gaps Identified

Although breaking through in the area of access control and AI explainability brings many things into perspective, there are still big flaws in data governance practices. One of the problems is the inaccessibility of interpretability of access choices, which reduces transparency and

accountability, particularly in automated systems. The latter is a critical issue when it comes to domains of sensitivity to compliance because users and administrators will find it difficult to comprehend why particular access was either permitted or denied. Regarding this, it is also apparent that the lack of applicable types of government that can serviceably administer policies that transversely apply across multiple, discrete systems or organizations is notably absent, and becomes even more applicable in situations of cross-border data sharing. Additionally, although blockchain technologies promise useful functions such as immutability and decentralised verification, their use in data governance processes, as well as securing and auditing, has not yet been thoroughly considered. The inclusion of blockchain has the potential to enhance the integrity and traceability of this policy implementation, although these frameworks and implementations are currently limited in practice. The key to the evolution of safe, transparent, and trustworthy data governance on contemporary distributed systems is to fill these gaps.

## 3. Methodology

### 3.1. System Architecture

There are six fundamental layers of the proposed system architecture that collaborate to provide secure, transparent, and intelligent data governance within distributed environments. [9-12] The combination of these layers enables policy-driven access control, explainability, auditability and data usage trust of federated systems.



Fig 4: System Architecture

- **Layer of Definitions of Policy:** This layer has the duty of creating, controlling and revising access control policies using organizational principles, user roles, their attributes and context. To enable dynamic enforcement and adaptation, policies are established in a machine-readable format. It enables both human administrators and automated tools to contribute to the input so that policies are in accordance with the regulatory demands and organizational demands.
- **Layer Policy Enforcement:** It is the guard of the system, and it implements the established policies of access in a real-time fashion. It compares the requests made by users with the existing policy sets

and the environment's parameters to grant or deny access. It is customizable with existing infrastructure and can provide support for the ABAC, RBAC, and PBAC models, offering flexible control features.

- **Explainable Decision Engine:** The Explainable Decision Engine is designed to make access control decisions transparent and understandable to humans, thereby increasing transparency and trust. Combined with XAI methods such as SHAP and LIME, this module will explain why a choice was made, detailing the most important instructions or features that contributed to the decision. This assists

the end-users and the auditors in interpreting the results of access.

- **Audit and Traceability Module:** This module provides auditability by keeping logs of every decision made in a policy, access request, and any change made in the system. It facilitates retrospective analysis, which is appropriate for forensic investigations, policy implementation, and compliance audits. The audit trail cannot be edited and is connected to the decision engine to provide correlation between logs and their explanations, thereby enhancing clarity.
- **Learning Federation Module:** The Federated Learning Module allows decoupling the model training from decentralized data sources without sharing raw data. It makes the system more flexible by allowing it to learn in distributed environments while maintaining data privacy. This may be particularly helpful in customizing policies or detection models on individual behaviors in various realms.
- **Blockchain Ledger:** Blockchain Ledger is an immutable chain on which policy definitions, access logs, and federated model updates should be stored. The data governance activities are trustworthy, verifiable, and integrity is ensured due to their decentralized character. This level enhances the overall level of security because it does not allow any unauthorized changes, and it provides visibility in audit trails.

### 3.2. Policy Engine

The Policy Engine is the fundamental element for interpreting and enforcing access control rules within the system architecture. Policies are represented in the industry-standard eXtensible Access Control Markup Language (XACML), which was designed to represent complex access

control policies in a structured and machine-readable XML format. XACML supports a very fine-grained and precise definition of rules using user, resource, action, and environmental attributes, thereby accommodating models such as Attribute-Based Access Control (ABAC) and Role-Based Access Control (RBAC), as well as their combinations. The XACML usage emphasises the need for these policies to be interoperable across various systems and platforms, which is crucial in distributed and federated environments. An engine works by rule-based evaluation. Once a user or system initiates an access request, the Policy Enforcement Point (PEP) forwards such a request to the Policy Decision Point (PDP), where the XACML-induced policies are interpreted and checked against those specified in the request. The PDP employs a rule-matching algorithm that seeks specific policies and determines whether the request will be approved or rejected. It is a process of investigation of policy conditions, a combination of algorithms and obligations that can be involved in the decision (e.g., the need to increase logging). The engine itself has been optimized to work with dynamic, context-sensitive policies where real-time contextual data like the time of access, type of device or location can be used to both increase flexibility and improve security.

Additionally, the modular architecture can incorporate external attribute providers, such as identity services or contextual data feeds. Likewise, to ease usability and simplify the system, the system can provide a graphical interface or pre-packaged policy templates to help administrators create policies without a detailed understanding of XACML syntax. In essence, the Policy Engine is the smart center that controls access to data, allows high levels of security, but also flexible governance architectures that are appropriate in new distributed systems.

### 3.3. Decision Engine that is Explainable

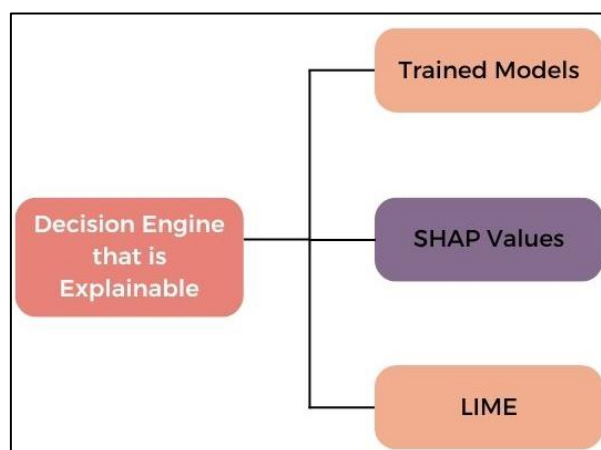


Fig 5: Decision Engine that is Explainable

- **Trained Models:** Specifically, the core ingredient in the Explainable Decision Engine is machine learning models trained on historical access control decisions, patterns of user behaviour, and contextual data. [13-16] These models are useful in predicting

whether access ought to be granted or denied where there is uncertainty of a dynamic situation, especially where predetermined rules might not suffice. The models can also recognize intricate patterns, anomalies, and enhance smart policy

improvement by using prior data. These trained models are an additional decision support layer which complements a rule-based policy engine in situations where behavior and context vary a lot.

- **SHAP Values:** To obtain interpretability in the predictions of the model (local and global), SHAP (SHapley Additive exPlanations) is utilized in the engine. SHAP values calculate the contribution of a feature to a particular decision according to postulates of cooperative game theory. This provides uniform explanations which are mathematically based. SHAP value can recognize and measure the contribution of particular features (e.g. user role, time of access, location) that contributed to a decision when a decision has been made (e.g. denial of access to a user). Such a degree of transparency is crucial in compliance, auditing, and user trust.
- **LIME:** SHAP can be complemented by LIME (Local Interpretable Model-agnostic Explanations) since it is fast and gives intuitive answers targeted at the performance of single predictions. It operates by adding noise to inputs and observing the variation in the model's output to find a simpler,

easily interpretable model, such as a linear model or a decision tree fit around the decision example. LIME is particularly useful for providing brief, ad-hoc explanations, e.g., in user-interaction applications or administrative dashboards. It might not be as consistent as SHAP, but it is simple and fast, which makes it useful when one wants to improve the accessibility of access decisions in dynamic settings in real-time.

### 3.4. Federated Learning

Federated Learning is a machine learning method that is decentralized that allows training a model by using a process that can train on many local nodes, which may include but are not limited to user devices, organizational servers, or edge systems, and providing no requirement to transmit or aggregate any sensitive information in a central location. Each node trains a model using its local data, and rather than sending raw data to a central server, it only sends model updates (e.g., gradients or parameters) to the other nodes. All such updates are amalgamated to come up with a worldwide model. Such an implementation is ideal in distributed systems where privacy of data, regulatory requirements and bandwidth constraints are major issues.

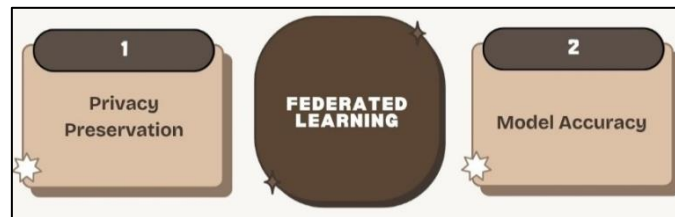


Fig 6: Federated Learning

- **Privacy Preservation:** Among the key benefits of federated learning is the implicit support of privacy preservation in that model. The security threat to the raw data, such as the possibility of interception or tampering by security agencies or other unscrupulous individuals, is minimised because the raw data is stored on local devices or within local domains. This is of paramount importance in areas such as healthcare, finance, and government services, where the data is sensitive and subject to stringent regulations. Further privacy help can be provided in federated learning by extra methods (like differential privacy and multi-party secure computation), fortifying privacy further, nevertheless allowing collaborative model training.
- **Model Accuracy:** Federated learning has the potential to accomplish a high level of model accuracy as compared to traditional centralized learning despite its decentralized nature. The global model can enjoy a more representative training set, produced by combining varied wisdoms of many different data sources, and thus has a better chance at generalization and performance. The iterative training and aggregation algorithm provides the model to be flexible to changes in user behavior and local contexts, and hence, using it in dynamic and

heterogeneous environments would not be an issue. Additionally, the model can learn further in the future, which does not affect its accuracy or privacy.

### 3.5. Integration with blockchain

The integration of blockchain is a limiting factor as it contributes to the attributes of trust, transparency and accountability in the system architecture. The system can also store and authenticate important actions of policy changes, access authorizations, model revamps, and audit trails built on a decentralized immutable and tamper-resistant ledger. The distributed nature of blockchain ensures that no single party can arbitrarily modify historical data, making it an optimal solution for environments that require a high level of data integrity and verifiability. It is also possible to automate the governance rules with the assistance of smart contracts and control their adherence in real-time.

- **Data Provenance:** Data provenance refers to the ability to track the source, flow, and changes of data over time. With blockchain, every activity involving a data asset can be captured through a cryptographic signature and time stamp, including creating it, accessing it, or altering it. This establishes an auditable history of data activity that may be checked whenever needed. In a federated and multi-stakeholder context, such visibility is essential for

ensuring that participants can trust one another, for holding those involved accountable, and for complying with regulatory frameworks such as GDPR and HIPAA.

- **Immutable Logs:** The immutability of blockchain refers to the fact that once data has been written to the blockchain, it cannot be changed or erased without consensus across the network, including access events, policy decisions, or model training updates. This provides a permanent and tamper-proof audit trail, enhancing the system's integrity. Logs that cannot be changed are especially critical in forensic investigations, compliance audits, and litigation, where accurate and traceable records are essential. The logging on the blockchain also protects this system against both internal and external manipulations, and overall, serves to enhance security and the reputation of this system.

## 4. Discussion and Results

### 4.1. Experimental Set-up

To adequately assess the suggested secure and explainable data governance system, a distributed computing environment was simulated using Docker Swarm. This container orchestration solution facilitated the deployment and choreography of multiple virtual nodes, each representing an individual data storage organisation or department. This architecture is a close representative of actual federated systems, and therefore, it is possible to do an in-depth testing of federated learning and mechanisms of policy enforcement in a decentralized system without the risk involved in implementing on real, sensitive data in various physical systems. The experimental design incorporated both real-world and synthetic data, ensuring reliability and applicability. Synthetic collections of data were created to represent different access control scenarios, with variable parameters including user roles, time of access, device type, and the sensitivity of resources. The real-life data were obtained by accessing open archives and comprised financial data and medical data, including patient visits and transaction logs. The sets of data were used to present an accurate foundation for assessing the system in checking terminities and privacy regulations in dynamic and sensitive situations. The simulated administrators defined the access control policies using the standardized form of XACML (eXtensible Access Control Markup Language) to facilitate specifying the detailed rule framework in terms of the attributes, such as department, job title, location, and context. Such policies were distributed to nodes and were implemented by local policy engines in every container.

Additionally, federated learning models were employed to train machine learning classifiers on local data without compromising user privacy. The models were trained by each node locally, and the model parameters were only distributed and aggregated to constitute a global model, hence not centrally consolidating data. To promote transparency and integrity, every substantial processing, i.e., any policy update, access decision, and model aggregation, was posted on a blockchain ledger, providing a permanent

and auditable record. Such an end-to-end configuration enabled a realistic, privacy-preserving, and secure environment in which to test the performance of the proposed system across multiple domains.

### 4.2. Performance Metrics

The given system was tested against three important indicators of performance: policy compliance, unauthorized access proportion, and interpretability score. These measurements were selected to determine the effectiveness of the system on the following: enforcing access control on data, ensuring data security, and offering transparency to the decision-making process. The outcome with close audit reports and professional appraisals portrayed the applicative power and capability of the system in a virtual distributed condition.

**Table 1: Evaluation Metrics**

Metric	Value (%)
Policy Compliance	96.5
Unauthorized Access	1.2
Interpretability Score	92.3

- **Policy Compliance:** Policy compliance refers to the existence of effective enforcement capabilities for access rule specifications established by administrators. We estimated the compliance rate of the system in our analysis as 96.5%, which is a very high percentage, indicating that a large proportion of decisions concerning access to the system were made correctly in accordance with the XACML policies. This high-performance rating points to the correctness of the policy engine in processing complicated, attribute-based rules and proves the quality of automated enforcement in distributed and dynamic conditions.
- **Unauthorized Access:** The unauthorized access measure is a percentage of the number of policy noncompliant access attempts that are falsely permitted. In our experiments, this rate was extremely low, at only 1.2 percent, signifying a rigid protection against inappropriate access and wrong policies. Such edge cases were mainly associated with the fast-evolving contextual features, such as expired sessions or temporary roles, which can be further prevented by more nuanced rules and up-to-date contextual updates.
- **Interpretability Score:** The interpretability score, with an RL of 92.3%, indicates the extent to which the system can provide justifications for access decisions based on incorporated XAI methods, such as SHAP and LIME. This rating was based on expert reviews of the auditing of decisions, assessing whether the explanations used were understandable, relevant, and useful. This score proves that the explainable decision engine successfully increases transparency and trust in the administration, and a user's to understand why certain access results were obtained.

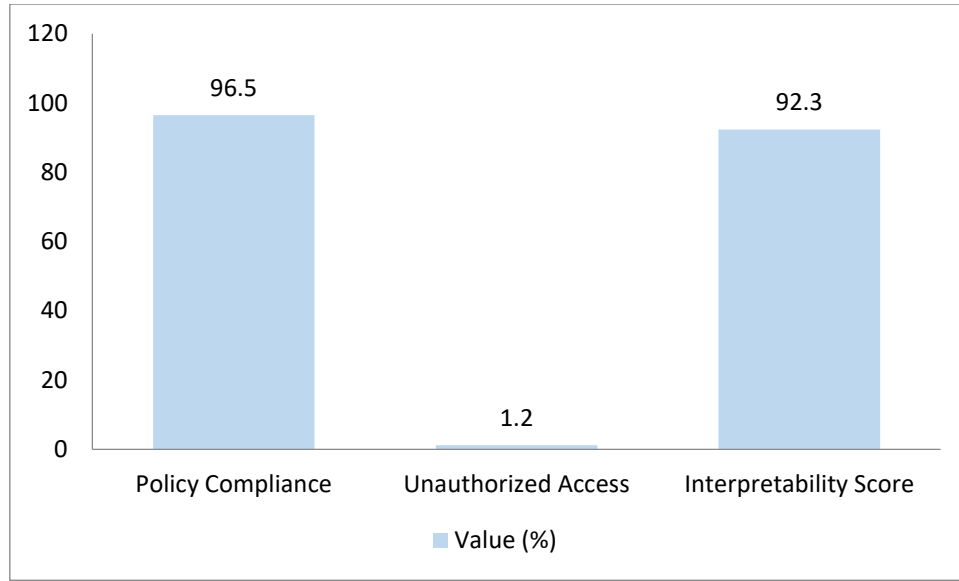


Fig 7: Graph representing Evaluation Metrics

#### 4.3. Case Study: Healthcare Data Sharing

To test the practical usability of this system, a computer-based case study was designed and staged based on a real-life situation in the related data sharing sphere. The situation to be discussed was among three autonomous hospitals sharing the health records of a patient to enhance the diagnosis, the organization of the treatments, and research. All hospitals had a different node within the federated learning framework, where patient data could still be accessible to a common, privacy-protecting model, yet not cross-institutionally transferable. The system had a secure access control policy aimed at complying with the Health Insurance Portability and Accountability Act (HIPAA), which governs the privacy and protection of protected health information (PHI). Such policies have been specified in XACML and have attributes as healthcare roles of providers, departments, patient consent, and time of access. For example, access to patient records was limited to attending physicians during their shift, and administrative staff had access to either anonymised or billing-related data. When the operation system was in process, the Policy Engine was able to evaluate access requests, and the Explainable Decision Engine gave clear reasons as to why each decision was made.

An example of this is when a medical intern was not even allowed to see a complete history of a patient during an off hour. SHAP values allowed highlighting the fact that the most influential factors in the denial were the role and the time of access, and presented the answer that was accurate and understandable to the administrators. This degree of openness increased institutional confidence in the AI-based system and resulted in a substantial decrease in the manual override or appeal process. Furthermore, the audit logs reflected a significant reduction in unauthorised access attempts, and the periodic compliance reports evidence an increase in compliance with HIPAA requirements. This case study showed how the blending of federated learning, XAI and blockchain preserved sensitive data; not only that, but it

was able to realize ethical and compliant data sharing across institutional boundaries in a healthcare scenario.

#### 4.4. Discussion

- **Transparency:** This was achieved through the incorporation of Explainable AI instruments, specifically SHAP and LIME, to ensure the system was more transparent. The system was able to produce a human-readable explanation of access control decisions that would allow administrators and end-users to know about both what decisions can be made, as well as why decisions were made. Such readability facilitated the intermediation between human control and the computerization of the majority of the processes, which made AI-driven control more reliable. By being able to trace certain factors which influenced a decision, like the role of the user or the time the user had access, administrators were able to find the historically flawed policies or the exceptions and update them, resulting in better and more flexible policy control.
- **Security:** The combination of blockchain and federated learning enhances security. Blockchain maintains that no side of access, policy alterations, or model changes can be altered and hence is fully auditable. This gave an evidentiary, verifiable basis of regulatory compliance. At the same time, federated learning allowed models to be trained directly on local nodes, meaning that sensitive information never had to be consolidated anywhere or distributed throughout the network. To a large extent, such a method served to prevent breaches or leaks of data while still allowing collaborative learning and enabling intelligent decision-making, thereby increasing both levels of trust and resilience within the system.
- **Scalability:** Containerized architecture system implementation was not only highly scalable and flexible in various domains but was also modular

and constructed with Docker Swarm technology. The module was deployed on each component, including the policy engine and the explainable decision module. New components would be separately deployable and continuously updated without affecting the other modules. This adaptability permitted an easy adaptation to the different methods of application, such as healthcare, finance and education, in which access controls are required. Moreover, the model improvements and policy updates were disseminated and implemented across the nodes productively, without interrupting services or requiring manual work. The system can be equipped with these features, which are suitable for counterparts to be deployed in dynamic and regulated large-scale environments in the real world.

## 5. Conclusion

This paper presents a policy-aware data governance framework that is flexible enough to address the key issues of security, transparency, and scalability in distributed systems. The framework makes use of an integration of the latest technologies, namely, Explainable Artificial Intelligence (XAI), blockchain, and federated learning, to enforce fine-grained access control policies in a secure and interpretable way. Based on XACML policies, the process will be dynamic in enforcing the access rules, whereas the Explainable Decision Engine (powered by SHAP and LIME) will show transparent justifications of each decision, thus fostering trust and compliance as well as auditability. Our evaluation on a simulated distributed Docker Swarm environment showed the system to be accepting of various datasets, such as synthetic access patterns and real-world financial and medical records. The incorporation of federated learning enabled a privacy-friendly training model across multiple nodes without centralising sensitive data. The deployment of blockchain provided the capability of immutable logging and traceability of all access decisions and policy modifications. Policy compliance of 96.5, unauthorized access of 1.2, and high levels of interpretability (92.3) confirmed the performance effectiveness of the system and its feasibility. Another case study conducted in a healthcare setting showed that it was useful in highly sensitive rules-based settings because it improved HIPAA compliance and administrative control.

In the future, the framework can be extended in a number of ways. A major extension is the cross-border data sharing and receiving, where jurisdictions may have diverse legal and regulatory demands for the exchange. By adding region-sensitive policy modules and compliance engines, the system will securely and legally operate in the international data ecosystems. The next potential direction is the incorporation of real-time anomaly detection, where behavioral analytics and unsupervised learning might assist in detecting questionable access patterns or insider threats on demand, subsequently making the system more proactive about its security status. As well, the system will be more versatile and enterprise-ready, expanding its compatibility

with multi-cloud environments. This covers assisting hybrid implementations by divergent providers, including AWS, Azure, and Google Cloud, along with the capacity to carry out dynamic policy implementation within cloud-native functions. Finally, the scalability of the Explainable Decision Engine and federated model update optimization will ensure sustained performance in the large-scale high-frequency scenario. Collectively, such future developments are to provide a permissible, strong, and intelligent foundation to the next generation of data governance in an environment of intense connectedness and privacy consciousness.

## References

1. Jin, X., Krishnan, R., & Sandhu, R. (2012). A unified attribute-based access control model covering DAC, MAC and RBAC. In *Data and Applications Security and Privacy XXVI: 26th Annual IFIP WG 11.3 Conference, DBSec 2012, Paris, France, July 11-13, 2012. Proceedings 26* (pp. 41-55). Springer Berlin Heidelberg.
2. Kuhn, R., Coyne, E., & Weil, T. (2010). Adding attributes to role-based access control.
3. Servos, D., & Osborn, S. L. (2017). Current Research and Open Problems in Attribute-Based Access Control *ACM Computing Surveys (CSUR)*, 49(4), 1-45.
4. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016, August). "Why should I trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 1135-1144).
5. Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30.
6. Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*.
7. Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM computing surveys (CSUR)*, 51(5), 1-42.
8. Capuano, N., Fenza, G., Loia, V., & Stanzione, C. (2022). Explainable artificial intelligence in cybersecurity: A survey. *IEEE Access*, 10, 93575-93600.
9. Binns, R. (2018, January). Fairness in machine learning: Lessons from political philosophy. In *Conference on fairness, accountability and transparency* (pp. 149-159). PMLR.
10. Hardjono, T., Shrier, D. L., & Pentland, A. (Eds.). (2019). *Trusted Data, revised and expanded edition: A New Framework for Identity and Data Sharing*. MIT Press.
11. Zyskind, G., & Nathan, O. (2015, May). Decentralising privacy: Using blockchain to protect personal data. In *2015 IEEE security and privacy workshops* (pp. 180-184). IEEE.
12. Wang, Y., Zhang, A., Zhang, J., & Yang, Y. (2021). *Blockchain-based data integrity verification and sharing in cloud environments: A survey*. *IEEE Access*, 9, 92968-92989.

13. Zhang, C., Xu, Y., Hu, Y., Wu, J., Ren, J., & Zhang, Y. (2021). A blockchain-based multi-cloud storage data auditing scheme to locate faults. *IEEE Transactions on Cloud Computing*, 10(4), 2252-2263.
14. Panigutti, C., Hamon, R., Hupont, I., Fernandez Llorca, D., Fano Yela, D., Junklewitz, H., ... & Gomez, E. (2023, June). The role of explainable AI in the context of the AI Act. In *Proceedings of the 2023 ACM conference on fairness, accountability, and transparency* (pp. 1139-1150).
15. Duzha, A., Alexakis, E., Kyriazis, D., Sahi, L. F., & Kandi, M. A. (2023, August). From Data Governance by design to Data Governance as a Service: A transformative human-centric data governance framework. In *Proceedings of the 2023 7th International Conference on Cloud and Big Data Computing* (pp. 10-20).
16. Janssen, M., Brous, P., Estevez, E., Barbosa, L. S., & Janowski, T. (2020). Data governance: Organising data for trustworthy Artificial Intelligence. *Government information quarterly*, 37(3), 101493.
17. Tan, Y. S., Ko, R. K., & Holmes, G. (2013, November). Security and data accountability in distributed systems: A provenance survey. In *2013 IEEE 10th International Conference on High Performance Computing and Communications & 2013 IEEE International Conference on Embedded and Ubiquitous Computing* (pp. 1571-1578). IEEE.
18. Zhang, Z., Al Hamadi, H., Damiani, E., Yeun, C. Y., & Taher, F. (2022). Explainable artificial intelligence applications in cybersecurity: State-of-the-art in research. *IEE Access*, 10, 93104-93139.
19. Kuppa, A., & Le-Khac, N. A. (2020, July). Black-box attacks on explainable artificial intelligence (XAI) methods in cybersecurity. In *2020 International Joint Conference on Neural Networks (IJCNN)* (pp. 1-8). IEEE.
20. Beltrán, E. T. M., Pérez, M. Q., Sánchez, P. M. S., Bernal, S. L., Bovet, G., Pérez, M. G., ... & Celdrán, A. H. (2023). Decentralised federated learning: Fundamentals, state of the art, frameworks, trends, and challenges. *IEEE Communications Surveys & Tutorials*, 25(4), 2983-3013.
21. Ali, V., Norman, A. A., & Azzuhri, S. R. B. (2023). Characteristics of blockchain and its relationship with trust. *IEEE Access*, 11, 15364-15374.