WCRG Global Research
WORLD COMET RESEARCH GROUP

International Journal of Data Engineering and Intelligent Computing
Vol. 2, No. 2, 2019
Doi: 10.67228/30715717/IJDEIC-2019PII5M8C
PP. 01-12

Original Article

Cross-Domain Policy Enforcement for Data Access and Governance in Data Mesh Systems

Dr. Tharindu Jayasinghe

Associate Professor, University of Peradeniya, Sri Lanka.

Received: 09-04-2019 Revised: 09-16-2019 Accepted: 09-27-2019 Published: 10-09-2019

ABSTRACT

Data Mesh is an emerging architectural paradigm aimed at addressing the complexities of scaling data management in large, decentralized organizations. With the rapid expansion of data and its increasing domain-specific nature, enforcing consistent data access policies across multiple domains has become one of the critical challenges for data governance. This paper explores the complexities and methodologies for implementing cross-domain policy enforcement within a Data Mesh system, where data ownership and access are distributed among different business units. We provide a detailed analysis of the issues surrounding data access, security, and compliance in such an architecture. Furthermore, we propose a novel framework for managing and enforcing cross-domain policies in a federated data governance context, leveraging existing access control models and modern policy enforcement technologies. Through case studies and real-world examples, we illustrate the challenges and solutions in maintaining security, transparency, and compliance in Data Mesh systems. The paper concludes with a discussion of future directions for research and development in cross-domain policy enforcement, emphasizing scalability, consistency, and the integration of new technologies.

KEYWORDS

Data Mesh; Cross-Domain Policy Enforcement; Data Governance; Data Access Control; Decentralized Data Architecture; Federated Governance; Role-Based Access Control (RBAC); Attribute-Based Access Control (ABAC); Policy Enforcement Framework; Data Security; Data Privacy; Compliance; Federated Identity Management.

Copyright © The Author(s) This article is distributed under a
Creative Commons Attribution-Share A like 4.0 International License.

ISSN: 3071-5717
PP: 01-12.

1. INTRODUCTION

1.1. Introduction to Data Mesh Architecture

Data Mesh is an innovative approach to handling large-scale data architectures that emerged as a response to the shortcomings of traditional monolithic data systems, which often fail to scale effectively in organizations with diverse and rapidly growing data needs. Unlike traditional data architectures that centralize data in large, monolithic data lakes or warehouses, Data Mesh decentralizes data ownership across various domains or business units. Each domain is responsible for managing its own data as a product, empowering teams to own the data they create and use. This approach emphasizes the principles of domain-oriented decentralized governance, self-serve data infrastructure, and the use of a federated model for governance, aiming to bring agility, scalability, and better alignment of data resources with business needs. In Data Mesh, the focus shifts from a centralized control model to a distributed, collaborative data management ecosystem where teams have autonomy and accountability over their own data.

1.2. Importance of Data Access and Governance

Data access and governance are critical components in the realm of modern data architecture. As organizations scale, particularly with the growth of data in distributed systems, managing who has access to what data and ensuring that data is handled appropriately becomes more complex. Proper data governance involves the creation of rules, policies, and processes that govern how data is accessed, shared, and utilized across an organization. In a Data Mesh architecture, where data is decentralized and spread across multiple domains, ensuring appropriate access and governance is crucial for maintaining consistency, security, and compliance. Without robust data access and governance mechanisms, an organization risks data breaches, compliance violations, and inefficiencies due to poorly managed or inconsistent data access controls. Data governance ensures that data is secure, available when needed, and meets the necessary regulatory and legal standards. It also ensures that data quality is maintained and that the right individuals and systems can access data in the proper context.

1.3. Need for Cross-Domain Policy Enforcement in Data Mesh

One of the key challenges with Data Mesh lies in enforcing consistent data access policies across multiple domains. Since Data Mesh distributes data ownership and responsibility across various business domains, each domain might have its own set of rules, requirements, and needs regarding data access and governance. However, this decentralization presents a challenge when attempting to ensure uniform policy enforcement across the entire organization. Cross-domain policy enforcement is necessary to maintain consistency in how data is accessed and used, regardless of where the data resides or which team owns it. Without proper enforcement, data could be accessed by unauthorized users, leading to security risks, legal issues, or breaches of privacy regulations. Additionally, the absence of consistent policy enforcement can lead to operational inefficiencies and make it challenging to audit data access across multiple domains. This necessitates the development of mechanisms that enable policies to be enforced consistently across different domains while maintaining the autonomy of each domain.

ISSN: 3071-5717
PP: 01-12.

1.4. Key Challenges and Opportunities

Enforcing cross-domain policies in Data Mesh systems comes with several challenges. One of the most significant challenges is maintaining the balance between decentralized ownership and centralized policy enforcement. Since each domain is responsible for its own data, it can be difficult to enforce common rules that span across domains. Additionally, the technical complexities of ensuring consistent, automated, and scalable policy enforcement across a federated environment can be daunting. Furthermore, cross-domain policy enforcement often needs to support a variety of data access models (such as Role-Based Access Control, Attribute-Based Access Control, etc.), which must be harmonized across different systems. However, despite these challenges, there are significant opportunities for improvement and innovation in this space. The development of unified policy enforcement frameworks, leveraging modern technologies such as AI and machine learning for policy automation, could streamline the process and offer scalability. Furthermore, advancements in data lineage and auditability can provide insights into how data is accessed and ensure compliance with legal and regulatory standards. The growing focus on federated identity management and self-service data access infrastructure in Data Mesh presents an opportunity to automate policy enforcement and simplify the process of data governance.

2. BACKGROUND AND RELATED WORK

2.1. Overview of Traditional Data Governance Approaches

Traditional data governance approaches were often based on a centralized model, where data was managed in large, monolithic data warehouses or lakes. In such models, data access was controlled from a single, centralized entity, typically an IT or data management team. Policies for data access, security, and compliance were defined and enforced by this central body, ensuring that all users across the organization adhered to the same set of rules. While this centralized approach worked well in smaller or less complex data environments, it became increasingly difficult to manage as organizations grew in size and data volumes exploded. These traditional systems often struggled to scale and could not accommodate the diverse needs of various business units. Moreover, the centralized control model hindered agility and made it harder to implement business-specific data strategies. The lack of ownership and accountability at the domain level led to inefficiencies, delays, and increased friction in the data access process.

2.2. Existing Research on Data Governance in Distributed Systems

As data systems became more distributed, researchers began exploring new methods for governing data in decentralized environments. Studies on distributed systems governance have focused on the challenges of maintaining consistency, security, and compliance across disparate systems. One key area of focus has been the development of distributed access control models that allow data access policies to be enforced across multiple, heterogeneous data sources. Research has also explored methods for ensuring data quality and lineage in decentralized environments, as well as techniques for managing metadata across domains. Another significant area of research has been on federated identity management, which allows for the consistent management of user identities and access permissions across multiple domains or systems. The growing importance of data privacy

ISSN: 3071-5717
PP: 01-12.

regulations (such as GDPR) has also driven research into ensuring compliance in distributed systems, especially in the context of cross-domain data access.

2.3. The Evolution from Monolithic Data Systems to Data Mesh

The shift from monolithic data systems to Data Mesh is part of the broader trend towards decentralization in technology and business operations. Monolithic systems were designed with a centralized architecture where data was stored in a single, large repository. However, as organizations grew and their data needs became more complex, these systems proved to be inflexible, slow, and difficult to scale. Data Mesh emerged as an alternative by decentralizing the data ownership across domains, treating data as a product owned by the respective business units. This approach aimed to address the scalability challenges of traditional data systems by enabling each domain to govern and manage its own data independently. As such, Data Mesh supports the creation of domain-specific data pipelines, data products, and access controls, giving individual teams more autonomy while maintaining an overarching governance framework.

2.4. Review of Relevant Policies and Technologies (e.g., RBAC, ABAC, XACML)

The evolution of data governance has seen the development of various policies and technologies for controlling access to data. Role-Based Access Control (RBAC) is one of the most commonly used models in traditional systems, where access permissions are assigned based on the roles users occupy within the organization. While RBAC simplifies administration, it lacks the flexibility required to address complex, context-dependent data access needs. Attribute-Based Access Control (ABAC) offers more granularity by evaluating attributes (such as user role, data sensitivity, and environment) to determine access rights, thus providing more fine-grained control. In a Data Mesh context, ABAC can offer a more flexible approach to governing access across diverse and distributed systems. Another key technology in this space is XACML (eXtensible Access Control Markup Language), a standard used for expressing access control policies. XACML is powerful in creating standardized, policy-driven access control systems that can be applied across distributed environments. These technologies, while powerful, need to be integrated and adapted to the decentralized, domain-driven nature of Data Mesh to ensure consistent policy enforcement.

3. DATA MESH ARCHITECTURE

3.1. Principles of Data Mesh

Data Mesh is built on four core principles: domain-oriented decentralized data ownership, data as a product, self-serve data infrastructure, and federated computational governance. The idea is to distribute data responsibility to the teams that generate and use it most, allowing those teams to treat data as a product that can be shared, accessed, and utilized with clear quality and accessibility standards. This contrasts with the traditional centralized approach to data management. A self-serve data infrastructure allows teams to easily access, process, and share data within the confines of governed boundaries, without needing to rely on a central IT team. Lastly, federated governance allows for consistent policies and standards across the entire organization while allowing domains to maintain autonomy in how they manage their data products.

ISSN: 3071-5717
PP: 01-12.

3.2. Domains and Domain-Oriented Decentralized Data Ownership

In Data Mesh, each domain is responsible for managing its own data, treating it as a product that is owned, maintained, and operated by the domain teams. This decentralization fosters agility and accountability, as the teams closest to the data have full control over its management, quality, and access. Each domain defines its own data products and ensures that these products are discoverable, accessible, and consumable by other domains. Data ownership is not limited to raw data storage but extends to ensuring the reliability, security, and compliance of the data, with each domain establishing its own set of policies.

3.3. Self-Serve Data Infrastructure

Self-serve data infrastructure allows teams within a domain to autonomously manage the ingestion, processing, and consumption of data. This removes the bottleneck of a centralized data engineering team and enables a more efficient flow of data through the organization. With a self-serve infrastructure, data engineers within domains can access standardized tools and resources to perform their work without needing complex integrations with central IT. This autonomy ensures that teams can act quickly in response to new business needs or changes in data requirements.

3.4. Federated Governance in Data Mesh

While Data Mesh emphasizes decentralization, federated governance ensures that overarching policies and standards are consistently applied across the entire organization. In a federated governance model, there is a central governing body responsible for setting high-level principles and guidelines for data security, privacy, and compliance, while each domain team is responsible for implementing these standards in a way that suits their specific needs. This allows for flexibility at the domain level while maintaining consistency across the entire data ecosystem.

4. CROSS-DOMAIN DATA ACCESS AND POLICY ENFORCEMENT

4.1. Defining Cross-Domain Data Access

Cross-domain data access refers to the ability to access data that resides in different domains across a Data Mesh architecture. Since each domain is responsible for its own data, users or systems in one domain may need to access data in another domain to perform their tasks. Cross-domain data access must be carefully managed to ensure that data is shared securely and that privacy and compliance standards are upheld. This can involve creating mechanisms for data discovery, access control, and secure data sharing between domains.

4.2. Challenges of Enforcing Policies across Multiple Domains

Enforcing policies across multiple domains in a Data Mesh environment is inherently complex due to the decentralized nature of the architecture. Each domain has its own set of policies and practices for data management, which may differ from other domains. Ensuring that access control policies, security standards, and compliance regulations are uniformly applied across the entire system without stifling the autonomy of individual domains is a significant challenge. Additionally, the lack of centralized control makes it difficult to monitor and audit access consistently across all domains.

ISSN: 3071-5717
PP: 01-12.

4.3. Technical Considerations for Cross-Domain Access Control

Technically, enforcing cross-domain access control requires the integration of multiple systems and technologies. This could involve leveraging federated identity management to ensure that users can be authenticated and authorized across different domains. Moreover, ensuring consistent access control models (such as RBAC or ABAC) across domains requires the integration of common policy enforcement tools and frameworks. Data catalogs, metadata management tools, and access control layers must be carefully integrated to enable seamless policy enforcement across the entire organization.

4.4. Role of Automation in Policy Enforcement

Given the complexity and scale of modern data systems, automation plays a crucial role in enforcing policies across domains. Automated policy enforcement frameworks can ensure that rules are consistently applied without requiring manual intervention, which is essential for maintaining the scalability of Data Mesh systems. Automation can also support continuous monitoring and auditing of data access to detect violations and provide insights into compliance. This automation is especially important in dynamic environments where data access needs may change rapidly.

5. DATA ACCESS MODELS

5.1. Role-Based Access Control (RBAC) vs. Attribute-Based Access Control (ABAC)

Role-Based Access Control (RBAC) and Attribute-Based Access Control (ABAC) are two prevalent models for managing access to resources in an organization. RBAC is traditionally used in centralized data management systems and defines access controls based on predefined roles within an organization. For example, an employee in a finance department may be granted access to financial data based on their role. While this model is simple and effective in environments where roles are well-defined, it becomes inflexible when access requirements change dynamically. In contrast, ABAC takes a more flexible, fine-grained approach by determining access based on attributes, such as user role, location, time of access, and the sensitivity of the data. In a Data Mesh environment, where data ownership is distributed across different domains, ABAC is often seen as more effective, as it allows for better context-based decision-making. ABAC allows for more nuanced control over who can access data, based on a broader range of factors, making it a better fit for decentralized data ecosystems like Data Mesh, where requirements can vary across domains.

5.2. Federated Identity Management

Federated Identity Management (FIM) refers to the technology and protocols that allow the secure sharing of identity information across different systems and organizations. In a Data Mesh architecture, where data access spans multiple domains, federated identity management enables users to authenticate once and gain access to data across different domains without needing to sign in repeatedly. FIM uses protocols like SAML (Security Assertion Markup Language) and OpenID Connect to facilitate Single Sign-On (SSO) and ensure that users' identities are consistently recognized across domains. This approach is particularly useful for managing user access to decentralized data systems, as it allows for a unified identity management framework that ensures

ISSN: 3071-5717
PP: 01-12.

seamless data access and enforces centralized authentication policies while maintaining the autonomy of individual domains.

5.3. Fine-Grained Data Access Controls for Heterogeneous Data Sources

In a Data Mesh system, data is often stored in diverse formats and resides in various systems that are owned and operated by different domains. To maintain data security and compliance, fine-grained access controls are required to govern how data is accessed and used across these heterogeneous systems. Fine-grained controls allow for specifying detailed rules about who can access specific parts of data, based on various attributes such as the user's role, the sensitivity of the data, the time of access, and more. For instance, in a health domain, access to patient records might be restricted not only based on the user's role (e.g., doctor, nurse) but also based on the type of data they need to access (e.g., medical history vs. prescription details). Implementing fine-grained access controls in a decentralized environment like Data Mesh ensures that each domain can manage its own data according to specific access needs, while still adhering to global governance and compliance standards.

5.4. Policy-Based Approaches and Standards (e.g., XACML)

Policy-based access control frameworks such as XACML (eXtensible Access Control Markup Language) are becoming essential in complex data environments like Data Mesh. XACML provides a standard for defining access control policies in a structured way, enabling centralized policy definitions while maintaining decentralized enforcement across different systems. By using XACML, organizations can define policies that determine how data is accessed across domains, based on various conditions and attributes. These policies can be fine-tuned to meet specific business needs and regulatory requirements. In a federated system like Data Mesh, the ability to enforce policies consistently across different domains and systems is crucial. XACML and other policy-based systems provide the flexibility to enforce access control rules in a distributed and scalable manner, while ensuring compliance with security, privacy, and governance standards.

6. DATA GOVERNANCE IN A DATA MESH CONTEXT

6.1. The Relationship between Data Governance and Data Access Policies

Data governance refers to the processes and rules that ensure data is properly managed, secured, and used responsibly across an organization. In the context of Data Mesh, data governance plays a critical role in ensuring that the decentralized, domain-based data products adhere to security, privacy, and compliance standards. Data access policies are an integral part of data governance, as they define the rules and guidelines on how data can be accessed, shared, and used within and across domains. In a Data Mesh architecture, where data is owned and managed by different business units, it is essential to establish governance frameworks that align access policies with organizational priorities. Data governance ensures that the right individuals or systems have access to the right data, ensuring that the data is used ethically, securely, and in compliance with legal regulations.

ISSN: 3071-5717
PP: 01-12.

6.2. Principles of Decentralized Data Governance

Decentralized data governance is a key principle of the Data Mesh architecture. It emphasizes distributing the responsibility of data ownership, stewardship, and governance to the domains that own and generate the data. In this model, each domain is responsible for managing the quality, security, and compliance of its own data products. This decentralized approach fosters greater autonomy and accountability within each domain, allowing them to implement governance practices that suit their specific needs. However, while domains maintain responsibility for their data, federated governance mechanisms ensure that overall organizational standards are met, ensuring a balance between domain autonomy and centralized oversight. Decentralized data governance also enables greater agility and faster decision-making, as each domain can tailor its governance practices to its unique requirements without relying on a centralized governing body.

6.3. Ensuring Data Security, Compliance, and Transparency in Data Mesh

In a decentralized environment like Data Mesh, ensuring data security, compliance, and transparency is critical. With data scattered across multiple domains, there are increased risks related to unauthorized access, data breaches, and non-compliance with regulations such as GDPR or HIPAA. To mitigate these risks, data security and compliance frameworks must be applied consistently across domains. Data encryption, identity management, and access controls should be implemented to protect data integrity and confidentiality. Transparency is another essential element of data governance, as organizations need to maintain detailed records of how data is accessed and used. This includes maintaining data lineage, which helps track the flow of data across systems, and logging all data access events for auditing purposes. A robust governance framework that incorporates these elements ensures that Data Mesh can scale while meeting security and compliance requirements.

6.4. Auditability and Traceability of Data Access across Domains

Auditability and traceability are key components of data governance, especially in distributed systems like Data Mesh. Auditability refers to the ability to track and record data access and changes, ensuring that organizations can monitor who accessed data, when, and for what purpose. Traceability goes a step further by enabling the tracking of data movement and transformations across domains. This is particularly important in a decentralized environment where data is spread across different systems, and the ownership and usage policies may vary. Ensuring that all data access events are logged and that the data lineage is transparent allows organizations to verify that their data access policies are being followed, and helps in identifying potential security breaches or policy violations. Furthermore, these audit and trace logs are critical for compliance, as regulatory frameworks often require organizations to maintain a clear record of data usage and access.

7. CASE STUDIES AND REAL-WORLD IMPLEMENTATIONS

7.1. Examples of Cross-Domain Policy Enforcement in Existing Data Mesh Systems

Several organizations have begun implementing Data Mesh principles, incorporating cross-domain policy enforcement frameworks to manage data access. For example, in a large financial institution, data ownership is distributed among various business units, such as retail banking,

ISSN: 3071-5717
PP: 01-12.

investment banking, and insurance. Each unit manages its own data but enforces access policies using a common federated identity management system, ensuring that users can access the data they need while maintaining compliance with financial regulations. These organizations typically use policy-based systems like XACML to ensure consistent enforcement across domains. Another example is in healthcare, where different departments manage patient data. Cross-domain policy enforcement ensures that only authorized personnel can access sensitive health records, and that the data is protected in compliance with health regulations such as HIPAA.

7.2. Lessons Learned from Industry Practices

One of the main lessons learned from the implementation of Data Mesh is the importance of establishing clear and consistent data governance frameworks across domains, while allowing flexibility within each domain. Organizations that have successfully implemented Data Mesh have emphasized the need for strong collaboration between business units and centralized governance teams. They have also realized the need to invest in automation tools to scale data access controls across multiple domains. Furthermore, organizations have found that clear data product definitions and metadata management are crucial to successful Data Mesh implementation. The integration of advanced identity management solutions and policy enforcement engines has also proven beneficial in overcoming the complexities of cross-domain data governance.

7.3. Potential Challenges and Their Solutions

Implementing cross-domain policy enforcement in Data Mesh is not without challenges. One of the main challenges is ensuring consistent policy enforcement across disparate systems and platforms. To address this, organizations have turned to standardized policy frameworks like XACML and the use of centralized policy engines. Another challenge is maintaining data quality and security while allowing for decentralized control. This can be addressed by incorporating automated data quality checks and centralized monitoring systems that help identify potential vulnerabilities. Lastly, privacy concerns and regulatory compliance can become more complicated in a decentralized model. However, these challenges can be mitigated by ensuring that all domains follow a unified governance framework and by using tools that automate compliance monitoring.

8. PROPOSED FRAMEWORK FOR CROSS-DOMAIN POLICY ENFORCEMENT

8.1. Design of a Cross-Domain Policy Enforcement Framework

A cross-domain policy enforcement framework in Data Mesh should provide centralized governance while respecting the autonomy of individual domains. This can be achieved by using policy engines that centralize policy definitions but enforce them within each domain's infrastructure. The framework should include automated tools for defining, enforcing, and auditing policies related to data access, security, and compliance across multiple domains. It should also support a flexible access control model like ABAC, ensuring that data access decisions can be made based on user attributes, data attributes, and the context of access. Furthermore, the framework should be extensible, allowing new domains and data sources to integrate seamlessly while maintaining consistent policy enforcement.

ISSN: 3071-5717
PP: 01-12.

8.2. Key Components and Technologies (e.g., Policy Engines, Data Catalogs)

Key components of a cross-domain policy enforcement framework include policy engines, which automate the decision-making process for data access control; data catalogs, which provide visibility into the data available in different domains; and identity and access management (IAM) systems, which ensure secure authentication and authorization across domains. Policy engines, such as those based on XACML, allow organizations to define policies that can be centrally managed but enforced across multiple domains. Data catalogs provide a unified view of all data products within the organization, helping to ensure that data access is governed and controlled effectively.

8.3. Mechanisms for Monitoring and Enforcing Policies across Domains

Monitoring and enforcing policies across domains in a Data Mesh requires the integration of centralized monitoring systems and decentralized enforcement mechanisms. Automated auditing tools can track data access events in real time, ensuring that policies are being followed and that any unauthorized access attempts are logged and addressed. Additionally, automated policy enforcement engines can ensure that data access decisions are made based on pre-defined rules, such as RBAC or ABAC, and that data is only accessed when appropriate. These monitoring systems should provide transparency and audit logs that can be accessed by both domain teams and centralized governance bodies.

8.4. Scalability and Performance Considerations

A key consideration in designing a cross-domain policy enforcement framework is scalability. As organizations grow and add more domains, the policy enforcement system must be able to scale without compromising performance. To achieve scalability, the framework should use distributed architectures for both policy enforcement and monitoring, enabling policies to be enforced at the domain level while still adhering to global governance standards. Performance can be enhanced by utilizing efficient policy evaluation engines, leveraging caching mechanisms to reduce the load on the system, and adopting decentralized data processing to avoid bottlenecks in centralized systems.

9. CHALLENGES AND OPEN ISSUES

9.1. Scalability of Cross-Domain Policy Enforcement

The scalability of cross-domain policy enforcement remains a major challenge in Data Mesh environments. As the number of domains and data products grows, ensuring that policies are applied consistently and efficiently becomes increasingly difficult. Additionally, scaling the monitoring and enforcement of policies across numerous systems without affecting performance can be complex.

9.2. Dynamic Policy Updates and Management

Dynamic policy updates are a challenge in environments like Data Mesh, where data access needs may evolve rapidly. Traditional approaches to policy enforcement often involve manual updates, which are time-consuming and error-prone. Automated, dynamic policy management systems are needed to ensure that policy changes are reflected across all domains in real time, without requiring manual intervention.

ISSN: 3071-5717
PP: 01-12.

9.3. Ensuring Consistency across Domains

Ensuring consistency in policy enforcement across diverse and autonomous domains is difficult, especially when different domains have different needs, priorities, and data access models. A solution to this issue is to establish a standardized set of baseline policies that all domains must adhere to, while still allowing for flexibility within each domain.

9.4. Addressing Privacy Concerns and Regulatory Compliance (GDPR, etc.)

Data privacy and compliance concerns are paramount in decentralized data environments like Data Mesh. Regulations such as GDPR require strict control over personal data, including where it is stored, how it is processed, and who can access it. Organizations must ensure that data access policies are aligned with these regulations, even as data is decentralized across domains.

10. FUTURE DIRECTIONS

10.1. Potential Improvements in Data Mesh Governance

As Data Mesh adoption grows, there will likely be improvements in governance tools and frameworks that better support scalability, compliance, and real-time policy enforcement. Future research and innovation will likely focus on creating more automated and intelligent policy enforcement systems that can scale across large, complex organizations.

10.2. Integration with Emerging Technologies (e.g., AI, Blockchain)

Emerging technologies such as AI and blockchain could play a key role in improving the efficiency and security of cross-domain policy enforcement. AI can help with automating policy decision-making and detecting anomalies in data access, while blockchain could provide a secure, immutable ledger for tracking data access events, improving auditability and traceability.

10.3. Research Opportunities in Cross-Domain Policy Enforcement

There are several opportunities for research in cross-domain policy enforcement. These include improving the scalability and performance of policy enforcement engines, exploring new access control models that are better suited for decentralized environments, and investigating how emerging technologies like AI and blockchain can be integrated into Data Mesh governance frameworks.

11. CONCLUSION

11.1. Recap of the Importance of Cross-Domain Policy Enforcement

Cross-domain policy enforcement is vital for ensuring data security, compliance, and transparency in Data Mesh systems. By ensuring that data access policies are consistently enforced across domains, organizations can maintain the integrity and reliability of their data while adhering to regulatory standards.

11.2. Summary of the Proposed Framework

This paper proposed a framework for cross-domain policy enforcement that combines centralized policy definitions with decentralized enforcement, leveraging modern policy engines,

ISSN: 3071-5717
PP: 01-12.

data catalogs, and identity management systems. This framework provides a scalable, flexible solution for governing data access in complex, distributed data environments.

11.3. Closing Thoughts on the Future of Data Governance in Data Mesh Systems

The future of data governance in Data Mesh systems will likely involve greater automation, better integration with emerging technologies, and more sophisticated models for policy enforcement. As organizations continue to decentralize their data management practices, ensuring effective governance will be key to unlocking the full potential of Data Mesh architectures.

REFERENCES

[1] Dreibelbis, A., Hechler, E., Milman, I., Oberhofer, M., van Run, P., & Wolfson, D. (2008). Enterprise master data management: An SOA approach to managing core information. Pearson Education.
[2] Eckerson, W. W. (2014). The data governance imperative. Technics Publications.
[3] Khatri, V., & Brown, C. V. (2010). Designing data governance. Communications of the ACM, 53(1), 148–152. https://doi.org/10.1145/1629175.1629210
[4] Loshin, D. (2009). Master data management. Morgan Kaufmann.
[5] Otto, B. (2011). A morphology of the organisation of data governance. ECIS 2011 Proceedings, 1–13.
[6] Weber, K., Otto, B., & Österle, H. (2009). One size does not fit all—A contingency approach to data governance. Journal of Data and Information Quality, 1(1), 1–27. https://doi.org/10.1145/1515693.1515696
[7] Zachman, J. A. (2008). The Zachman framework for enterprise architecture: Primer for enterprise engineering and manufacturing. Zachman International.