WCRG Global Research
WORLD COMET RESEARCH GROUP

International Journal of Data Engineering and Intelligent Computing
Vol. 1, No. 1, 2026
Doi: 10.XXXX/XXXXXXXX/IJDEIC-V1I1P104
PP. 31-40

Original Article

Comparative Study of Access Control Mechanisms in Data Mesh vs. Data Fabric Architectures

Niranjan[1], Sheetal Kumari[2]

[1,2]Assistant Professor, Department of IT, Acharya Narendra Dev College, New Delhi, India.

Received: 28-11-2025 Revised: 27-12-2025 Accepted: 04-01-2026 Published: 09-01-2026

ABSTRACT

The increasing complexity and volume of data in modern enterprises have led to the evolution of decentralized data architectures such as Data Mesh and Data Fabric. While both paradigms aim to democratize data access and improve scalability, their approach to data governance, particularly access control, varies significantly. This paper presents a comparative analysis of access control mechanisms in Data Mesh and Data Fabric architectures. It explores the foundational principles of both models, evaluates how they implement access policies, identity management, and compliance, and examines their adaptability to enterprise-scale requirements. By highlighting their strengths, limitations, and ideal use cases, this study aims to assist data architects and security professionals in making informed decisions when designing secure and scalable data infrastructure.

KEYWORDS

Data Mesh; Data Fabric; Access Control; Data Governance; Identity and Access Management (IAM); Zero Trust Architecture; Policy Enforcement; Decentralized Data Ownership; Metadata-Driven Security; Data Architecture.

WCRG
Global Research
ISSN: XXXX-XXXX
PP: 31-40.

1. INTRODUCTION

1.1. Background and Motivation

The exponential growth of data within modern enterprises has necessitated a paradigm shift in how data is stored, accessed, and governed. Traditional monolithic data architectures often fall short in handling distributed, large-scale data ecosystems. As businesses increasingly rely on data-driven decision-making, ensuring that data is both accessible and secure has become a strategic imperative. In this context, Data Mesh and Data Fabric have emerged as two prominent architectural approaches that aim to address the complexity of modern data landscapes. Each architecture offers a unique perspective on decentralization, integration, and automation, particularly with respect to access control—a critical component of any data governance strategy.

1.2. Importance of Access Control in Modern Data Architectures

Access control is a cornerstone of data security and governance, particularly in an era where data is distributed across hybrid and multi-cloud environments. Without robust access control mechanisms, organizations face heightened risks of data breaches, non-compliance with regulations, and internal misuse. Modern access control systems must go beyond simple authorization; they must enforce policies dynamically, support granular permissions, and integrate seamlessly with identity providers and audit systems. As data becomes more decentralized, the challenge is to maintain consistent security policies across domains without hindering data accessibility and innovation.

1.3. Overview of Data Mesh and Data Fabric

Data Mesh and Data Fabric represent two distinct yet complementary approaches to managing enterprise data. Data Mesh emphasizes decentralization, advocating for domain-oriented data ownership and self-serve infrastructure to empower cross-functional teams. It promotes federated governance, where access control is managed at the domain level, ensuring that data producers and consumers adhere to shared standards. In contrast, Data Fabric offers a centralized, technology-driven model that automates data integration and governance across disparate data sources. It leverages metadata, machine learning, and policy-based automation to create a unified data management layer. While Data Mesh focuses on organizational decentralization, Data Fabric is centered around technological unification.

1.4. Purpose and Scope of the Study

This paper aims to conduct a detailed comparative study of access control mechanisms within Data Mesh and Data Fabric architectures. It investigates how each approach handles authorization, policy enforcement, identity integration, and compliance. The study examines the strengths and weaknesses of both models, evaluates real-world tools and frameworks, and offers guidance on selecting an appropriate architecture based on access control requirements. The focus is not only on theoretical constructs but also on practical implications for enterprise-scale data systems.

ISSN: XXXX-XXXX
PP: 31-40.

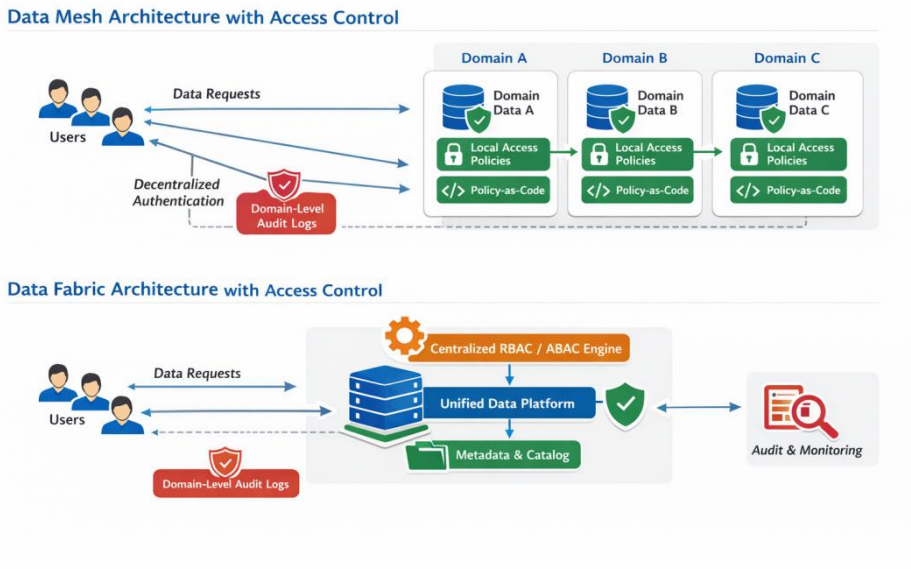

Fig 1 - Decentralized Access Control Mechanisms in Data Mesh Architecture

2. CONCEPTUAL FOUNDATIONS

2.1. What is Data Mesh?

2.1.1. Decentralized *Domain*-Driven Data Ownership

Data Mesh introduces a socio-technical paradigm that treats data as a product and distributes data ownership across domain-specific teams. Instead of centralizing data in a data lake or warehouse, it delegates data stewardship to the teams that are closest to the source of the data. These teams are responsible for maintaining the quality, availability, and security of their data products. By decentralizing control, Data Mesh aims to reduce bottlenecks and promote scalability, agility, and local autonomy—key attributes for organizations managing diverse and rapidly growing datasets.

2.1.2. Federated Governance

To maintain cohesion across distributed domains, Data Mesh employs federated governance. This model balances local autonomy with global standards by establishing a framework where governance policies such as those for access control, data quality, and interoperability are defined centrally but enforced locally. In this structure, access control policies are coordinated across domains through shared principles and interoperable tooling, allowing for consistent enforcement without undermining the independence of individual teams.

Table 1: Architecture Comparison – Data Mesh vs. Data Fabric

Feature / Aspect	Data Mesh	Data Fabric
Data Ownership	Domain-oriented (each team owns its data)	Centralized with integrated metadata layer
Governance	Federated governance per domain	Centralized governance and policies
Access Control	Decentralized, policy-as-code per domain	Centralized RBAC/ABAC enforcement

ISSN: XXXX-XXXX
PP: 31-40.

Scalability	Scales horizontally across domains	Scales through integration layers
Security Enforcement	Domain-level enforcement; may vary	Uniform enforcement across organization
Metadata Management	Domain-specific catalogs	Unified metadata layer
Data Access Auditing	Auditing per domain	Centralized audit logs

2.2. What is Data Fabric?

2.2.1. Metadata-Driven Integration

Data Fabric represents a technology-centric approach to data management that focuses on intelligent and automated data integration across the enterprise. It uses active metadata to create a contextual understanding of data across disparate systems and environments. This metadata enables dynamic discovery, classification, and policy application, including access control. By leveraging metadata, Data Fabric solutions can automate the enforcement of data security rules, track data lineage, and ensure compliance without manual intervention.

2.2.2. Centralized Orchestration and Automation

Unlike Data Mesh, Data Fabric centralizes the orchestration of data flows, governance, and security. It provides a unified control plane that manages data access, transformation, and delivery across the enterprise. Access control mechanisms in Data Fabric are often embedded within this orchestration layer, allowing organizations to define and enforce security policies uniformly. Automation capabilities further enhance security by proactively adjusting access rights based on context, user roles, or compliance requirements.

2.3. Key Differences in Architectural Philosophy

At the core, the philosophical divergence between Data Mesh and Data Fabric lies in their approach to decentralization and control. Data Mesh decentralizes data ownership and governance, emphasizing organizational transformation and domain-driven design. It relies on human-centric processes and cross-functional collaboration to scale data access securely. In contrast, Data Fabric centralizes control through intelligent systems that automate governance using metadata and AI. It aims for technical abstraction and simplification, enabling organizations to manage complex data landscapes through centralized automation. These contrasting philosophies significantly influence how access control is implemented and managed in each architecture.

3. ACCESS CONTROL IN DATA MESH

3.1. Decentralized Governance and Its Impact on Access Control

In a Data Mesh architecture, decentralized governance fundamentally reshapes how access control is implemented. Each domain team is responsible for the stewardship and security of its own data products. This distributed control model aligns with the principles of autonomy and scalability but introduces complexities in ensuring consistent access control across the enterprise. Without a central authority dictating security policies, organizations must establish clear guidelines and tools to

ISSN: XXXX-XXXX
PP: 31-40.

allow domain teams to implement their own access control mechanisms in accordance with global standards. This approach fosters agility and accountability, but it also requires strong coordination and trust among teams to avoid security gaps and policy inconsistencies.

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

Data Mesh environments often challenge traditional Role-Based Access Control (RBAC) models due to the dynamic and domain-specific nature of access requirements. While RBAC assigns permissions based on predefined roles, it can become unwieldy in a decentralized setting with many domains and evolving roles. Instead, Attribute-Based Access Control (ABAC) offers greater flexibility by allowing policies to evaluate user, resource, and environmental attributes at runtime. ABAC supports fine-grained, context-aware permissions that are well-suited to the autonomy of domain teams in a Data Mesh. For example, access can be granted based on a user's team, project, or the sensitivity of the data, allowing for a more adaptive and scalable approach to access control.

3.3. Policy Federation and Local Enforcement

To bridge the gap between centralized governance and decentralized execution, Data Mesh architectures often implement a federated policy model. In this model, enterprise-wide security principles such as data classification, compliance mandates, and identity standards—are defined centrally but interpreted and enforced locally by each domain. This requires robust tooling and governance frameworks that support the federation of policies while maintaining consistency. Tools like policy as code and decentralized policy engines (e.g., Open Policy Agent) enable domains to enforce access control locally in line with overarching enterprise directives. This blend of autonomy and standardization is key to ensuring that access control in a Data Mesh is both secure and scalable.

3.4. Challenges in Policy Standardization

Despite the benefits of decentralization, standardizing access control policies across domains is one of the most persistent challenges in Data Mesh. Domain teams may interpret governance guidelines differently, leading to inconsistent application of access controls. Furthermore, without a unified language or framework for defining policies, ensuring interoperability and auditability becomes difficult. Differences in tooling, expertise, and compliance awareness can result in vulnerabilities or inefficiencies. To address these issues, organizations must invest in education, standard templates, centralized registries for policy discovery, and shared tooling that supports consistent enforcement while respecting domain autonomy.

3.5. Case Examples or Frameworks (e.g., Open Policy Agent, Data Contracts)

Several tools and frameworks are emerging to support access control in Data Mesh environments. One prominent example is Open Policy Agent (OPA), a general-purpose policy engine that enables policy as code. OPA allows domain teams to write and enforce access control policies in a declarative language (Rego), making it easier to maintain and audit rules across distributed systems. Another concept gaining traction is data contracts, which formalize the interface between

ISSN: XXXX-XXXX
PP: 31-40.

data producers and consumers, including rules for access, schema consistency, and usage guarantees. These contracts serve as both a governance and enforcement mechanism, supporting security, quality, and compliance within a decentralized data ecosystem. By adopting such tools, organizations can achieve a practical balance between flexibility and control in their access governance models.

Table 2: Access Control Mechanisms Comparison

Mechanism	Description	Implementation in Data Mesh	Implementation in Data Fabric	Pros	Cons
Role-Based Access Control (RBAC)	Permissions assigned to roles	Domain-specific roles	Centralized roles	Simple, well-known	Limited flexibility for dynamic data
Attribute-Based Access Control (ABAC)	Permissions based on attributes/policies	Domain may define attributes independently	Centralized attribute policies	Highly flexible	Complexity in policy management
Policy-as-Code	Access policies coded as scripts	Common in domain-level governance	Less common, mostly centralized	Automatable & versioned	Requires DevOps integration
Data Tokenization & Masking	Data obfuscation for access control	Implemented per domain	Implemented centrally	Protects sensitive data	Overhead & performance impact

4. ACCESS CONTROL IN DATA FABRIC

4.1. Centralized Metadata-Driven Governance

In contrast to the decentralized nature of Data Mesh, Data Fabric emphasizes centralized governance powered by metadata. Metadata serves as the foundational layer that enables the Data Fabric to understand, classify, and manage data assets across the enterprise. Access control is governed through this metadata layer, which includes information such as data sensitivity levels, user roles, lineage, and regulatory tags. Because metadata is centrally managed, access policies can be uniformly defined and automatically applied across all connected systems. This approach not only streamlines compliance and auditability but also reduces the complexity of managing security in highly heterogeneous environments.

4.2. Automated Access Management Using Data Catalogs

Data catalogs in a Data Fabric architecture are more than just repositories of metadata—they are active governance tools that facilitate automated access control. These catalogs can integrate with identity management systems and use metadata to recommend, apply, or restrict access based on policies. For instance, if a dataset is tagged as containing personal identifiable information (PII), the catalog can automatically restrict access to users who do not have the necessary compliance training or authorization level. Furthermore, these catalogs often provide user-friendly interfaces for

ISSN: XXXX-XXXX
PP: 31-40.

requesting and approving access, reducing the administrative overhead and delays associated with manual security reviews.

4.3. Policy Enforcement Points and Control Planes

A key architectural feature of Data Fabric is the separation of the control plane (where governance policies are defined) and the data plane (where data flows and operations occur). Policy enforcement points (PEPs) are strategically embedded within the data pipeline to ensure that access controls are applied consistently, regardless of where the data resides or how it is processed. These enforcement mechanisms are centrally orchestrated, ensuring that policies are not only consistently deployed but also contextually aware—able to adjust enforcement based on factors like data sensitivity, user behavior, or external risk signals. This unified enforcement model improves both security posture and operational efficiency.

4.4. Integration with Enterprise IAM Systems

Data Fabric architectures are typically designed to integrate deeply with enterprise Identity and Access Management (IAM) solutions such as Active Directory, LDAP, or cloud-native IAM services. This integration enables fine-grained access control by mapping data access policies directly to organizational roles, departments, or user attributes managed within the IAM system. It also allows for the enforcement of advanced security protocols like multi-factor authentication (MFA), session timeouts, and dynamic access based on device or location. The close coupling of IAM systems with data governance platforms ensures that access control policies are enforceable, auditable, and adaptable to enterprise security standards.

4.5. Example Tools (e.g., IBM Cloud Pak, Informatica, Talend)

Several commercial Data Fabric solutions offer robust and automated access control capabilities out of the box. IBM Cloud Pak for Data leverages active metadata, data virtualization, and built-in governance features to enforce access policies dynamically. It integrates with enterprise IAM systems and offers centralized dashboards for monitoring policy compliance and access activities. Informatica Intelligent Data Management Cloud (IDMC) provides automated data discovery, classification, and access controls using AI-driven policy recommendations. Similarly, Talend Data Fabric combines data integration, quality, and governance under a unified platform that enforces access policies based on metadata and business rules. These tools exemplify the practical implementation of Data Fabric principles and highlight how access control can be operationalized at scale through centralized automation.

5. COMPARATIVE ANALYSIS

5.1. Governance Structure and Policy Management

The governance structures of Data Mesh and Data Fabric reflect their fundamental architectural philosophies. In Data Mesh, governance is federated, distributing responsibility to domain teams that implement access control according to shared standards. This structure promotes

ISSN: XXXX-XXXX
PP: 31-40.

autonomy and innovation but can make uniform policy enforcement difficult. In contrast, Data Fabric relies on centralized governance where policies are defined and enforced uniformly through metadata and orchestration platforms. While this centralization ensures consistency and simplifies compliance management, it may limit the agility of individual teams to adapt policies for specific needs. Therefore, organizations must weigh the benefits of flexibility versus uniformity when selecting between the two approaches.

5.2. Scalability and Flexibility

Data Mesh is inherently designed for scalability, particularly in organizations with complex, domain-driven structures. It allows data ownership and access control to scale with organizational growth, enabling teams to manage access independently. However, this flexibility can become a double-edged sword if not managed with strong federated standards. Data Fabric, while scalable in terms of technology, may face limitations when adapting access controls to rapidly changing organizational contexts, especially when policy changes require reconfiguration of centralized systems. Thus, while both are scalable, Data Mesh favors organizational flexibility, and Data Fabric favors technical consistency.

5.3. Security and Compliance Capabilities

Data Fabric typically excels in security and compliance due to its centralized enforcement of policies and integration with metadata-driven security frameworks. Automated policy application, audit trails, and integration with IAM tools make it particularly well-suited for highly regulated environments. Data Mesh, on the other hand, requires significant coordination to ensure that each domain adheres to security standards, and any non-conformity can create compliance risks. While frameworks like Open Policy Agent help mitigate this risk, the decentralized nature of Data Mesh means compliance is more reliant on process discipline than technical enforcement.

5.4. Ease of Implementation and Operational Complexity

Implementing a Data Mesh architecture demands cultural and organizational change, including new team roles, operating models, and governance practices. The complexity of managing domain-level access control mechanisms adds to the operational burden. Conversely, Data Fabric solutions are typically technology-driven and can be adopted with less organizational restructuring, although they may require significant up-front investment in metadata management and system integration. Operational complexity in Data Fabric lies more in technical orchestration, while in Data Mesh, it lies in coordinating decentralized teams and policies.

5.5. Use Case Suitability (Regulated Industries, Real-Time Analytics, etc.)

Data Fabric is generally better suited for industries where compliance, auditability, and centralized control are critical—such as finance, healthcare, and government. Its automation and standardization simplify governance and support regulatory requirements. On the other hand, Data Mesh is more appropriate for large, dynamic organizations such as tech companies or global

ISSN: XXXX-XXXX
PP: 31-40.

enterprises with diverse data needs. It enables faster innovation and localized control, which is ideal for real-time analytics, personalized user experiences, and product-focused data services.

6. CHALLENGES AND CONSIDERATIONS

6.1. Trade-offs in Decentralization vs Centralization

Choosing between decentralization and centralization in access control requires balancing autonomy and control. Data Mesh empowers domains but can lead to inconsistent policy enforcement. Data Fabric provides uniformity but may slow down localized innovation. Organizations must assess their maturity in data governance, risk tolerance, and need for agility to find the right balance between these models.

6.2. Data Sovereignty and Cross-Domain Access

Both architectures face challenges in addressing data sovereignty laws and enabling secure cross-domain access. Data Mesh must ensure that decentralized domains respect jurisdictional boundaries, while Data Fabric must incorporate legal compliance into its automated access controls. Effective identity management, data classification, and policy propagation mechanisms are essential to managing these complexities.

6.3. Evolving Standards and Interoperability

The lack of widely adopted standards in access control for modern architectures presents a significant challenge. Data Mesh often relies on evolving frameworks like data contracts and policy-as-code, which lack standardization. Data Fabric tools also face interoperability issues, especially when integrating with legacy systems. Advancing interoperability will be crucial for long-term success in either model.

6.4. Future Trends (e.g., AI-Driven Access Control, Confidential Computing)

Emerging technologies such as AI-driven access control, confidential computing, and decentralized identity (e.g., verifiable credentials) are likely to reshape access control strategies. AI can enhance dynamic policy enforcement by analyzing user behavior and risk context. Confidential computing can enable secure access to sensitive data without exposing it, benefiting both architectures. These trends point to a convergence of automation, intelligence, and security in next-generation data platforms.

7. CONCLUSION

7.1. Summary of Key Findings

Data Mesh and Data Fabric offer distinct approaches to access control, rooted in differing architectural philosophies. Data Mesh excels in decentralized, domain-oriented environments, fostering agility and scalability. Data Fabric, with its centralized, metadata-driven governance, provides strong compliance and uniform access policy enforcement.

ISSN: XXXX-XXXX
PP: 31-40.

7.2. Recommendations for Architecture Selection Based on Access Control Needs

Organizations seeking rapid innovation and distributed responsibility should consider Data Mesh but must invest in strong federated governance tools. Enterprises operating in heavily regulated environments or requiring consistent, automated control should favor Data Fabric. Hybrid approaches that combine centralized policy design with domain-level enforcement may offer the best of both worlds.

7.3. Directions for Future Research

Future research could explore hybrid governance models, AI-assisted policy management, and deeper integration of access control with data observability. Comparative performance studies of real-world implementations and longitudinal studies on the scalability and security of access control models across both architectures would also add significant value.

REFERENCES

[1] Dehghani, Zhamak. *Data Mesh: Delivering Data-Driven Value at Scale*. O'Reilly Media, 2022.
[2] IBM. "What is a Data Fabric?" IBM Cloud Learn Hub, 2023.
[3] Gartner. "Data Fabric Architecture is Key to Modernizing Data Management." Gartner Research, 2022.
[4] Open Policy Agent. "Policy-Based Control for Cloud-Native Environments." OpenPolicyAgent.org, 2023.
[5] Informatica. "Informatica Intelligent Data Management Cloud (IDMC)." Informatica Whitepaper, 2023.
[6] Talend. "Talend Data Fabric: Unified Platform for Data Integration and Governance." Talend Product Overview, 2023.
[7] Microsoft. "Decentralized Identity and Access Management." Microsoft Azure Docs, 2023.
[8] NIST. "Zero Trust Architecture." NIST Special Publication 800-207, 2020.
[9] AWS. "Best Practices for Implementing RBAC and ABAC in AWS Environments." AWS Whitepaper, 2022.
[10] Data Governance Professionals Organization (DGPO). "Federated Governance in Decentralized Data Architectures." DGPO Research Brief, 2022.
[11] Accenture. "AI in Data Governance: Automating Access Control and Classification." Accenture Insights, 2023.
[12] McKinsey & Company. "Building Trustworthy Data Platforms with AI-Driven Governance." McKinsey Digital, 2023.
[13] Google Cloud. "Implementing Data Mesh on Google Cloud." Google Cloud Architecture Center, 2023.
[14] Deloitte. "Navigating the Future of Data Governance in the Age of Decentralization." Deloitte Insights, 2023.
[15] Forrester. "Data Fabric vs. Data Mesh: Which One Is Right for You?" Forrester Report, 2023.