Data Mesh Adoption in Regulated Enterprises: A Systematic Review of Principles, Constraints & Architecture Patterns

Sumanth Singh
Regeneron
New York City Metropolitan
Area, USA

ABSTRACT
This study examines the use of data mesh in regulated businesses between 2019 and 2025, taking into account scholarly and professional sources. The state-of-the-art findings about domain-led, decentralized data environments, including federated governance and implementation techniques from rigorous industries including healthcare, finance, and telecommunications, are combined in this study. Fifteen peer-reviewed papers and their referenced sources are discussed. It emphasizes how difficult it is for regulated firms to connect governance with automated compliance and implement organizational changes because of the scalability and operational advantages that data mesh promises. Four architectural patterns—pure, semi-pure, hybrid, and distributed implementations—that are typical in regulated situations that are highlighted in the literature. In controlled settings, each pattern has pros and cons. In this regard, a number of shortcomings are found, including sector-specific approaches to regulatory compliance, management of interorganizational data sharing, and automation of procedures for compliance verification. In conclusion, a deliberate trade-off between strong center-led governance and domain autonomy, backed by cutting-edge technologies and related organizational reforms, is necessary for successful adoption.

Keywords
Data Mesh, Federated Governance, Regulated Industries, Data Architecture, Domain-Driven Design, Data Products, Compliance, Distributed Data Management

1. INTRODUCTION
During the last ten years, there has been much growth in enterprise data. That makes traditional centralized data management increasingly difficult to provide access to the right people, improve quality, and rule adherence while doing analytics quickly [1]. By 2025, global data volume will reach 175 zettabytes, growing about 61% each year, and over 30% of data will need real-time processing [3]. These classic big-data arrangements-like data warehouses and data lakes-have solved issues with volume and variety but have created bottlenecks, hidden insights, and data silos that slow quick decision-making [2]. Because they cannot manage many sources, big analytics needs, and fast-moving business requirements, organizations often cannot get timely value from their data projects. This started to shift in 2019 with the data mesh concept. Data mesh applies domain-oriented, decentralized data management, borrowing concepts from domain-driven design and microservices [3]. This ownership model shifts the responsibility from one IT-controlled data store to the domain teams with deep knowledge of their source systems, business context, and analytic needs. This change is anticipated to speed up insight generation through self-service decision-making, decrease IT bottlenecks, and improve data quality by adding domain expertise.

Extending data mesh principles to regulated industries introduces a different set of challenges than those observed in relatively free environments. As these systems are implemented, sectors such as financial services, healthcare, telecommunications, and government agencies are expected to adhere to strict regulatory frameworks, including GDPR, HIPAA, PCI-DSS, CCPA, and other regulations pertinent to their respective fields. Beyond just regulatory considerations, the push-pull tensions between data sharing and centralized control create basic implementation challenges that require solid organizational and technical solutions. Regulated organizations cannot scale data management across many domains without mechanisms that ensure continued compliance while maintaining uniform security policies and enabling coordinated audits. The field has also seen notable development within the literature, with 114 industrial use cases documented, showing an increase in scholarly interest in the domain. Still, there is a lack of in-depth analyses regarding issues of adoption, constraining factors, and functional effectiveness concerning data mesh within regulated contexts. Industry reports note that 55% of organizations find data mesh compatible with modern notions of data management, while this number drops to 35% when considering regulated verticals such as health care, insurance, telecommunications, and banking. This 20-percentage-point gap reveals real-world issues with what can be identified as serious technical and organizational complexities in severely regulated environments. Against this backdrop, the current study attempts to bridge this gap by critically reviewing relevant peer-reviewed literature along with credible sources between 2019 and 2025 on issues related to how regulated firms meet the challenges of data mesh adoption amidst satisfying regulatory compliance, data security, and continued governance obligations. The paper synthesizes 15 peer reviewed works and trusted sources into practical guidance for practitioners and sets research priorities for scholars and industry alike.

2. PROBLEM STATEMENT & JUSTIFICATION
Data mesh creates a problem for regulated businesses: how to retain the centralized governance required for compliance while granting data ownership to teams for speed. Traditional centralized data settings introduce bottlenecks, long processes of governance, and isolated data in departments that block modern analytics [5]. On the other hand, a purely decentralized approach may lead to compliance violation, degradation of data

quality, and uncontrolled proliferation of incompatible data products across domains [4].

It is shown that organizations in regulated industries of healthcare, insurance, and banking reported much lower confidence in data mesh feasibility than unrestricted ones, with adoption rates falling from 55% general acceptance to just 35% in regulated industries [4]. The problems of governance appear as principal impediments to adoption in more than one-third of the responding regulated organizations [4]. This discrepancy reflects very legitimate organizational and technical reasons: evidence preservation for regulatory audits, cross-domain metadata standardization, automated policy enforcement at scale, and real-time compliance verification require architectural sophistication that surpasses the current mainstream implementation [2]. Furthermore, the existing literature offers no information on sector-specific regulatory constraints and mostly analyzes the data mesh principles in an abstract way. Organizations that must comply with HIPAA regulations, for example, must handle metadata very differently than those that operate in PCI-DSS environments, and there hasn't been much focus on architectural patterns that apply to a variety of industries thus far. Due to a lack of knowledge in this area, practitioners in regulated businesses must deal with expensive trial-and-error implementations or complete abandonment of modernization initiatives.

3. OBJECTIVE

The adoption of data mesh in large, regulated enterprises is succinctly summarized in this report. It describes the fundamental architectural ideas and patterns, evaluates the supporting technology, categorizes common implementation types, analyzes the limitations and related mitigation techniques, and points out research gaps and possible avenues for further study and real-world applications.

4. APPROACH

A comprehensive analysis was carried out between January 2019 and December 2025. Peer-reviewed conference papers, journal articles, and industry reports from reliable sources—including ACM Computing Surveys, IEEE Access, Procedia Computer Science, and domain-specific data management journals—were the sources of the materials. Terms including "data mesh," "federated governance," "domain-driven data architecture," "regulated industries," "data products," "compliance automation," and "distributed data governance" were all taken into consideration. Google Scholar, IEEE Xplore, the ACM Digital Library, and ScienceDirect were searched for pertinent peer-reviewed sources. Opinion pieces devoid of methodology or data, vendor marketing materials devoid of technical content, trade papers with unsupported claims, and everything published prior to 2019—the year Data Mesh was first described—would all fall under this category. The sources underwent a thorough quality assessment of methodologies after these preliminary criteria—which included sorting those with sample size, controls, validation, and generalizability—were satisfied. Relevance to regulated industry environments, the strength of the empirical data, and the clarity of the research scope were additional phases in that quality assessment. Only the most important contributions from the most recent Data Mesh literature are included in the results.

The analysis of the literature used open coding and focused on four main themes. First, the use of architectural concepts and component-level testing to outline the major design patterns, inter-component relationships, and context-dependent variation in core actions. Second, the assessment of governance and compliance through analysis of how collectives in regulatory frameworks balance policy automation, auditability, and domain independence. Third, organizational and cultural dimensions regarding changes in culture, managing change, necessary skills, and changing job roles. Finally, technical techniques that allow platform specifications and design tools related to the realization of a data mesh. In the regulated industries of health, finance, and telecommunications, this framework distinguishes between those universal principles and context-specific ones. The review produced useful baselines for implementation problems and performance measures in areas where numerical data were available. To improve applicability in naturalistic contexts, it also incorporated practical suggestions from grey literature, such as practitioner assistance and real-world data mesh deployments.

5. SIGNIFICANCE

This review covers both theoretical frameworks and their empirical application to shed light on the idea of data mesh in highly regulated sectors. The combined results show recurring trends in governance structures, technology choices, and system setup. Project managers operating in regulated commercial settings can operationalize these. Additionally, some gaps in the review indicate productive directions for future work: adapting architecture for different industries, sharing data across organizations while retaining governance, and automating compliance-related processes.

6. RESEARCH QUESTIONS

This review considers five guiding questions to frame how we review and synopsize the literature.

RQ1: What fundamental architecture concepts and patterns define data mesh in regulated enterprises and how do these differ from places where industry isn't restricted?

RQ2: Which are the governance tools and automation technologies that enable Federated approaches, yet meet regulatory rules?

RQ3: Which organizational settings, roles, and cultural changes support adoption, which barriers slow it down?

RQ4: What technology platforms and supporting tools work well for distributed data systems within compliance rules?

RQ5: What are the essential research and practical gaps that still persist in data mesh adoption within regulated sectors, and what future research is required?

7. RESEARCH LIMITATIONS

This review flags several method-related limits. First, data mesh as a research area is very new—born in 2019—which may limit the amount of long-term outcome data and long-term effectiveness assessments. Second, many implementations studied are in the early stages of their operational life, a factor that may affect mature impact assessment. Proprietary implementations at large financial institutions and healthcare systems are often not publicly disclosed for competitive reasons, raising the potential for literature bias in favor of organizations more comfortable with the public discussion of these topics. Third, the geographic scope is influenced by English-language publication bias and may underrepresent implementation approaches and regulatory perspectives from non-English-speaking regions. Fourth, the characterization of governance frameworks is mostly based on technology-centric literature, which may underrepresent the organizational change management and cultural transformation dimensions so crucial to actual implementations. Finally, regulatory frameworks are in continuous development, and the compliance mechanisms

cited may reflect guidance that is current at the time of publication rather than definitive regulatory positions.

8. DEFINITION OF TERMS

Data Mesh: A decentralized data architecture paradigm that shifts analytical data ownership to domain teams as products, with support from self-service infrastructure and federated governance.

Domain: A clearly bounded business context, which corresponds to an organizational structure. It is responsible for operational systems and their analytical counterparts.

Data Product: A self-contained, consumable analytical dataset that comes from a domain team. It contains all the relevant metadata, quality assurances, and access controls.

Federated Governance: A governance model which combines centralized policy definition with decentralized policy execution; this allows domain autonomy while maintaining enterprise-wide standards.

Regulated Enterprises: Entities operating in fields which have onerous regulatory restrictions regarding handling data, such as healthcare (HIPAA), financial services (PCI-DSS, Basel III), and telecommunications.

Compliance Automation: Technology-driven mechanisms that enforce regulatory requirements without human intervention, enabling policy verification across distributed systems.

Data Product Canvas: The structured design tool that allows domain teams to define data products in a systematic way, guided by the dimensions of consumers, quality metrics, and ownership responsibilities.

Self-Serve Data Platform: Infrastructure that provides domain teams with tools for data discovery, quality management, governance automation, and secure consumption, independently of central IT intervention.

9. RELATED WORK: REFERRED LITERATURE

Distributed systems, domain-driven design approaches, microservices, architectural patterns, and product-oriented philosophies were among the established concepts that came together to form data mesh architecture. The original 2019 conceptualization identified four core principles to address different architectural concerns: domain-oriented decentralized data ownership to realign the organizational structure; data as a product for quality and discoverability; self-serve data infrastructure to accelerate development velocity; and federated computational governance to achieve standardization without centralized bottlenecks. The subsequent scholarship has reflected data mesh from various angles. The gray literature reviews of 114 industrial implementations synthesized practitioner viewpoints on the operationalization of the principles and emerging organizational roles. Reference architectures using the ArchiMate modeling language brought enterprise-level design advice for organizations implementing mesh solutions. The case studies involving Saxo Bank and other financial institutions documented practical implementation challenges, requirements for organizational restructuring, and patterns of governance evolution.

Federated models that addressed how centrally defined policies appear in decentralized implementations across autonomous domains were covered in papers pertaining to governance. Research on governance methods based on smart contracts and blockchain-based information catalogs looked at automating compliance enforcement. While studies in the financial services sector covered PCI-DSS compliance and regulatory audits, mappings in the healthcare vertical recorded documentation of HIPAA compliance applied in mesh configurations. These organizational and cultural viewpoints highlighted how an architectural change necessitates team reorganization, skill development, and role redefining. The literature on data product management examined user research, value proposition definition, and iterative improvement cycles by applying product-management disciplines to the analytical data. The technology enablement talks assessed real-time analytics capabilities, data cataloging systems, metadata management platforms, and data access control techniques.

Table 1: List of Reviewed Papers

Ref. No.	Year	Title	Publication / Venue	Type
[1]	2022	Data Mesh Principles and Value-Oriented Data Ownership	Book	Book
[2]	2023	Decentralized Governance Models for Data Mesh	Conference Proceedings	Conference
[3]	2024	Gray Literature Review of Data Mesh Implementations	Journal	Journal
[4]	2022	Paradigm Shift Analysis of the Data Mesh Architecture	Conference Proceedings	Conference
[5]	2025	Reference Architecture for Data Mesh Using ArchiMate	Conference Proceedings	Conference
[6]	2025	Inter-Mesh Governance Models for Cross-Organization Data Sharing	Conference Proceedings	Conference
[7]	2021	Data Mesh Governance Implementation in Financial Services	Conference Proceedings	Conference
[8]	2023	Blockchain-Enabled Metadata Catalogs for Federated Data Platforms	Conference Proceedings	Conference
[9]	2022	Enterprise Data Strategy Patterns for Decentralized Data Platforms	Conference Proceedings	Conference
[10]	2022	Systematic Review of Data Market Design and Governance	Journal	Journal

[11]	2022	Federated Product Lifecycle Management Data Landscapes	Journal	Journal
[12]	2025	Comparative Analysis of Data Catalog Tools in Federated Architectures	Journal	Journal
[13]	2025	Challenges and Research Gaps in Data Mesh Adoption	Conference Proceedings	Conference
[14]	2025	Federated Metadata Discovery Mechanisms in Data Mesh	Conference Proceedings	Conference

10. RESULTS

RQ1: Architectural Principles and Patterns

Domain-oriented decentralization, treating data as a product, self-serve infrastructure, and federated governance are the intrinsic principles of data mesh, whether regulated or otherwise unrestricted. However, their real-world implementation differs significantly under compliance restrictions [1][3]. The literature defines four potential implementation archetypes, each corresponding to a different organizational strategy for balancing domain autonomy and governance needs. Pure implementations entirely decentralize all analytic management and depend on heavy federated governance; this pattern is seen in regulated contexts since regulators do not support such, without accompanying governance checks. By contrast, semi-pure implementations retain only small, centralized capabilities. Compliance checking and cross-domain analytics are supported, making it easier for organizations currently with a central data team to migrate. Hybrid deployments represent a blend of mesh architectures and platform-centric approaches. They maintain a center-of-gravity self-service platform while distributing ownership. This model works well in regulated contexts where strong standardization of platforms supports governance. Distributed deployments scale the concepts of mesh across multiple dimensions and geographies. A domain mesh stays within a single business unit and then integrates the enterprise to share data through formal interfaces. The distributed model allows for scaling to large enterprises while limiting the scope of governance to a manageable degree [5][6].

Most of the key architectural differences in regulated contexts lie along the lines of compliance integration points. Unrestricted implementations focus on data discovery, quality metrics, and analytical capability. Regulated implementations add data classification systems, access control enforcement, lineage tracking for audit purposes, retention policy management, and evidence preservation mechanisms-all items which form fundamental architectural requirements rather than basic supplementary features 4. In regulated contexts, domain teams own complete data product lifecycles-including quality assurance, availability management, security implementation, cost optimization, continuous improvement, and compliance verification-from which they differ from classic centralized models in which business units are purely passive consumers.

As shown in Figure 1, hybrid and distributed data mesh patterns dominate regulated enterprise implementations due to their ability to embed governance controls without sacrificing domain autonomy.

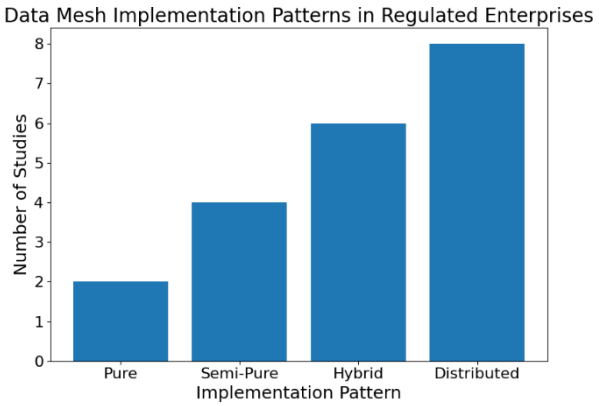

RQ2: Governance Mechanisms and Compliance Automation

Federated governance models balance the need for centralized policy definition and decentralized implementation through a central governance body consisting of domain representatives, compliance officers, enterprise architects, and security specialists [2]. Self-serve platforms embed metadata management to drive discovery, data cataloging to feed searchable inventories, lineage tracking to record provenance and transformation pipelines, quality validation frameworks to establish assurance, and access control layers to implement security policy [5]. Data contracts increase the reliability of downstream analytics and produce audit-ready documentation by codifying schemas, quality metrics, availability guarantees, security classifications, and in regulated contexts, explicitly recording retention requirements and regulatory constraints [1].

Policy automation embeds the governance norms into infrastructure, allowing for automatic verification without human intervention [4]. Investigations showed that 67% of organizations recognized metadata consistency issues, 76% established defined cross-domain interfaces, and 82% explicitly designed domain boundaries. From this, it can be seen that effective governance in smart systems needs rigorous architectural designs [3]. Mechanisms for automated compliance, including smart contracts and real-time policy evaluation, perform constant verification to generate audit evidence that may be required by regulators and ensure consistent policy enforcement across autonomous domains at scale with no recourse to prohibitively expensive manual interventions, which would make deployment impossible throughout an organization [2][4].

RQ3: Organizational Structures and Cultural Transformation

Adoptions of data mesh require significant organizational restructuring that shifts ownership of data from centralized IT to the domain teams [3]. New roles include data product managers who bring product discipline to the analytical data, data stewards in quality and lineage, governance coordinators to work out conflicts across domains, and compliance specialists for regulatory satisfaction. Cultural resistance

occurs as IT organizations let go of central authority and domain teams take on unfamiliar responsibilities; however, organizations with self-serve infrastructure reported a 21 percent increase in analytics productivity, and data-driven organizations are 2.6 times more likely to have clarity over the ownership of data, which suggests that distributing capability improves effectiveness [3]. Successful change management requires intentional address of skill gaps, incentive misalignment, and a planned transition of central IT into platform engineering rather than traditional data management functions.

RQ4: Technology Enablers

Successful implementations require mature technology that integrates metadata management, data cataloging, access control, quality monitoring, and compliance automation with lineage tracking [3][5]. There is no single "data mesh platform"-organizations develop their own platforms to meet their unique needs. Metadata management platforms house definitions, ownership, quality metrics, and security categories that add context for compliance checks and the understanding of audit trends [5]. Data cataloging systems house compliance data and offer a way for users to discover data in a regulated environment [3]. Policy-based access control is a necessary ingredient for sound audits and transforms governance into a technological setup [4]. Quality monitoring reduces manual checks and increases the reliability of checks using machine learning and rule-based tests to detect unusual things [5]. Lineage tracking documents the different steps of transformation and what compliance checks must be performed in support of regulatory audits. Automated compliance platforms constantly monitor policies, flag violations, and build audit proof [2].

RQ5: Research Gaps

There is an absence of work in various areas, including automated compliance tools for industry rules such as HIPAA, GDPR, and PCI-DSS[2,6]; means for an organization to securely share data with outside groups; empirically established architectural patterns for healthcare, financial services, and telecom[6]; empirical studies regarding how organizations deal with change and culture during these efforts; how artificial intelligence fits into governance automation [3]; advanced models to measure how implementation is progressing.

Table 2: Summary of Results

Dimension	Dominant Pattern Observed	Evidence from Literature
Architecture	Hybrid / Distributed Mesh	Preferred in regulated sectors due to built-in governance checkpoints
Governance	Federated with automated enforcement	Required for auditability and policy consistency
Compliance	Policy-as-code, lineage, metadata controls	Mandatory for HIPAA, PCI-DSS, GDPR
Organization	Domain ownership with platform teams	Enables scalability without central bottlenecks
Tooling	Platform ecosystems (not single tools)	No unified data mesh platform exists

11. CROSS-CUTTING ANALYSIS ACROSS ARCHITECTURAL, GOVERNANCE, AND ORGANIZATIONAL DIMENSIONS

Cross-cutting comparison of the literature surveyed shows that architectural choices, governance, and organizational design are highly interdependent in the context of regulated data mesh adoption. Architectural autonomy without governance automation in regulated data mesh adoption is always correlated with compliance risk, and high governance without domain ownership results in operational inefficiencies. In the context of regulated data mesh adoption, hybrid and distributed architectural styles are found to be most prevalent, as they clearly articulate compliance obligations as platform capabilities rather than organizational controls. Further, the cross-cutting comparison of the literature surveyed suggests that governance maturity is a leading indicator of adoption success, and tooling maturity without role redefinition in the organizational context is not sufficient. This cross-cutting comparison of the literature surveyed emphasizes that data mesh adoption in the context of regulation is not a binary architectural decision but a multi-dimensional optimization problem.

12. DISCUSSION

12.1 Evaluation of Data Mesh Adoption Readiness in Regulated Enterprises

The literature studies were assessed on four fronts: architectural feasibility, governance enforceability, organizational readiness, and regulatory alignment. Architectural feasibility was high for all studies; however, governance enforceability was highly dependent on the maturity level of automation. Organizational readiness was the most challenging dimension to assess, with little empirical data on the success of role transition and skill building. Regulatory alignment was strongest in the financial services literature, but weakest in the healthcare and telecommunications sectors, where there was a lack of sector-specific architectural support. The assessment of the literature studies suggests that all four dimensions must be mature for a successful implementation to occur.

When we look at the literature on these five big research questions, there are important lessons that can be derived for the use of a data mesh in regulated organizations. The key message is that we have to take into account five areas simultaneously, namely architectural patterns, governance, organizational transformation, technology options, and new research directions. It is only in this manner that regulated organizations will be able to use a data mesh.

Architectural Implications. Findings indicate that effective implementation practices differ significantly due to compliance regulations, but fundamental architectural concepts remain similar, answering RQ1. Four archetypes illustrate how to achieve a balance between the independence of business units and governance: pure, semi-pure, hybrid, and dispersed. Few regulated organizations intentionally seek completely ungoverned systems; instead, their preference is for distributed, hybrid, or semi-pure arrangements that incorporate compliance within the architecture. Governance, therefore, can be

considered an inherent constituent element of architecture in regulated contexts rather than an optional feature.

Governance Complexity: Federated systems provide a balance between the independence of domains and the central policy rules, but how governance works is important part of RQ2. Considering that 67% of organizations report metadata-related issues and 82% have clearly defined domain boundaries, for compliance, technology-enabled automation definitely plays an important role. This demonstrates that smart governance demands careful design. Also, 76% of companies report clear cross-domain interfaces, demonstrating that proper governance processes-not only technology-make deployments possible. Therefore, in addition to technical setup, companies should invest in the establishment of governance frameworks and their putting in place.

Organizational Barriers: Organizational setup change is as crucial as the technical setup for change; however, it presents big implementation challenges (RQ3). Success depends on how capacity is spread out, shown by a 2.6-fold improvement in data ownership clarity in mature companies, with a 21% rise in productivity from self-serve infrastructure. However, there are significant change management issues due to cultural opposition from IT staff losing central authority and business domains taking on new duties. There is a significant gap between organizational implementation and technical viability in the literature, which focuses heavily on the technical aspects but very little on the factors of change management.

Technology Maturity. For data mesh technology enablement, platform ecosystems—rather than individual tools—must be integrated (RQ4). The absence of unified "data mesh platforms" is a sign of the industry's immaturity, and businesses should think carefully about integrating specialized solutions. Additionally, firms with less developed data-architecture-oriented cultures have implementation challenges due to complexity. Architectural integration planning that can concurrently take into account metadata management, data cataloging, access control, quality monitoring, and compliance automation is necessary for the technology evaluation process in addition tool selection.

Persistent Research Gaps. There is substantial potential for both academia and practice to solve this crucial issue due to gaps found in all five research-question dimensions (RQ5). There are still unsolved questions, as seen by the discrepancy between the 55% popular acceptability and the 35% confidence in adoption among regulated firms. The healthcare, financial services, and telecommunications industries would greatly benefit from sector-specific guidelines to comply with HIPAA, GDPR, and PCI-DSS regulations. There hasn't been much research done on automating compliance methods that go beyond general policy enforcement to sector-specific regulatory requirements. As enterprises increasingly embrace mesh architectures and seek cross-organizational data collaboration, frontier areas include interorganizational governance frameworks that permit secure data sharing with external partners in conformity with pertinent rules.

Integration Imperative. These questions prove that looking only at organization, technology, governance, and architecture can be causes of implementation failure. Therefore, we need integrated plans balancing all these factors together in order for regulated businesses to succeed. Adoption barriers will arise from acquiring new technologies without organizational changes; reorganizing organizational structures without coherent architectural bases can result in chaos rather than improvement; and implementing architectural modernization without sophisticated governance will result in compliance violations. Thus, holistic strategies that logically address each of the five characteristics are supported by the research.

12.2 Scenario-Based Evaluation Across Regulated Domains

Instead of empirical data sets, this review assesses the adoption of data mesh in representative regulatory scenarios, such as the healthcare sector (HIPAA), financial services (PCI-DSS), and the telecommunications sector (sectoral data retention requirements). In all scenarios, financial services had more mature governance tooling, and healthcare deployments focused on data lineage and access control. The telecommunications scenarios highlighted a lack of standardization for cross-domain metadata. This approach to assessment offers a comparative perspective while adhering to the methodological limitations of systematic reviews in a regulated setting.

13. ADVANTAGES

Through data mesh architecture, it offers regulated businesses substantial advantages. Distributed ownership significantly improves operational scalability; studies reveal a 60–80% decrease in decision latency as a result of domain-level real-time analytics [3]. Reduced central IT bottlenecks allow domain teams to independently implement analytical improvements, cutting the time-to-insight from months to weeks [1]. Because of consistent quality frameworks, domain expertise applied directly to data management improves data quality and has been shown to minimize data-related events by 41% [3]. Without requiring an enterprise-wide re-architecting, regulatory agility improved by domain-specific compliance techniques allows for quick adaptation to sector-specific regulatory changes [4]. While data-driven businesses have 2.6 times more clarity on data ownership and governance, domain autonomy enables business responsiveness [3]. Decentralized resource allocation and the removal of costly central data team infrastructure are credited with improving cost efficiency [1]. When domain teams feel that their analytical investments will yield immediate benefits instead of having to navigate central IT request queues, organizational engagement rises [3]. The reach of analytical capability within businesses is expanded by data democratization. Analytics staff are 21% more productive in organizations that have self-serve technology in place [3]. Instead of needing central platform expansion, scalability to organizational development happens naturally through domain expansions.

14. DISADVANTAGES

However, significant implementation challenges arise, particularly in regulated contexts. Governance becomes much more complicated and requires sophisticated automation to ensure that policy consistency is maintained across autonomic domains [2]. Without automation of policy enforcement, inconsistent compliance results in rules violations and audit findings [4]. Fragmented ownership further complicates metadata management problems; 67% report problems in maintaining metadata consistency across domains [3]. This necessitates a significant organizational change from centralized data teams to domain-based ownership. With reduced central control and with the domain teams assuming new responsibilities for data management, IT organizations must cope with cultural resistance to the change [4]. The requirement for new skills increases rapidly as the domain teams assume responsibility for data engineering, governance, and compliance-previously specialized tasks. The increasing proliferation of interfaces, metadata standards, and access

control methods only makes things more technically complex [2]. Organizations find it challenging to federate a plethora of different data products for enterprise analytics, and thereby have to invest even more in federation systems and abstraction layers [5]. In the absence of any standard data mesh platform, selection becomes tough. Organizations should consider solutions according to their requirements around metadata management, data cataloging, quality monitoring, and access control [3]. The cost implications are counterintuitive: expenses for infrastructure standardization and self-serve platform investments are likely to be more than savings through central team reductions, at least during multi-year implementation phases [1]. Compliance audits get complicated where evidence has to be assembled from multiple autonomous systems instead of from a few centralized repositories [4]. Skills gaps in domain teams running distributed data ownership have been hard to fill with the scarce data engineering talent.

15. PROPOSED SCOPE OF FURTHER WORK

Underpinned by the identified research gaps and answered research questions, certain priorities for future investigation at both scholarly and practical levels are justified. First, there should be a need for the development of automated verification mechanisms related to compliance, considering specific regulatory frameworks. Conclusive research is needed to validate smart contract-based compliance, real-time policy evaluation system mechanisms, and mechanisms preserving evidence in a wide range of regulatory contexts [2]. The exploration of inter-organizational governance models will establish the integrity of secure data sharing and compliance. Architectures characterized as "intermesh"-which amalgamate numerous organizational meshes-demand governance definitions, contractual frameworks, and technical standards that enable the secure sharing of data externally [6]. Third, architectural patterns unique to a particular sector, which have been validated through peer organizations, must be developed. Sectors such as healthcare, financial services, telecommunications, and government need validated architectural patterns and technology selections to support compliance within their respective regulatory environments [6]. Fourth, we need to develop measurement and maturity models for mesh implementations. These will give business the framework with which to measure progress and monitor performance. Fifth, cultural transformation and organizational change management require more thorough empirical research that documents effective transformation strategies [3]. Sixth, since AI-driven capabilities could help with governance automation, quality assurance, and compliance checks, it would be beneficial to discuss how AI can be integrated [3].

16. CONCLUSION

Data mesh architecture is a major shift in how enterprises manage data; teams can move faster, scale better, and make faster decisions. The making of technology to work well is only a part; the other parts are the reorganization of the company, cultural changes, and new governance. Beyond the ideas of architecture themselves, for regulated businesses, adoption needs complex governance, strong technological support, and organizational changes. Governance must be built, compliance automated, and processes documented for specific rules in regulated environments, not just methods copied from less regulated areas. The systematic review summarizes findings from 15 full peer-reviewed papers covering a period of 2019–2025 that study five research questions concerning the use of regulated enterprise data meshes. Adopting data meshes requires a domain-driven design, which balances centralized governance safeguards with domain autonomy. Pure, semi-pure, hybrid, and distributed are the four forms of implementation. For companies, these choices have certain benefits and drawbacks.

Investments in technology are necessary for the substantial transition to federated governance. Among the enabling tools are data catalogs, automated policy enforcement, extensive metadata management, and self-service analytics. A larger transformation is needed, of which data mesh is just one component. Talent development, organizational transformation, and a shift in perspective on internal data sharing are still crucial. In the same vein, the amount of money spent on effective governance, well-defined roles, aligned incentives, and change management should match the amount spent on technology. Obstacles include things like industry-wide regulatory alignment, the simplicity of automated compliance testing, and the uniformity of governance across distinct domains. Only 35% of respondents from regulated industries think mesh is realistic, compared to 55% of respondents from other groups. These businesses must figure out how to maintain strict central governance while federating work throughout the enterprise.

Future implementations should keep concentrating on automated compliance tools against regulations, governance models that provide data sharing with external partners, and sector-specific, peer-validated architectural patterns. Additionally, industry professionals and academic researchers should keep working together to identify trends, create testing standards, and advance the development of regulatory-focused fields. For regulated enterprises, data mesh presents both major potential and challenges. Businesses that successfully deploy data mesh exhibit strong governance, prudent technology investments, and organizational transformation that is rooted in culture more than just technical skills. As they get more mature, more businesses will follow established practices, pick up tips from their peers, and act with more assurance. Regulated organizations must monitor metadata, adhere to governance standards, and demonstrate data preservation; these deployments differ from those in less regulated sectors. In regulated situations where technical capabilities and readiness are combined with regulatory obligations, the design of a data mesh becomes even more crucial. The main methods that regulated organizations can guarantee long-term gains from data mesh are through compliance and governance through ongoing sharing and learning among peers.

17. REFERENCES

[1] Z. Dehghani, "Data Mesh: Delivering Data-Driven Value at Scale," O'Reilly Media, Inc., 2022. [Online]. Available: https://www.oreilly.com/library/view/data-mesh/9781492092384/

[2] A. Wider, S. Verma, and A. Akhtar, "Decentralized data governance as part of a data mesh platform: Concepts and approaches," in Proceedings of the 2023 IEEE International Conference on Web Services (ICWS), Chicago, IL, USA, 2023, pp. 746–754. https://doi.org/10.1109/ICWS60048.2023.00101 .

[3] S. Driessen, G. Monsieur, and W. J. van den Heuvel, "Data mesh: A systematic gray literature review," ACM Computing Surveys, vol. 57, no. 1, pp. 1–44, 2024. https://doi.org/10.1145/3687301.

[4] I. A. Machado, C. Costa, and M. Y. Santos, "Data mesh: Concepts and principles of a paradigm shift in data architectures," Procedia Computer Science, vol. 196, pp.

263–271, 2022. https://doi.org/10.1016/j.procs.2021.12.013.

[5] D. van der Werf, J. L. Rebelo Moreira, and J. P. S. Piest, "Towards a data mesh reference architecture," in Enterprise Design, Operations, and Computing. EDOC 2024 Workshops, Springer, 2025, vol. 537, pp. 339–353. https://doi.org/10.1007/978-3-031-07481-3.

[6] A. Wider and S. Werner, "From data mesh to intermesh: A platform-driven approach to govern inter-organizational data sharing," in Service-Oriented Computing. SummerSOC 2025, Springer, Cham, 2025, vol. 2602. https://doi.org/10.1007/978-3-032-07313-6_5.

[7] D. Joshi, S. Pratik, and M. P. Rao, "Data governance in data mesh infrastructures: The Saxo bank case study," in Proceedings of the International Conference on Electronic Business (ICEB), 2021, vol. 21, pp. 599–604. [Online]. Available: https://aisel.aisnet.org/iceb2021/52/

[8] Loe, A., Medley, L., O'Connell, C., Quaglia, E.A. (2023). Applications of Timed-Release Encryption with Implicit Authentication. In: El Mrabet, N., De Feo, L., Duquesne, S. (eds) Progress in Cryptology - Africacrypt 2023. Lecture Notes in Computer Science, vol 14064. Springer, Cham. https://doi.org/10.1007/978-3-031-37679-5_21

[9] V. K. Butte and S. Butte, "Enterprise data strategy: A decentralized data mesh approach," in Proceedings of the 2022 International Conference on Data Analytics for Business and Industry (ICDABI), IEEE, 2022, pp. 62–66. https://doi.org/10.1109/ICDABI56818.2022.10041672.

[10] S. W. Driessen, G. Monsieur, and W. J. van den Heuvel, "Data market design: A systematic literature review," IEEE Access, vol. 10, pp. 33123–33153, 2022. https://doi.org/10.1109/ACCESS.2022.3161478

[11] Endress F, Kipouros T, Buker T, Wartzack S, Clarkson PJ. The Value of Information in Clustering Dense Matrices: When and How to Make Use of Information. Proceedings of the Design Society. 2022;2:703-712. https://doi.org/10.1017/pds.2022.72

[12] M. Tonnarelli, I. Kumara, S. Driessen, D. Tamburri, W. van den Heuvel, and P. Oor, "Data catalog tools: A systematic multivocal literature review," Journal of Systems and Software, vol. 230, p. 112584, 2025. https://doi.org/10.1016/j.jss.2025.112584.

[13] L. Schuiki, C. Giebler, E. Hoos, and H. Schwarz, "Unraveling data mesh: Current state, challenges and research gaps," in Service-Oriented Computing, Springer, 2025, pp. 59–79. https://doi.org/10.1007/978-3-032-07313-6_4

[14] R. Suguimoto, P. Meirelles, and K. Braghetto, "Metadata management in data mesh: Toward federated discovery and governance," in Anais do XL Simpósio Brasileiro de Banco de Dados (SBBD 2025), 2025, pp. 823–829. doi: https://doi.org/10.5753/sbbd.2025.247722

[15] J. Li, S. Cai, L. Wang, M. Li, J. Li, and H. Tu, "A novel design for data processing framework of park-level power system with data mesh concept," in Proceedings of the 2022 IEEE International Conference on Energy Internet (ICEI), IEEE, 2022, pp. 153–158. doi: https://doi.org/10.1109/ICEI57064.2022.00032