International Journal of AI, Big Data, Computational and Management Studies
Noble Scholar Research Group | Volume 6, Issue 2, 148-157, 2025
ISSN: 3050-9416 | https://doi.org/10.63282/3050-9416.IJAIBDCMS-V6I2P117

Original Article

Data Mesh Meets AI: Federated Ownership and ML-Enabled Contracts for Cross-Domain Data Collaboration

Sivadeep Katangoori[1], Diganto Ghosh[2]
[1]IT Specialist at Bank of America, USA.
[2]Senior Vice President at Bank of America, USA.

Received On: 04/04/2025 Revised On: 05/05/2025 Accepted On: 18/05/2025 Published On: 02/06/2025

Abstract: *In modern data-driven companies, the problem of fragmented data continues to stifle innovation, make it harder to make these decisions, and make it harder for people from different departments to work together. As companies grow, traditional centralized data architectures have trouble growing quickly & meeting the needs of various fields. This has led to the creation of the Data Mesh paradigm, which supports a decentralized approach to data management in which domain teams own, share, and utilize data as a product. However, enabling substantial cross-domain data collaboration within a distributed framework remains a considerable challenge, particularly regarding the establishment of trust, governance & interoperability across many other domains. This essay looks at how federated ownership and AI-driven smart contracts might change the way information is shared and managed in many other different fields. By adding ML to contract logic, businesses can automate access control, usage limitations, compliance checks & predictive alerts. This reduces friction and builds trust between teams. Through case studies and architectural analyses, we show how machine learning-enabled contracts allow for secure, scalable, and auditable data exchange while keeping autonomy and avoiding bottlenecks. Our study shows that adding AI to federated data governance makes data easier to get to and better quality, and it also gives domain teams more flexibility and responsibility to come up with new ideas. The paper offers a paradigm for businesses seeking to implement Data Mesh via intelligent automation, including practical insights into the design, implementation & measurable impacts of this technique in real-world contexts.*

Keywords: *Data Mesh, Federated Ownership, Smart Contracts, Machine Learning, Data Collaboration, AI Governance, Cross-Domain Data Sharing, Data as a Product, Federated Learning, Trust Automation, Data Contract Enforcement, Metadata Intelligence, Decentralized Data Infrastructure, Real-Time Compliance, AI-Driven Policy Enforcement.*

1. Introduction

1.1. Background

Organizations have had to deal with a lot of changes in how they handle information in the last several years. In the past, data was sent to centralized storage areas, which are often called data lakes or corporate data warehouses. The goal was to bring all the information together in one place so that it could be analyzed more easily & work more quickly. These centralized systems did offer better access & the visibility, but they typically failed to provide agility, scalability, & actual time collaboration, especially as businesses grew and data volumes rose.

The data mesh is a decentralized, domain-oriented way of working that sees data as a product & gives responsibility for it to the people who are closest to where it came from. Data mesh spreads ownership of all organizational data among teams that are unique to each domain, rather than relying on a single data team to handle it all. Each team is responsible for its own data, making sure that information can be found, is accurate & is available to others. This change is especially helpful for big companies with several other business divisions, where the needs & circumstances of each area are quite different.

The data mesh idea is more attractive, but it will only work if people work well together across domains. For data mesh to work well, teams need to be able to share and utilize data across these organizational boundaries. This includes not just technical factors but also cultural and operational ones. The main problem here is trust. Different teams typically have different goals, rules they have to follow & levels of risk they are willing to take. Sharing information might make people worry about misuse, quality, breaking the law, or losing control. This is where artificial intelligence (AI) really shines. It can not only analyze data, but it can also automate governance, enforce rules & predict problems before they happen.

Also, adding machine learning (ML)-enabled smart contracts gives data suppliers & customers an intelligent, rule-based way to automate their agreements. These contracts may enforce access rights, keep an eye on SLA violations & handle compensation procedures. This makes it possible for big groups of people to work together on their information with more freedom and accuracy than ever before. As businesses gather more and more data from more various places, such as IoT devices, CRM systems, social media & cloud apps, it has become very important to have seamless, reliable & automated data exchange across domains. The future of data infrastructure goes beyond just storage & pipelines. It will also involve governance, accountability & more complex interconnections.

Fig 1: Data Mesh Meets AI

1.2. Problem Statement

Despite progress in decentralizing data infrastructure, cross-domain data interchange still faces major challenges. Trust is very important. How can one team be sure that another team's data is accurate, up-to-date, and follows all rules, both within and outside the company? How can teams work out access in a way that is both safe and flexible?

Governance is another problem. Centralised governance typically makes it hard to adjust, whereas decentralised governance may lead to inconsistency, misalignment, or not following the rules. Without a standardised process for defining policies & keeping an eye on how data is used, working together may be quite too risky, especially in these regulated fields like healthcare, banking, or government. There is also the issue of who owns something. In a decentralised system, who gets access to the data when it is utilised in many other areas? Who is in charge if anything goes wrong? How can people who generate data be sure that their work is valued & not used for bad purposes?

In the end, compliance makes things much more complicated. GDPR, HIPAA, and CCPA are examples of laws that make it very hard to share, store & handle their information. Manual compliance checks are not only time-consuming, but they are also prone to their mistakes that might have many serious consequences. These problems cause friction, make it harder to get new ideas, and in the end, restrict the potential for innovation based on data.

1.3. Purpose and Scope

This research investigates the potential of integrating AI technology with the principles of data mesh and federated ownership as a feasible solution. Our goal is to show that companies can enable strong data collaboration across many other domains while keeping trust, compliance, and agility, as long as they have the right tools & governance structures.

This investigation will examine how federated ownership allows certain domains to retain control over their data assets while participating in a broader organizational data ecosystem. Adding AI-driven governance tools, such as smart contracts that use ML, anomaly detection, and automated data quality evaluations, might make data sharing more reliable & efficient.

We aim to ground the conversation in actual world examples while also giving a forward-looking view of what decentralized architectures & smart automation may do together. The creation of a data mesh that uses AI:

- The use of machine learning in establishing trust and binding contracts
- Case studies illustrating effective cross-domain uses
- Possible future paths for federated learning, AI that protects privacy, and automated policy enforcement

By the end of this essay, readers should have a better understanding of how new technologies might change how people think about data ownership, governance, and collaboration. This would make company data more scalable, intelligent, and focused on people.

2. Foundations of Data Mesh and AI Integration

2.1. What is Data Mesh?

Data Mesh is the latest way of thinking about data architecture that moves away from big data lakes and centralized governance. Data Mesh wants people to see data as a product instead of just a consequence of apps or limited access via centralized teams. This encourages teams throughout the business to actively manage, maintain & share their information.

A Data Mesh is based on four main ideas:

- Data Ownership Based on Domain: Data Mesh doesn't have a central data team in charge of all the information. Instead, it gives ownership to the teams who know the data best, which are generally the domain teams. For example, the sales department has sales data, the marketing department has marketing data, and so on. These teams are in charge of making sure that their datasets are of good quality, easy to get to, and well-documented.
- Data as a Product: Like any good product, data has to be easy to find, trust & utilize. This means putting money on SLAs, documentation, version control & interfaces. The goal is to make downstream clients feel confident that the data is accurate and useful.
- Self-Service Data Infrastructure: A self-service platform is important since it allows for decentralized ownership and makes things easier for each team by reducing the amount of technology

they have to deal with. This platform has capabilities for publishing their information, finding it, controlling access, tracking provenance, and many other things, all while hiding the infrastructure that makes it work.

- Federated Computational Governance: A Data Mesh doesn't need a central authority to keep a close eye on things. Instead, it is run by policy & is not centralized. Teams follow the same rules on privacy, security, naming standards, and quality, but they do it in different ways employing automation and standard practices.

At its core, Data Mesh is all about spreading data practices across a company without slowing things down. If done well, it makes data collaborative, responsible & useful.

2.2. Problems with Current Data Mesh Implementations

The idea of Data Mesh is interesting, however putting it into practice is a little hard. Companies typically run across a lot of problems when they try to make an operational mesh.

- Governance on a Large Scale: Federated governance seems to be empowering; nevertheless, without strong rules & automation, it might quickly turn into chaos. Different teams see regulations in these different ways, which leads to different ways of following them and security holes.
- Data Discoverability: Even if teams provide high-quality datasets, it is still hard to find the right data across many other different fields. Data is separated again, but this time in a more decentralized fashion, when there aren't good metadata and search features.
- Not all teams are mature or motivated enough to perceive their information as a product. This leads to data quality that is inconsistent, no version control & not enough documentation. Customers won't use the data if they don't trust it.
- Compliance and Security: When data ownership is decentralised, it is harder to follow standards like GDPR, HIPAA, or company security policies. It is harder to make sure that access restrictions and retention requirements are applied consistently when data is spread out over several domains.

These problems don't mean that Data Mesh is bad; instead, they show how important it is to have automation, better tools, and AI that can help fix problems.

2.3. A Look at Federated Ownership

Federated ownership is a key idea in both Data Mesh and the larger drive towards decentralisation. In a federated architecture, each domain, such HR, finance, engineering, or sales, is in charge of its own data assets. They decide how to publish, keep, and manage their data based on the rules that everyone in the company follows.

This strategy shows that specialist teams are good at what they do. In the end, people understand their data's context, details, and importance better than any central personnel could. Federated ownership also means being responsible at the same time. Teams must make sure that their data is safe, easy to get to, and of excellent quality. This is true not just for their own use, but for the entire business.

For federated ownership to work on a wide scale, it needs clear responsibilities and incentives.

- Help with limiting access and sharing data on the platform
- Standardised contracts or interfaces
- Ways to watch and provide comments on performance

Artificial intelligence is increasingly taking on a key role in making these operations more efficient and easier.

2.4. The Role of AI in Making Data Mesh Easier

Artificial Intelligence, especially machine learning, is an important feature of Data Mesh ecosystems because it automates and adds intelligence to them. Here are some other ways that AI is utilised to make the mesh better:

- Improvement of Metadata: Insufficient or missing information is a major barrier to finding and trusting data. AI can look over datasets on its own and provide you a lot of information, such as column descriptions, data types, lineage, and even PII identification. By looking at how people utilize data fields, natural language processing (NLP) algorithms can figure out what they mean.
- Data contracts in Data Mesh spell down what a dataset has to do, such as how it is structured, how fresh it is, and how available it is. AI can make sure that these contracts are followed. Anomaly detection algorithms can find quick modifications to a schema or missing values in a dataset. They will then let the owners know and stop further issues from happening.
- Usage Analytics: AI could help domain teams understand how their data is being used. Who is asking about which datasets? What fields are the most important? What kinds of questions are likely to fail? This information, which comes from machine learning-based pattern recognition, helps teams improve their data products and decide what improvements to make first.

AI ultimately serves as an intelligent layer inside the Data Mesh, discerning patterns, pinpointing risks, and uncovering insights that may otherwise stay hidden. It makes federated data systems more resilient, easier to find, and easier to use, while also cutting down on human work.

3. ML-Enabled Smart Contracts for Data Collaboration

3.1. What Are ML-Enabled Smart Contracts?

There is a breakthrough idea at the crossroads of AI, blockchain, and data governance: smart contracts that use Ml. These contracts go beyond regular blockchain contracts.

They are smart, active, and good at learning from data to make better decisions over time.

A smart contract is basically an agreement that is written down on a blockchain and enforced by itself. It works when certain conditions are met, like a "if this, then that" system. Your contract goes beyond just following the rules when it uses machine learning (ML). It looks at patterns, finds unusual behavior, and changes how it acts based on previous interactions.

3.1.1. An Overview of the Architecture

The standard architecture for a smart contract that uses machine learning has three main layers:

- Blockchain Layer: Provides an unchangeable, open ledger for putting contract logic into action. This makes sure that all data transfers and policy decisions can be checked and are safe from manipulation.
- AI/ML Layer: Has trained models that help people make decisions inside the contract. This layer may include an algorithm that finds unusual data access patterns or checks to see whether a request to share data goes against a privacy policy.
- Legal Logic Automation: Changes the language of policies, rules, or businesses into digital logic that smart contracts can understand. Think of it as adding GDPR rules or HIPAA compliance checks right into the software.

What happens? Contracts that not only follow rules but also learn, change & enforce them in ways that can be scaled up and are appropriate for the situation.

3.2. Smart Contracts for Enforcing Rules

Traditional smart contracts are great for making sure that basic rules are followed, like paying someone when they provide a product. But when it comes to data sharing, especially within different fields & companies, the rules are far more complicated. Imagine that two medical professionals wish to share patient information for study. There are a lot of things to think about, such as authorization, consent, jurisdiction, time limits & more. That's where smart contracts that use ML really shine.

3.2.1. Dynamic Policy Evaluation

AI may help figure out complicated conditional rules instead of utilizing rigid if-else logic. A contract might learn to spot requests for data access that seem legitimate but are not normal. It might compare how things are being used now with how they were used in the past, finding strange or suspicious activities in actual time.

3.2.2. Automated Enforcement

When an anomaly or violation is found, the smart contract may automatically take action, such as revoking access, informing stakeholders, or changing the terms of collaboration. This means that people don't have to become involved as often, which helps businesses stay in line with regulations that are continuously changing. AI brings a new level of contextual knowledge to policy enforcement, making sure that contracts are not just reactive but also proactive.

3.3. Federated Learning for Compliance with Rules and Finding Unusual Behavior

As data sharing increases, problems with privacy & compliance become worse. When training AI models in a central location, raw information typically has to be put together in one place. This might be a problem when dealing with sensitive or regulated material. In this case, federated learning is important.

3.3.1. What is Federated Learning?

Federated learning lets several parties train a ML model together while keeping their raw information private. Each participant trains a model on their own computer and only sends the modifications to the model (not the data) to a central coordinator. These changes are combined to make the global model better.

This technique is great for working with data from different domains since it keeps sensitive information where it belongs.

- Lessens dangers of not following the rules.
- Allows for shared intelligence while keeping anonymity.

3.3.2. Uses for Data Collaboration

Let's look at a lot of strong examples:

- Finding access to personally identifiable information: A number of financial institutions want to establish a mechanism for detecting illegal access to Personally Identifiable Information (PII). In federated learning, each company trains the model on its own data logs. The global method makes it better at finding access violations without needing to look at each other's data.
- Unauthorized Usage Detection: In a research partnership, federated learning can tell whether a data scientist is utilizing the data in a way that isn't allowed. If an ML model sees that a user is suddenly looking for too much data or getting information that isn't in their domain, it may send warnings or take action, all of which are controlled by the smart contract.

3.3.3. Scalable and Privacy-Conserving

Combining federated learning with smart contracts gives us a way to keep an eye on and control who may access information that is focused on their compliance. It looks like giving your data-sharing network a decentralized, communal immune system that can find and deal with threats in actual time while keeping your privacy.

3.4. Agreements for Sharing Data in Real Time

In traditional company structures, it takes a long time and a lot of work to negotiate and change data sharing agreements. It sometimes takes weeks or months for teams to agree on a contract, talk about the terms & then sign it. But this method doesn't work well in digital surroundings that change quickly.

3.4.1. Self-governing Negotiation

Companies may employ smart contracts with ML to create agreements that change on their own and negotiate terms in actual time. If an e-commerce platform and a logistics company share real-time supply chain data, for instance, their contract may automatically adjust the amount of bandwidth, the amount of time it takes for data to arrive, or the way payments are made based on the current load or market conditions. Imagine a contract that changes the price of hourly data access based on their actual time traffic or changes privacy requirements as more clients ask for them. Static, paper-based contracts wouldn't be able to make this type of change.

3.4.2. Adaptation to the Situation

A data collaboration agreement could include a clause that limits access to particular types of data based on their certain conditions. When a new privacy law goes into effect or a data breach is found, the smart contract can quickly change permissions, notify the right people, and even stop certain operations all on its own. This adaptable nature makes sure that partnerships for sharing data continue strong and are able to respond to changes in the law, technology, and business.

3.4.3. Built-in Trust and Openness

Both parties can see how the agreement is going over time since all interactions are stored on the blockchain. This encourages transparency and responsibility, which makes it easier to audit collaboration and show that you are following the rules to the authorities.

4. Architecture and System Design

Companies who want to be flexible and use AI in a fair way in a decentralized data ecosystem see a revolutionary opportunity in the combination of Data Mesh principles and AI-driven governance. This part describes the standard architecture and system design for a federated data mesh platform that uses AI. Smart contracts & ML are part of the system to make data exchange more trustworthy, open & independent across domains.

4.1. A Look at the Reference Architecture

The proposed architecture has four main tiers, each of which is important for enabling distributed ownership, enforced governance & AI-driven insights:

4.1.1. A layer of people who make data

IoT sensors, transactional databases, CRM tools, and SaaS platforms are just a few of the systems that make up this fundamental layer. These systems create, manage & share both raw & processed information. Each producer registers its schema, metadata, and rules with the appropriate domain catalogue to make sure that their information can be found and accessed in a safe way.

Main features:

- Helps with schema versioning & adding information
- Works with streaming and batch data pipelines like Apache NiFi and Kafka.
- Implements measurements for data quality that are particular to a domain

4.1.2. Federated Domains

Domain-centric data stewardship is the basic idea of Data Mesh & this layer shows that. Each domain works on its own, with its own databases, pipelines, rules & access to these controls. Domains are in charge of checking, storing, changing, and sharing information via APIs or marketplace interfaces.

Main tasks:

- Share data products together with Service Level Agreements & other paperwork: Access control lists (ACLs) are rules that decide who may access resources and what they can do with them. Lineage tracking is the process of following the path & history of information as it goes through different changes & procedures.
- Follow the rules for regional data governance: Corporate sectors (like Marketing and Finance), geographic areas, or specialized applications (like Risk Analytics and Customer Insights) may all be used to define their domains.

4.1.3. Smart Contract Engine

Smart contracts are programmable policy agents that make sure that rules governing access to data, prices, quality standards, compliance & lifecycle events are followed. These contracts are stored on a blockchain or decentralized ledger, which makes them impossible to change and lessens the need for them to be centralized by their administrators.

- Skills: Makes it easier to automate payments, stop access, or send messages if SLAs are broken.
- Includes licensing requirements and rules (such as GDPR and HIPAA)
- Makes it easier for consumers and producers to settle disputes via arbitration

4.1.4. A layer of trust for artificial intelligence

This smart layer always looks at, sorts & ranks datasets based on their quality, bias, compliance, and trust indicators. It works across several areas without having to combine raw information, which protects privacy.

- Completed tasks: Check datasets based on how often they are used, how recent they are, and how reliable they are.
- Use natural language processing (NLP) to find more fields that contain private information, such as personally identifying information.
- Give AI-generated ideas for changes to policies or contracts.

4.2. Basic Parts

4.2.1. Part of the Data Marketplace

This layer provides the mesh's public interface, which lets users, both within & outside the organization, search, request & exchange data items. It combines the activities of searching, visualizing, negotiating contracts, and onboarding.

Roles:

- Helps in faceted search by using metadata filters
- Includes preview panels and family trees
- Allows for flexible pricing plans, such as subscription, pay-per-use & barter.

The marketplace uses APIs and may be used as an external data monetization platform or for inside enterprise applications.

4.2.2. Framework for Trust and Compliance in Artificial Intelligence

This layer uses AI and ML algorithms to dynamically evaluate datasets & suggest ways to improve governance that go beyond just monitoring their routines.

The following submodules are included:

- Anomaly Detection Engine: Finds sudden changes in schema or strange data behavior
- Bias Detection Module: Looks at how fair the statistics are across different demographic groups
- Privacy Risk Classifier: Uses federated training to find data objects that are at high risk.

The AI layer is naturally explainable, which means that domain experts or compliance authorities may look at & check the results.

4.2.3. Managing the Life Cycle of a Contract

The Contract Lifecycle Manager (CLM) is an important tool for keeping track of contracts. It automates the whole process, from creating and negotiating contracts to versioning, expiry, and archiving.

Main features:

- Works with the smart contract engine to enforce rules automatically.
- Gives AI ideas for each clause based on how often it is used and how many times it is broken.
- Includes dashboards for legal & business stakeholders to keep an eye on these contract statuses and key performance indicators (KPIs).

This part makes it easier to share B2B data in a safe & scalable way and reduces the amount of work that people have to do to evaluate contracts.

4.2.4. Framework for Data Observability

This layer lets people who use information and those who create data assess its health, quality, and reliability in actual time.

Skills:

- Keeps an eye on changes, freshness & completeness in the schema
- Gives a temporal trend analysis of how data products are used
- Notifies users of issues based on their limits they set or problems that AI finds

This observability stack works well with the federated domains, making it easier to keep service levels up & fix many problems before they happen.

4.3. Technological Framework

This ecosystem of architectural technologies is modular, cloud-based, and designed to be easy to add to. Here are the main types of technology, along with some instances of each:

4.3.1. Basic Infrastructure Kubernetes:

- Manages microservices in containers across domains to make them more scalable and reliable.
- Istio/Linkerd: Makes communication between services safer & easier to see.

4.3.2. Managing and cataloguing data OpenMetadata, Amundsen, and DataHub:

- These tools let you organize & keep track of data products, metadata, and lineage across many other different areas.
- Apache Kafka / Pulsar: Makes it possible to stream events at a fast rate, which lets data spread quickly.

4.3.3. AI and Machine Learning TensorFlow Federated / PySyft / Flower:

- Makes it easier to do federated learning, which is a way to train models in a decentralized way.
- Scikit-learn and XGBoost make it easier to score trust, find more anomalies, and categorize things in a way that makes sense.
- LLM APIs (like OpenAI and Cohere) make it easier to explain rules & laws in plain terms.

4.3.4. Blockchain and Smart Contracts

- Ethereum and Hyperledger Fabric set up and run smart contracts that control data transactions.
- Chainlink uses outside data sources (oracles) to make these conditional contracts work.
- IPFS/Filecoin: Decentralised storage that you may choose to use for metadata or consent documents.

4.3.5. Safety and Compliance

- The Open Policy Agent (OPA) controls access in a precise, policy-driven way.
- HashiCorp Vault manages secrets & encryption keys in a lot of many other different places.
- Differential Privacy Libraries: Used in the AI layer for analytics that protect privacy.

5. Case Study: Cross-Domain Data Collaboration in Healthcare

5.1. Background and Context

In today's networked healthcare system, data moves across many other areas, such as hospitals, outpatient clinics, insurance firms, research institutions, and government authorities. Each organization is in charge of a different part of the patient's experience, such as appointments with doctors, lab testing, insurance claims & compliance audits. But even though they all want to improve patient outcomes, these groups generally work alone. Think about a regional healthcare network in Europe that has five main hospitals, a few clinics, two big health insurers, and a national regulatory body. Every company generates significant data but stores it in disparate systems, each defined by distinct standards, rights, and governance rules. They recognize the potential of integrated data, especially for predictive diagnoses and chronic disease research, but technological and organizational problems have made it hard for them to work together in the past.

5.2. The Problem

There were three main issues that were important:

- Data Silos: Each hospital kept its own patient information, usually in proprietary electronic health record (EHR) systems. This fragmentation made it very impossible to conduct long-term studies of patients or thorough health evaluations of a single person among providers.
- Regulatory Complexity: Laws like the General Data Protection Regulation (GDPR) and HIPAA put strong limits on how personal health information may be used. These rules were necessary, but they made it very hard to follow them, especially when trying to share data between public and commercial groups or across international borders.
- Absence of Trust: Even if it was possible to share information using technology, the institutions didn't trust each other enough. People were afraid of abuse, unfair competition, and liability, which made it hard for people to work together. All of the stakeholders thought they would have less control.

The executives of the hospital network realized that they needed the latest plan that balanced data sharing with privacy & governance with the latest ideas.

5.3. Using AI-Enabled Contracts to Implement Data Mesh

Instead of putting all the data in one big data warehouse, the network adopted a data mesh architecture. This meant using federated ownership, where each institution managed its own data as a "domain product" and gave others the right to find, query, and use it.

5.3.1. Domains that are federated

Each hospital and clinic put together a separate team to handle their datasets, which included quality, documentation, access controls & availability. Insurers and the regulatory authority set up their own areas where they kept records of things like claims history, pharmaceutical reimbursements, and compliance data. The discovery and access layer was provided by a centralized platform, which let analysts and researchers find datasets from many other different areas without having to move data around.

5.3.2. Contracts Made Better by AI

The system employed AI-powered smart data contracts to control who could get into this complex web of businesses. These contracts weren't set legal documents; instead, they were dynamic regulations that the machine followed on its own. They used AI models that were made to respect GDPR, HIPAA, and privacy laws in their own nation to do the following:

- Enforcement of policies in actual time: If a researcher from Hospital A wants to see de-identified lab results from Clinic C, the AI would look at the request and decide whether or not to provide it to them based on the law and rules for that field.
- AI systems automatically hide sensitive information before it leaves the source, which is how they provide privacy by design.
- Ongoing education: The AI models automatically altered the contracts when laws changed or new company rules were put in place, which cut down on the amount of work that had to be done by hand.

This AI-powered technology turned the complicated legal issues around data sharing into a method that was easy to follow, scalable, and compliant.

5.4. Outcomes

The switch to a data mesh with AI-implemented contracts made a big difference for everyone.

5.4.1. Numbers that can be counted

- How well the query works: In the past, it took weeks to get legal and regulatory approval for requests to access data. Using smart contracts for automatic approvals and enforcement has cut the time it takes to respond to enquiries by 87%, from an average of 10 working days to only 12 hours.
- Compliance Violations: Before the policy was put into place, audits commonly found unintended violations of the rules, particularly in overseas research investigations. In the first year following implementation, policy violations went down by 64% because of automated enforcement and records that could be checked.
- Research Collaboration: Better and safer access to data has led to a 250% increase in collaborative research projects at universities. These include multi-center trials for cancer medicines and AI-based risk assessment models for chronic diseases.

5.4.2. Trust and Culture

The most important outcome was a shift in viewpoint. Institutions that were reluctant to provide data began to make proactive contributions. Contracts that were clear, enforceable, and flexible created a new level of trust.

Patients also felt more comfortable consenting to the research usage of their data because they knew that strong, automatic safeguards were in place.

5.5. Lessons Learned Bringing together performance and openness

One problem was that the AI enforcement layer produced a delay. For analytics queries, the delay was not a big problem, but it was clear in real-time healthcare apps like emergency triage. The team had to set up different levels of data access: one that was intended for speed and supervised by strict pre-approved rules, and another that was meant for operations that needed a lot of transparency and compliance.

- Domain maturity is very important: The data mesh approach works best when domain teams are skilled and can work on their own. Several smaller clinics had trouble at beginning with handling the data they produced. To get everyone ready for the minimum level of readiness, we needed concentrated training, changes to staff, and systems that let people work together.
- The Importance of Being Able to Change Contracts: Healthcare is a field that is heavily regulated and changes quickly. Static rule sets quickly get out of date. The AI-driven features of the contracts that may change with the law were significant. Still, both the legal and data science teams had to keep an eye on the models to make sure they were fair and properly calibrated.

6. Conclusion

As data ecosystems get bigger & more complicated, traditional centralized data architectures have begun to fail when it comes to speed and cross-functional usage. This is where the strong synergy between Data Mesh and Artificial Intelligence (AI) comes into play. It gives businesses a more decentralized, smart, and scalable way to get the most out of their information. Data Mesh essentially fosters federated ownership, conceptualizing data as a product and promoting accountability among the teams that create and grasp the data most efficiently. This change in architecture and culture lets domain teams take charge of their data pipelines, governance & lifetime, which makes it easier for them to work together in a more natural and responsive way. AI, particularly ML, adds automation, forecasting, and adaptability to this decentralized architecture at the same time. Adding machine learning to contracts creates a new way to build trust and openness. These smart contracts not only set clear expectations for data producers and consumers, but they also learn, adapt, and enforce data quality, use rules, and compliance in real time. This flexible strategy is better than static service-level agreements because it lets teams build trust without having to be watched over by people all the time. As these smart contracts evolve based on their use patterns and quality metrics, they provide consistent and dependable data flows across more than several domains.

Federated ownership and AI-driven data governance work together to provide a foundation for long-term & scalable data collaboration. Companies may stay flexible & creative without giving up quality or following the rules. Machine-imposed conventions provide teams greater freedom while keeping them on the same page. People who use data may rely on the outputs they utilize, while people who create data are held accountable in a way that is useful rather than burdensome. This paradigm has a lot of promise to help businesses of all sizes do more with their data. To be successful, you need more than simply technology; you need to modify how you think. It is important to adopt a data product culture, buy advanced technologies & encourage empathy amongst people from different departments. The combination of Data Mesh with AI gives us a way to build strong, smart & cooperative data ecosystems on a solid foundation. This goes beyond just new building designs or another use of AI. It is a complete rethinking of how we work with data across teams, silos, and domains, which will help the whole organisation make better decisions, get insights faster, and build trust. Companies who embrace this vision today will lead the way to a future based on data.

References

1. Hechler, E., Weihrauch, M., & Wu, Y. (2023). Data fabric and data mesh approaches with AI. *Apress Berkeley, Berkeley, CA, USA*.
2. Kumar Doodala, A. N. (2024). Validating UX consistency Across Omnichannel Platform. American International Journal of Computer Science and Technology, 6(6), 87-97. https://doi.org/10.63282/3117-5481/AIJCST-V6I6P109
3. Dey, D., Slomka, P. J., Leeson, P., Comaniciu, D., Shrestha, S., Sengupta, P. P., & Marwick, T. H. (2019). Artificial intelligence in cardiovascular imaging: JACC state-of-the-art review. *Journal of the American College of Cardiology*, *73*(11), 1317-1335.
4. Srigadde, B. R., & Devaraju, J. M. (2024). Building a Reusable AI Connection Utility Class. International Journal of Emerging Research in Engineering and Technology, 5(2), 188-200. https://doi.org/10.63282/3050-922X.IJERET-V5I2P119
5. McDermott, D. (1976). Artificial intelligence meets natural stupidity. *ACM Sigart Bulletin*, (57), 4-9.
6. Shiramalla, R. (2024). Secure Multi-Cloud API Orchestration between Salesforce, Oracle CPQ, and Azure. American International Journal of Computer Science and Technology, 6(3), 102-113. https://doi.org/10.63282/3117-5481/AIJCST-V6I3P108
7. Botsch, M., Kobbelt, L., Pauly, M., Alliez, P., & Lévy, B. (2010). *Polygon mesh processing*. CRC press.
8. Muppaneni, R. K. (2022). From Legacy ERP to Cloud-First: A Transformation Story with Dynamics 365. International Journal of Emerging Research in Engineering and Technology, 3(4), 153-164. https://doi.org/10.63282/3050-922X.IJERET-V3I4P117
9. Muppaneni, K., & Vejella, M. (2023). Security and Data Privacy in Redux Stores. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 4(4), 153-162. https://doi.org/10.63282/3050-9262.IJAIDSML-V4I4P117
10. Farahani, B., & Monsefi, A. K. (2023). Smart and collaborative industrial IoT: A federated learning and

data space approach. *Digital Communications and Networks*, *9*(2), 436-447.
11. Suryadevara, S. S. K., & Shaik, K. (2023). Real-Time Anomaly Detection and Attack Mitigation for Cloud-Based Content Delivery Paths Using AI. International Journal of Emerging Research in Engineering and Technology, 4(1), 175-185. https://doi.org/10.63282/3050-922X.IJERET-V4I1P119
12. Kishor, K. (2022). Communication-efficient federated learning. In *Federated Learning for IoT Applications* (pp. 135-156). Cham: Springer International Publishing.
13. Kumar Doodala, A. N. (2024). Service Virtualization for API-First development: A shift-Left Testing Strategy. American International Journal of Computer Science and Technology, 6(4), 50-58. https://doi.org/10.63282/3117-5481/AIJCST-V6I4P105
14. Parakala, A. (2024). Agentic Automation: What's next for Jobs. American International Journal of Computer Science and Technology, 6(6), 25-35. https://doi.org/10.63282/3117-5481/AIJCST-V6I6P103
15. Ghimire, B., & Rawat, D. B. (2022). Recent advances on federated learning for cybersecurity and cybersecurity for federated learning for internet of things. *IEEE Internet of Things Journal*, *9*(11), 8229-8249.
16. Allenki, S. S. (2022). Securing Databases in the Cloud with RBAC and Encryption Best Practices. International Journal of Emerging Research in Engineering and Technology, 3(3), 173-182. https://doi.org/10.63282/3050-922X.IJERET-V3I3P117
17. Vppalapati, M. (2022). The Storage Stack Nobody Draws: Cabling, Panels, and the Illusion of Isolation. International Journal of Emerging Research in Engineering and Technology, 3(2), 211-220. https://doi.org/10.63282/3050-922X.IJERET-V3I2P121
18. Manias, D. M., Shaer, I., Yang, L., & Shami, A. (2021, December). Concept drift detection in federated networked systems. In *2021 IEEE global communications conference (GLOBECOM)* (pp. 1-6). IEEE.
19. Muppaneni, K. (2023). Virtual DOM vs Real DOM: Performance Benchmarks. International Journal of AI, BigData, Computational and Management Studies, 4(4), 180-189. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V4I4P118
20. Shiramalla, R. (2023). Optimizing Cross-Platform Enterprise Integrations Using Workato: A Case Study of Salesforce and Oracle SaaS Applications. International Journal of Emerging Trends in Computer Science and Information Technology, 4(1), 232-243. https://doi.org/10.63282/3050-9246.IJETCSIT-V4I1P124
21. De Martino, V., & Palomba, F. (2023). Classification, challenges, and automated approaches to handle non-functional requirements in ml-enabled systems: A systematic literature review. *arXiv preprint arXiv:2311.17483*.
22. Gaddam, R. R. (2022). Advanced Data & Model Drift Detection at Scale. International Journal of AI, BigData, Computational and Management Studies, 3(2), 124-136. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V3I2P113
23. Takkalapally, D. (2024). ShiftLeft-AI: Machine Learning Framework for Proactive Performance Assurance in CI/CD Pipelines. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 5(4), 285-296. https://doi.org/10.63282/3050-9262.IJAIDSML-V5I4P126
24. Issa, W., Moustafa, N., Turnbull, B., Sohrabi, N., & Tari, Z. (2023). Blockchain-based federated learning for securing internet of things: A comprehensive survey. *ACM Computing Surveys*, *55*(9), 1-43.
25. Muppaneni, R. K. (2022). Data Privacy in the Age of AI: How Dynamics 365 Handles Regulatory Challenges. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 3(4), 159-170. https://doi.org/10.63282/3050-9262.IJAIDSML-V3I4P117
26. Srigadde, B. R. (2024). Agents, LLMs, and Salesforce with Multi-Cloud Provider (MCP). International Journal of Artificial Intelligence, Data Science, and Machine Learning, 5(3), 277-288. https://doi.org/10.63282/3050-9262.IJAIDSML-V5I3P127
27. Radhika, K., Alla, B. S. R., Mamidi, S. B., & Puppala, N. (2022, December). A Review on Federated Learning-Based Network Intrusion Detection System. In *International Conference on Information and Management Engineering* (pp. 707-713). Singapore: Springer Nature Singapore.
28. Suryadevara, S. S. K., & Nakirikanti, S. (2023). Privacy-Preserving Personalization Using Federated Learning in AEM. International Journal of AI, BigData, Computational and Management Studies, 4(4), 190-199. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V4I4P119
29. Parakala, A. (2024). Self-Learning Bots & Cloud-Native Platforms. International Journal of Emerging Trends in Computer Science and Information Technology, 5(4), 132-141. https://doi.org/10.63282/3050-9246.IJETCSIT-V5I4P114
30. Hong, C. S., Khan, L. U., Chen, M., Chen, D., Saad, W., & Han, Z. (2021). *Federated learning for wireless networks* (pp. 3-239). Singapore: Springer.
31. Allenki, S. S., & Lee, N. (2022). Performance Tuning Cloud-Hosted Databases: Resource Allocation & Query Optimization. International Journal of AI, BigData, Computational and Management Studies, 3(4), 152-163. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V3I4P116
32. Khan, R. S., Sirazy, M. R. M., Das, R., & Rahman, S. (2022). An ai and ml-enabled framework for proactive risk mitigation and resilience optimization in global supply chains during national emergencies. *Sage Science Review of Applied Machine Learning*, *5*(2), 127-144.
33. Takkalapally, D., & Takkellapally, M. R. (2024). AI-SynPerf: Synthetic Data Intelligence Framework for 5G Mobile Performance Simulation. International Journal of Emerging Trends in Computer Science and Information Technology, 5(1), 182-194.

https://doi.org/10.63282/3050-9246.IJETCSIT-V5I1P118

34. De Martino, V., & Palomba, F. (2023). Classification, challenges, and automated approaches to handle non-functional requirements in ml-enabled systems: A systematic literature review. *arXiv preprint arXiv:2311.17483*.
35. Vppalapati, M., & Talasila, P. K. (2022). Correlated Independence: Why Redundant Storage Systems Share the Same Fate. International Journal of Emerging Trends in Computer Science and Information Technology, 3(1), 169-179. https://doi.org/10.63282/3050-9246.IJETCSIT-V3I1P119
36. Gaddam, R. R. (2022). Cost-Aware Autoscaling for Batch vs. Online Inference. International Journal of Emerging Trends in Computer Science and Information Technology, 3(4), 134-143. https://doi.org/10.63282/3050-9246.IJETCSIT-V3I4P113
37. Mitra, A., Bera, B., Das, A. K., Jamal, S. S., & You, I. (2023). Impact on blockchain-based AI/ML-enabled big data analytics for Cognitive Internet of Things environment. *Computer Communications*, *197*, 173-185.
38. Taluri, R. (2024). A Hybrid Data Engineering and Generative AI Architecture for Intelligent Data Governance, Metadata Management, and Automated Data Quality Assessment on AWS. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 5(2), 230-240. https://doi.org/10.63282/3050-9262.IJAIDSML-V5I2P126