




# Data Mesh Architecture

How Leading Companies  
**Win with Data**

# Table of Contents

Introduction: The Modern Data Deluge	3
What Is Data Mesh?	8
How Data Mesh Creates Better Data Products	14
Your Path to Data Mesh: A Structured Framework for Maturity	20
Tips for Success: Do's & Don'ts	25
Summary	28
How Marlabs Can Help	29

# Introduction: The Modern Data Deluge





Data has multiplied at a relentless pace over the last two decades. The sheer volume of data that businesses can access has vastly outgrown the paradigms and architectures originally created to derive organizational insights.


Not only has the amount of data we can accumulate multiplied, but so has its complexity. Unstructured and semi-structured data introduced previously unimagined opportunities for innovation – and new challenges. When companies struggle to convert this raw data into actionable insights, they lose revenue, fall behind their competitors, and fail to meet customer standards.

The traditional data management model no longer works for modern organizations. At its core, this model simply wasn't built for the scale, speed, and complexity of modern data demands. Now, we face systems where:

- Centralized architectures are overly rigid.
- Overworked IT teams can't keep up with requests.
- Customized reports and data access require more waiting than business users can afford.
- Cross-functional projects require such complex data gathering that collaboration is stifled.
- The availability of data history and documentation is limited and inconsistent.
- Trust in organizations' data erodes.

For more about the history of the traditional model and how we ended up here, check out our companion whitepaper, [\*Data-as-a-Product Paradigm: How Leading Companies Unlock Their Data.\*](#)





# What Makes Data Mesh Architecture Different?





Over the years, many solutions promised to fix persistent data problems, but most fell short. Why? Because almost all of these solutions simply tried to apply more and more complex technology to the issues.

Data mesh works differently. While it sometimes requires technological changes, data mesh isn't a technology shift; it's a fundamental change in how teams own, manage, and deliver data. It's a paradigm shift that addresses not only technology but also people, processes, systems, and culture.

Zhamak Dehghani developed the framework for data mesh architecture to address the recurring data challenges, bottlenecks, and slow delivery of insights happening across industries. Data mesh decentralizes data ownership to domain teams (like sales, marketing, and finance) while also embedding strong, unified governance into the data product lifecycle across domains. This reduces delays; aligns data with business context; and ensures the data products conform to consistent quality, security, and compliance rules.

A data mesh architecture adopts a data-as-a-product (DaaP) paradigm and rejects centralized pipelines controlled by IT. That means instead of funneling all requests through a central data team, the organization empowers different domain teams to manage and maintain their own data products. This puts data in the hands of the people closest to the data and increases agility, accountability, and alignment with the problems they're trying to solve. Governance then ensures every data product is trustworthy, discoverable, and interoperable — without slowing progress or innovation.





# Preview What's Possible

Rather than waiting weeks for IT to build new reports, imagine generating them instantly via self-service access to accurate, pre-packaged data that's ready to use.

What if your sales teams no longer had to chase spreadsheets and your data analysts no longer had to waste valuable time reconciling conflicting data sources for every single request? How would it change your company to have up-to-date, accurate data on customers, inventory, and operations available at any time via a self-service data marketplace?



With data mesh, all of that is possible.


The shift to data mesh isn't about buying another tool (though that might sometimes be necessary). It's about rethinking organizational mindsets around data and unlocking powerful new ways to uncover value from the data you already have.

## Data Mesh Produces Results

Data mesh goes beyond theory. It's a proven architecture that has gained traction with CIOs and CDOs worldwide.

Early adopters of data mesh solve long-standing data challenges while gaining a serious competitive edge. At its core is the practice of treating data like a product, which we explore in more detail in our companion whitepaper, *Data-as-a-Product Paradigm: How Leading Companies Unlock Their Data*. Harvard Business Review reports, "[Companies that treat data like a product can reduce the time it takes to implement it in new use cases by as much as 90%.](#)"





# Who Is Data Mesh Best For?

Not every organization is the ideal candidate for data mesh right away, given how extensive of a transformation is required for the mindset and practice shifts. However, data mesh and data-as-a-product thinking are the future for data-driven organizations.

So, even if you're not ready to take on the transformation now, learning about it today will give you insights into how to make smarter, more informed decisions about what's ahead.


A data mesh architecture works best for mid-sized to large organizations, particularly those struggling to manage data and discover timely insights. If your business is small, then the traditional, reactive model will likely keep working for you for a while. Another major consideration would be if all your systems are on-premises, in which case moving to a data mesh would require a larger commitment in technological change.

That said, as your business grows, the model that works today likely won't be able to scale (meaning it won't be able to adapt to growth, expanding data volume, and increasing business needs). For now, implementing a data mesh at a small company with limited systems might be overkill, since you likely aren't facing as many of the issues, or at least not at the same intensity, that more mature organizations have run into. You don't need it...yet.

Whether you're considering making the switch to data mesh soon or just want to understand its potential for your future, we'll explain data mesh principles and how to develop high-quality data products clearly to help cut through some of the confusion caused by misleading information written on the topic.

Unfortunately, sometimes even organizations who say they've shifted to a data mesh architecture have still implemented the paradigm and practices incorrectly. Often, that's through a lack of understanding of the core data mesh principles as something that "just happens" as part of a technology change. However, data mesh is about much more than technology; it's an organizationally transformational initiative that foundationally changes how your teams think about, approach, and share data.

So, if you're a CIO, CDO, or other business leader looking for clarity and support, you're in the right place.





# By the time you finish reading this whitepaper, you'll understand:

- What data mesh is
- The four pillars of data mesh
- How a data mesh creates better data products
- How to follow a structured framework for data mesh maturity
- Strategies for success with implementing DaaP thinking and data mesh
- How Marlabs' experienced data consultants can help



## Let's hop in.

# What Is Data Mesh?

Data mesh is a **sociotechnical, architectural, and organizational model** that decentralizes data ownership and delivery. It's a logical, domain-driven view of data, not a single platform or technology.

Thoughtworks, the company where Zhamak Dehghani first defined the concept of data mesh, explains:

*"Data mesh is a decentralized organizational and technical approach in sharing, accessing, and managing data for analytics and ML. Its objective is to create a sociotechnical approach that scales out getting value from data as the organization's complexity grows and as the use cases for data proliferate and the sources of data diversify. Essentially, it creates a responsible data-sharing model that is in step with organizational growth and continuous change."*


*-[Thoughtworks Technology Radar](#)*







# Data mesh is based on four pillars:

<b>Domain-oriented ownership:</b>	<b>Data-as-a-product approach:</b>	<b>Self-serve data infrastructure:</b>	<b>Federated computational governance:</b>
Domains own the data they know best, which reduces bottlenecks while improving quality and relevance.	Every dataset is treated like a product, where product owners provide clear documentation, meet quality standards, and ensure discoverability.	Domains receive the tools and platforms they need to publish, access, and use data without heavy IT intervention.	Shared policies and automated checks ensure compliance, security, and interoperability while also preserving agility.






# What Is a Domain, and Why Does It Matter?

Domains are business functions (like marketing, finance, HR, or ops). Domain-driven software development is a concept started by Martin Fowler, aligned with domain-driven design (DDD) principles. As part of data mesh, Zhamak Dehghani applied his ideas to data management.

In traditional data management, business functions depend on IT to pull, clean, and prepare data, often causing delays and missed opportunities. With data mesh, each domain acts as a data product team that owns its data end-to-end (from ingesting that data to serving it to internal and external teams). Domain-level ownership drives and enforces accountability, data quality, consistency, and responsiveness while reducing reliance on centralized gatekeepers. The domain team turns data into discoverable, governed, and accessible products that others can discover, trust, use, and even reuse as a part of a new data product.

For example, the marketing team might create data products to compare campaign performance across channels, and the finance team might create data products that combine data on revenue, expenses, and forecasting. At the same time, production teams may develop a data product for customer subscription churn across product types. All three teams may reference or incorporate data from the other teams' products to make informed decisions about where marketing, production, and finance interact. They may combine, compare, and reference each other's data products to determine which product types should receive the most funding for research and marketing. Or they may build an AI machine learning model to suggest which type of campaigns to use for a new product offering based on what has historically been most effective — and cost-effective — in reaching audiences who both purchased similar products and renewed subscriptions for them.



# Treating Data Like a Product

The data-as-a-product (DaaP) paradigm is about introducing product management thinking to data management. Instead of viewing data as the unintentional byproduct of digital activities, we intentionally design data products with purpose, audience, and iteration in mind.



A major part of the shift to product thinking is defining end users, business goals, and feedback loops for every data asset or product. Domain teams that treat data like a product view the products they create as valuable, purposeful, and iterative. They design their products to meet specific audiences' needs and with specific use cases in mind.

At the same time, they also ensure their products are accessible to other audiences and for other use cases, keeping in mind that the products will be available across the organization via a self-service platform. The DaaP paradigm is what ensures that when users find relevant data in that self-service platform, they can trust that is clean, accurate, and up-to-date, ready to be used for their own needs.

DaaP leads to payoff like:

- Faster, more reliable decisions
- Customer retention
- Improved data quality and trust
- Easier, more effective collaboration
- Improved user experiences
- Accelerated analytics, AI, and innovation
- Lower operational costs and more efficient operations
- Reusable data – at scale

For a deeper dive on how treating data like a product benefits businesses, check out our companion whitepaper, [\*Data-as-a-Product Paradigm: How Leading Companies Unlock Their Data.\*](#)





# Goodbye, Centralized Warehouse. Hello, Data-Platform-as-a-Service.

A common concern about making the transition to a data mesh is the question of, “What happens to the centralized IT team?” It’s important to know that **you aren’t replacing the IT team**, just changing their responsibilities—specifically, changing their responsibilities to be more manageable.


Over time, the IT team transitions to focus more on data-platform-as-a-service and governance efforts, rather than filling every data request across an organization.


As you replace the centralized warehouse with a data-platform-as-a-service framework, your IT team will focus on stewarding the platform and providing governance. The technology team builds a self-service platform that enables decentralized, domain-oriented data. Then, they manage the platform’s design, CI/CD, version control, traceability, lineage, data as code, security and privacy, metadata management, automation, observability, costs, life cycle ownership, policies, and data contracts.

Not only does this eliminate bottlenecks and ensure compliance, but it ensures consistency, maintenance, and adaptability. The IT team shifts to give the business simple access to findable, well-described, reusable data products through a self-service platform – but they don’t have to also create all the products themselves.

Instead of IT leading product discovery, the domain teams build the data product, identifying customer needs and developing the use cases, vision, roadmap, objectives, and key results. They collect and validate the data, implement the solutions, maintain the product, and drive adoption. It’s also their responsibility to enrich, cleanse, and apply the rules (provided by the IT team) to the data.

When the teams’ efforts come together in a data marketplace (e.g., Snowflake), this facilitates seamless integration and collaboration, providing access to ready-to-use and AI-ready data products.






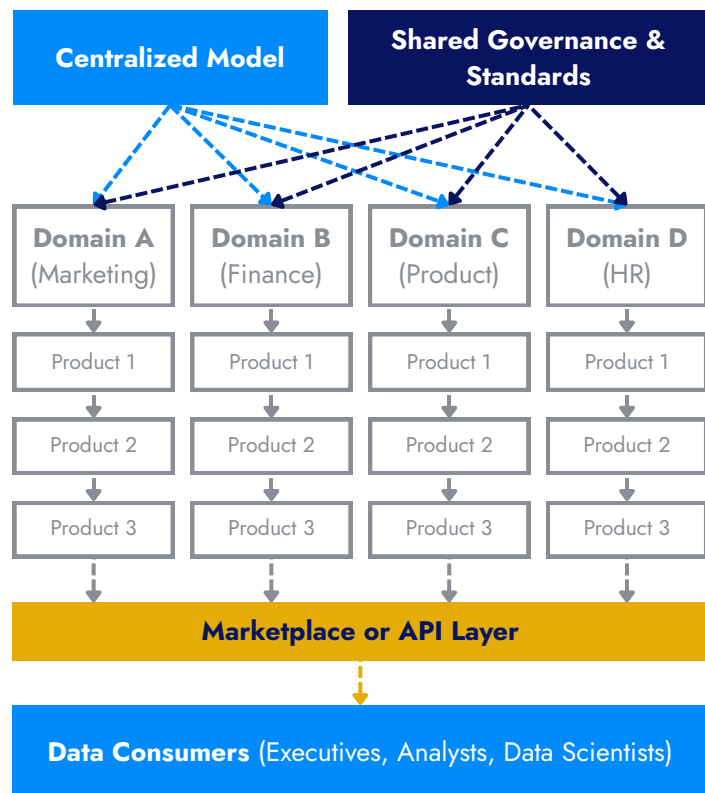
# How It All Comes Together


1. Data ownership moves from a single IT team to **decentralized domain owners**.
2. Those owners follow a **data-as-a-product paradigm** and apply product management thinking to create trustworthy data products.
3. They then distribute those data products, including their data contracts, through a **self-service data access platform**, such as data port, data marketplace, or API.
4. The IT or data team ensures that platform is secure and interoperable through **federated governance**, shared policies, and automated testing and maintenance.

Through this approach, organizations make data scalable, business-aligned, reusable, and valuable across the enterprise. It looks something like this:

Delays disappear when you're no longer waiting for a single team to serve everyone. Your organization transforms into a network of data creators and empowered consumers working from the same trusted data products designed for action.

That's how you make data work like the rest of your business: **intentionally, efficiently, and with purpose**. Even if your functions are otherwise siloed, a data mesh architecture can start to break down those walls so teams can effectively work together and build on one another's achievements, which is essential to compete in the modern digital world.





# How Data Mesh Creates Better Data Products

## What's a Data Product?

A data product is not the same as a data-as-a-product paradigm but rather the result of developing data assets with such an approach.


A data product is a consumable, standalone product or asset that you design, build, and maintain to fill the needs of internal or external users. A data product isn't just a raw table or even report of data. It's a managed, documented, versioned, and discoverable dataset engineered for reuse, one that incorporates data from multiple sources.

By applying the components of data mesh, we build more useful, sustainable, and accurate data products. You build a data product with such clear specifications and understanding and quality that someone can **repurpose it in ways you never imagined**. This intentional design turns raw data into a strategic, reusable, measurable asset and shifts your data culture from reactive reporting to proactive, product-driven data enablement.

Every data product consists of code, data sets, metadata, and infrastructure. Different layers of data products are built for different audiences with various levels of technical need. A data product could simply be a conglomerated data set that compares customer social media interactions across platforms. Or it can be more complex, like a 360° customer profile developed from data from both in-person and digital interactions across the organization.

### Other data product use cases include:

- Enterprise risk scoring
- Equipment failure probability
- Performance management data
- Inventory, demand, & stock forecasting
- Spend analysis cube
- Vendor performance & spend
- Workforce demographics & turnover
- Marketing channels attribution



# Key Characteristics of a Well-Built Data Product

A well-designed data product created within a data mesh architecture has certain characteristics that make it more useful. A simplified version of the most important of these characteristics is an acronym commonly used in data-as-a-product thinking, **FAIR: findable, accessible, interoperable, and reusable**.

Sometimes it can help to be a bit more specific, as we've done in the following list. Ultimately, the important takeaway is that these characteristics ensure a data product is well designed, purpose-built, and developed with the consumer in mind.

Whether it's a sales manager reviewing performance, a data scientist training a model, or a compliance officer preparing for an audit, it's important that each consumer has access to reliable, actionable data that's ready to use. This reduces the friction involved in acquiring data through the traditional "request-and-wait" process.

## A well-built data product is...



Discoverable



Understandable and self-describing



Addressable



Trustworthy



Secure and governed



Independent and interoperable




Natively accessible



Observable



A clearly defined owner



Some sources may call certain characteristics by other names or combine some of the characteristics we list below. In general, though, a well-built data product is:



### **Discoverable**

Users can easily find data products that are available to them through an intuitive, searchable data product marketplace.



### **Understandable and self-describing**

The data product includes rich metadata about what it is, who owns it, and how to use it. Typically through a data contract (which we discuss more in a later section), a data product clearly communicates the purpose, structure, organization, quality, and lineage.



### **Addressable**

Data products come with a fixed and reliable way for users to retrieve them with a unique, consistent address, like through a standardized URL or API endpoint.




### **Trustworthy**


In a data mesh architecture, platform and domain teams verify the trustworthiness of data and continually check for updates. They also ensure data products come with documentation, service level objectives (SLOs) that are dynamically updated, and quality metrics to guarantee freshness and trust. They provide understandable metadata that lends to trust, particularly by explaining data lineage: when the product was produced, its timeliness, how often it's updated, how complete it is, and its statistical shape.



### **Secure and governed**

The data platform includes and enforces role-based access controls and complies with governance policies by design. It also protects privacy and blocks unauthorized users. Access levels can be controlled even down to granular access within a specific data product based on business needs. Data products and platforms both include protections that comply with laws to protect sensitive personal information.





## Independent and interoperable

Data products are designed to be autonomous and meaningful on their own. At the same time, owners build their data products to work seamlessly with other data products by following shared standards like naming conventions, schemas, and access protocols, all of which enable integration. They deliver richer insights when joined and correlated with other data products or when otherwise used as a foundation for creating a new data product.



## Natively accessible

Data products are available in a variety of formats (with multi-modal access) to meet different needs. In all formats, users access data products through a user-friendly, self-service platform, such as a data marketplace, API, or data port. This way, business users don't need to learn a new tool and skill for each product type or business domain.




## Observable


The data platform provides visibility into data products' performance and usage with a documented service level (e.g., update frequency, uptime, query statistics). The platform also supports monitoring and alerting to ensure reliability.



## A clearly defined owner

Each data product is owned and managed from end to end by a cross-functional domain team closest to the data and its business context. The team includes a roadmap, a lifecycle, and feedback loops for the product, and the product owner reports to the business, not IT. Because the owner understands the data and its business purposes, they maintain the product, iterate it, document it, and share it with the rest of the organization through a data platform where others can see and access the data.





# Why Data Characteristics Matter


If a data product is missing any of the characteristics we listed, it becomes stale and needs to either be reworked or be pruned out of the ecosystem.

Following a data-as-a-product approach as part of data mesh means continuously iterating products. This keeps products both high-quality and actively useful. Ensuring these qualities is how a data-as-a-product approach creates data that people trust. Without this approach, organizations often end up using flawed data and can't vouch for its quality, trustworthiness, and timeliness. **In a data mesh architecture, it's non-negotiable that every product be trustworthy.**

The actual characteristics of the data product make the difference between raw datasets that are hard to use versus data assets that can be understood and repurposed across a variety of applications and teams. It's the difference between traditional, hard-to-use data and modern, accessible data that gives you a strategic advantage.

Data mesh principles ensure that your data product is easy to find, understand, and access; that it's accurate, secure, and interoperable; and that someone is actively managing and updating its quality.

The goal is to make the right data available to the right people, at the right time, in the right format.





# The Right Format: The Benefits of a Data Contract

Data products should be structured to be engaging, intuitive, and easy to access. You also want to design a data product with such clear specifications and understanding of the quality of the data that someone can repurpose it for a use that you haven't even imagined.

Within a data mesh architecture, you'll package all those specifications into something called a data contract, which travels with your data product to explain all its ins and outs. A data contract details each data product's attributes (size, shape, etc.), quality, and usage, so others can combine or repurpose it on their own. Data contracts are key to ensuring your data products have the characteristics that make them trustworthy across teams.


By using data to create purposeful data products, we create something tangible from often bulky data that might otherwise be difficult to imagine, grasp, and utilize. Then, by pairing each data product with a data contract, we turn that tangible product into a reusable and reliable asset across domains, rather than a siloed asset that might be "available" to many teams but is only truly functional for the team that created it.

Data contracts also set up the building blocks for everyone in the organization to develop a common language, an impact that is often underestimated. A shared language bridges gaps between IT, business, and analytics teams and breaks down misalignments and misunderstandings. Instead of each team working sequentially until their part of the project is complete (and throwing the initiative over the wall to another team), the data contract helps teams work more collaboratively.

## Other aspects that fall into place more easily because of data contracts include:

- Collaborating across domains on use cases
- Breaking down communication barriers
- Ensuring clarity with models
- Managing project and data complexity
- Implementing data products
- Searching for relevant data products








# Your Path to Data Mesh: A Structured Framework for Maturity

## Reaching Organizational Data Maturity with Data Mesh

Every organization is different, so there's no single path to adopt a data mesh architecture. However, it can be helpful to consider what organizational data maturity looks like while transitioning to a well-integrated data mesh architecture.

Achieving organization-wide data maturity happens in stages. It's a journey to move from fragmented, reactive data practices to a proactive, product-driven ecosystem that fuels growth and innovation. Each phase builds the foundation for the next. In the book *Implementing Data Mesh* by Jean-Georges Perrin and Eric Broda, the phases are described in 5 phases, which we've adapted below and represented as evolving modes of transportation. Like with the progress in transportation methods, there will be new developments **within** the stages of organizational maturity outlined below, and development may not be completely linear. However, there's a much more highly noticeable difference **between** the stages.





Moving to each new level takes work. Moving to a data mesh architecture requires **major organizational transformation**, which means proper change management and team enablement are key. Here, we'll describe how we help organizations develop and mature their data mesh architecture through each phase, taking them from an initial, siloed structure to an architecture where data products create foundations for innovation.


## Structured Framework for Maturity




### Phase 1: Initial Siloed Data

In the earliest phase of organizational data maturity, organizations face non-standardized data products, often due to a problematic adoption of any consistent approach to handling data. At this point, teams mostly use data only reactively for reporting. Typically, these issues arise from:

- Siloed data with manual management
- A lack of clear standards
- Ad-hoc and replicated creation of data products and requests across teams
- Data costs that outweigh their return on investment





At this stage, teams recognize the need for change, but their data initiatives often lack clear ownership and stall out. This is where most organizations start when they reach out to data consulting firms like Marlabs for support.

The first step we then take towards creating a data mesh foundation is to create a pilot project or proof of concept (PoC) with a well-defined use case. In doing so, the project demonstrates tangible results and value to stakeholders – thus driving internal adoption – and builds a foundation for future use cases.

To complete a proof of concept, we first identify an engaged business partner with a clear understanding of their data and a specific, actionable goal or KPI. Then, we manually run a small DaaP pilot, applying “FAIR” (which stands for “findable, accessible, interoperable, and reusable”) principles to the project. Once the pilot is complete, we communicate the results to the organization to demonstrate the value of the approach.

## Phase 2: Treating Data as a Product within Domains

In the second phase, we use the foundations developed in the proof of concept to establish data-as-a-product concepts throughout the organization. To do this well, we also start to implement another important factor of data mesh architecture: the establishment of domains. From these domains, we then establish dedicated product owners.

As domain teams and product owners identify use cases and input them into a demand management framework, we manage the creation of complex data products with reliable data contracts and self-service interfaces. These ensure the accessibility, quality, and reusability of the products. Because teams begin thinking of data a product (instead of a byproduct), **the culture shifts towards a structured and strategic approach to aligning with business objectives.**

With DaaP foundations established, enterprises begin to consider how to move past proofs of concept and sustainably scale enterprise-wide data product production. That typically includes:

- Expanding pilots across multiple domains
- Creating standardized and reusable components and templates during manual pilots
- Identifying manual processes that can be automated or otherwise replaced

By spreading product-thinking mindsets and continuing organizational conversations about the change, the organization prepares to build efficient, repeatable processes.







## Phase 3: Automated, Manufactured Products

Once the organization scales the foundational DaaP paradigm across domains, those domains create data factories to **streamline the development of data products**. They implement the standardized components and templates they created through their pilot projects to automate processes and develop data products on demand.

This automation keeps production:

- Efficient
- High quality
- High velocity
- Reliable
- Consistent

At the same time, leadership, domain teams, and consultants work to continue spreading data product management thinking through the organization, especially as the data factories make it simpler and smoother to “manufacture” high-quality products that measurably impact KPIs. That creates a firm foundation of mature domain ownership and decentralized accountability, which paves the way for effective standardized governance and self-service capabilities across the organization.





## Phase 4: Governance & Self-Service Enablement

Along with strong domain ownership and accountability in managing data from end-to-end, the maturing data-as-a-product approach makes it possible to add the final two components of a data mesh architecture: a decentralized approach to governance and self-service capabilities. Rather than top-down IT management, organizations embed governance, standards (including FAIR principles), and monitoring into workflows. They also implement standardized data product certification processes to ensure compliance, quality, and security.


These processes, alongside self-service enablement across domains, ensure that data products are interoperable. That way, products in such a diverse and dynamic data ecosystem can be understood, combined, or built upon by anyone with appropriate access. Consumers can access data products on demand via self-service, and because of the standardized governance procedures and data contracts, they know they can trust the data products they find.

**Enterprise-level, decentralized governance and interoperability maximize the impact and ROI of data assets across the enterprise.** Not only this, but when paired with strategic change management, they also allow organizations to begin leveraging data for strategic decision-making, AI-powered initiatives, and other innovations.

## Phase 5: Innovation

**At their most mature, DaaP-driven organizations use data products to drive investments, innovation, and returns on investments.** Domain leaders explore new opportunities, improve decision-making, and gain competitive advantage not only through their own domains' data products but also by collaborating together and building off of the products created by other domains. These mature organizations show a strong sense of ownership and accountability for treating data as a product across domains according to data mesh principles.

Domain leaders also fully integrate products as key strategic assets, and product owners embed continuous improvement into every data product to ensure its usefulness grows and adapts with the organization. At the same time, the IT team manages the platform and ecosystem to ensure data products have all the qualities needed for compliance, interoperability, and intelligent decision-making.





# Tips for Success: Do's & Don'ts

Becoming an organization with a mature data mesh architecture and foundation is not an overnight operation.

It's a **major transformation** impacting people, processes, and technology – one that requires a major commitment of time and effort up-front. With such a large change, it's easy to make a misstep. Even otherwise-mature organizations make mistakes implementing data mesh architecture. Here, we've gathered the most important insights we've seen from helping our clients succeed – and from watching other companies struggle.

## **Don't underestimate what it takes.**


Don't underestimate the work involved in changing to a new paradigm. Just like the movement from waterfall to Agile was a huge pivot, this is going to be a substantial shift. Assess what must be done, gain buy-in from leadership and influential people at all levels, and then be prepared to run into obstacles.


## **Do have leadership champion the change.**

Executives play a huge role. Leaders have to champion and support the establishment of the technical infrastructure to make data mesh work. Empower the domain leaders and team by promoting the value of a data-as-a-product paradigm, reiterating value creation, making systematic changes to prioritize the initiative, and celebrating successes of change agents during the pilot program and throughout the entire journey. Bring people along with you as a team.

## **Don't rely on IT to drive the change.**

Data mesh principles must be driven by the business, not IT. This is critical. IT should stand up the platform (e.g., Snowflake or Immuta), but it's the business users who should own the data products, manage resources, and create data. Give the business a dedicated architectural support team of technologists and data engineers who report to and are accountable to the business. Limit IT's role to platform maintenance, development, and governance.





### **Do empower business users.**

Choose motivated, committed domain experts and empower them as product owners. You'll need to prepare the people closest to the business domain for a major shift and make sure they're willing to learn new paradigms and processes in order to be engaged champions.

The product owner needs to be someone who is accountable to the business. Technical people should be accountable to the business, too, not to IT. Communicate this upfront to set expectations for this cultural transition. It might feel like a lot of extra work at first, but in the long run, it helps everyone stop doing time-consuming tasks that feel like a waste. Help teams think of it as a “pay more now or pay way more later” decision.

### **Do plan for organizational change management.**

Since this is a completely different way of approaching data, companies need to embrace organizational change management throughout the journey to ensure adoption. As you dig into the weeds of implementation, it might be all too easy to be consumed by the effort and wind up sidetracked from ongoing communication and change management. It's important in those times to remind yourself — and your change agents — that the best systems in the world will fail if you don't prepare your people. Help them feel like they will be a part of the new order, not left behind.


### **Do prepare to make technology changes.**


Technology isn't everything — not by a long shot. But it can be a showstopper, so you may need to commit upfront to bringing on new technologies if needed. For example, if your organization has some systems on-premises, moving to the cloud may be essential for data mesh to succeed. Make sure you gain buy-in early for the possibility of investing in new technology to support a modern data ecosystem.

Whether you need to make major technology changes or not, it's always vital to take the time to put in place the technical framework you'll need. You'll also want to prepare associated technical processes in place for managing code, versioning, continuous improvement and development, and sprints.

### **Do follow data mesh principles closely.**

It may even be helpful to read Dehghani's book [Data Mesh: Delivering Data-Driven Value at Scale](#) before you start. Just like incorrect Agile adoption in software, it's easy to veer off course if you don't understand the data-as-a-product paradigm and the other components of data mesh. Even notable brands mistakenly say they're treating data as a product in a data mesh, yet create data assets in a centralized architecture — which will lead to problems down the road.





### **Don't ignore governance and accountability.**

Data product ownership without clear roles and responsibilities for how to create products will inevitably result in low-quality, inconsistent data products – which defeats the whole purpose of a data mesh. Without data contracts, standardized templates, secure access, and other governance necessities, products not only grow stale, inconsistent, and untrustworthy, but they also open the organization to compliance and regulatory risks.

### **Don't quit when friction sets in.**

As tempting as it may be to go back to what's familiar when change starts making teams tense, don't give up hope on the data mesh principles. Don't centralize data products under IT or relinquish control of owning the data products to IT. Know that organizational tension with such a major shift is normal; understanding, competency, and adoption of the new architecture will take time. Compromising the paradigm principles actually only prolongs the time under tension, and it could jeopardize your investment and compromise the ability to scale.

### **Don't jump into data mesh if you're not ready.**

Implementing a data mesh architecture takes a major investment of time, energy, and commitment to change. It's for organizations with genuine business needs and pain points that would be solved by a move to domain-level data product ownership. Ensure that your organization needs a decentralized model and that you're ready to give it the time and energy it requires.





# Summary


As organizational data continues to exponentially multiply, traditional data management models aren't able to keep up.

Overworked IT teams bottleneck processes, collaboration stalls, and teams mostly receive data that's simply the unintentional byproduct of digital activities. Data mesh architecture changes all of that.

More than just adding another technology to the stack, data mesh is a sociotechnical, architectural, and organizational model that introduces new ways of thinking about and managing data. Data mesh:

- Puts data in the hands of the people and domains closest to its context.
- Empowers those people to intentionally design data products.
- Provides self-service access throughout the organization to those data products.
- Governs the platform and products so that they're useful, trustworthy, and interoperable.

Implementing data mesh requires openness to change, strong leadership, and a deep commitment to closely following data mesh tenets. But the payoff is substantial.



# How Marlabs Can Help

At Marlabs, our team is here to help you succeed. As you read about earlier in this paper, we've guided major organizations in their data mesh journeys.

Our approach stays true to data mesh principles, and we offer deep, hands-on experience in data, analytics, cloud, machine learning, and AI. Across industries, we've successfully delivered scalable platforms and business-ready data products. Our team will provide you with the business, technology, and industry acumen to holistically assess whether treating data as a product is the right move for your organization.

If you choose to move forward, our process is simple. We:

**Assess your readiness.** We take time to understand your people, processes, and technology; to identify gaps; and to align data mesh strategy with your business goals and objectives.


**Build an actionable roadmap.** We show you how to move from your current state to data mesh thinking, and we include checkpoints and continual feedback so you can mitigate risk and avoid common challenges.

**Build the foundation.** Whether that means migrating systems to the cloud or designing a platform that meets all the governance, compliance, and observability criteria you need, we help you put all the vital components in place.

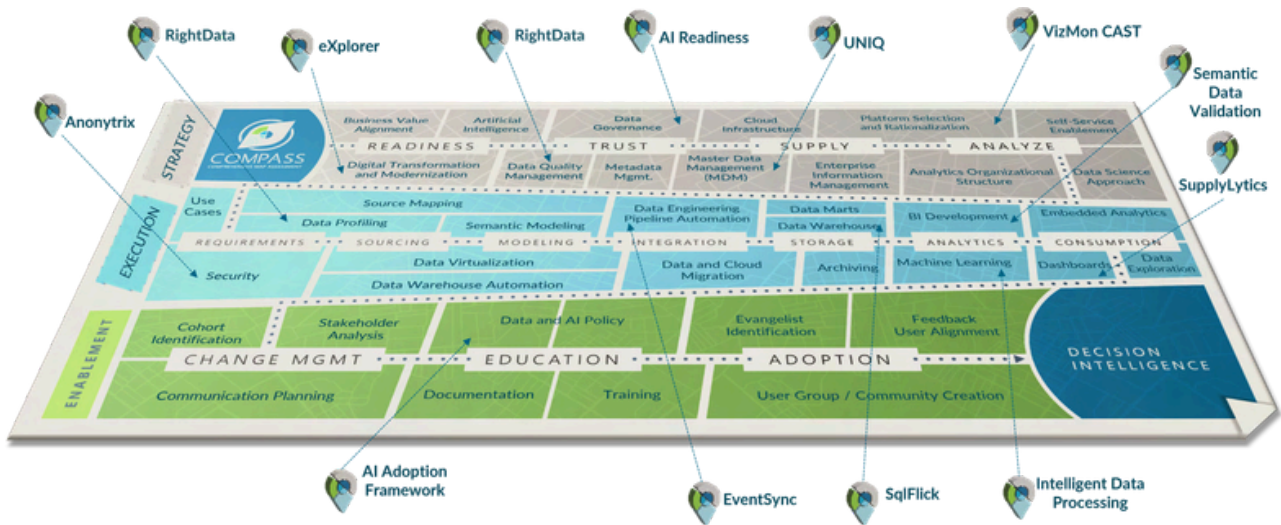
**Empower domain teams.** We help with change management, tools, training, and frameworks to help your teams own and deliver their own data products.

**Deliver quick value.** We'll help you choose pilots to start implementing a data mesh architecture that demonstrate measurable impact and set the stage for the organization to scale the initiative.

**Accelerate adoption.** Our team deeply believes in ensuring enablement throughout our engagements. We implement change management a part of how we help our clients, guiding you with best practices and executive support to speed up adoption.



In short, we'll meet you where you are, work with you in the way that helps you best, and ultimately help you unlock the full value of your data. Of course, it's much more complicated than that. That's why we've built our proprietary MAP methodology, a roadmap of which is shown in image the below. We use this methodology to take clients from **strategy** through **execution** and into **enablement**.



In a world where hype and misinformation sell, we focus on providing honest expertise. It's how we've built client loyalty that spans decades. If you're interested in discussing how a data mesh architecture might transform your company, we're ready to help.

