



Point of View /

Data Mesh: Beyond the Hype



CONTENT

02	Introducing the Data Disruption of Tomorrow
03	Data Mesh - Not (All) About Technology
04	Four Principles of Data Mesh
05	Changes in Status Quo
08	Why Data Mesh?
09	Are You Ready For Mesh?
12	Conclusion
13	Authors

Introducing the Data Disruption of Tomorrow

“An idea is like a virus –, resilient and, highly contagious. The smallest seed of an idea can grow. It can grow to define or destroy you.”
— Christopher Nolan

Disruption too can define or destroy you. As a key part of a business strategy, disruptions have brought about drastic changes in the corporate ecosystem. However, today, it alone may amount to nothing. For modern and successful businesses, speed is critical while scaling innovations and adapting to disruptions.

To achieve sustained success, innovation must be continuous, measurable, and accountable. Traditional R&D setups are insufficient to meet the new demands of innovation. Two metrics that can ascertain its success in an organization are maturity and scale.

Mature innovations entail the use of data and analytics to drive both the discovery and execution of the innovation. And scaling innovations require constant experimentation with data in a decentralized way. This combination would ensure a better quality of ideas while engaging employees across the organization.

The current competitive innovation environment demands organizations to shift toward a decentralized and managed data strategy. Data mesh has emerged as a possible solution. It was conceptualized in 2019, and in just three years, it has taken the data world by storm.

Let's delve deeper and understand Data Mesh and when to use it.

Data Mesh - Not (All) About Technology

While existing data strategies and paradigms delivered many organizations a victory yesterday, the continuously changing playing field will not produce the same results in the near future.

Simply put, to ensure a win in the long term, an organization must:

- ⋮ Disrupt with value, scale, and pace.
- ⋮ Prioritize constant, data-driven, and managed innovation.

Delivering sustained and reproducible victories on the business front calls for evolving our data strategies.

When we look from this angle, data mesh doesn't look like a technological solution. It doesn't look like a ground-breaking overhaul of the status quo either. Rather, it seems like a logical next step to compete in this competitive playing field. A paradigm shift that feels rational and necessary.

If you find yourself struggling to compete, if you feel that your current organizational structure is finding it difficult to cope with the disruptions in the market, data mesh might be the remedy to your problem.

At its core, data mesh is a socio-technical paradigm that enables organizations to extract the maximum value from their data to drive progress with pace at scale. This approach places heavy emphasis on sharing, managing, and accessing analytical data within and across organizations.

When implemented properly, data mesh can enable your organization to:

- ⋮ Stay prepared for disruption
- ⋮ Quickly adapt to change
- ⋮ Innovate with scale and maturity
- ⋮ Deliver at scale with continuous agility

Prior to constructing a data mesh, one must be cognizant of the four interacting principles that enable an organization to achieve the objectives of a data mesh.

Four Principles of Data Mesh

Domain Ownership: Data mesh has an architecture that arranges the analytical data by domain. In this kind of arrangement, the domain's interface to the rest of the organization includes the operational capabilities and access to the analytical data that the domain serves.

⋮ **Data-as-a-Product:** In the Data Data-as-a-Product model, an organization's data is viewed as a product, and the data team's role is to provide that data to the organization in ways that aid good decision-making and application building.

⋮ **Self-Serve Data Platform:** It is a new generation of data platform services that empower cross-functional teams to share data. The platform services concentrate on removing friction from the journey of shared data from source to consumption.

⋮ **Federated Computational Governance:** A data mesh strategy requires a governance model that adopts decentralization, domain ownership and automated execution of decisions by the platform. So, this governance model, which is very different from the traditional governance model, is called the federated computational governance model.

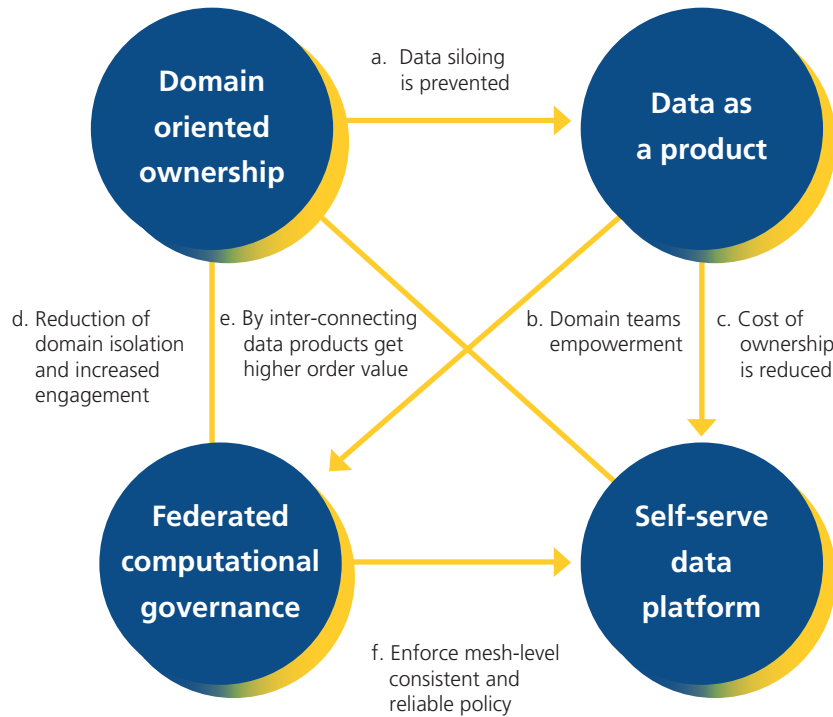



Figure 1: Four principles that describe a data mesh architecture


Changes in Status Quo

Implementing the data mesh transforms the structure of data management within an organization. Any company that intends to implement the mesh architecture must be aware of these structural changes and commit to them:

- **From an organizational standpoint**, data mesh calls for ownership and accountability of data to be placed with business domains in place of ownership and accountability being taken care of by specialized data platform teams.



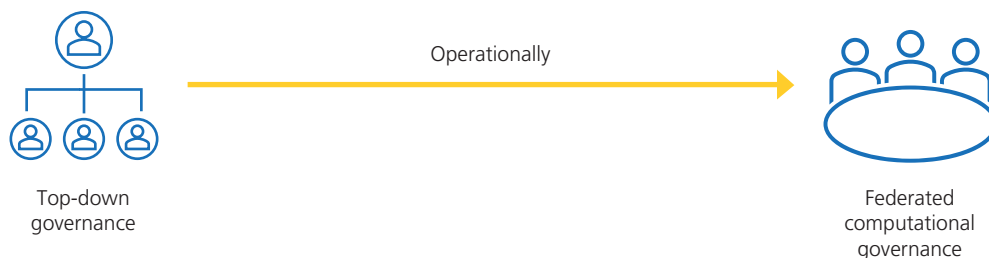
⋮ **From an architectural standpoint,** data mesh calls for data to be distributed across the mesh as data products as opposed to hoarding of data within monolithic warehouses and lakes.



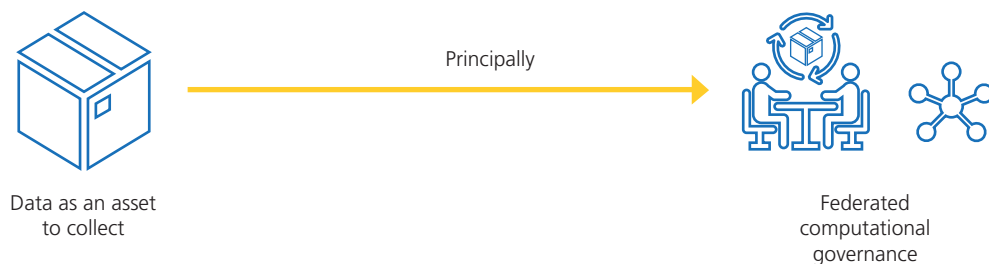
⋮ **From a technological standpoint,** data mesh necessitates viewing data and code as two equal parts that contribute to/is obtained from a solution as opposed to viewing data as a by-product of code.



⋮ **From an operational standpoint,** it calls for data governance to be federated and automated through computational policies as opposed to a centralized top-down approach to governance.



⋮ **From a principle standpoint,** data is no longer viewed as an asset but as a product.



⋮ **From an Infrastructure standpoint,** a fragmented platform comprising of analytical and operational components is replaced by a well-integrated platform.



Why Data Mesh?

“In business, what’s dangerous is not to evolve”
— Jeff Bezos

Naturally, these changes are not easy to implement, and one must commit to them to obtain favorable results. This may seem daunting, but one must keep in mind that the benefits far outweigh the cost of the initial setup. We are at a point in time where it is critical to evolve our understanding of data to maximize its value. Persisting with the existing system will eventually lead to a dead end. Here are some of those lanes you would do well to avoid, lest you find yourself staring at a brick wall.

Sr No.	Current Lane	Eventual Dead End	Bypass With Mesh
1.	Centralizing data into one platform	As organizations scale into larger entities, it becomes impossible to prevent data silos from forming.	Data is viewed as a product and the ownership and management of data are handled by business domains.
2.	Centralized solving of data requests	Disconnected teams will not understand the needs of a business holistically. This time-consuming activity will severely compromise an organization’s ability to innovate with agility.	Data products and their quality is maintained by data owners who are incentivized to ensure quality. This is in turn visible to the rest of the organization and the data is ready to use.

Sr No.	Current Lane	Eventual Dead End	Bypass With Mesh
3.	Data experts in a centralized platform become entrenched in a specific domain.	Will inevitably create platform-related bottlenecks.	Data mesh incentivizes all stakeholders to understand and appreciate other parts of the pipeline by building cross-functional domain teams around the data product.

Table 1: Potential dead ends that current system will lead to and how mesh avoids these issues.

Are You Ready For Mesh?

“Readiness is an oft mentioned, but frequently forgotten necessity”
— Erik Prince

Building a mesh even before assessing your organization’s preparedness to commit to one will inevitably result in being caught in the web. Thus, conducting a **Mesh Readiness Assessment** before taking the leap will help organizations to gauge their readiness.

A good data mesh readiness assessment platform should cover the following dimensions:

Organization Complexity, Data Landscape Exploration, Data-Driven Mission, Executive Support, Biz-Tech Collaboration, Initial Adopter, Contemporary Engineering and Commitment.

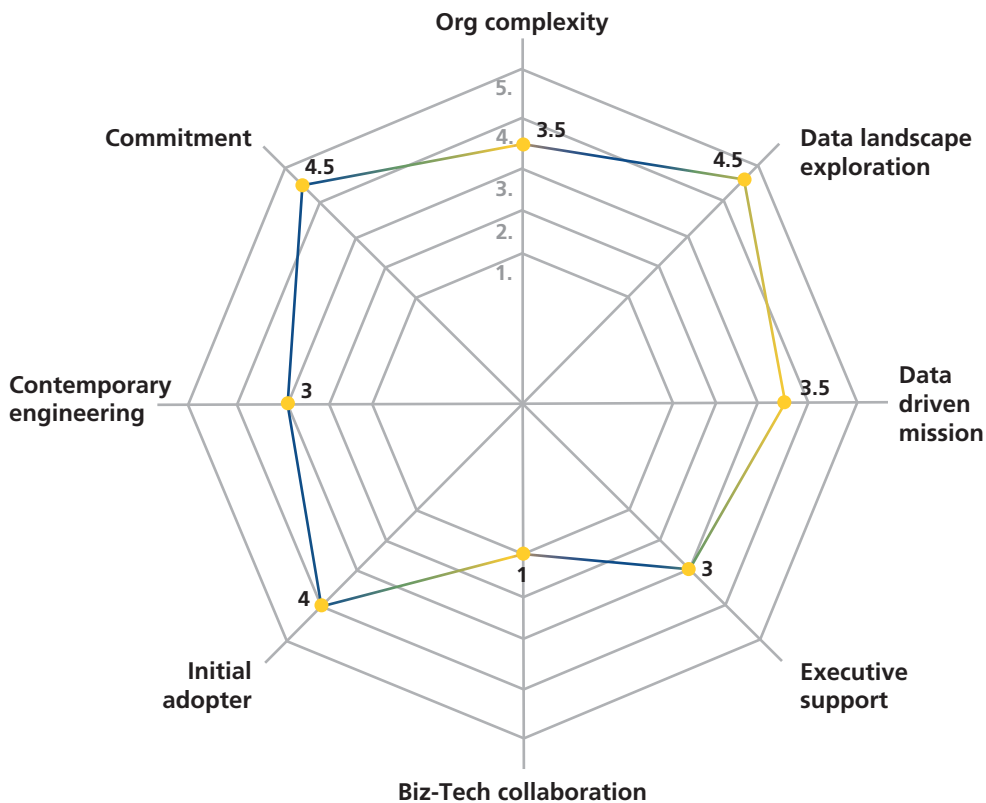


Figure 3: Sample Data Mesh Readiness Assessment Result:
Organizations scoring medium or high on each of these dimensions
will be ready to adopt data mesh successfully.

⋮ **Organization Complexity:** This dimension measures the primary criteria for adopting data mesh. These are:

- Intrinsic business complexity
- Proliferation of data sources and data use cases.

Thus, organizations with proliferating data sources and use cases that get a moderate to a high score on data complexity will benefit from a data mesh.

- ⋮ **Data Landscape Exploration:** This dimension measures the nature of the data landscape. To adopt data mesh and make full use of it, it takes more than just great technology and quality data. Understanding the scale of data growth, cloud factor, and resource availability becomes very crucial.
- ⋮ **Data-Driven Mission:** This dimension measures the commitment of product teams and business units of the organization in using intelligent decision-making and actions in their applications and services. This kind of commitment is possible only when the organization identifies data-enabling ML and analytics as a business strategic differentiator.
- ⋮ **Executive Support:** This dimension measures the organization's "resistance to change". An organization that is ready to embrace data mesh should be open to change. Data mesh demands the organization to change how the people work. Thus, this is where executive support and top-down engagement of leaders come into the picture.
- ⋮ **Biz-Tech Collaboration:** This dimension measures the relationship between the business and technology. For example, organizations that consider technology as a supporter of the business and not at the core often externalize technical capabilities to external vendors and will be willing to buy and plug in a ready-made solution for their business needs. Such organizations are not ready to adopt data mesh.

- ❖ **Initial Adopter:** This dimension measures the organization's risk-taking nature, and their willingness to adapt to new technologies because data mesh demands the spirit of experimentation, taking risks, learning, and evolving. On the contrary, organizations that are not willing to take risks and like to only adopt well-tested, refined paradigms may need to wait for some time.
- ❖ **Contemporary Engineering:** This dimension measures the alignment of the organization with modern software engineering practices because data mesh works with modern technology stacks that are API-driven, easy to integrate, and require smaller teams, as opposed to one centralized team.
- ❖ **Commitment:** This dimension measures an organization's commitment towards data mesh principles and the adoption of data mesh transformation journey. A successful data mesh implementation necessitates a high level of organizational commitment toward creating distributed leadership and, revamping the working structure, among others.

Conclusion

"A journey of a thousand miles begins with a single step"

— Lao Tzu

Despite all the hype around data mesh, we must concede that it is still nascent. While organizations have adopted Data Mesh at the proof-of-concept (PoC) stage, few have deployed it in their production environment. Without a doubt, it is a promising idea that could potentially solve various business problems and drive success. We are still learning, however, and as we learn, the concept of data mesh grows and becomes more mature.

Authors



Kavya Ganesh

Data Engineer, Snowflake Center of Excellence, LTIMindtree

Kavya strongly believes in the concept that learning is a continuous process and that it does not have any end. She is an avid learner and with digital world's rapid pace, she feels there is always something new to learn! Kavya loves to spend her free time in learning new skills or watching travel logs. She loves travelling and spending time in nature. She is a trained Carnatic classical singer who has successfully cleared junior level and is practicing senior level. When she isn't working, she can be found listening to her favourite songs. Dancing is also one of her passions. She believes yoga followed by some good dance moves make a good fitness-mantra! She loves to explore new places and different cuisines around the world. Kavya is currently working as a Data Engineer at LTIMindtree's Snowflake Center of Excellence where she is currently involved in building accelerator tools related to data migration and utilities for governance features of a cloud data platform.



Srinivasaraghavan Sundar

Senior Data Engineer, Snowflake Data Center of Excellence, LTIMindtree

A jack of all trades, striving to become the master of some, Srinivas usually spends his time either picking up new skills or honing and deep diving into his areas of interest. He is currently a part of LTIMindtree's Snowflake Data COE in the capacity of a Senior Data Engineer with a real penchant for Data Science. When he isn't working, he can be found reading books, learning (or at least trying to) new languages and practicing the keyboard.

Srinivas has a fiery passion for theatre; aside from presiding over his college's theatre club, he was part of Crea-Shakthi's Malgudi Days troupe which toured South India in 2018.

Today, he channels his creativity into finding unique solutions for technical problems, and, well, writing blogs like the one above.

LTIMindtree

LTIMindtree is a global technology consulting and digital solutions company that enables enterprises across industries to reimagine business models, accelerate innovation, and maximize growth by harnessing digital technologies. As a digital transformation partner to more than 700 clients, LTIMindtree brings extensive domain and technology expertise to help drive superior competitive differentiation, customer experiences, and business outcomes in a converging world. Powered by 81,000+ talented and entrepreneurial professionals across more than 30 countries, LTIMindtree — a Larsen & Toubro Group company — combines the industry-acclaimed strengths of erstwhile Larsen and Toubro Infotech and Mindtree in solving the most complex business challenges and delivering transformation at scale. For more information, please visit www.ltimindtree.com.