Reimagining manufacturing: From centralised platforms to domain-driven intelligence

Contents

01 Executive summary

Manufacturing industry is under immense pressure to improve productivity, reduce downtime, ensure quality, and respond faster to market demand due to increasing global competition, rising operational costs and fast evolving technology landscape. While data exists across plants, machines, and enterprise systems, most manufacturers struggle to convert the vast pool of data into timely and trusted insights. Centralised data platforms, though well-intentioned, often become bottlenecks, slowing innovation and distancing data teams from operational realities.

Implementing a data mesh architecture on cloud platforms enables manufacturing organisations to convert siloed OT and IT data into discoverable, governed data products owned by domain teams, delivering faster time to insight, higher data quality, and repeatable business outcomes.

By combining key principles such as domain ownership, self-serve platform capabilities, and federated governance, manufacturers can operationalise use cases such as predictive maintenance, quality control, and supply chain visibility.

02 Key challenges

Many manufacturers continue to rely on manual data entry processes which not only impact operational efficiency but also increase the operational costs due to rework and human errors. The rise of AI, particularly GenAI, emphasises the necessity of high-quality data, as manufacturers realise that successful AI deployments rely on data readiness. Edge computing is also gaining traction among manufacturers as it addresses challenges related to processing large volumes of unstructured data efficiently and supports applications such as predictive maintenance and production quality control.

The manufacturing industry is also pivoting towards customer-centric, tech-enabled platforms, moving away from traditional models to embrace Industry 4.0 technologies like IoT, AI, and predictive analytics for efficiency and ESG-linked value creation. At the same time sustainability and regulatory requirements are also influencing business strategies, demanding new skills and governance models to comply with ESG expectations.

All these challenges lead to adoption of modern data architectures by the manufacturers enabling them to integrate diverse data sources, domain-oriented ownership, real-time analytics, and consistent governance. Data mesh supports distributed, domain-specific data consumers and views data-as-a-product with each domain handling their own data pipelines resulting in accelerated speed to market, improved data quality, reduced data and analytics delivery timelines and stronger alignment between business needs and data ownership.

03 Modern data architecture for the manufacturing industry

Modern data architecture in the manufacturing industry focuses on unifying IT and OT data to enable real time visibility, predictive insights, and cross plant optimisation. It typically combines a lakehouse for scalable analytics and ML, data fabric for governance across diverse OT/IT systems, and data mesh principles for domain ownership across plants, quality, operations, supply chain, and maintenance. The goal is to break data silos, standardise KPIs (e.g. OEE, FPY, OTIF), and empower frontline teams with faster, more reliable data for improved productivity, quality, and decision making. Table 1 describes various kinds of architecture and their uses cases.

Table 1: Overview of modern data architectures and their key use cases

	Description	Use case
Data warehouse	Centralised system that stores structured, processed, and cleaned data from various sources for reporting, analytics, and business intelligence. It follows a schema-on-write approach and is optimised for fast queries.	Suitable for highly structured, centralised reporting where metric control and consistency are more important than flexibility or scale of raw data.
Data lake	Storage repository that holds raw, unprocessed data in any format—structured, semi structured, or unstructured. It uses schema on read, making it flexible for big data, ML, and advanced analytics.	Used for large-scale raw storage and diverse data storage for exploratory analytics and advanced processing but requires additional governance layers for enterprise reporting.
Lakehouse	Combines the strengths of data lakes and data warehouses. It stores raw data like a lake but also provides data management, governance, ACID transactions, and performance features similar to a warehouse. It supports both BI and ML workloads.	Suitable when enterprise wants a unified platform for BI, ML and large-scale data processing with central engineering control.
Data mesh	Organisational and architectural approach where data is treated as a product, and ownership is distributed across domain teams. Instead of a central data team, each domain manages its own pipelines, quality, and governance—promoting scalability and self service analytics.	Suitable when multiple business domains produce and consume data and centralised teams cannot scale, requiring domain-owned data products with federated governance.
Data fabric	Architecture that provides a unified and intelligent data management layer across all data sources—on premises, cloud, hybrid, and multi cloud. It uses automation (AI/ML) to seamlessly integrate, govern, and deliver data wherever it's needed. The goal is to make data accessible, consistent, and reusable across the entire organisation.	Suitable when data is distributed across many systems and the organisation needs metadata-driven integration, lineage, policy, enforcement, and unified access without heavy data movements.

Source: https://www.pwc.com/gx/en/issues/technology/tech-translated-data-mesh-data-fabric.html

From silos to smart systems: Data mesh for manufacturing excellence

Data mesh is a sociotechnical model that decentralises data ownership to domain teams, treats datasets as discoverable, supported products, and uses a self-serve platform and an automated, federated governance and a common interoperability layer (metadata, contracts, APIs) to enable trusted, scalable data sharing and analytics across the enterprise. Key principles of data mesh are:

- **Domain ownership:** Each domain team is responsible for its own data, aligning data ownership with team boundaries.
- **Data as a product:** Data is treated as a product where domain teams must provide high-quality data that meets the needs of other domains.
- **Self-serve data infrastructure:** A dedicated platform enables domain teams to create and manage interoperable data products independently.
- **Federated governance:** A governance group ensures standardisation and compliance across the data ecosystem, promoting interoperability.

Data mesh in manufacturing

Data mesh is reshaping how manufacturing teams work with data by moving responsibility away from a single central IT group and placing it directly within operational areas like production, quality, and logistics. Instead of relying on a large, centralised data lake, each domain manages its own data and treats it as a valuable product. This reduces delays and gives engineers and plant leaders immediate access to the IoT and machine data they need. With data managed closer to where it's generated, manufacturers can roll out AI-driven initiatives such as predictive maintenance or supply chain improvements much faster. Some of its use cases are listed in Table 2.

Table 2: Decision matrix for evaluating data mesh against typical use cases

	Cross domain High – Multiple domain involved Low – Single domain	**Require domain expertise** High – Strong domain knowledge required Low – Data structures are standardised and can be modeled centrally	**Require domain ownership of data** High- Domain ownership of pipelines Low - Central ownership of data pipelines	**Multi-plant/ multi site** High - Enterprise scale Low - Local optimisation	**Real-time / IoT heavy** High -Streaming Low - Batch	**Data scale (volume, velocity)** High- Large and high frequency data Low – Aggregated transaction data	**Enterprise KPI standardisation** High – Across organization Low – Local operational measures	**Data mesh to be considered**
Supplier comparison/ procurement analytics	High	High	High	High	Low (Mostly Batch)	Medium	High (Consistent Metrics)	**Yes** (highly cross-domain, requires strong business ownership, and depends on unified KPIs)
Inventory visibility across locations	High	Medium	High	High	Medium (near real time inv updates)	Medium	High (Consistent Analytics across)	**Yes** (multi-site coordination and cross-domain ownership)
Logistics load optimisation	High	High	Medium (Data from multiple systems but operational)	High	Medium (near-real-time data)	Medium	Medium (Algorithm-driven)	**Can be considered** (multiple domain-owned inputs but optimisation logic is centrally owned logistics domain)
Real-time production monitoring	Medium	High	High	High	High	High	Low (Measured IOT data)	**Lakehouse** better option as this needs real-time IoT/SCADA streaming, high-volume time-series storage, and ML-based fault detection
Improve returns processing and customer support	High	High	High	Medium	Medium (near real time)	Medium	High	**Yes** (require strong cross-domain ownership, consistent KPI definitions to ensure trust and accountability)

Source: https://www.pwc.in/industries/financial-services/fintech/point-of-view/financial-services-data-and-analytics/may-2022.html

How data mesh accelerates AI use cases in manufacturing

Traditional centralised architectures make it difficult to maintain data readiness, often resulting in monolithic data lakes, semantic layers bolted on top, and analytical datasets stripped of the rich business context needed for intelligent systems to succeed. These patterns create context loss, information dilution, and silent drift—all of which directly impact model accuracy.

Data mesh architecture addresses these challenges by decentralising data ownership into domain teams (production, quality, supply chain, maintenance), who curate data as interoperable, governed products. These domain-specific data products serve as high-quality training and inference inputs for AI workloads.

Table 3: Key AI use cases in the manufacturing industry

Material selection and management	Identify optimal sourcing opportunities based on different parameters and recommend sourcing strategies that minimise costs while meeting all other requirements.
Production planning	Improve production planning and execution by continuously analysing sales demand and selecting optimal production volumes, configuration and sequencing to drive higher efficiency and/or lower costs.
Predictive maintenance	Carry out analysis and augmentation of high volumes of sensor data, prioritisation of maintenance activities and workplan creation to fully optimise allocation of activities within the businesses.
Enhanced pricing	Process large pools of historical and competitor benchmarking data, accounting for demand/capacity and forecasting trends to suggest and optimise pricing. Customer profiling towards pricing can also be done to increase engagement and conversion rates.
Technical support	A training platform integrated with GenAI chatbot can provide on-the-job support to workers to augment and personalise training on SOPsA10 and RIMA11/ERPA12 platforms, thus ensuring compliance and quality assurance in supply chain functions.

Source: The Impact of Generative AI in Industrial Manufacturing 2023, PwC

04 Operationalising data mesh in manufacturing: Architecture and implementation roadmap

The following reference architecture shows how a manufacturing data mesh enables domain-owned data products, federated governance, and self-service consumption across manufacturing domains.

Figure 1: Data mesh reference architecture

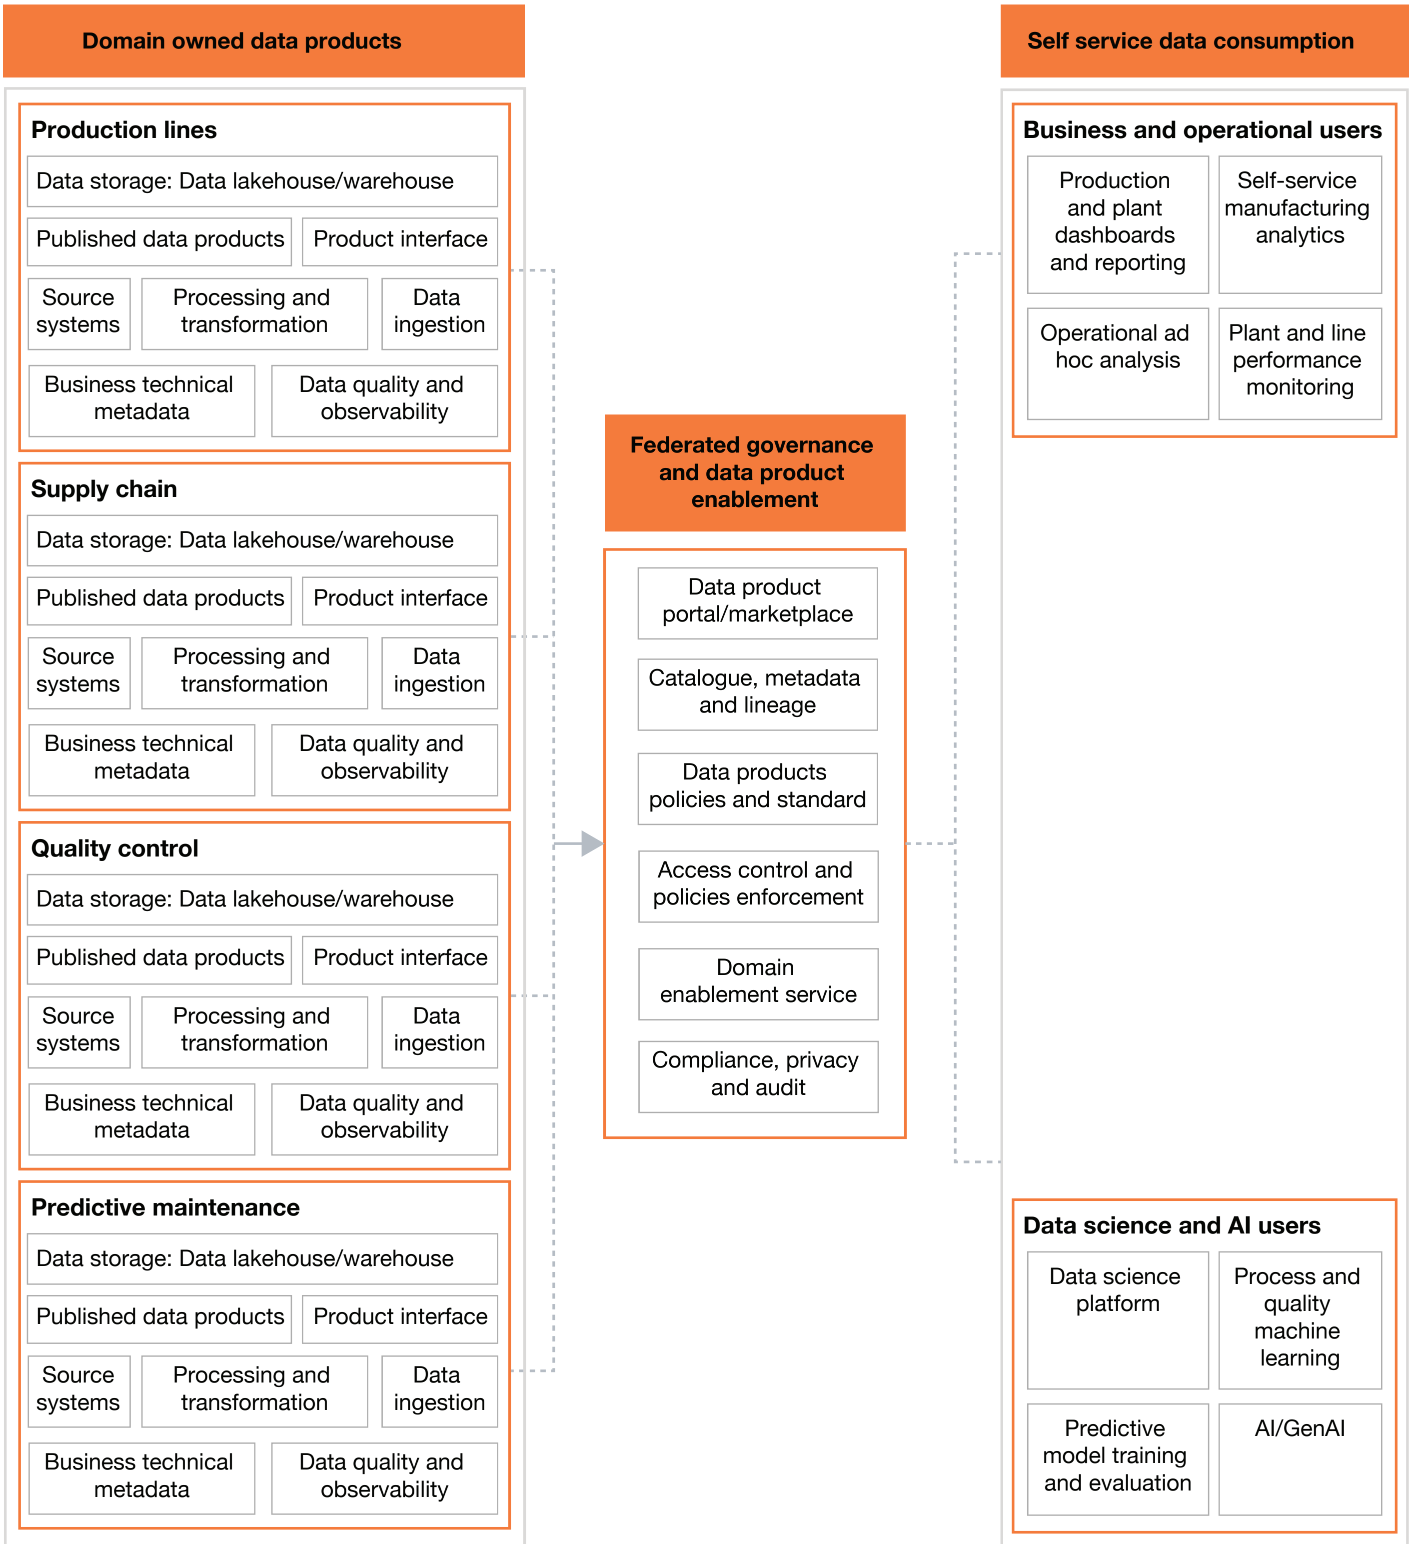

Layers considered:

- **Domain owned data products:** Organise around manufacturing domains, i.e. production lines, supply chain, quality control and predictive maintenance (non-exhaustive), where each domain owns data end to end.
- **Federated governance and data product enablement:** Provides centralised governance guardrails while enabling domain autonomy through a federated operating model. This enables data product discovery, business metadata management and subscription workflows, along with capturing business context, ownership and usage intent.
- **Self-service data consumption:** Enables consumers to discover, request and consume trusted data products in a governed and self-service manner. Data mesh adoption is a phased transformation journey which includes operating models, governance, ownership and ways of working. Manufacturers should adopt an incremental approach rather than a large-scale transformation. The following phased approach provides a roadmap for manufacturers to implement data mesh within their organisation:

Figure 2: Data mesh implementation phased approach

1	2	3	4
Leadership and business sponsorship alignment	**Establish manufacturing data foundation**	**Industrialise and scale**	**Democratise and virtualise**
• Successful adoption of data mesh requires executive sponsorship across functions such as supply chain, quality, engineering, and IT with clear alignment on business outcomes. • Prioritise high-value use cases, identify required investments, and define a phased transformation roadmap to guide adoption and execution.	• Define business-aligned data domains and introduce federated ownership across manufacturing functions. • Deliver an initial set of high-value use cases through cross-functional teams, while the central platform establishes shared capabilities such as cataloging, metadata, and governance guardrails.	• Move from pilot phase to repeated delivery while DataOps, automation, monitoring, and a clear operating model across domain, platform, and governance teams. • Scale adoption by assessing product usage, prioritising high-value assets, and improving metadata, traceability, and interoperability across plants and enterprise.	• Enable broader adoption through a data product marketplace that allows users to discover and access trusted data products. • Expand coverage across plants and functions over time, while using virtualisation and distributed access patterns to reduce unnecessary data movement.

05 Conclusion

In an environment where operational continuity, quality, and compliance are non negotiable, data mesh enables faster and safer decision making across the enterprise. Plant teams gain ownership of trusted operational telemetry, quality and analytics teams consume SLA backed data products with confidence, and legal and compliance functions rely on policy driven controls rather than labour- intensive reviews.

When implemented on a modern cloud or hybrid data platform, data mesh can further accelerate time to insight, reduce mean time to resolution for data incidents, and make regulatory reporting and forensic analysis viable at scale. In an era of smart manufacturing, data mesh could become not just a data architecture, but a strategic foundation for resilient, data driven operations.

06 Abbreviations

KPI	Key performance indicator
OT	Operational technology
IT	Information technology
OEE	Overall equipment effectiveness
FPY	First pass yield
OTIF	On time in full
ACID	Atomicity, consistency, isolation, durability (database transaction properties)
BI	Business intelligence
ML	Machine learning
SOP	Standard operating procedures
RIM	Resource information management
ERP	Enterprise resource planning
IoT	Internet of Things
SCADA	Supervisory control and data acquisition

About PwC

We help you build trust so you can boldly reinvent

At PwC, we help clients build trust and reinvent so they can turn complexity into competitive advantage. We're a tech-forward, people-empowered network with more than 364,000 people in 136 countries and 137 territories. Across audit and assurance, tax and legal, deals and consulting, we help clients build, accelerate, and sustain momentum. Find out more at www.pwc.com.

PwC refers to the PwC network and/or one or more of its member firms, each of which is a separate legal entity. Please see www.pwc.com/structure for further details.

Contact us

Prakash Suman
Managing Director, One Consulting
PwC India
prakash.suman@pwc.com

Somoj Chowdhury
Director, One Consulting
PwC India
somoj.chowdhury@pwc.com

Author

Sushant Prabhakar

Editor

Rubina Malhotra

Design

Kirtika Saxena

pwc.in

KS/July2026 - M&C 53854