LTIMindtree

iNXT
From LTIMindtree

Whitepaper

# Connected Data Smarter Decisions

## Industrial Data Fabric

# Contents

# Executive Summary

**The volume of data is growing rapidly. Customer preferences are evolving just as fast. Many organizations still struggle with disconnected operational and IT systems. These gaps create challenges in sustaining digital transformation. Without a unified operational view, AI-driven initiatives become harder to implement. Industrial Data Fabric now plays a vital role. It helps manufacturing ecosystems accelerate innovation and maintain long-term digital momentum.**

Industrial Data Fabric (IDF) solutions serve as the cornerstone for digital transformation in manufacturing and industrial sectors. By harmonizing diverse data types, ranging from time-series machine data to maintenance logs from Enterprise Asset Management (EAM), production orders from Manufacturing Execution Systems (MES), and inputs from ERP, scheduling, and material management systems, IDF enables a unified, contextualized view of operations. This alignment across the time domain and semantic layers reveals critical relationships between data types and their content, unlocking actionable insights for business process optimization.

At its core, IDF is not merely a data integration framework; it is a strategic enabler that bridges operational technology (OT) and information technology (IT), forming a unified namespace and governance model. Through contextualization, transformation, and activation of data, IDF supports advanced capabilities such as Digital Twins, Industrial AI, and intelligent decision systems.

IDF acts as a catalyst for innovation—streamlining processes, optimizing manufacturing operations, and enhancing productivity, availability, and quality. It supports ROI realization, unlocks new revenue opportunities, and ensures robust data security across business functions and data pipelines.

IDF empowers key industries with target capabilities:

- Manufacturing: Enables predictive maintenance, production optimization, quality control and assurance, and supply chain efficiency.
- Energy and Utilities: Supports energy demand forecasting, predictive maintenance, and energy optimization.
- Oil & Gas: Drives production optimization and enhances operational efficiency in drilling and exploration.

# Introduction

Industrial Data Fabric is a new and evolving database management design that leads to the creation of strategic applications with cloud and data technologies, purpose-built for industries such as manufacturing, energy, and utilities. It enables organizations to harness real-time operational data from assets—such as industrial equipment, sensors, and control systems—to unlock deeper insights and drive efficiency far beyond the capabilities of traditional IT infrastructure.

Driven by emerging industry trends, manufacturing enterprises are increasingly prioritizing investments in data and cloud platforms to unlock greater business value. This shift reflects a strategic response to the growing need for scalable, secure, and intelligent infrastructure that supports digital transformation, operational efficiency, and innovation on a larger scale.

**By 2027 – 2032**

- **17.5%** CAGR increase estimated for Industrial Cloud Platform Market by 2032
- **$3 trillion** McKinsey estimates that cloud adoption could generate in global value by 2030
- **$200 billion** Cloud Manufacturing Market size is estimated to reach, growing at a CAGR of 16.1% by 2030
- **70%** of enterprises will use industry cloud platforms to accelerate their business initiatives by 2027

**Key Segments contributing to this opportunities..**

- Increased adoption of IoT devices and sensors
- Digital Transformation
- Asset Performance Management
- Real Time Data

- IaaS – Infra Optimization
- PaaS – Digital Transformation
- SaaS – Applications across Processes
- Accelerate Digital Transformation

Figure 1: Industry metrics [1] [2] [3] [4]

**Key capabilities include:**

- Real-time data collection from OT systems
- Data standardization and transformation for consistency and usability
- Unified namespace formation to organize data across devices and environments
- Advanced analytics and visualization for operational intelligence
- AI enablement through high-quality, governed data pipelines

By leveraging these capabilities, organizations can enhance operational performance, accelerate AI adoption, and make informed, data-driven decisions that optimize both strategy and execution.

The true promise of Industry Data Fabric lies in its ability to eliminate manual efforts in data integration, standardization, and contextualization once implemented. While some vendors offer end-to-end Industry Data Fabric solutions, it's important to recognize that a composable architecture—built from modular technology components—is essential to adapt to evolving use cases and operational demands. This approach ensures flexibility, scalability, and long-term relevance as business needs and data sources continue to grow and diversify.

Artificial Intelligence (AI) in manufacturing must meet the stringent standards of industrial-grade environments—capable of processing vast volumes of machine data and identifying complex operational patterns. Industrial-grade AI empowers organizations to scale digital transformation, achieve operational goals, and evolve into resilient, sustainable digital enterprises.

Industrial AI, when combined with digital twins, enables the creation of intelligent virtual models that analyse data, predict outcomes, and automate processes for real-time optimization, predictive maintenance, and quality assurance. This synergy enhances operational efficiency by streamlining workflows, reducing waste and energy consumption, and accelerating product development through virtual testing and simulation—ensuring changes are validated before physical implementation.

# Challenges

The list of industry challenges includes data silos, overwhelming data volumes across hybrid environments, poor data quality and governance, the complexity of integrating diverse data sources, and the need for real-time insights for better decision-making.

To fully leverage data across IT, OT, and enterprise domains, it is essential to understand and address the inherent challenges that hinder seamless integration and value realization.

**Key Data Challenges in Industrial Enterprises**

**Fragmented Data Silos**
Data is often isolated across departments and platforms, limiting visibility and making it difficult to access consistent, enterprise-wide information for decision-making.

**Complex Hybrid Environments**
Organizations operate across on-premises systems, multiple cloud platforms, and edge devices—creating challenges in managing diverse data formats and access protocols.

**Exponential Data Growth**
The surge in structured, semi-structured, and unstructured data overwhelms traditional data management systems, resulting in vast amounts of underutilized “dark data.”

**Lack of Unified Data Integration**
Extracting and harmonizing data from IoT sensors, cloud services, and legacy systems is complex, hindering the creation of a single, comprehensive view of enterprise data assets.

**Data Quality, Governance, and Security**
Ensuring high-quality data, enforcing governance policies, and maintaining security across distributed environments is difficult, leading to compliance risks and operational errors.

**Demand for Real-Time Insights**
In fast-paced industrial settings, businesses require immediate, actionable insights to optimize operations and maintain a competitive advantage.

**Operational Overhead and Cost Burden**
Traditional data architectures, such as warehouses and data lakes, involve heavy data lifting, which increases development time and costs, and perpetuates siloed systems.

**Need for Self-Service & Agility**
There’s growing pressure to deliver faster ROI and empower business users with intuitive, self-service access to data. Thereby, reducing reliance on IT teams.

# Enterprise-Wide Challenges

Most organizations lack a comprehensive data management framework that can effectively collect, connect, integrate, and deliver distributed data from diverse OT systems and applications. The shortage of skilled personnel capable of managing and extracting value from this data further compounds the challenge.

Moreover, the growing shift toward cloud-based data platforms introduces additional complexity. In many cases, OT systems operate in environments with limited or unreliable network connectivity, making real-time cloud integration difficult.

# Brief Solution Outline

Industry Data Fabric is a foundational enabler for digital transformation in the manufacturing, energy, and utilities sectors. It empowers these industries to modernize product operations by integrating automation, IoT, AI, and robotics in a more human-centric and sustainable manner. This ensures continuous monitoring, improved operational efficiency, and consistent product quality—key principles of Industry 5.0.

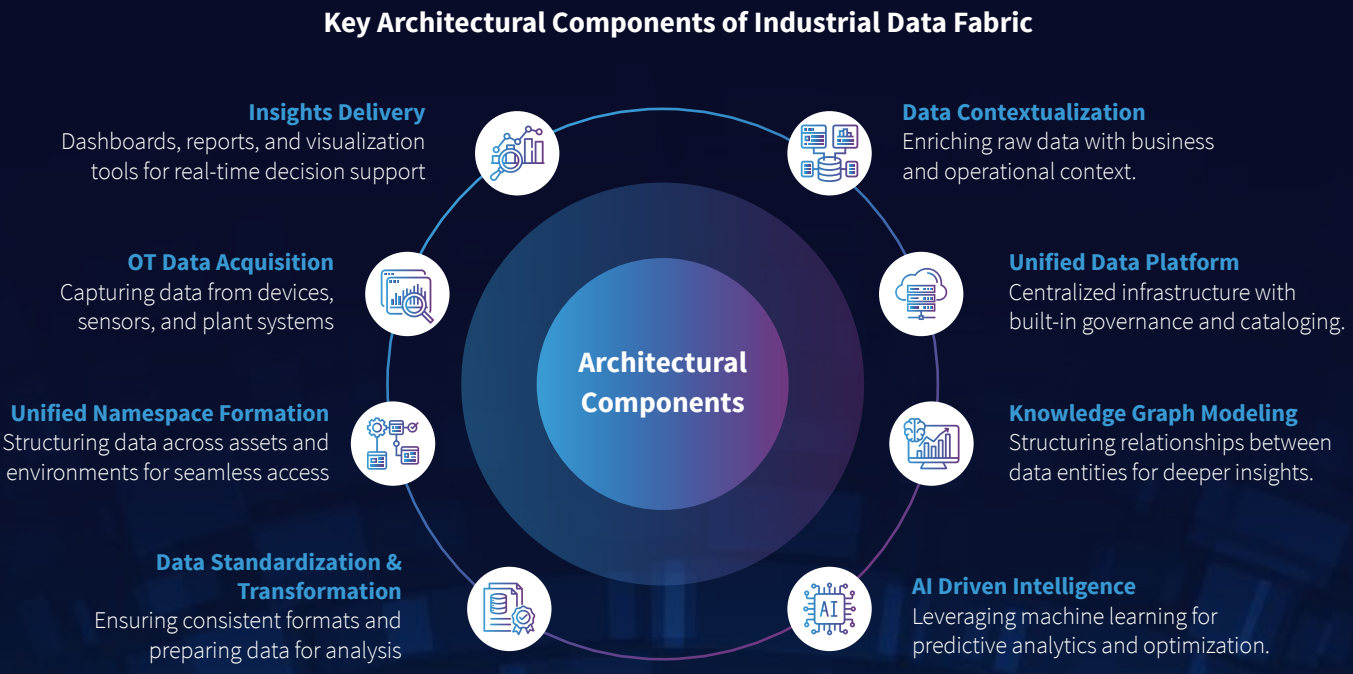


Figure 2: Key Architectural Components of Industrial Data Fabric

## Key Design Considerations for a Future-Ready Solution

- **Seamless IT-OT Integration**
  Addressing long-standing integration challenges between information and operational technologies.
- **Unified Namespace Architecture**
  Establishing a single source of truth for industrial data across devices, systems, and sites.
- **Risk Visibility**
  Enabling identification and mitigation of upstream and downstream operational risks.
- **Connected Industrial Ecosystem**
  Providing a holistic view of operations by linking data across the entire industrial value chain.
- **Cost-Efficient Performance**
  Balancing cost control with increased production throughput and quality assurance.

## What Industry Data Fabric Enables.

An Industry Data Fabric acts as an interconnected network of data, applications, and services, designed to scale with evolving industrial use cases. It:

- Contextualizes siloed data to create a unified, intelligent data layer.
- Connects ecosystems—linking manufacturers with customers, suppliers, and partners.
- Digitalizes supply chains and production processes, overcoming the rigidity of legacy systems.
- Consolidates, analyzes, and democratizes data from the factory floor to the end of the supply chain.
- Provides a command center view of operations, enhancing visibility and decision-making across the enterprise.

# Detailed Solution

Realizing an Industry Data Fabric requires building a secure, scalable, and flexible interconnected platform that enables centralized data governance and insightful data sharing across industries. The solution architecture includes the following components:

### Industrial Edge

**Edge Engineering**

- Establishes a seamless network connecting OT and IT systems
- Integrate sensors and actuators for real-time data capture

**Industrial Agents**

- Context-aware and goal-oriented agents that support autonomous decision-making

### OT Layer

**MES, SCADA & Historian Systems**

- Serve as the memory center for production operations

**Industrial Standardization**

- Implements Unified Namespace (UNS) and standardizes industrial data models

### IT Layer

**Digital Twin & AI Integration**

- Enables enhanced scenario simulations and closed-loop optimization

**Cloud Data Platforms**

- Utilizes data Lakehouse and warehouse architecture for industrial data management

**Unified Command Center View**

- Provides a single-pane-of-glass interface for proactive decision-making across delivery, cost, performance, and maintenance metrics

Let's look at conceptual architecture for Industry Data Fabric:

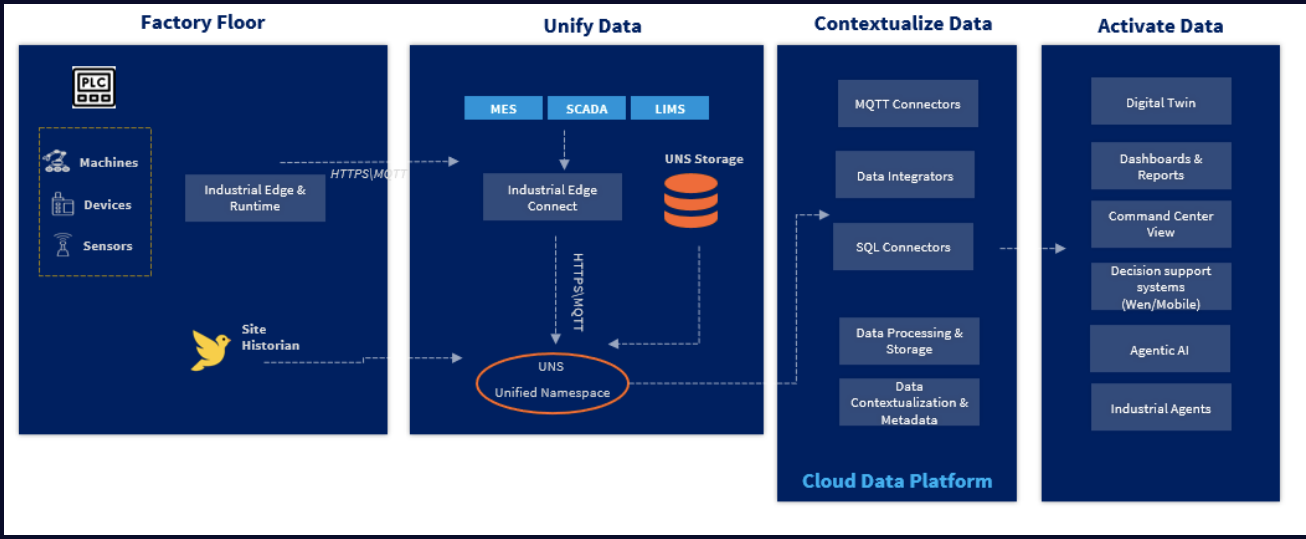


Figure 3: Conceptual Architecture for Industry Data Fabric

### Industrial Edge

Edge Engineering: Establishes a robust network connecting OT and IT systems, integrating real-time data from PLCs, machines, devices, and sensors, which is captured in site-specific historian systems across multiple enterprise locations.

Industrial Agents: Context-aware and goal-oriented agents that support autonomous decision-making through runtime AI model execution.

### Data Integration via MQTT

The Industrial Edge connects using the MQTT protocol, aggregating data from MES, SCADA, and LIMS systems alongside historian and sensor data.

### Unified Namespace (UNS) at the OT Layer

All industrial data originating from machines, sensors, devices, historian systems, MES, SCADA, LIMS, and enterprise sources is unified using the Unified Namespace concept and securely stored in UNS Storage at the OT layer. These data at the OT layer are integrated with the IT layer, primarily cloud, with different data integrators, MQTT connectors, and SQL connectors. The data is transformed, standardized, and processed. It is then contextualized with master data and metadata and stored in a centralized Data warehouse or Data Lakehouse.

The contextualized data is then activated with consumers. A digital twin is implemented using near-real-time cloud services, providing a real-time view of product operations. Dashboards and reports are created with this contextualized data from Data warehouses. A command centre view is also created, providing a single pane of view of complete operations, enabling data-driven decision-making. Along with decision support systems, there are systems created using web and mobile technologies that help field engineers and managers make informed decisions with intelligent industrial data.

By leveraging both real-time and historical industrial data, Agentic AI and autonomous industry agents are developed to harness the full potential of industrial insights and execute predictive actions.

Throughout this lifecycle, the formation and standardization of a Unified Namespace (UNS) is critical to establishing a single source of truth. The UNS serves as the golden record for Industrial IoT data, structured according to ISA-95 model standards and enriched through contextualization with master data and metadata—typically stored in a centralized data warehouse. This architecture enables a hierarchical and comprehensive view of the Industrial IoT ecosystem, supporting scalable and intelligent operations.

## 5.1. Implementation Steps

The implementation of an Industry Data Fabric involves a structured, multi-step approach:


Figure 4: Data Fabric Implementation

**Identify Industrial Data Sources**

- Capture data from sensors, assets, enterprise systems, and other operational technologies.

**Unify the Data**

- Understand and integrate OT, IT, and enterprise data using Industrial Edge connectors and integrators.
- Standardize the data into a Unified Namespace (UNS) for consistency and interoperability.

**Contextualize the Data**

- Process, transform, and enrich the data within a cloud data platform, aligning it with master data and metadata.

**Activate the Data**

- Implement Digital Twin models for real-time simulation.
- Develop Industrial AI and Agentic AI for predictive and autonomous decision-making.
- Build decision support systems accessible via web and mobile interfaces for field operations.
- Create dashboards and reports to enable rapid, data-driven decisions.
- Establish alerting mechanisms for proactive issue resolution.

**Democratize the Data**

- Govern and securely share data across teams and partners to foster collaboration and transparency.

## 5.2. Key Benefits

**Scalable Industrial Platform**
Enables rapid onboarding across diverse industrial assets, sites, and use cases, with flexibility powered by a Unified Namespace (UNS) architecture.

**High Performance and Availability**
Leverages AWS's robust infrastructure to ensure consistent performance and reliability for industrial workloads.

**Future-Ready Architecture**
Supports self-service capabilities, insight-driven planning, and intelligent decision-making through advanced analytics and visualization tools.

**Operational Efficiency and Cost Optimization**
Facilitates automation at scale, reducing the cost of generating insights while continuously delivering business value.

## Use Cases

The Industry Data Fabric adopted by our customers offers comprehensive data management capabilities, including:

- An integrated command center view that consolidates industrial data across silos, enabling rapid and informed decision-making.
- Robust data management features such as standardization, transformation, cataloging, and governance, along with advanced integration with third-party tools and cloud PaaS services.
- A Unified Namespace (UNS) that securely stores sensor and device data as industry-grade data products, enriched with discoverable data models and accessible via a knowledge graph.
- Self-service analytics, including customizable reports and dashboards.
- Secure data sharing mechanisms to support seamless collaboration across teams and partners.

## Capabilities

LTIMindtree provides end-to-end Industry Data Fabric design, development, and implementation using industry-standard platforms such as Cognite and Palantir, as well as custom solutions leveraging Azure, AWS, and GCP data services. One of our most extensive implementations was delivered for a global leader in flavours, fragrances, food ingredients, and health & biosciences, spanning 110 manufacturing plants and 92 laboratories. The objective was to enhance visibility into manufacturing and maintenance operations, logistics costs, and drive improvements in productivity and operational efficiency. We developed a unified Data Fabric platform that integrates IT and OT data to generate actionable insights, enabling predictive and prescriptive analytics for intelligent decision-making and optimization. A command centre dashboard was designed to provide a comprehensive view of delivery performance, cost metrics, manufacturing efficiency, freight and warehouse spend, maintenance, and quality performance.

## Conclusion

Ongoing discussions highlight the need to improve data storage, management, and usage for changing business needs. By modernizing data management practices, organizations can enhance data security, unlock actionable insights, and drive operational efficiency—ultimately enabling smarter business processes and more responsive customer service.

# References


1. *Industrial Cloud Platform Market Size, By Deployment (Public Cloud, Private Cloud, Hybrid Cloud), By Type (Infrastructure as a Service (IaaS), Platform as a Service (PaaS), Software as a Service (SaaS)), By Component (Solution, Services) By Application, By End-user, Forecast 2024 – 2032,* gminsights, May 2024: https://www.gminsights.com/industry-analysis/industrial-cloud-platform-market

2. *Projecting the global value of cloud: $3 trillion is up for grabs for companies that go beyond adoption,* McKinsey, November 28, 2022: https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/projecting-the-global-value-of-cloud-3-trillion-is-up-for-grabs-for-companies-that-go-beyond-adoption

3. *Cloud Manufacturing Market - Forecast(2025 - 2031),* industry ARC: https://www.industryarc.com/Research/cloud-manufacturing-market-research-800364

4. *What Are Industry Cloud Platforms?, Lori Perri,* Gartner, November 16, 2023: https://www.gartner.com/en/articles/what-are-industry-cloud-platforms

5. *How Data Fabric Can Optimize Data Delivery,* Gartner: https://www.gartner.com/en/data-analytics/topics/data-fabric

6. *Enterprise Data Readiness for the era of generative AI,* IBM: https://www.ibm.com/solutions/data-fabric

7. *Transforming Manufacturing with AWS,* AWS: https://aws.amazon.com/solutions/industrial/

8. *Industrial Data Fabric solution best practices – Part 1, organizational readiness and alignment,* Shaun Kirby, Ramesh Chinnasamy, AWS, June 13, 2024: https://aws.amazon.com/blogs/industries/industrial-data-fabric-solution-best-practices-part1/


To learn more, write to us at
**iNXT.IndustryDataFabric@ltimindtree.com**

# About the Authors

## Meenakshi Muralidharan

Principal Director of Architecture, LTIMindtree

Meena is the practice head of industrial middleware and data engineering at LTIMindtree's iNXT service line that focuses on building large scale Cloud, IoT and data platforms. She is a renowned multi-cloud expert professionally certified in Azure, AWS, GCP and Databricks. She helps our customers with technical advisory, design, architecture, and implementation of cloud-based IoT and data platforms. She is LTIM's AWS technical Ambassador and has co-authored books on Microsoft Synapse Analytics.

## Mohamed Imran

Senior Solutions Architect, LTIMindtree

Mohamed Imran is a highly experienced senior solution architect with over two decades of expertise in designing and architecting innovative data, connected IoT and cloud solutions. With specialization in Microsoft .NET, Data, Microsoft Azure, AWS, and GCP, Imran has led successful large-scale customer implementations and pre-sales engagements across global markets, consistently driving impactful business outcomes. Known for bridging business goals with technology strategies, Imran excels in crafting compelling proposals with technology solutions, RFP responses, technical advisory and documentation. A passionate advocate for emerging technologies and client success, Imran continues to shape enterprise digital transformation initiatives worldwide