Operationalizing *Adaptive AI* Across the *Manufacturing Shopfloor*

Capgemini

Foreword

In today's manufacturing landscape, digital ambition is not a constraint – execution is. The vision of the "lighthouse" or even "lights-out" factory has been discussed for years. Yet, when we move beyond the headlines and truly "peel the onion," a more fundamental question emerges: what are the core capabilities required to make this vision real, scalable, and repeatable?

This paper addresses this question with a pragmatic lens. Rather than revisiting the well-understood benefits of digital transformation, we focus on the technical imperatives that underpin it. At the heart of this shift lies a decisive move toward real-time, event-driven architecture an approach long theorized, but only now becoming tangible at scale.

Manufacturers today are navigating increasing complexity: smaller batch sizes, compressed innovation cycles, and the need for greater resilience across global operations. Meeting these demands requires more than incremental change. It calls for solutions that are inherently modular, context-aware, and capable of orchestrating intelligence across the entire value chain from engineering to shop floor to supply network.

What is particularly compelling is that what is long used to solely be a vision is no longer abstract. By combining advances in edge computing, cloud scalability, and time-series intelligence, we are now able to anchor this vision in concrete, deployable building blocks.
This creates a pathway from concept to implementation bridging the gap that has historically limited progress.

We are pleased to bring to you this Everest Group report as an invitation to move beyond aspiration. It is about translating a long-standing vision into actionable steps, grounded in technology that is ready today, and designed for what comes next.

Nicolas Rousseau
Head of Engineering
Technologies and Transformation
Capgemini Engineering

April 2026

Operationalizing Adaptive AI across the Manufacturing Shop Floor

Contents

Introduction

Manufacturing enterprises are at a critical inflection point. Persistent macroeconomic volatility, supply chain disruptions, shifting demand patterns, labor shortages, and rising cost pressures are exposing the limits of traditional, rigid operating models and incremental digitalization. As a result, there is growing urgency to move beyond incremental improvements toward more intelligent, responsive, and adaptive operations.

This need is giving rise to the concept of adaptive manufacturing systems – environments in which production systems can continuously learn, self-optimize, and respond dynamically to changing conditions. This may include adjusting process parameters to maintain quality, reconfiguring production schedules in response to disruptions, or optimizing energy consumption in real time. The shift toward adaptive manufacturing systems promises not only productivity gains but also fundamentally more agile, resilient, and future-ready manufacturing operations. This transition is increasingly enabled by advances in AI, particularly in real-time analytics, agentic systems, and data-driven decision-making, allowing manufacturers to move beyond static automation toward more adaptive operations.

However, while the vision is compelling, the path to achieving it is far from straightforward. Most enterprises continue to operate in complex brownfield environments characterized by legacy systems, fragmented data architectures, limited IT-OT interoperability, and tightly coupled processes. Scaling AI beyond pilots remains a persistent challenge, constrained by process rigidity, technology silos, talent shortages, and accumulated data and technical debt. As a result, tangible, enterprise-wide impact in terms of productivity, quality, and responsiveness remains limited for many manufacturers.

In this Viewpoint, we examine how manufacturers can bridge this gap and transition toward adaptive manufacturing systems enabled by AI-driven operations. Key themes covered include:

- The current state of manufacturing and the need for adaptive manufacturing systems
- Barriers to scaling adaptive AI across plants and shop floors
- A technology-led framework to operationalize adaptive AI
- The role of the ecosystem, including technology providers, hyperscalers, and system integrators, in enabling this transformation

Enterprise leaders across manufacturing, operations, engineering, and digital transformation can use this report to understand the architectural, data, and operating model shifts required to operationalize AI across shop floor environments.

Adaptive AI in manufacturing

Today, manufacturing enterprises operate in an environment defined by increasing complexity, volatility, and pressure on operational performance. They are also contending with structural challenges that are reshaping how production systems must function, as illustrated in Exhibit 1.

Exhibit 1: Key challenges shaping manufacturing transformation

Source: Everest Group (2026)

Challenge severity Low → High

Fragmented data and legacy system architectures

Siloed IT and OT systems limit real-time data access and interoperability and as a result, data remains underutilized and difficult to scale across plants.

Supply chain disruptions and network complexity

The need for improved coordination across distributed production networks is challenged by limited real-time visibility. Everest Group survey data reveals that over 50% of enterprises cite supply chain disruptions as a major operational challenge.

Workforce and talent constraints

Over reliance on experienced operators and domain experts, exacerbated by labor shortages. Nearly half the manufacturers Everest Group surveyed identified talent gaps as a key barrier to transformation

R&D to manufacturing disconnect

The need for lifecycle continuity and traceability from engineering to the aftermarket has been identified as a top priority for manufacturing organizations. This is based on an Everest Group survey of 234 CXOs from leading US and European engineering and manufacturing organizations.

Proving potential ROI from AI investments

Fragmented deployments and a lack of integration limit enterprise-wide impact. As per Everest Group data, more than 80% of manufacturers report difficulty in scaling AI beyond pilot use cases.

Demand variability and customer-centric production

Variable demand and shorter product cycles require greater flexibility in how production systems are configured and operated

.

Everest Group data indicates that over 80% of manufacturers report one or more of these challenges as a major operational constraint. As a result, decision-making across manufacturing operations remains largely reactive. Data is primarily used for retrospective analysis rather than real-time control, and insights generated through analytics are not consistently integrated into execution systems. This disconnect between insight generation and operational response limits the impact of digital investments and reduces manufacturers' ability to respond effectively to changing conditions.

The need for adaptive manufacturing systems

Taken together, these factors place manufacturing enterprises in a state of partial transformation. While significant progress has been made in digitizing operations, most organizations have yet to achieve the responsiveness, flexibility, and coordination required to operate in today's dynamic environment. Bridging this gap requires moving beyond static, rule-based systems toward adaptive, intelligence-driven approaches that can respond to variability in real time.

In this context, adaptive manufacturing systems refer to manufacturing environments where systems continuously sense operational conditions, analyze data in real time, and dynamically adjust processes across production, quality, maintenance, and supply chain functions. These environments enable a shift from reactive operations to systems that can learn, adapt, and optimize continuously. Intelligence is embedded directly into operational workflows, allowing decisions to be executed within the flow of production rather than through manual intervention or delayed analysis.

At a system level, adaptive manufacturing is characterized by three key capabilities:

- Continuous sensing and operational data, contextualized through semantic models, enabling real-time visibility across the shop floor
- Event-driven decision-making, where actions are triggered dynamically based on process conditions, demand signals, or equipment performance
- Closed-loop integration of decision-making and execution, enabling systems to detect and respond to issues automatically or with minimal human intervention

The adoption of adaptive manufacturing systems can deliver measurable improvements across manufacturing performance, as outlined in Exhibit 2.

Exhibit 2: Adaptive factories' impact on manufacturing systems

Source: Everest Group (2026)

Dimension	Key metrics (not exhaustive)	Key considerations for manufacturing leaders	Adaptive factories' potential impact
Downtime	Mean time to repair, unplanned downtime percentage	Reduce downtime and repair times to improve efficiencies and optimize costs	30-50% reduction in unplanned downtime
Quality	Defect rate, Right First Time (RFT)	Enhance quality to reduce waste, rework, and customer dissatisfaction	25-35% improvement in RFT
Productivity	Operator efficiency, throughput per shift	Enhance productivity, increase labor effectiveness, and factory utilization	20-30% gain in productivity
Decision speed	Time to replan/reschedule	Accelerate agility in responding to disruptions	3-5x acceleration in decision-to-action turnaround
Cost efficiency	Overtime costs, scrap/waste ratio	Leverage technology to create a direct bottom-line impact	15-25% reduction in waste and unplanned overtime

Beyond operational efficiency, adaptive manufacturing systems also play a critical role in improving resilience and sustainability. By enabling systems to respond dynamically to disruptions, enterprises can maintain operational continuity even in volatile environments. At the same time, real-time process optimization supports more efficient use of energy and materials, advancing sustainability goals and regulatory compliance.

Early examples of this shift are already visible in leading manufacturing environments, as illustrated below.

Case in point: Siemens

At Siemens' Erlangen PCB factory, industrial-grade AI is deployed alongside digital twins and advanced robotics to enable autonomous, real-time operations. AI-enabled robotic systems can identify and pick unsorted components, orient them correctly, and prepare them for assembly without predefined configurations – replicating tasks that previously required human dexterity. In parallel, a comprehensive digital twin allows AI models to be trained and validated in a virtual environment using synthetic data, significantly reducing development time and risk while improving accuracy. These capabilities extend to production processes such as automated quality inspection and precision assembly, where AI-driven systems detect defects with greater accuracy.

Together, these implementations demonstrate how AI is moving beyond isolated pilots to enable autonomous, adaptive, and continuously optimized manufacturing operations.

> "There was a world before AI. We are now transitioning to a world that makes full use of it – across factories, buildings, grids, and transportation."
>
> – Roland Busch, CEO, Siemens

Key challenges in scaling adaptive AI in manufacturing

While factories such as Siemens' Erlangen facility highlight what is possible, they also underscore the gap between leading-edge implementations and the broader industry. Although most manufacturers have begun investing in AI, Everest Group data indicates that over 75% of AI pilots failed to move from isolated pilots to consistent, enterprise-wide deployment across plants, production lines, and core operational workflows, reflecting a gap between experimentation and enterprise-wide adoption.

Most manufacturing enterprises continue to operate with fragmented systems and limited integration of intelligence into core workflows. The challenge lies not in deploying and maintaining individual AI use cases, but in building a parallel, data-driven layer on top of existing systems, processes, and operating models required to enable continuous, coordinated decision-making at scale.

These challenges are deeply embedded in manufacturing environments and can be broadly categorized across four dimensions: process, technology, data, and skills, as demonstrated in Exhibit 3.

Exhibit 3: Structural barriers limiting the scalability of adaptive manufacturing systems

Source: Everest Group (2026)

	Key challenges
Process rigidity and operational silos Most manufacturers continue to rely on manual or semi-automated decision-making processes, limiting their ability to transition to real-time, event-driven operations	Manual, rule-based, siloed, and sequential workflows limit the ability to respond dynamically to real-time events and disruptions Concerns around data privacy, safety, and security Uncertainty around model and agent governance
Legacy technology environments Nearly 50% of manufacturers report challenges in integrating IT and OT systems across plants, highlighting the scale of legacy constraints	Legacy IT-OT systems using different data formats and communication protocols operating in isolation Lack of interoperability and standardization across platforms and plants
Data fragmentation and limited real-time accessibility Over 60% of manufacturers surveyed by Everest Group lacked AI-ready data environments and cited data quality and accessibility as primary barriers to AI adoption	Data silos across systems and functions limit the effective use of the large volumes of data generated in manufacturing Incomplete, inaccurate, or limited availability of real-time data
Skill gaps for continuous improvement Talent/Labor shortages and a lack of cross-functional expertise have been identified as critical barriers to scaling AI initiatives	Shortage of skilled operators Disconnect between manufacturing engineers (with deep domain knowledge) and data scientists Change management issues

These constraints are particularly pronounced in brownfield environments, where new capabilities must be integrated into existing infrastructure. This adds complexity, as enterprises must balance innovation with operational continuity. Ultimately, scaling adaptive AI requires coordinated changes across technology architecture, data management, system integration, process design, and decision-making models.

Building adaptive manufacturing systems

While the need for adaptive AI in manufacturing is increasingly well understood, the path to implementation remains complex. The challenge is not merely deploying AI models, but enabling a connected architecture that links the shop floor, real-time decision layers, and enterprise intelligence. This ensures that decisions are triggered in real time and actions are coordinated across the production environment. Achieving this requires a shift from fragmented digital initiatives to a structured, scalable technology architecture – one that enables a human-in-the-loop, autonomous, and self-learning shop floor seamlessly integrated with enterprise systems. Exhibit 4 outlines a roadmap for enterprises to adopt and implement such an adaptive architecture.

Exhibit 4: Roadmap to operationalize adaptive manufacturing systems

Source: Everest Group (2026)

Define the foundational principles for an adaptive AI architecture

Establish the core design tenets that will govern the system to ensure that the adaptive AI architecture/platform remains scalable, resilient, and capable of evolving with changing manufacturing and business needs.

Engineer a converged, intelligence-driven architecture

Design a unified architecture that can link enterprise IT, shop floor OT, data layers, and agentic systems to enable a human-in-the-loop, self-optimizing, adaptive manufacturing system.

Build the partner ecosystem to operationalize the vision

Identify and align the right mix of partners (technology, platform, and service providers) to help bring this architecture to life.

The following sections expand on the key steps highlighted in Exhibit 4.

Foundational principles for an adaptive AI architecture

Adaptive manufacturing systems are increasingly built on AI-native architectures that enable real-time, data-driven decision-making. At its core, an adaptive manufacturing system is underpinned by five design principles that enable scalable, resilient, and intelligent operations, as depicted in Exhibit 5.

Exhibit 5: Key tenets of an adaptive AI architecture

Source: Everest Group (2026)

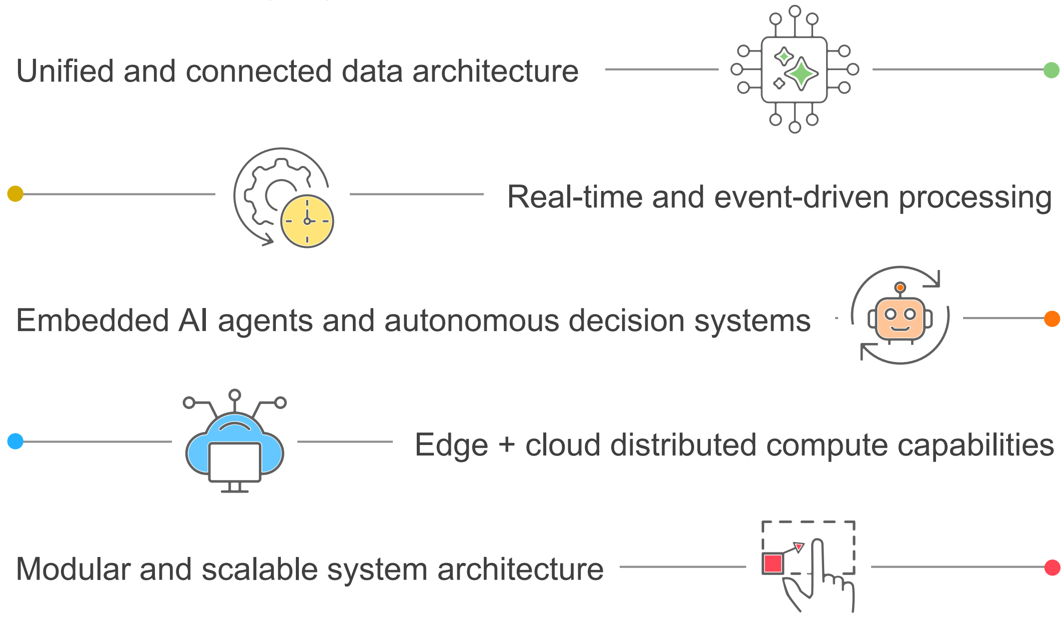

Unified and connected data architecture

Adaptive manufacturing requires a data-centric architecture where information is not only generated but also structured, contextualized, and accessible across systems. This brings with it the need for interoperability to enable cross-functional visibility, coordination, making operational data the foundation for decision-making across IT, OT, and enterprise layers. This also includes leveraging semantic data models and digital representations of assets and processes to provide manufacturing and business context, enabling more meaningful and accurate decisions.

Real-time and event-driven processing

For a truly adaptive system, manufacturing environments must support real-time processing and an event-driven architecture, where systems respond to events as they occur rather than relying on predefined schedules or periodic updates. In traditional environments, decisions are often based on static thresholds or batch processing. In contrast, an event-driven architecture enables systems to detect and act on signals such as shop floor disruptions in real time, triggering appropriate responses immediately. This is typically enabled through a Unified Namespace (UNS) and associated messaging protocols such as MQTT, which facilitate continuous data exchange and help break down silos across IT-OT environments.

However, as manufacturers extend these architectures beyond site-level integration to enterprise-wide, event-driven coordination, lightweight messaging alone may not be sufficient, creating a need for complementary event-streaming or pub-sub capabilities at the broader enterprise layer.

Embedded AI agents and autonomous decision systems

Embedding intelligence directly into workflows is essential to enable autonomy and continuous learning. In advanced environments, this extends to agent-based systems that dynamically orchestrate decisions and task execution across workflows, while still allowing for human oversight where needed. In practice, however, this is still limited to integrating AI-driven insights with execution layers such as MES, quality systems, and maintenance workflows, with real-time insights automatically triggering actions such as process adjustments, rescheduling, or maintenance interventions.

Edge + cloud distributed compute capabilities

Adaptive manufacturing requires a balance between centralized and decentralized decision-making. A hybrid edge + cloud architecture enables processing both locally (at the edge) and centrally (in the cloud), based on latency, scalability, and business needs. To scale this model across plants, enterprises increasingly need a containerized, microservices-based architecture that allows AI models, applications, and orchestration logic to be deployed, updated, and monitored consistently across distributed environments. This creates a converged edge-cloud operating fabric, reducing the need for site-by-site intervention while enabling continuous rollout of new capabilities.

Modular and scalable system architecture

Scalability demands modular, flexible, and reusable system components that can be deployed across plants and production lines with minimal reengineering. Standards-based architectures (for example, MQTT, UNS, Asset Administration Shell, and ISA-95) enhance interoperability and replicability, particularly in brownfield environments.

In addition to technology capabilities, adaptive manufacturing systems must meet industrial-grade requirements for reliability, security, performance, observability, and maintainability. Embedding strong governance and monitoring into manufacturing processes and system workflows, such as production execution, quality management, and maintenance, ensures the architecture remains resilient, secure, and continuously optimized as it scales.

Engineering an intelligence-driven architecture to scale adaptive AI

With these foundational principles in place, the next step is to design a unified architecture that integrates AI across IT, OT, and data layers to enable a self-learning, autonomous manufacturing system. Exhibit 6 illustrates a representative view of such an architecture.

Exhibit 6: Adaptive AI architecture in manufacturing

Source: Everest Group (2026)

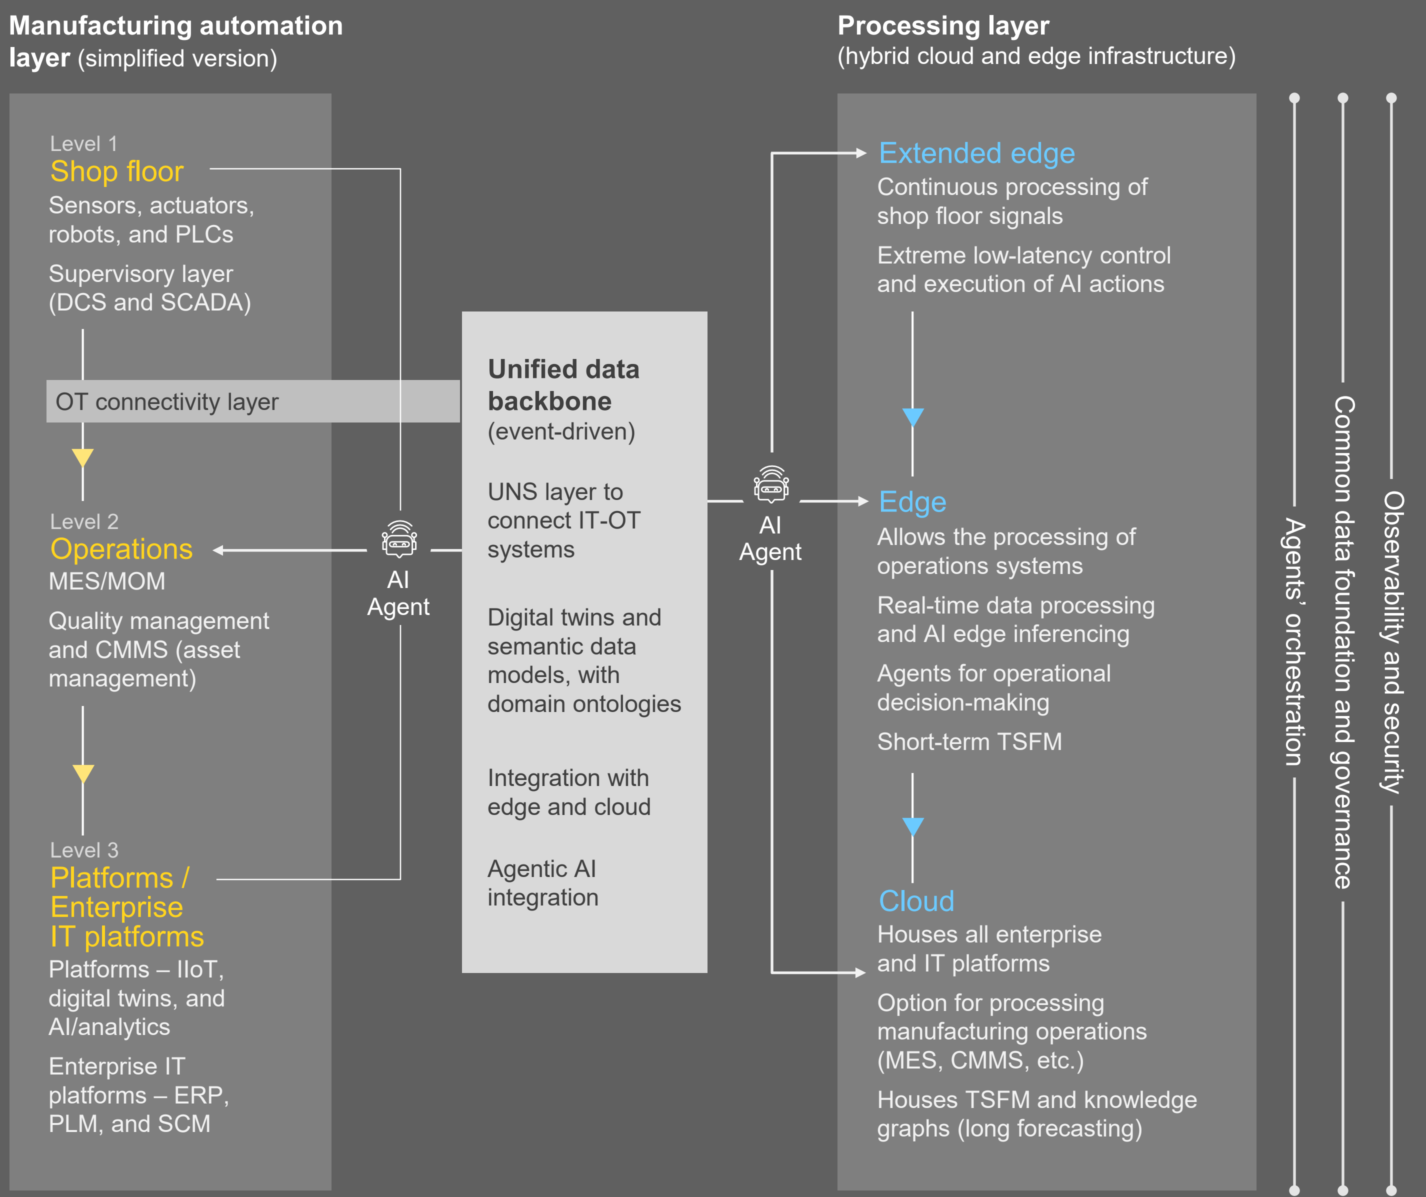

The architecture comprises three key components.

Layered manufacturing stack

The architecture builds on a simplified manufacturing automation pyramid:

- The shop floor layer includes machines, sensors, actuators, PLCs, and control systems such as DCS and SCADA that generate real-time operational signals
- The operations layer comprises systems such as MES/MOM, QMS, and CMMS, which orchestrate workflows across production, quality, and maintenance
- The enterprise layer includes ERP, PLM, IoT platforms, digital twins, and analytics systems that support enterprise-wide planning and optimization

Hybrid cloud-edge infrastructure

To enable real-time responsiveness and scalable intelligence, the architecture distributes workloads across cloud and edge environments:

- The cloud layer hosts enterprise platforms, AI/analytics systems, and advanced capabilities such as Time-series Foundation Models (TSFMs), knowledge graphs, and AI agents. It supports mid- to long-term forecasting, planning, scheduling, and scenario analysis, leveraging large-scale historical and cross-plant data to train and run foundation-scale time-series models. As techniques proven in other time-series-intensive domains are adapted to industrial contexts, TSFMs are increasingly entering manufacturing, expanding their role in forecasting, prediction, and broader decision support across production and supply networks
- The edge layer processes and acts on data locally within operational systems, enabling low-latency, near real-time decision-making. It includes historian systems, edge AI models, and short-cycle time-series inference to support local prediction and response, particularly where latency and operational continuity are vital
- The expanded edge layer brings compute closer to the shop floor, allowing AI systems to directly interact with machines, sensors, and process conditions, enabling continuous adaptation

In practice, this architecture is most scalable when cloud and edge environments are treated as a single, orchestrated infrastructure layer, with containerized microservices enabling centralized deployment, remote updates, and lifecycle management of AI models and applications across sites. This hybrid model ensures optimal distribution of workloads based on latency, scalability, and operational requirements, while also supporting plant-level data sovereignty.

Unified event-driven data backbone

At the core of the architecture is a unified, event-driven data backbone that connects manufacturing systems with processing layers. It integrates IT-OT data into a contextualized structure and enables real-time decision-making. Key elements include:

- A UNS that enables seamless integration across shop floor, operations, and enterprise systems by organizing and contextualizing data across layers, while serving as the logical structure for event-driven coordination across manufacturing workflows
- Digital twins and semantic models that provide contextual representations of assets, processes, and configurations, enabling more accurate, predictive, and adaptive decision-making
- A standardized OT connectivity layer (for example, OPC-UA and MQTT) ensures reliable communication and real-time data access at the operational level. However, enterprise-scale event orchestration may require complementary streaming and pub-sub mechanisms to handle higher data volumes, richer contextualization, and IT-OT convergence across sites

In practice, many manufacturers will require a hybrid event backbone: lightweight protocols such as MQTT support local shop floor data exchange, while more robust event-streaming platforms, such as Kafka-based architectures, enable enterprise-wide coordination, contextualization, and orchestration. Another vital component is agent-driven orchestration, which enables seamless coordination across systems and workflows. As manufacturing environments become more interconnected, AI-driven agents can orchestrate actions across layers, dynamically managing workflows, triggering decisions, and enabling closed-loop execution. These capabilities support key use cases such as predictive maintenance, quality optimization, and dynamic production scheduling.

In addition, the architecture embeds governance, observability, and security as cross-cutting layers to ensure system reliability, performance monitoring, traceability, and regulatory compliance in industrial environments.

Together, these components enable a shift from static, reactive operations to dynamic, adaptive systems. Data flows continuously across the enterprise, decisions are triggered in real time, and actions are coordinated across functions. This allows manufacturers to respond to variability more effectively, improve efficiency, and optimize performance across the production system.

While establishing foundational principles and engineering a converged architecture are critical first steps, operationalizing adaptive AI at scale requires a strong partner ecosystem. Enterprises must align technology and service providers to translate this architectural vision into deployable, scalable solutions across manufacturing environments.

Ecosystem's role in enabling adaptive AI

Realizing the vision of an adaptive manufacturing architecture is inherently complex, particularly for large enterprises with distributed operations. External partners such as technology and service providers play a vital role, and effective ecosystem collaboration is essential for successful implementation at scale.

Technology providers' role

As outlined earlier, adaptive manufacturing systems rely on real-time data processing, event-driven decision-making, and integrated intelligence across the production environment. Delivering these capabilities requires software platforms, especially AI-enabled applications, to align with event-driven architecture standards and interoperability guardrails. Enterprises must therefore engage with a coordinated set of technology providers, including industrial platforms, hyperscalers, and full-stack AI providers, to realize their adaptive manufacturing vision.

Each of these ecosystem participants serves a distinct yet complementary role, as illustrated in Exhibit 7.

Exhibit 7: Key technology providers enabling adaptive manufacturing systems

Source: Everest Group (2026)

Ecosystem layer	Illustrative logos	Illustrative solutions	Role in enabling adaptive manufacturing
Industrial platforms and automation providers	SIEMENS	• Industrial Edge and Insights Hub • Digital Twin Composer • Industrial AI OS and industrial copilots	• Enable digital twins and validation environments • Support edge AI and real-time control • Enable digital thread and IT-OT integration
Cloud and hyperscaler platforms	aws, Azure, Google Cloud	• Azure IOT Operations • Industrial Data Fabric (AWS) and Manufacturing Data Fabric (Microsoft) • AI Factories (AWS) and Microsoft Foundry	• Enable TSFM • Provide IIoT platforms and enable sovereign AI and predictive supply chain analytics • Provide linkages to other industrial platforms for data continuity and enable the unified data backbone
Full-stack AI technology providers	NVIDIA	• NVIDIA Omniverse • NVIDIA GPU stack and Physical AI	• Enable real-time inference and simulation • Support digital twins, robotics, and AI compute • Integrate with industrial platforms and hyperscalers for data continuity and enable the unified data backbone

Industrial technology platforms, such as Siemens, form the foundation of the ecosystem. They provide core capabilities across automation, MES, digital twins, and edge AI. These platforms operationalize data-centric architectures by enabling digital threads that connect physical assets with digital systems, ensuring consistent data flow across production environments. Their digital twin capabilities allow enterprises to simulate and validate changes before deployment, improving decision accuracy and reducing risk. Increasingly, these platforms are embedding AI directly into their offerings, including industrial AI operating systems spanning engineering and shop floor environments.

Hyperscalers such as Microsoft Azure and Amazon Web Services (AWS) provide the hybrid edge-cloud infrastructure needed to scale adaptive manufacturing systems. They offer IIoT platforms that unify IT-OT data through common semantic models and embed AI capabilities for use cases such as anomaly detection and predictive analytics. They are also advancing TSFMs, which are emerging as a key enabler in manufacturing. These models help capture demand patterns, asset behavior, process variability, and other dynamic operational signals at scale. Building on techniques proven in other time-series-intensive domains, TSFMs enhance forecasting, prediction, and planning, particularly as enterprises move beyond narrow, single-use case models toward more reusable intelligence layers. Additionally, hyperscalers support low-latency and sovereign AI deployments and collaborate closely with industrial and AI technology providers to ensure seamless interoperability.

Full-stack AI providers, such as NVIDIA, contribute advanced capabilities such as edge computing platforms, physical AI model development tools, and high-performance simulation environments. These capabilities enable accelerated processing for real-time optimization and digital twin applications, supporting advanced use cases such as predictive optimization and virtual factory modeling.

While these participants collectively enable a layered ecosystem aligned with adaptive manufacturing principles, their effectiveness ultimately depends on how well their capabilities are integrated into a cohesive architecture.

Service providers' role

While technology platforms provide the foundational building blocks, service providers play a vital role in designing, operationalizing, and scaling adaptive manufacturing architectures. Their role has evolved beyond traditional system integration to encompass technology consulting, architecture design, and ecosystem orchestration.

Leading providers enable the unified data backbone and event-driven architecture at the core of adaptive systems. This includes defining data models, implementing UNS frameworks, and integrating IT-OT systems to establish seamless, real-time data flows across the shop floor, operations, and enterprise layers. Increasingly, this also involves designing hybrid event architectures, where lightweight protocols such as MQTT support local operational exchange, while more robust event-streaming mechanisms enable enterprise-scale orchestration, cross-site contextualization, and broader IT-OT convergence.

They also design and implement the processing layer, including hybrid edge-cloud architectures tailored to manufacturing environments. This involves determining the optimal distribution of workloads across cloud, edge, and extended edge layers based on latency, scalability, and operational requirements. It also includes defining containerized, microservices-based deployment patterns and orchestration mechanisms that allow manufacturers to manage edge environments remotely, standardize deployments across plants, and continuously roll out new AI models and applications without extensive on-site reconfiguration.

In parallel, service providers are strengthening their AI and agentic capabilities through in-house investments and ecosystem partnerships. Leveraging domain expertise, reusable frameworks, and accelerators, they embed agent-driven workflows into operations and scale AI-enabled use cases such as predictive maintenance, quality inspection, and dynamic production scheduling.

Critically, service providers act as ecosystem orchestrators, coordinating across industrial platforms, hyperscalers, and specialized technology providers to ensure interoperability and alignment. This orchestration is essential to enable data-centric, event-driven, and scalable architectures. Their partnerships also facilitate the deployment of digital twins and semantic data models, which are key to contextual and predictive decision-making.

Overall, while integration and change management remain important, the role of service providers is increasingly centered on architecture-led transformation, defining how systems connect, how data flows, and how intelligence is embedded into operations at scale. They also play a key role in operationalizing governance, observability, and security frameworks to ensure reliable, compliant, and secure deployment of AI-driven systems.

Conclusion

Manufacturing is entering a new phase, shifting from isolated digital initiatives to the operationalization of intelligence across the shop floor. This transformation is enabled by adaptive AI, where systems continuously sense conditions, analyze data in real time, and dynamically adjust processes across production environments. In these adaptive manufacturing systems, intelligence is embedded directly into workflows, enabling continuous learning, adaptation, and optimization.

The journey toward adaptive manufacturing will depend on enterprises' ability to modernize data foundations, integrate AI into core workflows, and establish robust governance structures. At the same time, ecosystem collaboration will be critical, as no single organization can deliver the full range of capabilities required to scale adaptive systems.

Looking ahead, advances in agentic AI, digital twins, real-time optimization, and TSFMs are expected to accelerate this transformation. In particular, TSFMs are likely to enhance manufacturers' ability to forecast, anticipate, and respond to variability across production and supply networks, making them a key enabler in the short to mid-term. As these technologies mature, the boundary between digital and physical operations will continue to blur, giving rise to manufacturing environments that are not just automated, but also intelligent, responsive, and increasingly self-optimizing.

In this evolving landscape, manufacturers that successfully embed AI into their operational fabric will be best positioned to achieve sustained competitive advantage.

Everest Group®
With you on the journey

Everest Group is a leading research firm helping business leaders make confident decisions. We guide clients through today's market challenges and strengthen their strategies by applying contextualized problem-solving to their unique situations. This drives maximized operational and financial performance and transformative experiences. Our deep expertise and tenacious research focused on technology, business processes, and engineering through the lenses of talent, sustainability, and sourcing delivers precise and action-oriented guidance. Find further details and in-depth content at [www.everestgrp.com.](http://www.everestgrp.com)

This study was funded, in part, by
Capgemini

For more information about
Everest Group, please contact:

+1-214-451-3000

info@everestgrp.com

For more information about
this topic please contact the author(s):

Nishant Udupa, Vice President
nishant.udupa@everestgrp.com

Anurag Mukherjee, Senior Analyst
anurag.mukherjee@everestgrp.com

Kartikeya Panja, Analyst
kartikeya.panja@everestgrp.com

Notice and Disclaimers

Make
it real.

Capgemini is an AI-powered global business and technology transformation partner, delivering tangible business value. We imagine the future of organizations and make it real with AI, technology and people. With our strong heritage of nearly 60 years, we are a responsible and diverse group of over 420,000 team members in more than 50 countries. We deliver end-to-end services and solutions with our deep industry expertise and strong partner ecosystem, leveraging our capabilities across strategy, technology, design, engineering and business operations. The Group reported 2025 global revenues of €22.5 billion.

Make it real. | www.capgemini.com