

Invented for life



# Redefining Manufacturing Automation with Agentic AI

The Promise of Manufacturing Co-Intelligence

A white paper from Bosch Connected Industry



# Table of Contents

<b>1. Executive Summary</b>	<b>3</b>
<b>2. Entering a New Era of Manufacturing Intelligence</b>	<b>4</b>
<b>3. The Strategic Imperative of Agentic AI in Manufacturing</b>	<b>6</b>
<b>4. Bosch Case Studies: Agentic AI in Action</b>	<b>9</b>
<b>5. The Critical Role of Data and Semantics</b>	<b>11</b>
<b>6. The Context of Manufacturing Data: Why Domain Knowledge Matters</b>	<b>13</b>
<b>7. Navigating the Limits of Agentic AI</b>	<b>15</b>
<b>8. Bridging the Gap: Real-World Barriers to Agentic AI in Manufacturing</b>	<b>18</b>
<b>9. The Bosch Approach in Manufacturing Co-Intelligence</b>	<b>19</b>
<b>10. Stepping Into Manufacturing Co-Intelligence</b>	<b>21</b>
<b>11. Navigating the Age of Agentic AI</b>	<b>22</b>

# 1.

## Executive Summary

Manufacturers stand at an inflection point where the convergence of traditional industrial systems and cutting-edge AI technologies is beginning to reshape the global manufacturing landscape. The integration of Agentic AI introduces new strategic possibilities that extend far beyond the incremental gains of past industrial revolutions. While automation and digitalization brought connectivity and efficiency, Agentic AI enables a new level of intelligent, context-aware operations.

This white paper explores the promise of Agentic AI as a cornerstone for manufacturing transformation. It introduces the concept of co-intelligence – an approach that unites humans and machines in seamless collaboration. Through real-world case studies, including Bosch’s pioneering applications, it outlines how manufacturers can approach the adoption of Agentic AI not as a single leap but as a strategic journey. The paper also offers a practical roadmap, from early-stage readiness and pilot use cases to semantic modeling, hybrid system integration, and responsible deployment.

Agentic AI is not intended to replace human intelligence but rather to enhance decision-making capabilities, enable proactive optimization and support agile responses to disruption. Manufacturers who recognize this shift and take decisive action will unlock unprecedented levels of productivity, adaptability, and innovation in manufacturing automation powered by artificial intelligence.

# 2.

## Entering a New Era of Manufacturing Intelligence

For years, manufacturers have sought to unlock the full potential of a fully connected and integrated factory. While grappling with data silos, legacy system integration, and cybersecurity concerns, they have made significant strides: achieving documented efficiency gains of up to 20% in some areas, reducing downtime, and improving overall equipment effectiveness. Throughout this journey, the vision of autonomous factories, capable of self-optimization, proactive problem-solving, and enhanced human-machine collaboration, was a distant aspiration, a mere concept with no clear understanding of the precise form this embedded intelligence might take.

That was until now. The sudden emergence and rapid dissemination of Generative AI has emerged as a powerful catalyst for the manufacturing landscape. For the first time, we are witnessing AI that can understand complex engineering challenges, generating novel solutions and reasoning across diverse datasets with unprecedented speed and accuracy. This breakthrough is not merely another incremental improvement but a major stride to fulfill the vision of intelligent, self-optimizing manufacturing systems closer than ever before.

While machine learning has excelled at automating specific, repetitive tasks, Agentic AI is prepared to unlock a new level of efficiency, adaptability, and innovation across the entire manufacturing value chain. Instead of isolated robots and machines following pre-programmed routines, Agentic AI creates a sophisticated orchestrator of interconnected, intelligent agents. Each of them can sense, reason and act independently, while still collaborating to achieve common goals. These agents can dynamically optimize workflows, predict equipment failures, and proactively address unforeseen challenges.



Rather than relying on static rules and human intervention for every decision, Agentic AI empowers manufacturers to create truly responsive and resilient operations. Imagine a system where a sudden surge in demand is automatically addressed by reallocating resources, optimizing production schedules, and even anticipating potential supply chain bottlenecks – all without direct human intervention.

However, the journey to realizing this vision requires careful planning and execution. Successfully integrating Agentic AI into manufacturing demands a strategic approach that considers not only the technological challenges but also the crucial human element. Embracing this transformative technology means rethinking existing processes, fostering collaboration between humans and AI, and building a future where intelligent machines augment, not replace, the expertise of the workforce. By carefully navigating this complex landscape, manufacturers can unlock unknown levels of productivity, agility, and competitive advantage.

# 3.

## The Strategic Imperative of Agentic AI in Manufacturing

After years of transformative effort already invested in automation and digitalization, the prospect of another major overhaul can be daunting. However, in light of the fundamental shift in competitive dynamics initiated by Agentic AI, this is no longer a “nice-to-have” innovation. It has become a strategic imperative for manufacturers seeking to thrive in the coming decade. Agentic AI will likely redefine entire industries, including manufacturing. Companies that proactively integrate Agentic AI into their core operations will gain a decisive advantage in several key areas:



### **Accelerated Innovation Cycles:**

AI can significantly compress innovation timelines. By leveraging Agentic AI to simulate, analyze, and optimize designs, materials, and processes, manufacturers can drastically reduce the time to market for new products and services. This speed and agility are essential in rapidly evolving modern markets.

### **Mass Customization and Flexible Production Systems:**

The ability to deliver highly individualized products at scale is becoming a strategic differentiator. But this shift toward mass customization requires production systems that are not only responsive but also highly flexible and reconfigurable. Agentic AI enables manufacturers to dynamically adapt production lines, manage complex configurations, and optimize resource use in real-time – making customized production efficient, scalable, and economically viable.





### **Data-Driven Operations and Predictive Decision-Making:**

Manufacturing is increasingly a real-time, data-driven discipline. Agentic AI transforms vast streams of data in actionable insights—and even autonomous decisions—across production, logistics, and quality control. Predictive maintenance, demand forecasting, and real-time optimization are just a few of the areas where this shift delivers significant value.

### **Resilient and Adaptive Supply Chains**

In an increasingly volatile world, global supply chains face constant disruption. Agentic AI enhances resilience by predicting potential bottlenecks, optimizing inventory levels, and identifying alternative sourcing options instantly. These adaptive capabilities help manufacturers maintain continuity and minimize risk.



### **Sustainability and Resource Efficiency:**

As pressure mounts for greener operations, Agentic AI can play a key role in optimizing energy consumption, minimizing waste, and enabling circular manufacturing practices. It not only supports compliance with environmental regulations but also opens new pathways for sustainable innovation.

### **Workforce Transformation and Human-AI Collaboration:**

Agentic AI will transform how people work—not just what they do. Human-AI collaboration will become central to factory operations, requiring new skills, redefined roles, and a culture of continuous learning. Manufacturers that invest in upskilling their workforce and redesigning workflows for AI-human collaboration will unlock higher productivity and associate engagement.



In essence, Agentic AI is not just about automating tasks. It is about augmenting human capabilities, empowering associates to make better decisions, and transforming the entire manufacturing organization into a more intelligent, agile, and resilient entity. The manufacturers who recognize this strategic imperative and proactively invest in Agentic AI are likely to lead the next wave of industrial innovation.

As with any investment, however, the choice of technology partners and suppliers plays a crucial role. Given the current geopolitical landscape, manufacturers must carefully assess not only what they invest in but also where it comes from. For European manufacturers, choosing providers within the EU may offer greater long-term stability and alignment with regional values and regulations.

Those who hesitate risk falling behind their more forward-thinking competitors. The future of manufacturing belongs to those who embrace the power of Agentic AI—and do so with strategic clarity about their partners.



# 4.

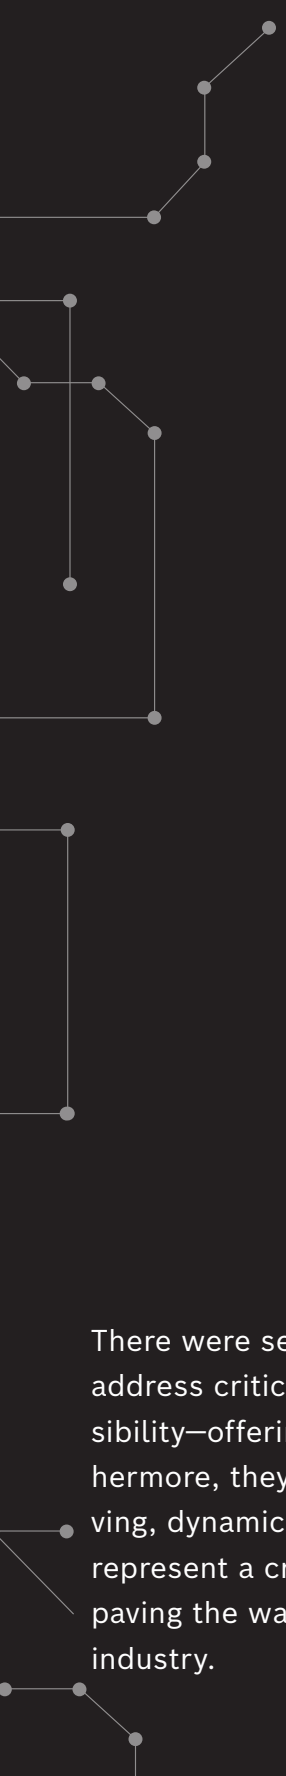
## Bosch Case Studies: Agentic AI in Action

While the field is still in its early stages, green shoots are already emerging,— with a few pioneering companies demonstrating early successes in experimenting with Agentic AI across various applications. These initial breakthroughs signal that its transformative potential is no longer just theoretical.

For companies already invested in data-driven maintenance and quality control, Agentic AI offers a powerful multiplier effect. The strong foundation laid by earlier systems—established data pipelines, domain expertise, and proven ROI—provides fertile ground for intelligent agents to flourish. Instead of merely predicting when a machine will fail, Agentic AI can now proactively optimize maintenance schedules based on constantly evolving conditions, autonomously adjusting parameters to minimize downtime and maximize efficiency. In quality control for process industries, these agents move beyond simply identifying defects, delving into root cause analysis in real time, and adjusting the production line based on autonomous reasoning.



Bosch, a pioneer in leveraging new technologies, is enhancing its manufacturing processes with AI agents. Having generated numerous ideas for potential use cases, the three most promising are now laying the foundational stones for further Agentic AI integration:



The Shopfloor AI Agent addresses the persistent challenge of unplanned equipment failures and the resulting downtime. By leveraging natural language processing and drawing upon cross-plant knowledge bases, this agent can rapidly diagnose issues, prioritize maintenance tasks, and even propose solutions, which effectively bridging communication gaps between multilingual shopfloor teams. This real-time responsiveness minimizes downtime, leads to faster issue resolution, and delivers significant cost savings. In a plant with 2,000 workers, this translates to estimated annual savings of €0.85 million.

The Workforce Scheduling AI Agent tackles the complexities of dynamic workforce allocation in a rapidly changing manufacturing environment. It integrates real-time data on worker availability, skill sets, and production line demands. This agent autonomously optimizes shift schedules, minimizes disruptions caused by unplanned absences, and ensures the right talent is precisely where and when it is needed. This adaptive approach unlocks greater workforce agility and efficiency than traditional static scheduling methods, generating estimated annual savings of €2 million in a 2,000-worker plant.

The Talk to Your Data AI Agent democratizes access to valuable insights by breaking down the barriers of complex data analysis. It empowers production workers to interact with data using natural language, enabling them to solve specific problems with low entry barriers, such as retrieving OEE figures from the last two months in real-time. By putting data directly into the hands of a broader range of personnel, this agent fosters a data-driven culture, empowering informed decision-making and driving continuous improvement. Estimated savings for this application are €0.5 million annually in a 2,000-worker plant.

There were several strategic reasons behind Bosch's selection of these particular use cases. They address critical pain points in manufacturing—downtime, workforce optimization, and data accessibility—offering immediate opportunities for improvement and demonstrating tangible ROI. Furthermore, they showcase the diverse capabilities of Agentic AI, including real-time problem-solving, dynamic resource allocation, and simplified data interaction. These focused applications represent a crucial first step: not only delivering practical value in real-world settings but also paving the way for the broader adoption of Agentic AI across Bosch's operations and the wider industry.

# 5.

## The Critical Role of Data and Semantics

Like many companies, Bosch began its exploration of Agentic AI in areas where it had already made significant investments and gained substantial experience in working with data. Indeed, high-quality data is a critical enabler for building intelligent systems. Data, even more than during the initial Industry 4.0 and digitalization waves, has become a strategic asset—one that demands rigorous management, architecture, and contextualization.

High-quality data – accurate, complete, consistent, and timely – is crucial for Agentic AI systems to function reliably and deliver meaningful outcomes. Poor data leads to flawed decisions, undermining performance and trust. However, building a perfect data architecture from the outset is not essential for initial pilot projects. A systematic approach to identifying gaps within the current IT/OT landscape and incrementally working toward a target architecture may be a more compelling and cost-effective strategy.

But data quality alone is not enough. To truly harness the value of data, companies need a

robust, future-ready data architecture, one that supports not only storage and access but also scalability, integration, and security. Scalability ensures that systems can grow without performance bottlenecks, especially as data volumes expand. A flexible architecture that supports integration offers seamless interoperability between systems such as MES, ERP, and SCADA, facilitating efficient data flow across structured and unstructured formats. Lastly, safeguarding sensitive operational data with encryption, access controls, and governance protocols ensures data integrity and security.

However, there is another essential layer to consider when preparing data for Agentic AI: semantic understanding. While data quality ensures accuracy and consistency, and a well-designed architecture enables scalability and integration, neither provides meaning on its own.

Agentic AI systems go beyond traditional automation. They must be able to perceive, interpret, reason, and act autonomously within complex environments. To do this effectively, much like humans, these systems must comprehend the relationships and meaning embedded in data—something only semantic modeling can enable. Semantic modeling provides structured meaning through ontologies, taxonomies, and knowledge graphs. It allows data to be machine-interpretable, not just machine-readable.



Unlike systems that see “12345” as merely a number, a semantically aware system recognizes it as a part ID linked to a specific machine, production order, or quality standard. This contextual understanding allows AI agents to reason about their environment, make dynamic decisions, and adapt automatically.

Without semantic modeling, even a high-quality dataset within a modern architecture remains just raw input. Agents may analyze patterns but lack the contextual intelligence to understand the significance or act meaningfully across systems. This limitation results in narrow, rigid applications that struggle with interoperability and are painful to scale. Semantic models also act as a common language between MES, ERP, SCADA, and AI engines. They preserve meaning during data exchange, which is essential for collaborative decision-making, traceability, and explainability—especially in regulated industries.

Ultimately, semantic modeling transforms data from a passive asset into an active participant in decision-making. It enables Agentic AI systems to operate with a shared understanding of the environment, unlocking capabilities such as predictive diagnostics, self-optimization, digital twins, and human-agent collaboration. Without it, autonomy remains superficial.

# 6.

## The Context of Manufacturing Data: Why Domain Knowledge Matters

While semantic modeling is an essential step in preparing data for the use of Agentic AI, there is an equally critical dimension in contextualizing manufacturing data for autonomous systems: the incorporation of domain knowledge—specifically, the logic of the production system and the tacit know-how of the people who operate within it.

Manufacturing environments are complex, dynamic systems where decisions are shaped not only by real-time data but also by long-standing operational practices, implicit priorities, and domain-specific heuristics. Operators, line supervisors, and process engineers often possess a nuanced understanding that is not visible in structured data: how to respond to a particular machine's behavior, when to intervene manually, or how to balance throughput against quality during a shift change. This tacit knowledge is vital for configuring and training AI agents as well as for validating the output of agentic workflows. So they are not only technically functional but also operationally effective.

Agentic AI workflows are not abstract algorithms—but are designed to act within a production system, influence its behavior, and make decisions aligned with its operational goals. As such, their design must be deeply grounded in the logic of that system: its material flow, constraints, interdependencies between machines, and the temporal rhythms of production. This logic is not always evident from data alone—it requires translating the real-world complexity of manufacturing into forms that autonomous agents can reason with responsibly.

For example, an AI agent tasked with optimizing machine scheduling must understand far more than just machine availability and cycle times. It must also take into account changeover rules, operator shift patterns, downstream bottlenecks, and preventive maintenance schedules. These critical considerations are rooted in domain expertise—often undocumented—and must be visible, structured, and embedded into the agent's decision logic.



Most Agentic AI frameworks aim to democratize agent creation through intuitive no-code or low-code platforms, reducing the need for intensive collaboration across AI engineers, data architects, and domain experts—except in more complex enterprise scenarios. However, domain knowledge remains indispensable even for initial deployments. Setting up effective Agentic AI workflows in manufacturing requires direct engagement with production teams to elicit unwritten rules, recognize failure modes, and capture informal workarounds. These insights must be rooted in models, ontologies, decision policies, and human-machine interactions to support the vital concept of “human in the loop.” Only then can Agentic AI become a reliable and trusted participant in the production environment.

In this way, domain knowledge serves as the interpretive lens that turns raw data into operational intelligence. It ensures that Agentic AI workflows are not only aligned with technical possibilities but also with the real, lived logic of the shop floor. Disregarding this key factor increases the risk of developing autonomous solutions that are detached from the operational contexts they are designed to enhance.

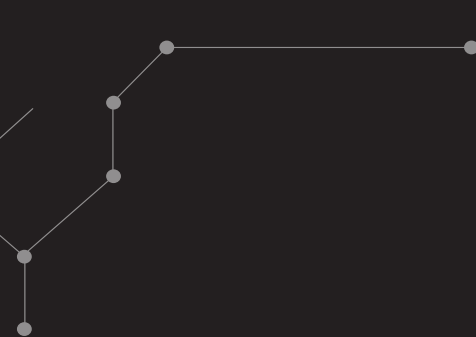
# 7.

## Navigating the Limits of Agentic AI

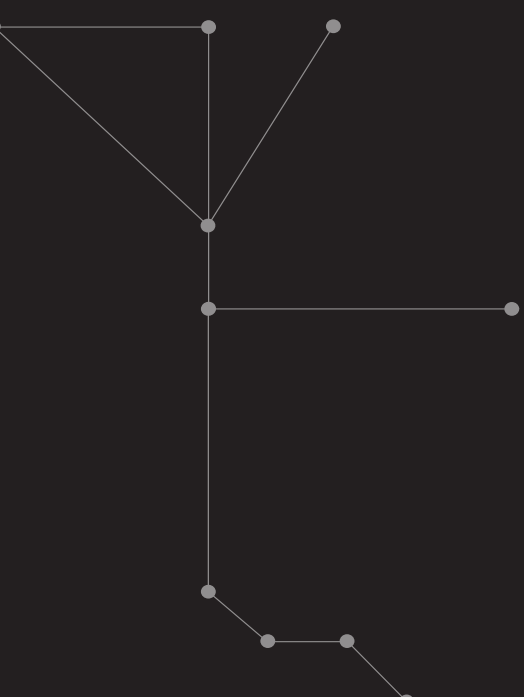
Despite its transformative potential, Agentic AI faces critical limitations that must be understood to ensure responsible and effective deployment within manufacturing environments. While intelligent agents offer new possibilities for decision-making, optimization, and adaptability, there are three key dimensions in which current capabilities fall short.

First, the probabilistic nature of Agentic AI introduces unpredictability that undermines consistency and safety. Unlike rule-based systems, agents may not produce the same outcome—even when presented with identical inputs. This non-determinism poses real risks in safety-critical areas such as quality assurance, where inconsistent decisions could result in defective products passing inspection—or safe ones being wrongly rejected. Such variability also opens the door to manipulation and creates a lack of traceability, which is unacceptable in regulated manufacturing environments. Moreover, for a few tasks—such as precise calculations or deterministic logic—Agentic approaches are simply unsuitable. However, this does not preclude agents from effectively orchestrating such tasks by delegating them to specialized, reliable tools. The key is knowing when to rely on the agent's autonomy and when to embed safeguards or structured hand-offs.

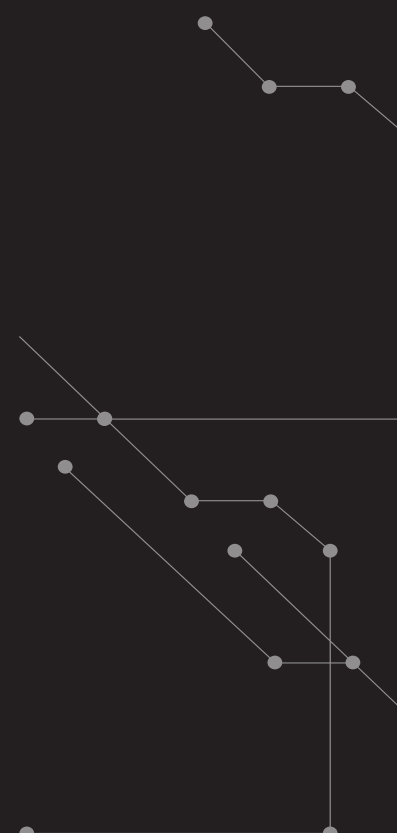





Second, Agentic AI lacks a comprehensive understanding of human, environmental, and ethical dimensions. While an agent may optimize for production efficiency, it might simultaneously increase emissions, violate safety protocols, or prioritize short-term gains over long-term sustainability. This inability to evaluate broader contextual impacts reinforces the need for human oversight and hard-coded safeguards—capabilities that AI alone cannot yet replace.



Third, the economic cost of running agents at scale remains prohibitive. Utilizing agents for every task, including those typically managed by Manufacturing Execution Systems (MES)—is not currently cost-effective. Advanced Agentic systems rely on large-scale language models that consume significant computational resources and token volumes. For the foreseeable future, this cost barrier will persist, making it impractical to replace deterministic systems such as MES for routine operational tasks.



Taken together, these limitations make it clear that a fully autonomous, agent-only production system is neither realistic nor advisable in the near term. Instead, the path forward lies in hybrid architectures—systems that combine the structure, control, and reliability of MES with the adaptability and intelligence of Agentic AI.



In this hybrid model, MES serves as the operational backbone, maintaining real-time execution, compliance, and data integrity, while interfacing with agents that handle higher-level reasoning, optimization, or exception management. For example, an agent might generate a revised production schedule in response to a supply chain disruption, but the actual execution and enforcement would still flow through the MES.

Looking ahead, MES platforms will increasingly integrate AI capabilities to enable predictive analytics, anomaly detection, and intelligent automation, while remaining the trusted interface between shopfloor operations and digital intelligence. This symbiotic relationship ensures that manufacturing remains not only innovative and efficient but safe, ethical, and economically viable.

To harness Agentic AI's full value, these constraints must be tackled in a structured manner. Emerging agent frameworks are actively working on first-stage solutions, including:

- Mechanisms for non-deterministic quality control
- Implementation of guardrails and multi-level permissions
- Human-in-the-loop design models
- Configurable autonomy, audit trails, and explainability
- Source citation and transparency

These foundational elements will be essential to ensuring that Agentic AI is deployed safely and responsibly in the structured, high-stakes environment of modern manufacturing.

# 8.

## Bridging the Gap: Real-World Barriers to Agentic AI in Manufacturing

Building hybrid models that enhance the possibilities of intelligent automation introduces a new level of complexity for manufacturers. The practical road to implementing Agentic AI is lined with substantial challenges that span technical, operational, and cultural dimensions. A key barrier lies in the data landscape. MES systems provide structure and visibility, but the data they generate is often fragmented, siloed, and lacking in contextual richness. It presents a significant challenge for Agentic systems, which rely not only on accurate inputs but also on data that reflects the logic, relationships, and intent behind operations. Without semantic modeling to define how entities and processes connect, agents are unable to interpret their environment meaningfully. The absence of semantic modeling hampers AI's potential to move beyond surface-level automation.

Compounding this issue is the fundamental mismatch in how MES and Agentic AI operate. MES systems are deterministic, rule-based, and tightly controlled. In contrast, Agentic AI is probabilistic and adaptive and often produces non-reproducible outcomes. Bridging this gap requires more than just integration—it demands a shared operational logic. Without it, AI-generated actions may conflict with the rigid structure of MES, leading to execution friction or outright failure. The risk increases in regulated industries, where traceability and repeatability are essential.

Beyond technology, the human factor presents its own set of challenges. Production staff often possess tacit knowledge—informal workarounds, judgment-based decisions, and context-sensitive practices—that are invisible to digital systems. Agentic workflows that fail to respect or incorporate this knowledge risk causing operational errors and generating resistance from the workforce. Trust in AI must be earned, and that begins with involving operators in designing workflows that enhance—rather than replace—their roles.

Even when these challenges are acknowledged, economic feasibility remains a pressing concern. Current Agentic AI models are computationally intensive and not yet scalable for real-time, shop-floor-wide deployment. As a result, manufacturers must be strategic, focusing on use cases where the return on investment justifies the cost. Overcoming these barriers is no easy task—made even more difficult by the lack of widely accepted reference architectures or standards.

# 9.

## The Bosch Approach in Manufacturing Co-Intelligence

A hybrid model that combines the best of both worlds—where MES continues to serve as the structured foundation for execution and compliance, while Agentic AI operates as a complementary layer—is still not fully available as a role-model implementation. Bosch, having invested in both the development of MES systems and the advancement of digital twins and semantic models, is now well-positioned to bridge these domains. Drawing on insights from both areas—and leveraging the deep domain knowledge accumulated through years of operational practice—Bosch is currently refining its approach to manufacturing co-intelligence, built on three interconnected components:

### Agentic AI Framework



1

A sophisticated AI engine serves as a powerful enabler of operational excellence, allowing autonomous systems to make context-aware decisions and dynamically adapt to complex production environments. By embedding Bosch's deep manufacturing expertise into agentic workflows—through training and validation in real customer settings—these systems can act with precision, reliability, and alignment to real-world constraints. This fusion of expert knowledge and autonomous capability leads to significant improvements in productivity, quality, and responsiveness, resulting in a new level of performance that amplifies the gains already achieved through traditional automation and digitalization.

## 2

### Composable MES

The Nexeed Industrial Application System (Nexeed IAS) represents Bosch's high level of proficiency in optimizing industrial operations. Born out of a long history of standardization, Nexeed IAS has proven to enhance transparency, efficiency, and adaptability across production lines. Designed with openness and interoperability in mind, it integrates with a wide array of systems and data sources. Its scalable architecture allows it to be deployed from single-machine setups to global manufacturing networks, flexibly adapting to diverse production needs. Nexeed IAS continues to serve as the structured backbone for execution and compliance.

## 3

### Bosch Semantic Stack

Supporting the entire architecture, the Semantic Stack acts as the common language that facilitates seamless data exchange and interpretation across all components—Agentic AI, Nexeed IAS, and the existing IT infrastructure. It establishes a single source of truth, promoting transparency, traceability, and robust data-driven decision-making throughout the organization. This unified data layer powers advanced analytics, enables the use of digital twins for simulation and optimization, and enhances predictive capabilities, thereby maximizing the value of data and enabling proactive optimization throughout the manufacturing lifecycle.

Bosch's approach directly addresses key barriers identified in earlier sections. The Semantic Stack helps dismantle data silos by offering a shared interpretive layer, enabling Agentic AI to understand its environment in context. The Composable MES bridges the gap between deterministic MES systems and probabilistic Agentic AI while upholding the standards of compliance and traceability. Integrating deep domain knowledge into its Agentic AI engine helps Bosch to preserve critical know-how and strengthen operator confidence. Moreover, Bosch acknowledges that the broader manufacturing ecosystem extends beyond its own architecture, hence the emphasis on interoperability as a critical success factor for enterprise-wide value creation.

# 10.

## Stepping Into Manufacturing Co-Intelligence

Before diving into the specifics of manufacturing co-intelligence, manufacturers must lay a solid foundation for AI adoption. It starts with prioritizing high-value use cases—specifically, identifying where GenAI and Agentic AI can deliver the most significant impact on operations.

One of the key advantages of starting with agentic workflows is the ability to achieve tangible results quickly, without the need for time-consuming system integration. Today, standalone use cases are implementable with agent integration through semi-automated onboarding processes, offering a high degree of flexibility. Notably, semantic data modeling is not a prerequisite to begin early-stage use cases in Agentic AI. It enables organizations to explore its potential using a low-barrier, test-and-learn approach.

Naturally, the benefits are even greater when a robust, scalable data and technology foundation is already in place—especially one that supports semantic modeling. In the long term, manufacturers need to systematically identify and address gaps in data quality, architecture, and governance. Doing so creates a fertile environment for scaling agentic workflows and ensures their sustained impact across the AI journey.

In parallel, the way operations are structured requires a fundamental rethink: What are the concrete and redefined workflows in which AI creates value? How will human-AI collaboration materialize in daily operations? Preparing and training personnel is another critical step on the path to co-intelligence. It means cultivating AI expertise and fostering a culture of innovation throughout the organization.

Lastly, having a trusted partner—one that brings not only technical expertise and practical solutions but also a proven track record of real-world implementation—can be a valuable cornerstone as organizations take their first steps toward manufacturing co-intelligence.

# 11.

## Navigating the Age of Agentic AI

The path forward for manufacturers in the age of Agentic AI hinges on strategic integration and a nuanced understanding of both its opportunities and limitations. While the vision of fully autonomous factories remains aspirational in the near term, tangible benefits are within reach through the incorporation of intelligent agents.

A balanced approach is key—favoring hybrid architectures that leverage existing MES systems while gradually integrating the adaptive intelligence of Agentic AI. It enables manufacturers to maximize the value of current investments while laying the foundation for more sophisticated autonomy.

Focusing on high-impact use cases with measurable ROI—such as predictive maintenance, quality control, and resource optimization—is essential for demonstrating early value and building momentum for broader adoption.

Critically, investing in robust data infrastructure and semantic modeling is foundational. Agentic AI thrives on high-quality, context-rich data—not just raw information. It means not only capturing data but also imbuing it with meaning and ensuring seamless interoperability across systems.

Furthermore, unlocking the transformative potential of Agentic AI depends on cultivating human-AI collaboration. Fostering partnerships between AI engineers, data scientists, and domain experts is vital for designing effective workflows that seamlessly blend human expertise with machine intelligence.

Finally, manufacturers must also navigate the ethical and societal implications of increasingly integrated AI systems—ensuring data security, promoting responsible use, and proactively addressing algorithmic bias.

The transition to Agentic Manufacturing demands a strategic mindset, a commitment to continuous learning, and a willingness to adapt. By addressing these factors thoughtfully and responsibly, manufacturers can unlock new levels of efficiency, agility, and competitive advantage—ultimately shaping a new era of industrial innovation.



# Make the first move!

## **The Next Leap in Industrial Productivity**

Bosch isn't just adapting to the future of manufacturing—we're actively shaping it. We're pioneering manufacturing co-intelligence, a transformative approach that enables manufacturing companies to achieve unprecedented levels of autonomy, efficiency, and resilience.

Traditional connected solutions deliver insights. Bosch delivers action. Our agent-based AI solutions go far beyond data analytics, proactively optimizing processes, predicting and preventing disruptions, and driving continuous improvement. This is more than just an incremental change; it's a fundamental shift in the manufacturing industry.

Find out more and contact us.

[www.manufacturing-co-intelligence.com](http://www.manufacturing-co-intelligence.com)

### **Robert Bosch Manufacturing Solutions GmbH**

PO Box 30 02 20  
70442 Stuttgart  
Germany  
+49 (711) 811 10 900

[www.bosch-connected-industry.com](http://www.bosch-connected-industry.com)  
© Robert Bosch Manufacturing Solutions GmbH  
June 2025

Bosch and the device mark are registered trademarks of Robert Bosch GmbH, Germany.

This document is a schematic representation and not an operating manual.