WHITE PAPER

Welcome to the Thinking Factory

How AI is Giving MES a Brain

By Francisco Almada Lobo
CEO at Critical Manufacturing

Critical manufacturing
an ASMPT company

Index

Executive Summary

Manufacturing is entering a new era, one where intelligence is no longer confined to isolated algorithms or analytics dashboards but infused throughout operational systems themselves. In this context, traditional boundaries between automation, MES (Manufacturing Execution Systems), and data platforms are dissolving. In their place emerges a unified, AI-ready architecture that is not only capable of orchestrating production but of reasoning, adapting, and improving autonomously.

This white paper explores the technological and organizational shifts necessary to enable this transformation. It introduces the concept of the "Thinking Factory", a digitally integrated, self-learning environment built on a core synergy between automation, MES, and data architecture. It also outlines how the next generation of AI, from classical machine learning to large language models (LLMs) to agent-based systems, depends not only on smarter models, but on smarter infrastructure.

Without a foundational data model that links operations, execution, and context, even the most advanced AI solutions will remain underutilized and fragmented. The goal of this paper is to educate decision-makers and digital transformation leaders on why this foundation is critical, what it looks like, and how to approach it strategically, positioning AI not as a bolt-on capability but as a native layer of the manufacturing brain.

> "By 2029, computers will have emotional intelligence and be convincing as people."
>
> **Ray Kurzweil**

The Age of Intelligent Manufacturing

Ray Kurzweil's vision of exponential technological change is no longer theoretical. In the world of manufacturing, this change is manifesting through rapid advances in automation, artificial intelligence, and connectivity.

While the industrial revolution once unfolded over decades, we are now witnessing paradigm shifts in just a few years, driven by the convergence of data, algorithms, and computational power.

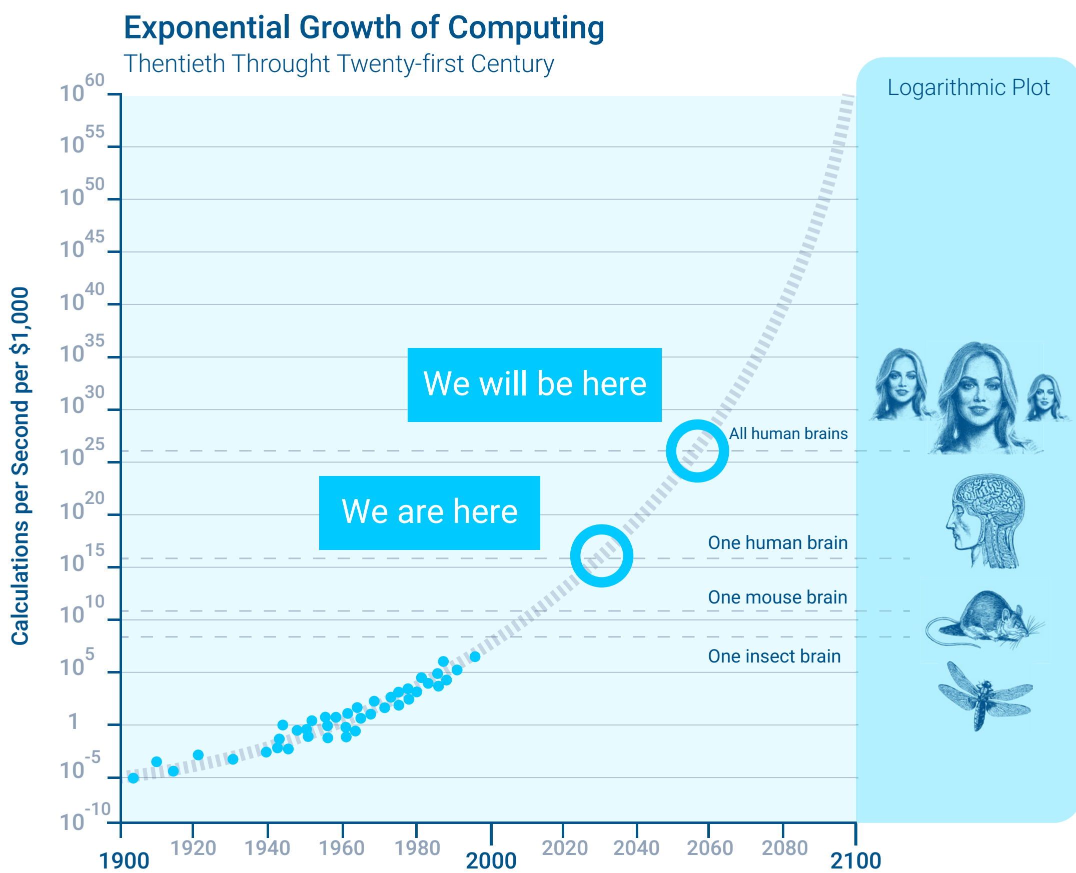

Figure 1: Exponential Growth of Computing

Source: The Singularity is Near, Ray Kurzweil, 2005

The industrial sector stands at the inflection point of what can be called the "Thinking Factory". Machines are not only executing commands, but they are also starting to understand context, learn from patterns, and adapt their behavior in real-time. LLMs (Large language models) such as GPT-4 and Claude 3 are demonstrating complex reasoning capabilities once thought exclusive to human operators. Multi-agent systems are allowing distributed systems to collaborate toward goals. And yet, the reality on many factory floors is still governed by brittle rules, siloed data, and narrow automation logic.

The message is clear: transformation is not optional. In a world where AI is becoming a first-class operational citizen, the ability to adapt rapidly is no longer a competitive advantage; it is a survival imperative.

The Manufacturing Adaptation Gap

Despite the technological possibilities, most manufacturers are not ready for the future. This is due to what Scott Brinker termed Martec's Law: technology changes exponentially, but organizations change logarithmically.

The result is an adaptation gap, a growing distance between what is technologically possible and what is operationally implemented.

Martec's Law and the Adaptation Gap

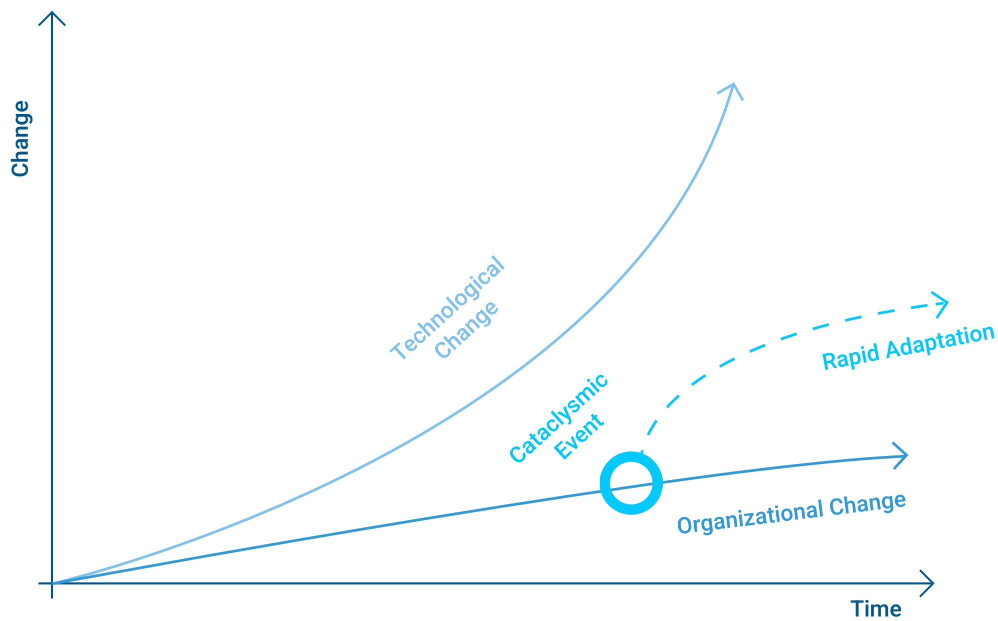

Figure 2: Martec's Law and the Adaptation Gap

So we won't experience 100 years of progress in the 21st century - it will be more like 20,000 years of progress (at today's rate). Ray Kurzweil (2001) "The Law of Accelerating Returns"

Legacy MES systems were never designed for intelligent, real-time reasoning. They operate on fixed logic, require manual configuration, and cannot natively ingest unstructured data like logs, images, or text. AI and data science teams often work in silos, disconnected from the systems that run the shop floor. The tooling is fragmented. The data is incomplete or inconsistent. The result? AI pilots may show promise in isolation, but they fail to scale in production.

This is not just a technology issue; it is a matter of architecture. It is a problem of thinking in layers when the opportunity demands systems thinking. If MES, automation, and data platforms remain loosely coupled, AI will remain peripheral, a sideshow rather than a central nervous system.

The Foundation of AI in Manufacturing: Data, MES & Automation

Before manufacturers can capitalize on advanced AI models, they must first modernize their digital core. That begins with integrating three foundational components: automation, MES, and data platforms. Automation generates the real-time signals that reflect the physical state of production. MES provides the operational context and governs process logic. Data platforms store and expose both structured and unstructured data for analytical use.

Together, these form the “Holy Trinity of Smart Manufacturing”. Only when they are harmonized around a shared data model can downstream AI capabilities, like predictive quality, intelligent scheduling, or AI agents, operate reliably. A well-architected MES and data platform stack provides not only the operational data but also the semantic context required for AI to reason effectively. Without this context, classical machine learning models remain shallow and narrow in scope, and generative AI models, such as LLMs, become disconnected from the real-world environments they are meant to support. In these cases, LLMs may hallucinate, misinterpret inputs, or make unreliable decisions due to the lack of grounding in structured operational data.

The Holy Trinity of Smart Manufacturing

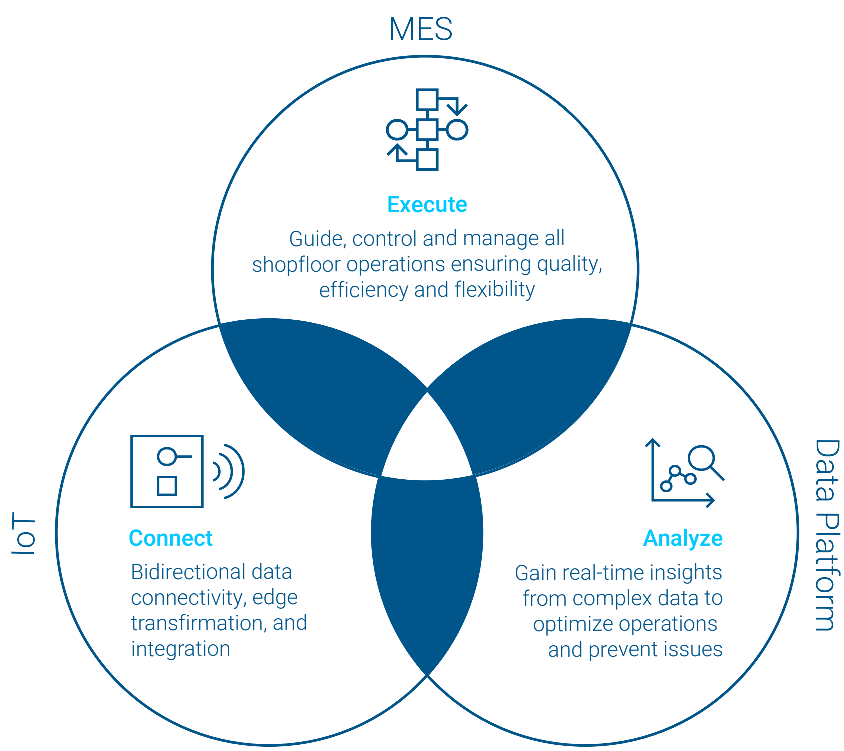

Figure 3: The Holy Trinity of Smart Manufacturing

From ML to LLMs to Agents: The Three AI Waves

Artificial Intelligence in manufacturing is not a monolith; it is an evolving spectrum. Understanding how it has progressed helps explain why past efforts fell short, and why new models hold promise. The trajectory can be understood in three overlapping but distinct waves of AI: classical machine learning, LLMs, and AI agents.

The Three Waves of AI

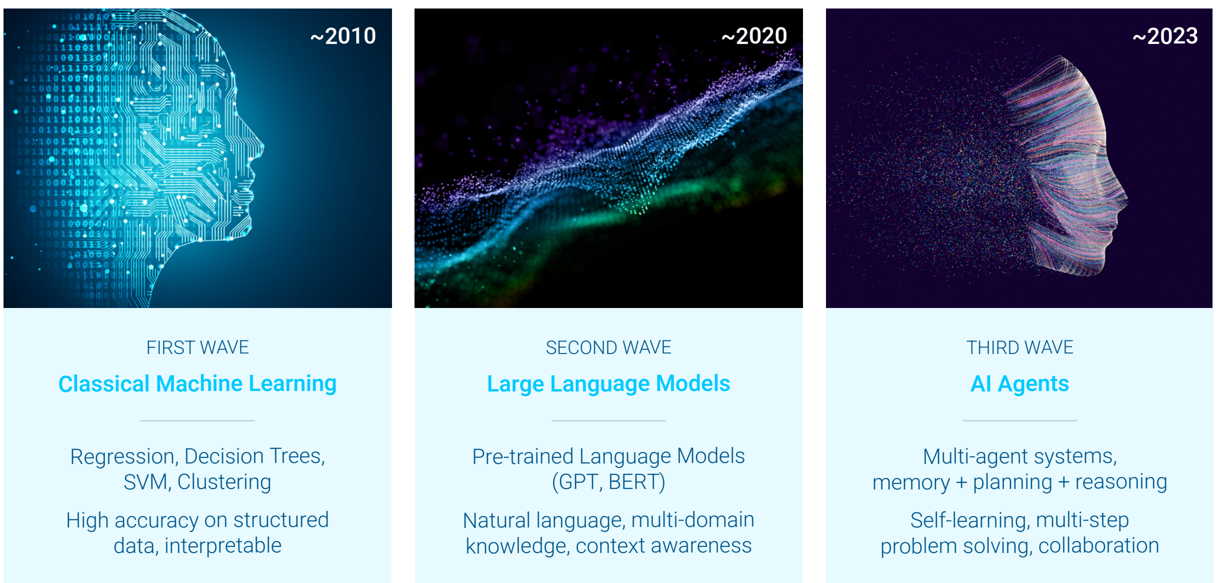

Figure 4: The Three Waves of AI

The first wave, classical ML, enabled predictive models based on structured data and statistical learning. These systems could forecast demand, predict maintenance needs, or detect anomalies, but they required clean inputs and extensive human supervision.

The second wave, LLMs, ushered in general-purpose intelligence across natural language and unstructured domains. These models, trained on massive corpora, could summarize logs, draft SOPs, and engage with operators in plain language. However, they struggled with domain grounding, hallucination, and operational reliability.

The third wave, AI agents, combines LLM capabilities with memory, planning, and tool use. These agents can take actions, adapt to feedback, and reason across steps. They do not just answer questions; they solve problems, initiate workflows, and collaborate across systems. In manufacturing, this evolution brings intelligence to the edge of operations, where it is most needed.

The Promise of LLMs on the Factory Floor

LLMs, when applied carefully, hold immense promise for industrial applications. Their language understanding capabilities allow them to bridge traditionally siloed functions, engineering, operations, maintenance, and IT, through intuitive, conversational interfaces.

Key Promises

- **Natural Language Interfaces:** Operators or engineers can ask questions like “What caused the yield drop yesterday?” or “Show me maintenance trends for Line 3,” and receive useful insights.
- **Automated Documentation:** SOPs, work instructions, shift handover notes, and exception reports can be generated or summarized on the fly.
- **Root Cause Exploration:** When paired with structured logs or time-series data, LLMs can assist in initial diagnostics or anomaly triaging.
- **Cross-System Reasoning:** By ingesting knowledge across MES, ERP, equipment data sources, and PLM systems, they can unify insights without the need for bespoke dashboards or custom analytics tools.

LLMs do not require retraining to operate across domains, making them faster and cheaper to deploy than classical ML. They enable experimentation and iteration at a scale previously impractical. However, with power comes risk, especially when applied in safety or compliance-critical environments.

The Pitfalls: Why LLMs Struggle Without Structure

Despite their capabilities, LLMs are not turnkey solutions. Their behavior is driven by probability, not certainty. When disconnected from real-time, structured data or deployed naively, LLMs can lead to dangerous or misleading outputs.

Hallucination Syndrome: LLMs may generate plausible-sounding but entirely false statements. In manufacturing, this can result in incorrect diagnostics, unsafe suggestions, or inaccurate documentation, especially when outputs are not traceable to source data.

Goldfish Syndrome: LLMs have no long-term memory. They cannot retain context across time, meaning insights, anomalies, or operator preferences from one shift, day, or week are lost unless manually reintroduced. This makes continuity difficult, particularly in troubleshooting recurring or evolving issues.

Inverted Assistant Syndrome: LLMs can suggest actions like "reschedule this job" or "adjust inspection frequency," but they do not execute these actions, nor do they specify how to carry them out in operational systems. This results in a disconnect between intelligence and action, creating bottlenecks rather than automation.

These limitations make LLMs unreliable when deployed without a structured framework. Worse, their fluency often masks their fallibility, making it difficult for users to judge the reliability of the outputs without deep domain knowledge or validation mechanisms.

LLM Syndromes

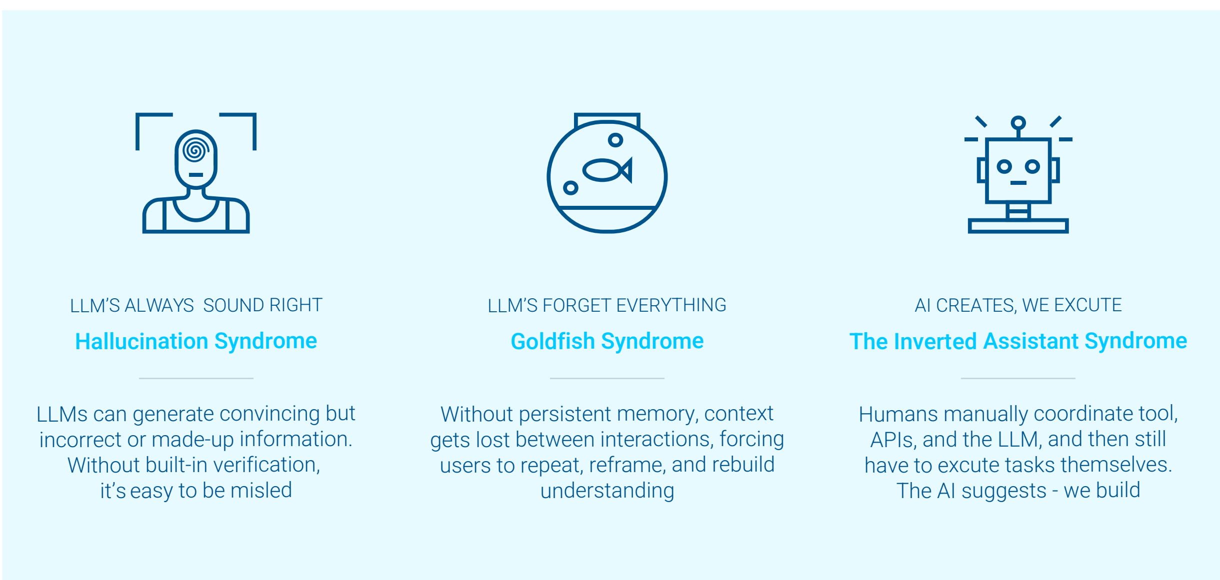

Figure 5: LLM Syndromes

Techniques to Specialize and Stabilize LLMs

To overcome these pitfalls, several complementary techniques have emerged. They help tailor LLM behavior to domain-specific needs and reduce risk by grounding the model in verified context.

Retrieval-Augmented Generation (RAG): RAG connects an LLM to a trusted external knowledge base, often a vector database built from manuals, logs, or structured data sources. For MES-specific applications, this can include MES documentation, historical MES data, contextual production rules, or configuration files.

Fine-Tuning: Fine-tuning involves training an LLM on curated, domain-specific data such as MES event logs, configuration data, or annotated operator interventions. This improves the model's performance on repetitive tasks and technical language, offering higher reliability for manufacturing operations.

Prompt Engineering: Prompt engineering involves crafting precise instructions, examples, or templates to guide the LLM's behavior without altering the model itself. This can include templates for SOPs, structured query prompts, or MES-specific Q&A examples.

Together, these techniques bring LLMs closer to operational usefulness in manufacturing. However, true transformation requires systems that initiate, plan, act, and adapt: AI agents.

Methods for Specializing LLMs

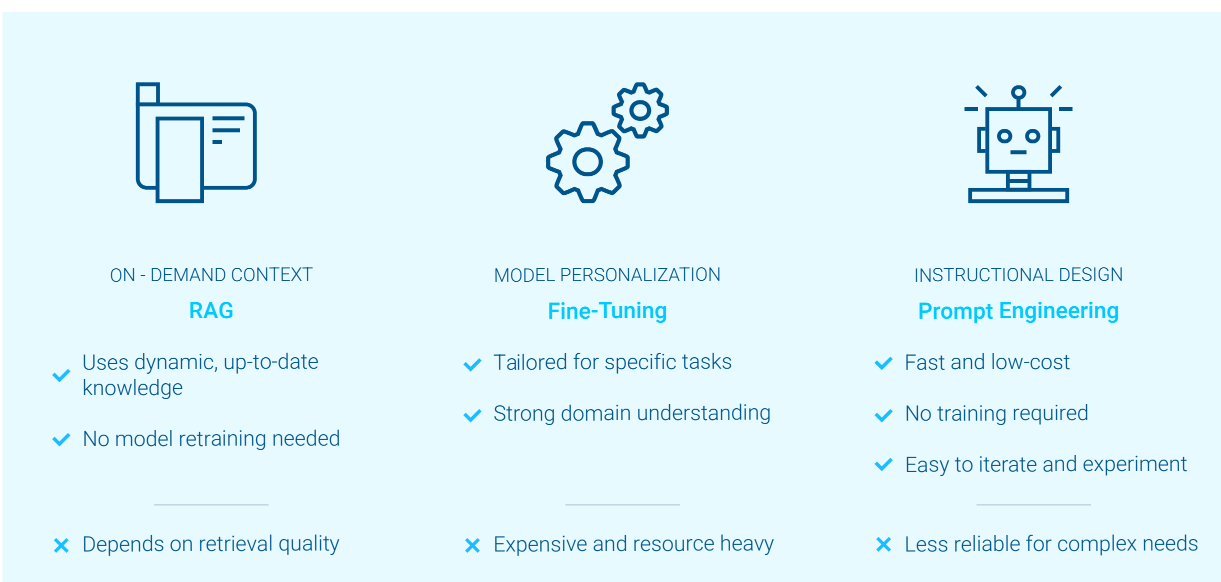

Figure 6: Methods for Specializing LLMs

The Rise of AI Agents and Agentic Workflows

While LLMs offer intelligent suggestions, they remain fundamentally stateless and passive.They wait for a prompt, provide a response, and forget what just happened. In manufacturing, this is not enough.

What is needed are systems that do not just interpret, but also plan, act, learn, and collaborate. That is the leap from LLMs to AI Agents.

Brains, Memory, and Tools: Inside an AI Agent

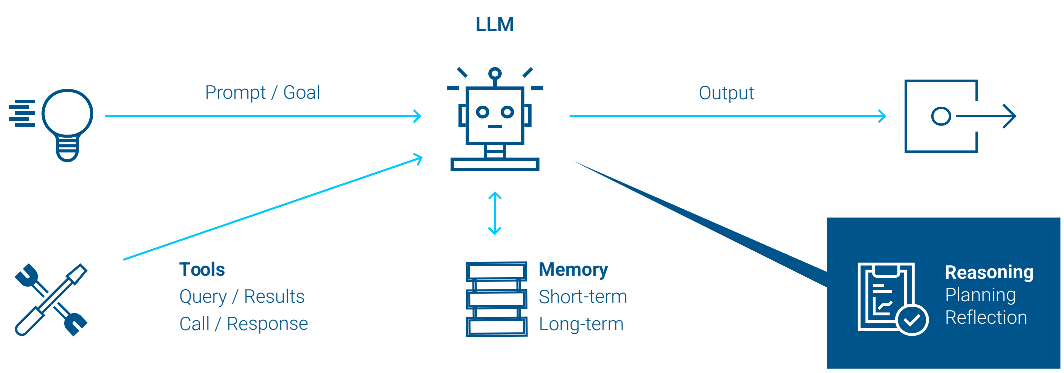

Figure 7: Brains, Memory and Tools: Inside an AI Agent

An AI agent is a system that has a goal, uses tools to interact with the world, retains memory, and applies reasoning to achieve its objective.

In MES, agents enable workflows that can respond to dynamic inputs, such as changing production conditions or equipment failures. They can reschedule operations, reroute materials, and optimize decisions without relying on hardcoded rules.

Agentic Workflow

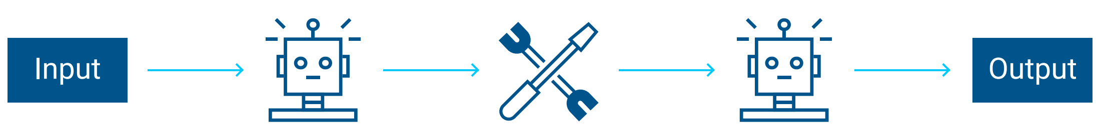

Figure 8: Agentic Workflow

AI Agent

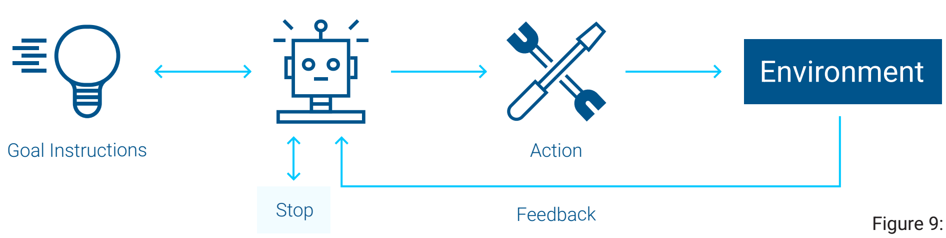

Figure 9: AI Agent

Unlike traditional MES logic, which follows linear, deterministic flows, agentic workflows operate with conditional logic and reasoning.

This allows for greater flexibility and resilience, key characteristics of next-generation MES systems.

Model Context Protocol (MCP): The Infrastructure for AI in MES

As agents multiply and specialize, they need a consistent way to communicate with MES systems, access data, and share context. Model Context Protocol (MCP) is the infrastructure that supports this.

MCP provides a framework that allows agents to discover available APIs, access shared memory, and communicate securely with MES and data platforms. It includes MCP Clients embedded in each agent and an MCP Server that exposes MES objects like materials, schedules, and quality parameters.

With MCP, agents can operate in a modular fashion, without custom integrations for every function. It standardizes interaction and enables agents to be rapidly developed, deployed, and scaled.

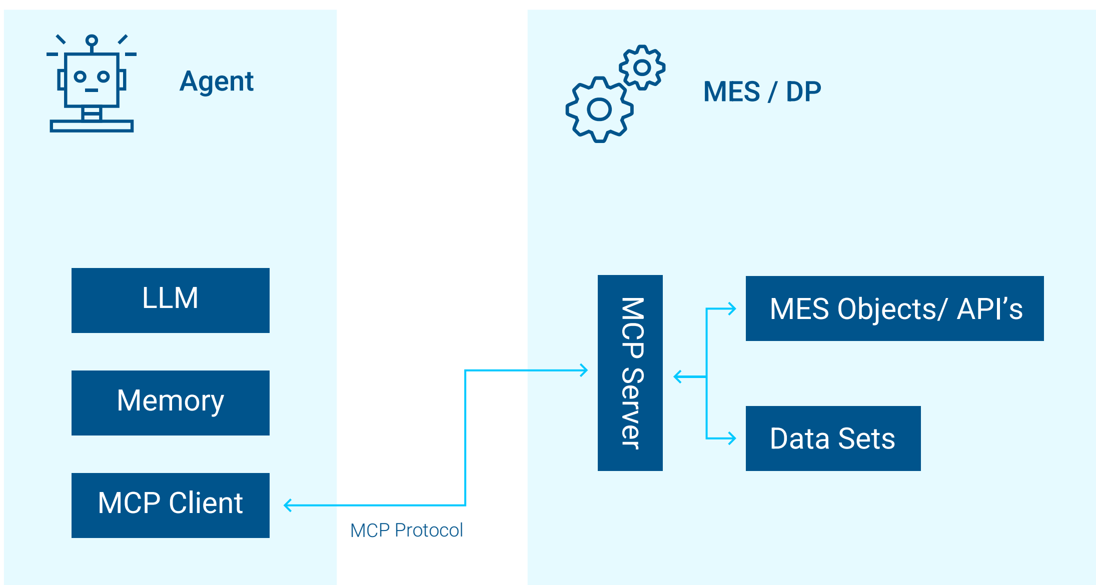

Figure 10: Communication between agents and MES/DP

Agent-to-Agent Communication (A2A): How Intelligence Scales

In complex manufacturing environments, no single agent can do everything. Intelligence must be distributed. Agent-to-Agent Communication (A2A) enables this.

With A2A, specialized agents, such as scheduling, material flow, or quality agents, can discover and collaborate. For example, if a material flow agent detects a bottleneck, it can notify the scheduling agent. If a maintenance agent anticipates a breakdown, it can signal the quality agent to adjust inspection frequencies.

These interactions happen through structured protocols built on MCP. They are logged, observable, and governed by shared policies. This enables factories to evolve into intelligent ecosystems where decisions are distributed, coordinated, and continuously optimized.

A24 - Agent to Agent Protocol

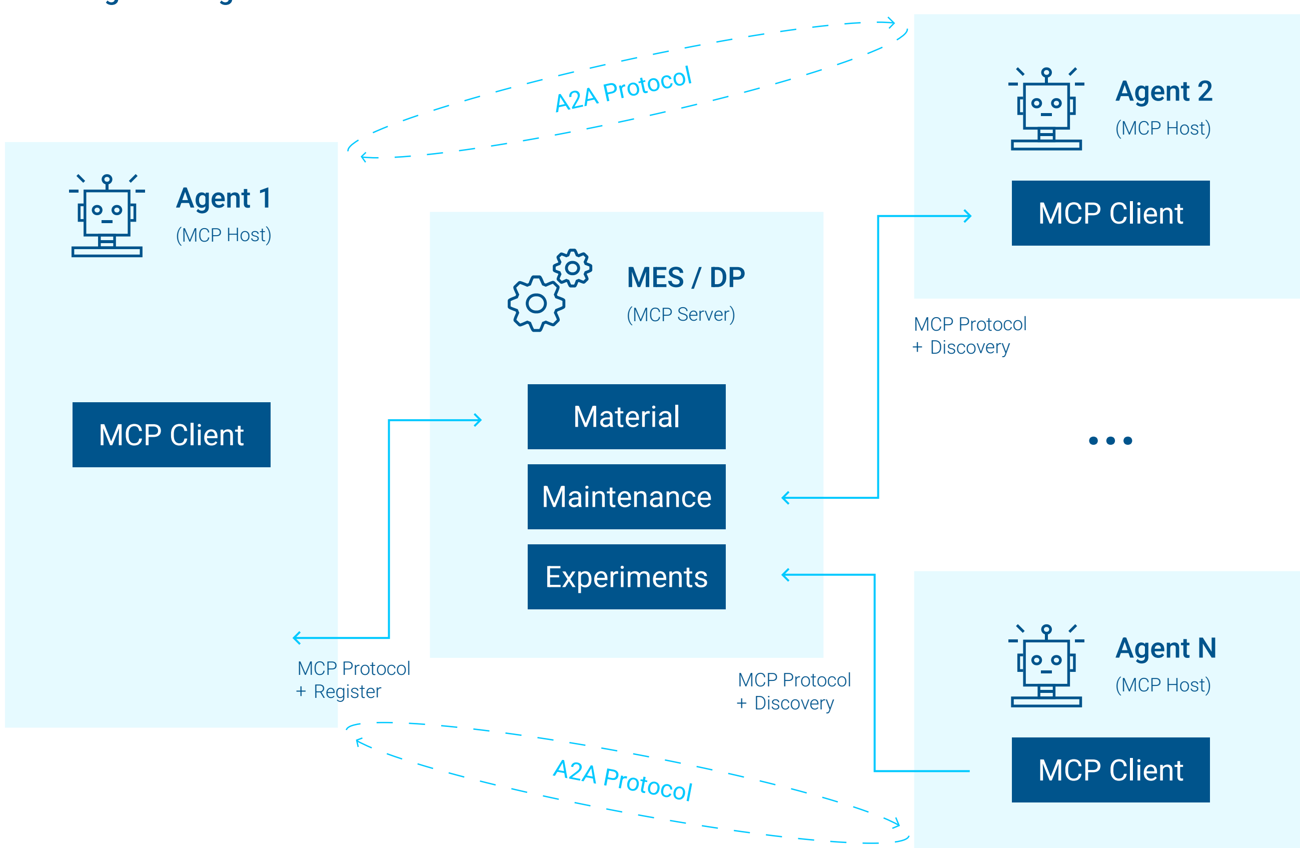

Figure 11: A24 Agent to Agent Protocol

MES Agency Levels: From Automation to Autonomy

MES has historically been rule-driven. The shift to agent-based systems introduces aspectrum of intelligence, or agency.

At the lowest level is static automation: fixed rules, no decision-making. Agentic workflows offer greater flexibility, adapting to context while still operating within guardrails. Autonomous agents go further, using reasoning, memory, and goals to make decisions without predefined logic. At the highest level, orchestrator agents oversee and coordinate other agents to achieve systemic optimization.

This evolution allows companies to move gradually from deterministic control to adaptive intelligence, gaining value at each step while maintaining operational safety and stability.

MES Agency levels

Figure 10: MES Agency Levels

The Learning Flywheel of MES AI Agents

AI agents differentiate themselves through learning. They execute actions, observe results, and adjust behavior based on outcomes. This creates a flywheel of continuous improvement.

For example, a scheduling agent may reroute jobs to reduce cycle time. If the change improves throughput, the agent reinforces that behavior. If it causes a delay, it adapts. Human feedback, like approvals, overrides, or operator comments, feeds into the agent's memory and helps refine future decisions.

This loop enables MES to evolve dynamically, improving over time without manual reprogramming. The system does not just automate tasks; it becomes increasingly competent at executing them

The learning flywheel of MES AI agents

Memory & Behavior Update

Learns from feedback and results to refine future decisions (adapts rules, ranking, or confidence thresholds)

Agent Takes Action

Executes a decision in MES (e.g., reschedules job, rereoutes material)

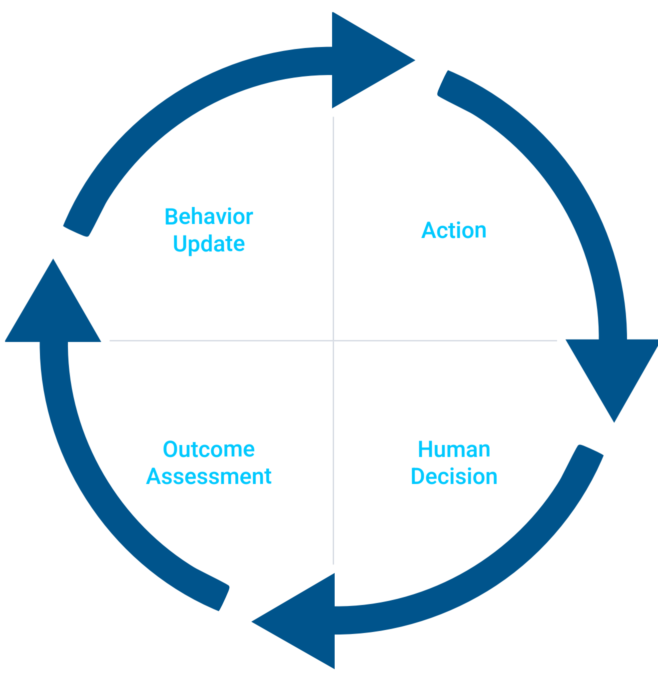

Outcome Assessment

Agent evaluates effectiveness of its decision based on KPIs (e.g., yield, downtime, throughput impact)

Human Overrides or Confirms

Operators review, override, or approve agent decisions

Figure 13: The Learning Flywheel of MES AI Agents

Safety Mechanisms for AI Agents in Manufacturing

Autonomy in manufacturing must be earned. AI agents must operate within strict safety and governance constraints.

Policy enforcement ensures agents stay within predefined boundaries, for example, never bypassing a safety check or altering validated procedures. Human-in-the-loop oversight allows operators to approve or override agent decisions, especially in critical scenarios. Decision transparency means every agent action is traceable: what data it used, why it acted, and what it expected to achieve.

Together, these mechanisms allow AI agents to act with autonomy while remaining accountable and aligned with operational goals and regulatory requirements.

AI agents safety mechanisms

Figure 14: AI Agents Safety Mechanisms

Strategic Takeaways for Manufacturing Executives

The shift to agent-based MES is not a question of possibility; it is a question of readiness. Success requires a foundational investment in data architecture, integration, and governance.

Executives must start by aligning automation, MES, and data platforms around a shared semantic model. From there, pilot agentic workflows within contained, high-value use cases. Build feedback loops early, especially those involving operators and managers, and embed safety controls from day one.

This is not an IT upgrade. It is an operating model shift, from control to collaboration, from configuration to cognition. Those who act early will lead the next era of industrial intelligence.

The transition from traditional MES to agent-based, intelligent systems is more than a technical evolution; it is the beginning of a new manufacturing paradigm. As AI agents become more capable, more connected, and more embedded into operational logic, factories will no longer require rigid orchestration from the top down. Instead, intelligence will begin to emerge from the system itself.

These agents will anticipate failures, optimize performance across objectives, and learn from every shift. They will not replace people, but they will change what people need to focus on. Operations will move from intervention to oversight, from reaction to improvement.

This is the Thinking Factory: a system that not only runs, but reasons. That adapts. That improves. And that unlocks the full potential of AI not as a tool, but as a native feature of manufacturing itself.

From business rules to agents: evolving the MES stack

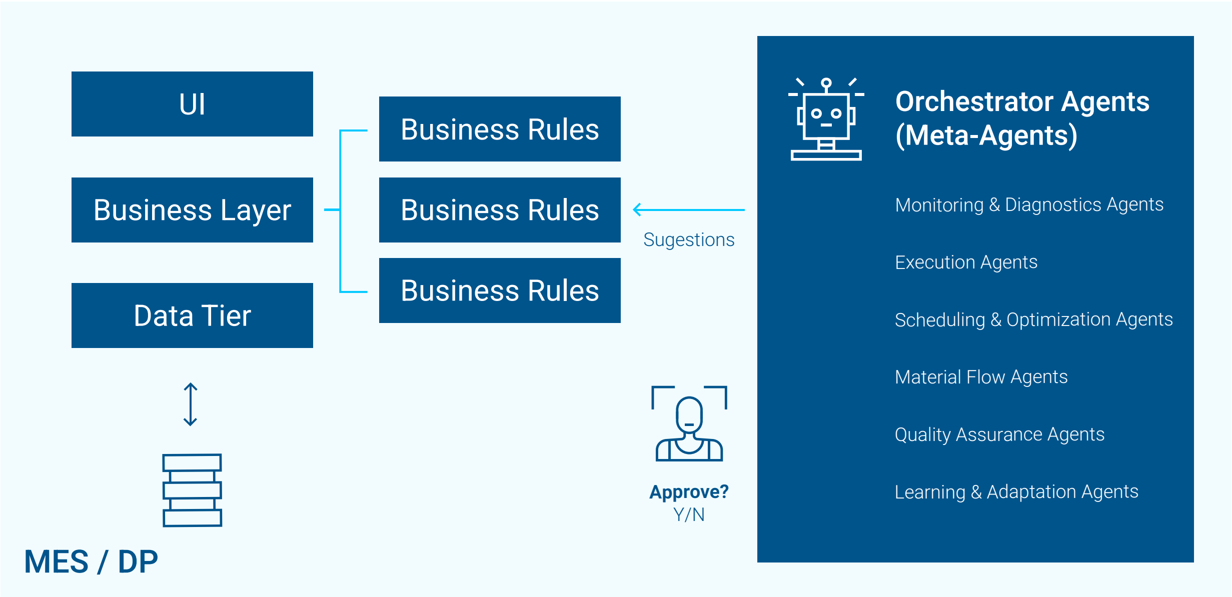

Figure 15: From Business Rulesto Agents: Evolving the MES Stack

ABOUT THE AUTHOR

Francisco Almada Lobo is recognized as a top strategic thought leader and evangelist on digital transformation, specifically Industry 4.0, manufacturing operations and the factories of the future. Francisco co-founded Critical Manufacturing in 2009 and has been CEO since 2010.

Francisco started his career in a CIM R&D Institute and joined Siemens Semiconductor in 1997. Throughout his tenures at Siemens, Infineon and Qimonda, he specialized in optimizing highly complex, discrete manufacturing operations. In 2004, he led the first migration of an MES system in a running high-volume facility.

Francisco holds various positions within the smart manufacturing and venture capital industries, including being a Member of the 200M Fund's Investment Committee, Executive Committee Member of SEMI Smart Manufacturing Technology, Member of the Forbes Technology Council and Advisor to many Industry 4.0 startups.

ABOUT CRITICAL MANUFACTURING

Critical Manufacturing provides the most modern, flexible and configurable manufacturing execution system (MES) available. Critical Manufacturing MES helps manufacturers stay ahead of stringent product traceability and compliance requirements; reduce risk with inherent closed-loop quality, integrate seamlessly with enterprise systems and factory automation and provide deep intelligence and visibility of global production operations. As a result, customers are Industry 4.0-ready. They can compete effectively and profitably by easily adapting their operations to changes in demand, opportunity or requirements, anywhere, and at any time.

To learn more visit: www.criticalmanufacturing.com