

From Enterprise Knowledge to Agentic AI: A Framework for AI Investment Decisions in Manufacturing

Takeaways

Not all AI delivers the same value. The way AI is implemented determines both what it can do and the investment required. Understanding the architectural differences between single-system AI, cross-system knowledge retrieval, and agentic workflows is the starting point for any informed investment decision.

The fundamental trade-off is effort vs. data scope. Out-of-the-box AI from software providers is limited to one system's data. Cross-system knowledge retrieval opens the AI's aperture to the full enterprise, and that is where the real productivity gains begin.

The knowledge foundation is the step most organizations underestimate. Cross-system knowledge retrieval is the prerequisite for anything truly agentic. Organizations that skip this step and jump directly to agentic AI typically find themselves in pilot purgatory.

Trustworthy AI requires more than good models. Security, reliability, and governance become increasingly critical as AI gains autonomy. When an AI agent can plan, decide, and act independently, the consequences of getting trustworthiness wrong escalate.

AI adoption is a progression, not a single step. Most organizations will run multiple AI implementation approaches simultaneously. Sequencing investments wisely, starting with the knowledge foundation, maximizes value and minimizes risk.

Introduction

The promise of AI in manufacturing and engineering has never been louder, or more confusing. Nearly every PLM and ERP provider has introduced AI capabilities in the last 18 months, each using slightly different language to describe what they are delivering. One calls it a copilot, and another an agent. A third says they have integrated retrieval-augmented generation (RAG) into their software platform. For business leaders trying to make investment decisions, the result is a market where everything sounds the same but can mean something quite different.¹

CIMdata works with industrial organizations across the spectrum of AI maturity, and we see four friction points come up in nearly every conversation (Figure 1).

¹ Research for this paper was supported by ChapsVision.

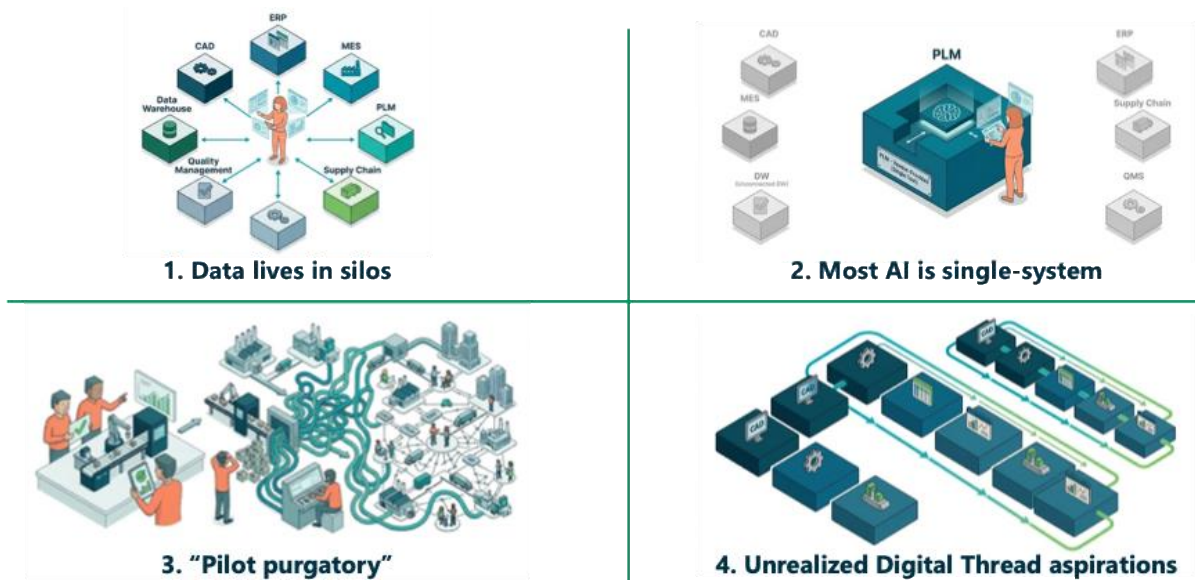


Figure 1: The Reality Gap: Four Friction Points CIMdata Observes Across Manufacturing Organizations Adopting AI

First, data lives in silos. The average engineer’s work touches five to ten different systems on any given day. The AI available through each of those systems can only see one. A survey of 300 engineering leaders conducted by SimScale and Global Surveyz found that 93% expect AI to deliver meaningful productivity gains, but only 3% report achieving significant impact to date.² The top barrier, cited by more than half of respondents, was siloed data and legacy tools not built to share information.

Second, most deployed AI remains single system. Your PLM solution provider may offer a chatbot. Your ERP provider may have a copilot. Each is useful within its own boundaries, but none can see across the boundaries that define an engineer’s actual workday.

Third, organizations are getting stuck between pilot and production. The proof of concept works in the lab but scaling it across the enterprise is where initiatives stall. An MIT study analyzing 300 public AI deployments found that 95% of enterprise generative AI pilots fail to deliver measurable business impact, and the core issue the report identifies is not model quality but poor data integration.³

Fourth, the digital thread remains unrealized for most organizations. The aspiration of a connected information backbone spanning design through manufacturing through service has been on roadmaps for a few decades now. Most organizations have made real progress, but the full vision remains elusive. AI does not make this aspiration go away. It makes it more urgent, because AI is now both a consumer of connected enterprise data and, increasingly, a potential enabler of the digital thread itself.

The confusion facing decision-makers is not a knowledge gap. It is a categorization gap. The industry has not given them a consistent vocabulary for distinguishing between fundamentally different AI approaches. CIMdata has identified five distinct implementation patterns for AI in PLM and manufacturing environments to address precisely this problem. This commentary focuses on the three patterns most relevant to the current enterprise AI conversation: single-system AI (Pattern 1), cross-system knowledge retrieval (Pattern 2), and agentic AI workflows (Pattern 5). For the complete framework, including Patterns 3 and 4, see CIMdata’s educational webinar, “The Five AI Implementation Patterns in PLM: A Framework for Informed Decision-Making,” [available on CIMdata.com](https://www.cimdata.com).

² [The State of Engineering AI 2025, SimScale/Global Surveyz](#)

³ [The GenAI Divide: State of AI in Business 2025, MIT NANDA Initiative](#)

The AI Implementation Landscape

The most common AI experience for engineers and manufacturing professionals today is what CIMdata classifies as Pattern 1: AI delivered by existing software providers. In this pattern, the provider handles the entire AI pipeline, including data preparation, retrieval, and model interaction, all within the boundary of their system. A company configures and uses it. It works well for searching within their delivered system, summarizing documents, and getting recommendations based on the data it manages. It is genuinely useful, and it is rapidly becoming table stakes among leading PLM and ERP providers.

But the boundary matters. Everything the AI can see is inside one software provider's perimeter. An engineer's work, however, does not respect those perimeters. On any given day they may need information from PLM, ERP, quality systems, supplier portals, test databases, service records, and numerous legacy repositories that have been around for decades. Asking single-system AI for help is like asking a colleague to assist but telling them they can only look in one filing cabinet. They will give a good answer based on what they can see, but they are likely to miss most of the picture. This is one of the main architectural reasons AI pilots fail to scale (Figure 2).

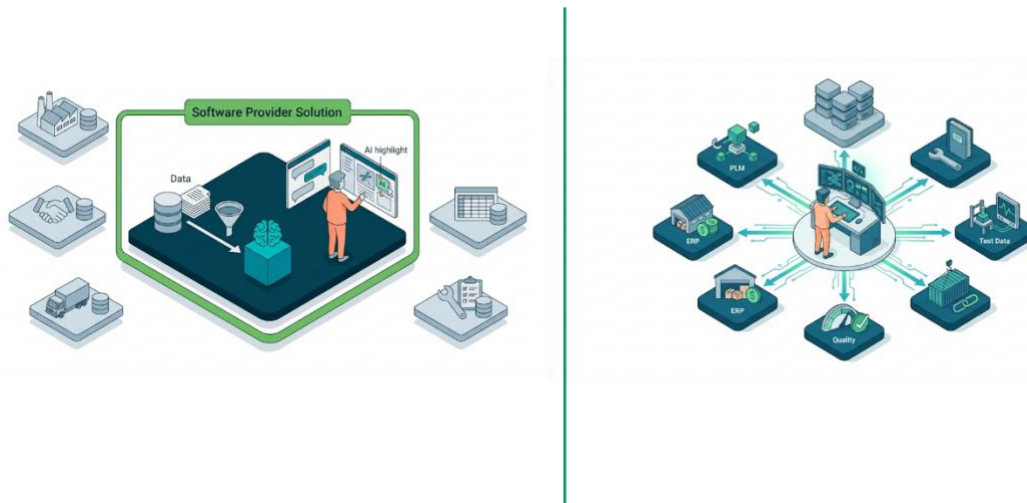


Figure 2: Single-System AI vs. Cross-System Knowledge Retrieval: Breaking Through the Data Boundary

Cross-system knowledge retrieval (Pattern 2) breaks through this boundary. Instead of one software provider owning the AI pipeline inside their system, the organization takes ownership of a retrieval layer that spans the entire enterprise: PLM, ERP, supply chain, quality systems, shared drives, and potentially every system in between. The underlying mechanism is Retrieval-Augmented Generation, or RAG. In practical terms, RAG automates the manual process most knowledge workers already perform (i.e., finding relevant documents across multiple systems, extracting the key passages, and feeding them to an AI model to get a grounded, cited answer). The difference is that the system does this in seconds across the full enterprise (i.e., all those systems under the RAG umbrella), rather than requiring an engineer to spend hours searching across disconnected systems and their associated databases and repositories.

The term "RAG" covers a wide range of sophistication—not all RAG implementations are equal. A basic, or "naive," RAG implementation splits documents into chunks, creates mathematical embeddings, and retrieves based on similarity search. This works at small scale but struggles with large, heterogeneous enterprise document collections. More advanced implementations combine vector search with keyword retrieval, natural language processing, and multimodal analysis to improve accuracy and comprehensiveness. The distinction between naive and advanced RAG is not academic—it is often the difference between a successful enterprise deployment and a stalled pilot.

For most organizations, cross-system knowledge retrieval is where the highest-impact, near-term AI investment lies. When an engineer can ask a question and get an answer drawing from test data, design specifications, supplier documentation, and service records simultaneously, the impact on decision quality and speed is substantial. But it is important to be precise about what Pattern 2 does and does not do. Cross-system retrieval is still fundamentally a question-and-answer interaction. The human asks, and the AI finds, synthesizes, and then responds.

Agentic AI (Pattern 5) refers to systems that can plan, decide, and act autonomously toward a goal. These systems operate on a fundamentally different model. The user provides a goal, not a question. The AI orchestrator dynamically determines what steps to take, what tools to call, what data to retrieve, and what sequence to follow, adjusting based on what it finds along the way. This is the core distinction between retrieval and agency—with retrieval, you ask and the AI answers, and with agentic AI, you direct and the AI works.

Two points are critical. First, agentic AI requires a cross-system knowledge foundation. An agent that needs to reason across the enterprise needs access to enterprise knowledge. Organizations that skip Pattern 2 and jump directly to Pattern 5 are the ones CIMdata sees stuck in pilot purgatory six months later. Second, trustworthiness becomes exponentially more important as AI gains autonomy. Traditional software follows a script. An agent makes decisions, and this introduces non-deterministic behavior that enterprises have not had to manage in software before. Even a small per-step error rate compounds rapidly across multi-step agentic workflows. Governance for agentic AI must itself be automated. This is because manual oversight cannot keep pace with agents operating at machine speed.

These patterns connect directly to the digital thread. AI raises the stakes for connected enterprise data because it is no longer just an IT aspiration—it is the substrate on which AI reasons. The shift from Pattern 1 to Pattern 2 changes who performs integration. In Pattern 1, the user manually connects information across systems. In Pattern 2, AI takes on that role. When mature, the retrieval layer becomes a persistent, governed knowledge fabric, and that is the point at which the digital thread shifts from aspiration to operational reality. Pattern 5 takes this further: agents do not just retrieve from the knowledge fabric, they act on it, executing multi-step tasks that span systems and disciplines. At that point, the digital thread is no longer just connected information; it is the operational backbone of autonomous work.

Practical Guidance for AI Investment

When evaluating AI solutions for manufacturing and engineering environments, CIMdata recommends four key questions as a starting framework:

1. **Knowledge scope:** does the AI see across the enterprise or only within one system? This immediately tells you whether you are looking at single-system AI or something broader.
2. **Architectural clarity:** does the AI retrieve and summarize, or does it plan, decide, and act? Both have value, but they are fundamentally different capabilities with different prerequisites and risk profiles. The market is rife with what has been called “agent washing,” where solutions that are essentially chatbots or assistants get marketed as agentic AI.
3. **Trustworthiness:** are security, accuracy, and auditability enforced? The more autonomy the AI has, the more these dimensions matter.
4. **Organizational readiness:** are data governance, change management, skills, and clear policies on AI authorization in place? This is the dimension that does not show up in product demos but determines whether the implementing organization receives value at scale.

On sequencing, the practical recommendation is straightforward. Most organizations already have some version of Pattern 1 through their existing software providers. The highest-impact next step is Pattern 2: building the cross-system knowledge retrieval layer. Organizations that have done this well have typically stood up working systems in months, not years. The heavy lifting is in data preparation and integration, not the technology itself. Agentic AI (Pattern 5) is more involved, and its requirements are cumulative. Ultimately, the knowledge foundation has to come first.

On the build-vs.-buy spectrum, organizations face a range of options. At one end, you can build RAG pipelines from open-source libraries if you have the engineering talent, time, and appetite for ongoing maintenance. In the middle, cloud platforms such as AWS Bedrock, Google Vertex AI, and Azure AI Foundry provide configurable infrastructures that require less custom development but still demand hands-on setup, parameter tuning, and interface development. At the other end, purpose-built platforms handle the retrieval, security, and agent orchestration layers as integrated capabilities, often with optimizations (such as hybrid retrieval strategies) that are difficult to replicate independently. Each position on the spectrum involves trade-offs between control, speed to production, and maintenance burden, and the right answer depends on the organization's specific situation.

ChapsVision in Context

ChapsVision's capabilities map to the two dimensions that determine whether enterprise AI agents succeed—what the agent knows and how the agent behaves.

On the knowledge side, Sinequa, ChapsVision's enterprise knowledge platform, provides the cross-system retrieval foundation (Pattern 2). The platform connects to more than 200 enterprise applications, including PLM solutions such as Windchill and Teamcenter, and employs a hybrid retrieval approach that goes beyond simple vector search to combine statistical retrieval, natural language processing, semantic embeddings, and multimodal understanding. Security and access control enforcement happen at the retrieval layer itself, ensuring that agents only access information they are authorized to see. This positions Sinequa as the knowledge foundation on which agentic capabilities depend.

On the behavior side, ChapsAgents provides the orchestration, deployment, and governance layer for enterprise AI agents (Pattern 5). The platform includes a no-code agent builder, an agent catalog for discovery, a workflow builder for cases where agents should follow predefined processes, and governance capabilities spanning monitoring, guardrails, and human-in-the-loop checkpoints. The architecture is large language model (LLM)-agnostic and deployment-agnostic and supports emerging interoperability protocols including Model Context Protocol (MCP) and Agent-to-Agent (A2A) (Figure 3).

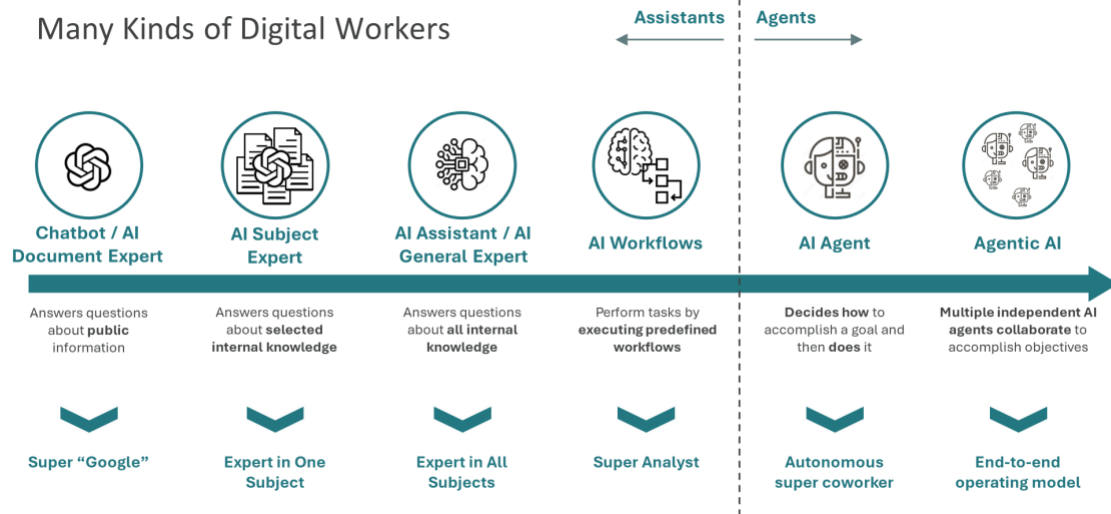


Figure 3: The Digital Workers Spectrum: Matching Architecture to Use Cases
(Courtesy of ChapsVision)

As an example, a global automotive powertrain manufacturer has embarked on an enterprise-scale agentic AI deployment using ChapsVision’s platform, and the approach illustrates several principles discussed in this commentary. They began with the knowledge foundation, integrating multiple source systems (i.e., engineering, manufacturing, and operations) into a unified enterprise knowledge base built on the digital thread concept. They deployed Sinequa’s hybrid RAG, going beyond naive vector search to combine multiple retrieval methods for more accurate and comprehensive context. They matched architecture to use cases, recognizing that not every problem requires a full agent. Initially, the organization identified approximately 50 use cases and grouped them by required capability level, from simple document experts to fully agentic workflows. Importantly, they built governance in from the start, not as an afterthought.

Among the agentic use cases were a drawing compliance agent that cross-references technical drawings against applicable standards and identifies non-compliance, a design review agent that compiles supporting evidence from across enterprise systems into a comprehensive review package, and a test plan agent that identifies applicable requirements and generates efficient, non-redundant test plans informed by historical data and part provenance.

Conclusion and Recommendations

The AI implementation landscape in manufacturing is noisy, and the terminology is inconsistent enough that informed decision-making requires a deliberate framework. These are CIMdata’s independent recommendations for organizations navigating this landscape:

- **Invest in the cross-system knowledge foundation now.** For organizations still operating primarily with single-system AI from existing software providers, building the enterprise knowledge retrieval layer is the highest-impact next step and the prerequisite for anything truly agentic.
- **Be precise about what you want to buy.** Understand whether a solution retrieves and summarizes or plans and acts. Misaligned expectations between retrieval and agency are a common source of failed implementations.
- **Treat trustworthiness as a prerequisite, not an afterthought.** Security, reliability, and governance become more critical with every increment of autonomy. Automated governance is essential for agentic deployments.

- **Assess organizational readiness alongside technology readiness.** Data governance, change management, and clear AI authorization policies determine whether technology investments deliver value at scale.
- **Include ChapsVision in your evaluation** for enterprise knowledge and agentic AI. Their platform addresses both the knowledge foundation (through Sinequa’s enterprise search and hybrid RAG) and the agentic layer (through ChapsAgents’ orchestration and governance), making them a relevant option for organizations pursuing the implementation path described in this commentary.

The winners in enterprise AI will not be those with the best models, but those with the best-connected knowledge.

About CIMdata

CIMdata, a global strategic management consulting firm, provides services designed to maximize an enterprise’s ability to design, deliver, and support innovative products and services. For more than forty years, CIMdata has provided industrial organizations, providers of digital technologies and services, and investment firms with world-class insight, expertise, and best-practice methods on a broad set of product lifecycle management (PLM) topics and the digital transformation they enable. CIMdata also offers research, subscription services, publications, and education through certificate programs and international conferences. To learn more, visit www.CIMdata.com or email info@CIMdata.com.