

Bridging the great AI agent and ERP divide to unlock value at scale

Resources focused on AI are coming at the expense of enabling ERP to provide the system capabilities AI needs to thrive.

*by Bjørnar Jensen, Darwin Deano, and Michael Allison
with Talha Bin Asad*



Investments are flowing into AI at unprecedented levels. The headlines tout massive investments in AI and AI-adjacent companies. This pace of activities is also manifesting at the business level, as companies shift budgets from IT functions into AI. Our latest research shows that almost half of all IT organizations surveyed are planning to invest in gen AI initiatives, with levels of investment dropping significantly for core IT capabilities such as infrastructure and architecture.¹

At one level, this shift makes sense as companies move resources to AI to take advantage of the incredible opportunities. But at another level, this move is creating a “great divide” in which companies are focused on AI at the expense of the supporting enterprise resource planning (ERP) system capabilities, many of which are critical in enabling AI.

This emerging divide is having negative consequences. The experimentation with AI (and gen AI more specifically) has led to a proliferation of use cases and experiments that are unsupported by the underlying end-to-end processes, data, people, and technologies that enable these cases to scale (in other words, “pilot purgatory”). The numbers bear this out: Only about 40 percent of companies report any enterprise-level EBIT impact from their AI initiatives.²

With the ongoing excitement around AI agents, ERP applications are often treated as an afterthought and considered an unwieldy legacy technology. This “ugly stepchild” attitude toward ERP dangerously undervalues its importance in the AI conversation. Not only do AI use cases rely on much of the data and many of the applications that are housed in ERP systems, the end-to-end nature of workflow transformation that drives much of the potential value of AI agents requires thoughtful integration with the ecosystem of ERP capabilities (Exhibit 1).

ERP’s value as a core enabler for scaling agents is particularly the case in highly scaled transactions in which the workflow logic is well defined. AI agents become an extension of the ERP, especially in long-tail (the extensive set of traditionally lower-value opportunities) and exceptional cases where high manual effort is required.

This reality is not to imply that ERP systems can evolve along traditional “business-as-usual” pathways. While it is unlikely that AI agents will replace the ERP in the near or medium term because of system complexity, companies should consider not only how AI agents will disrupt ERP operations but also how they provide a powerful capability for evolving and modernizing ERP itself.

The ERP foundation to help unlock AI value at scale

AI is projected to generate between \$17 trillion and \$26 trillion in global economic impact (Exhibit 2).³ The majority of CEOs and executives recognize this value potential—about 80 percent of companies report using gen AI in at least one function.⁴ Our latest global AI survey shows that about 40 percent of organizations report EBIT impact from AI, though most attribute less than 5 percent EBIT impact. Those achieving 5 percent or more EBIT impact from AI report

¹ See the forthcoming article, “The intelligent enterprise: How tech leaders can rewrite the rules of growth,” publishing in February 2026.

² “The state of AI in 2025: Agents, innovation, and transformation,” McKinsey, November 5, 2025.

³ “The economic potential of generative AI: The next productivity frontier,” McKinsey, June 14, 2023.

⁴ “The state of AI in 2025: Agents, innovation, and transformation,” McKinsey, November 5, 2025.

Exhibit 1

Most business domains enable key processes either directly or partially with enterprise resource planning as a core system.

Enterprise resource planning system as core

Core system	Partial	Not core		
Finance and accounting Record to report Order to cash Procure to pay Plan to perform	Supply chain and logistics Demand to supply Deliver to cash Logistics execution	Operations Plan to produce Make to deliver Deliver to customer	Customer service and success Request to resolution Return to refund	Business development Identify to contract Partner to revenue
Procurement and purchasing Source to settle Procure to pay Supplier management	Human resources Hire to retire Recruit to onboard Develop to reward	Project management Initiate to close Plan to deliver	Health, safety, and environment Incident to prevention Risk to mitigation	Quality management Detect to correct Nonconformance to action
Learning and development Plan to learn Train to certify	Information technology Request to resolve Incident to resolution Change to deploy	Product management and R&D Initiate to close Plan to deliver	Sustainability and environmental, social, and governance Measure to report Monitor to improve	Sales Quote to cash Customer management Lead to order
Marketing Lead to conversion Campaign to lead Market to opportunity	Legal and compliance Risk to resolution Audit to report Contract to close	Data and analytics and business intelligence Incident to prevention Risk to mitigation	Investor relations Plan to communicate Report to market	Corporate strategy and planning Plan to execute Project to benefit Initiative tracking
Risk management and internal audit Control to certify Risk to mitigation	Change management and organization development Plan to adopt Change to sustain	Innovation and digital transformation Idea to execution Prototype to scale	Corporate communications and PR Message to market Crisis to recovery	Corporate social responsibility Plan to impact Engage to report

McKinsey & Company

pushing for redesigning workflows, scaling faster, and investing more aggressively in transformative AI.⁵ The AI potential is vast, but most companies struggle to translate it into tangible business outcomes.

Our latest research shows that high performers are much more likely than others to have taken AI agents to the scaling phase.⁶ One of the key actions they take, as highlighted in our book *Rewired*, is focusing the transformation at the domain level—in other words, a function or journey (Exhibit 3). Focusing on the domain allows companies to address all the interrelated use cases so holistic change and scale can happen.

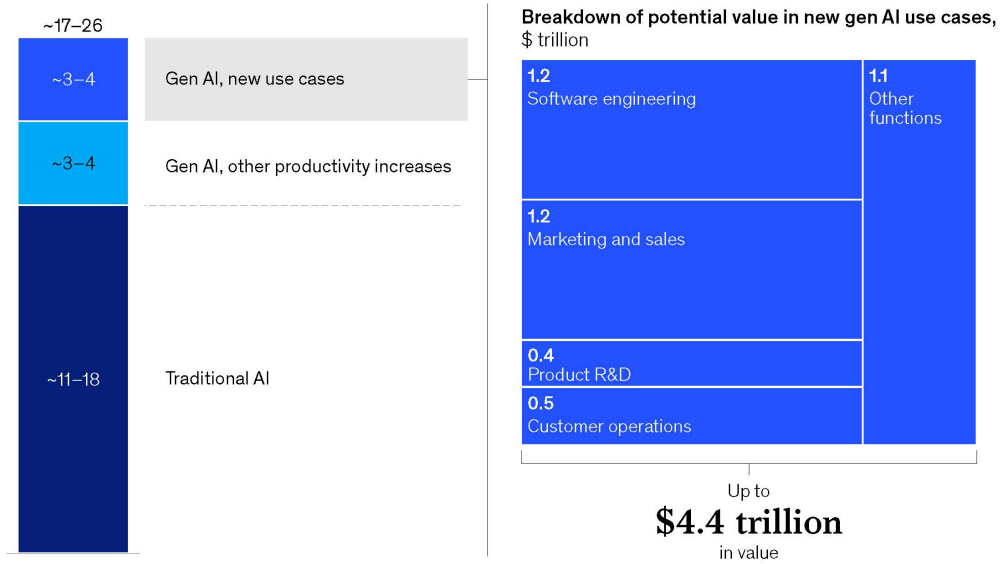
⁵ “The state of AI in 2025: Agents, innovation, and transformation,” McKinsey, November 5, 2025.

⁶ “The state of AI in 2025: Agents, innovation, and transformation,” McKinsey, November 5, 2025.

Exhibit 2

AI offers significant value potential.

AI's potential impact on global economy, \$ trillion



Source: "The economic potential of generative AI: The next productivity frontier," McKinsey, June 14, 2023

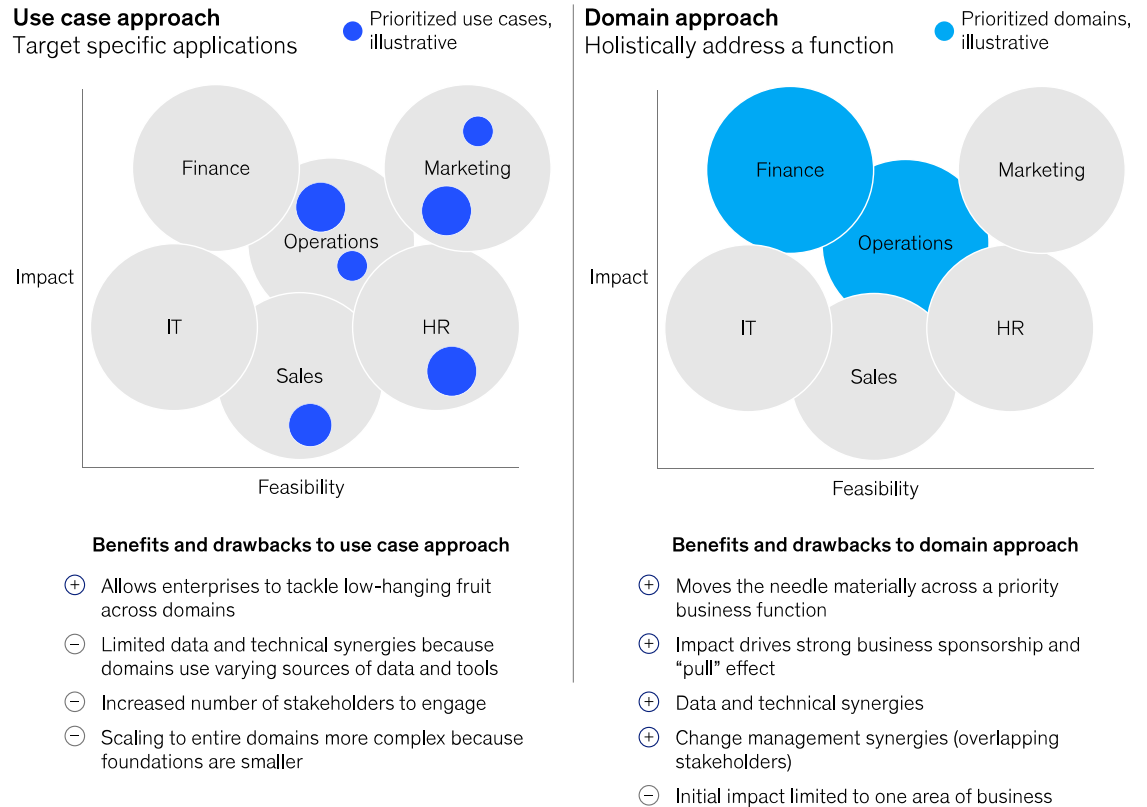
McKinsey & Company

Business domains bring together people, processes, and data in the flows that determine how a company actually operates. These domains cannot run or transform without ERP, which defines how value moves through the business: It structures the data, sets the rules, and organizes the workflows that power daily operations.

Many leaders focus only on ERP technical debt, forgetting about ERP's "equity"—the deep process knowledge, clean data structures, and built-in business logic that represent the company's operating DNA. These capabilities are the fuel that powers AI in business.

Exhibit 3

An AI domain transformation approach ensures a cohesive AI glide path and helps avoid throwaway or unnecessary work.



McKinsey & Company

How to harness ERP to help power AI

ERP’s role in enabling AI transformations will vary by company and company strategy, as might be expected. But in our experience, successful cases broadly follow the same playbook.

Clarify the value at the workflow level

Identify the most important problems in a business domain and tie them to clear, measurable outcomes—such as margin, cost, service levels, or working capital. As our colleagues have highlighted in other articles,⁷ addressing these opportunities for AI requires them to drill down to

⁷See, for example, “The change agent: Goals, decisions, and implications for CEOs in the agentic age,” *McKinsey Quarterly*, October 1, 2025; Lareina Yee, Michael Chui, Roger Roberts, and Stephen Xu, “One year of agentic AI: Six lessons from the people doing the work,” *McKinsey*, September 12, 2025.

the workflow level, not at the use case or tool level. For each priority AI workflow (such as dynamic inventory allocation, intelligent sourcing, or AI-assisted production planning), work backward from the decision the AI should make and list the specific ERP elements that it depends on: which primary data (materials, plants, customers, suppliers), which transactions (orders, deliveries, purchase orders, production orders), which events (stock changes, delays, confirmations), and which configuration or business rules (lead times, lot sizes, approval limits).

Buy or build?

As organizations embed AI deeper into their organization, they face a strategic question: What should be bought versus built? In a space evolving this quickly, waiting for a perfect off-the-shelf solution or attempting to custom-build everything can both become high risk. AI changes too fast for multiyear build cycles, and re-creating foundational layers can drain resources without creating differentiation.

A more resilient approach is to buy standardized capabilities—such as embedded approval agents, predefined data products, ERP-integrated orchestration frameworks—and reserve custom development for the select areas where domain-specific logic or proprietary workflows create real competitive advantage. The decision is also not one and done; unlike the traditional software as a service (SaaS) world, AI requires a continuous, deliberate reassessment of what to create and what to consume. Enterprises that get this balance right avoid fragmentation, reduce long-term maintenance burdens, and accelerate value realization.

The most effective way to do this is to bring together domain experts, ERP functional experts, and AI practitioners in short, structured working sessions. These sessions are practical and specific; the team “walks” along the target AI workflow step by step and explicitly marks which ERP tables, fields, and processes must be accurate, available, and exposed to AI to run and scale (see sidebar, “Buy or build?,” to review how to think through resourcing agentic work).

Define your terms to ensure accuracy

A shared ontology (basically a shared map of how your business defines technologies) grounded in ERP is essential because it gives AI one consistent set of data definitions, process logic, and business rules to operate on. This is how to ensure AI decisions are accurate, aligned with how the business operates, and scalable across the enterprise. Rather than creating every component of this ontology from scratch, enterprises can leverage existing, well-defined ERP data products with custom extensions that can accelerate ontology development (for example,

SAP Business Data Cloud platform). Focusing on ontology also implies less focus on traditional ERP data warehousing and data replication, one way that AI will drive the evolution of ERP.

Embed agentic capabilities inside workflows

Placing AI directly inside the steps where work gets done—approvals, planning, recommendations, forecasting, and exception handling—is crucial for adoption. It also helps AI agents to perform faster, smarter, and more reliably because AI is applied exactly where decisions are made and work gets done rather than sitting off to the side as a separate tool.

Modern ERP platforms make this easier: For example, SAP provides ready-made AI agents that can slot into cross-functional “agent squads.” By integrating AI and ERP, your agentic workflows are grounded in best practices for process and data, effectively safeguarding process and data integrity while unleashing the power of agents. At the same time, the integration of AI represents a departure from the traditional software-as-a-service (SaaS) ERP solution approach.

Balance flexibility with stability in architecture

The goal is to balance flexibility with stability. Use open components where customization or innovation is needed, and use your ERP and cloud platforms where scale, reliability, and security matter most. A thoughtful architecture prevents fragmentation, one of the most common reasons AI programs stall.

To connect ERP systems to AI-enabled workflows, begin by making sure the AI can access the right data and processes at the right moment. This means exposing clean, structured ERP data—such as orders, inventory, supplier information, or production schedules—through standard data services or APIs. It also means giving the AI a controlled way to feed decisions back into ERP by using existing action interfaces or workflow triggers. In this way, when the AI recommends a change, it can update the operational system directly and safely.

It is important to wrap the end-to-end steps into a single workflow that links ERP events, AI logic, and business actions. An orchestration layer sequences the flow—pulling ERP data, sending it to the AI, receiving the recommendation, and writing the result back into ERP—and event triggers from the ERP ensure the AI runs only when something meaningful happens, such as a stock change or a supplier delay. Together, this creates a smooth, responsive workflow in which ERP and AI operate as one system, with intelligence applied exactly where work happens.

A key decision is whether you need a separate data platform or whether you can leverage an existing one. Many ERP and cloud platforms now offer built-in data services and integration frameworks, which can simplify integration with major cloud providers and often reduce cost and complexity.

Measure and refine the system

Organizations need a way to measure and refine the system continuously. Consider setting up a “value mission control,” a small team and dashboarding setup that continuously tracks how AI-enabled workflows are performing, links process metrics back to business value, and quickly flags where tuning or fixes are needed so the impact keeps growing instead of fading.

Many process-mining and ERP analytics platforms offer catalogs of operational performance indicators that can help map to higher-level value levers. These can serve as a starting point for building value traceability, if organizations tailor them to the specific metrics and outcomes they care about (for example, SAP Signavio or Celonis). This approach enables organizations to maintain momentum in a fast-changing landscape, starting with the value thesis and anchoring AI initiatives in business logic and measurable outcomes rather than isolated proofs of concept.

Key considerations for moving forward

The targeted modernization of the ERP to support AI transformations is complex, and as such, leaders should look to do the following:

- **Make ERP a core part of the AI conversation.** If managing the ERP for an AI domain transformation is delegated to IT and forgotten, it will struggle. CIOs and CTOs must elevate ERP from a back-office system to a strategic enabler. A key element in pulling that off requires them to explicitly tie all ERP initiatives to not only AI initiatives but also the resulting value opportunity for the business. This point is crucial in shifting the mindsets that ERPs are “old-school back-office systems” to important enablers of successful AI transformations. Make sure that ERP is part of meaningful strategic, planning, and execution conversations.
- **Develop a risk management strategy now.** To manage risk in ERP-enabled agentic AI transformations, organizations must address both traditional system-integration risks and new AI-specific challenges (for example, autonomous decision-making, data-quality sensitivity, and model drift). The most important mitigations are establishing tight human-in-the-loop governance for high-impact decisions—identify who those people are and what the high-impact decisions will be. Spend sufficient time implementing robust data controls and logging that track every AI-initiated action within the ERP environment. Given the unpredictability of agentic solutions for the foreseeable future, it’s important to properly resource testing capabilities.
- **Be clear on the unit economics of change and rigorously track P&L impact.** With rising pressure on margins and productivity, AI investments must show direct, traceable impact on the P&L. Ensure ERP elements of the work have specific KPIs that link to measurable impact. Be sure to account for the higher change management costs associated with AI programs than would normally be included in an ERP business case. As a rule of thumb, our experience has shown that for every \$1 cost in developing a model, you need to spend \$3 in change management.⁸

⁸ “Moving past gen AI’s honeymoon phase: Seven hard truths for CIOs to get from pilot to scale,” McKinsey, May 13, 2024.

Find more content like this on the
McKinsey Insights App



Scan • Download • Personalize



Unlocking AI's full potential requires treating ERP not as legacy baggage but as a key enabler that makes intelligence scalable, safe, and valuable. Companies that close the divide between AI ambition and ERP readiness will move fastest from experimentation to real, defensible P&L impact. Those that don't will continue to watch AI promise outpace AI performance.

Bjørnar Jensen is a senior partner in McKinsey's Zurich office; **Darwin Deano** is a partner in the Chicago office, where **Talha Bin Asad** is an associate partner; and **Michael Allison** is a partner in the Montréal office.

Copyright © 2026 McKinsey & Company. All rights reserved.