

Agentic AI in ERP Systems: What Academic Research Tells Us (and What Vendors Don't)

A practitioner-researcher's dual reading of the agentic wave in Workday, SAP, and Oracle

Gabriel Aubert Desjardins

2026-04-24

Table of contents

When the ERP Stops Being a Tool and Becomes a Colleague	2
The Distinction That Changes Everything: <i>AI Agent</i> vs. <i>Agentic AI</i>	2
Workday, SAP, Oracle: Three Strategic Bets, Not One Race	3
Workday: The Unified Front Door of Work	3
SAP: Process Orchestration Above All	4
Oracle: The Industrialization of the Agentic Application	4
What This Divergence Means for You	4
What Vendors Won't Tell You (But Research Reveals)	5
The Cognitive Cost Is Real and Measurable	5
Not All Employees Respond the Same Way	5
Long-Term Effects Remain a Black Box	6
What Academia Doesn't See (But the Field Reveals)	6
Agentic Commercial Mechanics Are Invisible in Research	6
Sovereignty Does Not Exist in Academic Journals	7
The Devil Is in the Process	7
Four Questions to Ask Before You Sign	7
The Practitioner-Researcher Stance	8
Scientific References	9

Cet article est aussi disponible en français — [Lire la version française](#).

Workday launches Sana. SAP deploys Joule and announces 30+ agents for Q1 2026. Oracle introduces its Fusion Agentic Applications. They all talk about the same thing, but none of

them tell the same story. Here is what eighteen months of academic literature review and fifteen years of field practice taught me to hear behind the announcements.

When the ERP Stops Being a Tool and Becomes a Colleague

Two years ago, talking about an “AI assistant” in an ERP meant, at best, a chatbot capable of answering “how many vacation days do I have left?” Today, the three market leaders — Workday, SAP, Oracle — promise agents capable of generating a requisition, orchestrating an HR-IT onboarding, closing a financial month, or producing a budget variance report with minimal human supervision.

This is a genuine qualitative leap. And it is also a minefield.

On one side, the anticipated benefits are concrete and measurable. On the other, the blind spots in vendor discourse are systematic, and recent scientific literature is only beginning to name them. For an executive who needs to decide in 2026 what to do with this wave — embrace it, defer it, frame it — clarity comes neither from marketing decks nor from academic reviews in isolation. It comes from the intersection of both.

This blog post offers precisely that intersection. I will give you neither vendor hype nor academic skepticism. I offer a dual reading: that of a practitioner who deploys these platforms with clients, and that of a doctoral student working on this very subject.

The Distinction That Changes Everything: *AI Agent* vs. *Agentic AI*

Before going further, a conceptual detour is necessary. The words used by vendors today are vague enough to muddle the entire debate.

A [taxonomy published in *Information Fusion* in 2026 by Sapkota, Roumeliotis, and Karkee](#) proposes a structuring three-generation distinction.

Generation 1 — Generative AI. These are large language models like GPT-4, Claude, or Gemini. They produce text, code, summaries. They do not “do” anything inside your system.

Generation 2 — AI agents. Specialized modules capable of executing a specific task via an API call or integration. An agent that creates a leave request. An agent that reconciles an invoice. An agent that answers an HR question. Useful, but essentially reactive and task-specific.

Generation 3 — Agentic AI. Multi-agent systems capable of dynamically decomposing a complex task, retaining context from one day to the next, coordinating multiple specialized agents, and operating with significant autonomy in evolving environments.

This distinction is not an academic quirk. It has a direct operational implication. An AI agent that makes a single well-defined step is governed like an integration. An agentic system that orchestrates multiple subtasks, negotiates with other agents, and learns over time is governed like an employee. The nature of control, responsibility, and risk shift to a different level entirely.

When Workday presents Sana as “the new front door of work,” when SAP formally distinguishes its deterministic *Joule Skills* from its autonomous *Joule Agents*, or when Oracle introduces its *Fusion Agentic Applications* as “a new class of outcome-driven applications” — they are all describing, in their own vocabulary, this passage from Generation 2 to Generation 3.

Keep that in mind. It is the lens that makes the rest readable.

Workday, SAP, Oracle: Three Strategic Bets, Not One Race

Once this taxonomy is in place, you stop comparing features to features. You compare theories of change. And here, the three vendors diverge radically.

Methodological note: the comparative points below are based on public vendor communications and field observations at the time of writing. This landscape is moving quickly.

Workday: The Unified Front Door of Work

Workday has made two connected bets. The first: user experience takes priority over functional depth. With Sana, deployed globally in March 2026 as the platform’s new conversational front door, Workday no longer wants you navigating application screens. You talk to Sana, and Sana orchestrates the Workday, partner, and third-party agents needed to fulfill your request.

The second bet is more original. Workday treats its agents as entities to be managed in a dedicated system of record: the **Agent System of Record (ASOR)**. Lifecycle, identity, security, observability, cost, and impact. In other words: if Sana is the interface, ASOR is the HRIS of agents. This is an approach unlike either of the other two vendors, and it is probably the most aligned with recent scientific literature on the governance of agentic systems.

On the commercial side, Workday has also invented its own vocabulary with **Flex Credits** — a fungible annual credit consumption model that directly challenges the traditional per-seat licensing logic.

SAP: Process Orchestration Above All

SAP takes almost the opposite approach. Where Workday sells an experience, SAP sells **process depth**. Joule is anchored in the Business Suite and draws its strength from what SAP calls *process grounding* — the detailed business process knowledge accumulated over decades.

SAP's public architecture is also the most legible of the three. AI Foundation for models, GenAI Hub for LLM orchestration, Knowledge Graph for business semantics, HANA Cloud for data, and Joule Studio for building custom agents. All integrated into SAP BTP with centralized identity logic and permission propagation down to application-level calls.

SAP is also presented as highly open to third-party models — GPT-4, Claude, and Gemini are accessible alongside SAP's proprietary models like SAP-RPT-1 and SAP-ABAP-1. In its public communications, SAP announces more than 30 specialized agents and 2,500 *Joule Skills* for Q1 2026.

The commercial model blends **Base AI** included in the subscription and **Premium AI** billed in *AI Units* — a mixed logic that demands careful attention when calculating total cost.

Oracle: The Industrialization of the Agentic Application

Oracle pushes the product logic furthest. The **Fusion Agentic Applications** introduced in March 2026 are not agents. They are applications of an entirely new type, natively designed to execute business outcomes by mobilizing teams of specialized agents.

The vendor publishes the densest catalog, covering ERP, Risk, HCM, SCM, Procurement, Product Lifecycle, Quality, and CX. AI Agent Studio allows organizations to customize or build their own agent teams. OCI Generative AI provides the model layer with access to Oracle LLMs, OpenAI via Azure, Anthropic, Gemini, and others.

Oracle's unique argument is **deployment flexibility**: from public cloud to sovereign on-premises datacenters, through hybrid architectures. To the best of our knowledge at the time of writing, it is the only vendor explicitly positioned on this complete continuum.

Commercially, Oracle includes AI Agent Studio in the Fusion subscription for several use cases, but charges an additional *Custom AI Agent subscription* for advanced customization or the use of third-party LLMs.

What This Divergence Means for You

These three bets are not three ways of doing the same thing. They are three different theories about what creates agentic value in the enterprise. Workday bets on adoption through experience. SAP bets on business process mastery. Oracle bets on productization and infrastructure flexibility.

The right choice depends less on which product is “best” than on your organization’s center of gravity and the maturity of your processes.

What Vendors Won’t Tell You (But Research Reveals)

This is where my dual perspective matters most. Vendor decks do an admirable job covering architecture, technical governance, and projected ROI. They are remarkably silent on what recent academic research identifies as the true success or failure factors in an agentic deployment.

The Cognitive Cost Is Real and Measurable

A study published in the *Journal of Management* (FT50) by [Shao, Shao, and Huang in 2024](#) tracked employees exposed to AI augmentation tools on a daily basis. The results are both encouraging and sobering.

On one side, frequent AI tool use is associated with significant knowledge gains and better end-of-day performance. This is what vendors emphasize.

On the other side — and this is what vendors do not emphasize — frequent use also generates **information overload** that degrades performance and compromises post-work recovery. The same tool simultaneously produces both effects, on the same person, on the same day.

This duality has a concrete implication. The typical vendor ROI metrics — time saved, tickets reduced, adoption rate — mask an invisible cognitive cost that only surfaces over time, in the form of burnout, staff turnover, or disengagement. None of the three vendors measure this today.

Not All Employees Respond the Same Way

According to Shao et al., two individual factors strongly moderate the effects of AI exposure. **Openness to experience** as a dispositional trait, and **positive affect** as a momentary state. Concretely, two employees in the same role with the same training can experience the arrival of an AI agent in diametrically opposite ways.

Vendors and system integrators tend to talk about “users” as a homogeneous category. Research says the opposite. An adoption strategy that ignores individual moderators ends up producing flattering average adoption rates alongside invisible pockets of resistance or burnout in the KPI dashboards.

Long-Term Effects Remain a Black Box

A panel published in 2024 in *Communications of the Association for Information Systems* by [Haase, Kremser, Leopold, Mendling, Onnasch, and Plattfaut](#) — an interdisciplinary group of researchers in IS, automation, and human-automation interaction — explicitly calls for research on the **long-term psychological impact** of the RPA-LLM combination in business processes. To my knowledge, this call has not yet received any serious empirical response.

This means that no one — neither vendors nor researchers — can say with confidence today what happens to an HR, finance, or IT team exposed to an agentic environment for 24 months. That is not a reason to do nothing. It is a reason to do it with discernment.

In the same vein, [Brynjolfsson, Chandar, and Chen \(2025\)](#) report differentiated labor-market effects from AI adoption, including stronger pressure on some entry-level roles in exposed occupations. The source is still a working paper, but it is useful for framing short-term labor adjustment risk.

In parallel, multiple Gartner market notes on enterprise GenAI point to a similar operational pattern: adoption remains uneven, use cases are scaling faster than governance standards, and many organizations still struggle to industrialize usage beyond early phases. These market signals are useful for managerial framing, but they remain market-read evidence and are not substitutes for academic proof.

What Academia Doesn't See (But the Field Reveals)

The symmetry is necessary. If vendors have their blind spots, so does the academic literature — and some of them are massive.

Agentic Commercial Mechanics Are Invisible in Research

Workday's *Flex Credits*, SAP's *AI Units*, Oracle's *Custom AI Agent subscription*. Not one academic article reviewed seriously examines the organizational, behavioral, and strategic implications of these consumption-based models.

Yet this is a paradigm shift. Moving from per-seat licensing to consumption-based pricing potentially changes usage behavior, organizational trade-offs, and even the perception of AI as a scarce or abundant resource. In the field, I already see organizations rationing their use of agents out of fear of budget overruns, and others over-consuming without measurement. No academic theory currently describes this phenomenon.

Sovereignty Does Not Exist in Academic Journals

For a Quebec, Canadian, or more broadly sovereignty-constrained organization — Loi 25, public sector, regulated financial services, healthcare — the question of where data and the AI runtime are placed is central. The strategic analysis I draw from explicitly identifies Oracle as the only vendor offering a complete continuum from public cloud to sovereign datacenter.

This dimension is entirely absent from the academic literature on AI in ERP systems. It is a blind spot that should make us cautious about the “universal” frameworks proposed in research, which often implicitly assume public SaaS deployment.

The Devil Is in the Process

Academic AI-ERP integration frameworks, such as the one proposed by [Sarferaz in *IEEE Access*](#), remain at a high level of abstraction. They describe generic methodological steps. The field, however, lives in the exceptions.

A Quebec payroll cycle with its RL-1, QPIP, CNESST, and ministerial orders cannot be delegated to an agent in the same way as a standardized American payroll process. A procurement approval matrix that varies by cost center, amount, and expense type requires modeling that generic frameworks do not capture. An onboarding path conditioned on remote work status and provincial tax obligations introduces branching that no universal prompt can anticipate. It is at this level of granularity — regulatory, sectoral, organizational — that real agentic deployments are won or lost.

Four Questions to Ask Before You Sign

If you are a CHRO, CIO, VP Finance, or VP Transformation and need to decide in 2026 what to do with agentic AI in your ERP, here are the four questions I recommend asking before signing anything.

1. Which processes actually warrant controlled delegation? Do not start with “which agent should we buy.” Start by mapping your processes along three axes — frequency, friction, and risk. The best candidates for the first wave are high-frequency, high-friction, moderate-risk processes. HR self-service, expenses, helpdesk, requisitions, onboarding, minor cash management. Not the monthly financial close. Not payroll. Not high-regulatory-impact transactions.

2. What lived indicators will you track, in addition to management KPIs? ROI, time saved, and tickets reduced are necessary but insufficient. Add at minimum three lived indicators — perceived cognitive load, sense of control over one’s work, and the coping strategies employees mobilize. These measures are not exotic. They have existed in HRM and IS research

for years. They are absent from vendor dashboards. You will need to build them yourself — and that is precisely what will distinguish a successful deployment from a costly one.

3. What agent governance beyond the purchase? If your agents are digital colleagues, who supervises them? Who validates their access rights? Who decides on their decommissioning? Who measures their budget footprint? Workday with ASOR, SAP with Cloud Identity Services, and Oracle with its RBAC frameworks offer different mechanisms. The right choice depends on the maturity of your IT function and your tolerance for governance industrialization.

4. What is the real economics beyond the proof of concept? Consumption-based models make pilots attractive and large-scale deployments unpredictable. Before signing, require a realistic simulation of your full-load annual cost — including premium LLMs, third-party connectors, API calls, and customization overages. Ask explicitly for the list of situations that trigger additional billing. You will be surprised.

The Practitioner-Researcher Stance

To close, a word about the lens that allowed me to write this article.

I co-lead a Workday consulting firm. Every week, I see Quebec and Canadian organizations facing exactly the choices I have just described. I participate in real deployments, with their successes and their failures.

I am also a doctoral student working on the transformation of knowledge work by generative AI in an Industry 5.0 context. I spend a significant part of my time reading recent scientific research and confronting theoretical frameworks with field observations.

This dual stance is not comfortable. It forces you to acknowledge the limits of both worlds. Field practice alone produces compelling narratives that do not generalize. Science alone produces rigorous frameworks that are often inoperable. It is the back-and-forth between the two that creates value — for clients, for teams, and, I believe, for the discipline itself.

If you have read this far, you probably share that intuition. You know that the next ERP wave will be won neither in marketing nor in theory. It will be won in the quality of the intersection between what research teaches us and what the field forces us to understand.

I will always have more questions than answers. But it is that back-and-forth between theory and field that makes agentic deployments succeed.

Scientific References

- Sapkota, R., Roumeliotis, K. I., & Karkee, M. (2026). AI Agents vs. Agentic AI: A Conceptual Taxonomy, Applications and Challenges. *Information Fusion*, 126, 103599. <https://doi.org/10.1016/j.inffus.2025.103599>
- Shao, Y., Huang, Z., Song, Y., Wang, M., Song, Y. H., & Shao, R. (2025). Using Augmentation-Based AI Tool at Work: A Daily Investigation of Learning-Based Benefit and Challenge. *Journal of Management*, 51(8), 3352–3390. <https://doi.org/10.1177/01492063241266503>
- Haase, J., Kremser, W., Leopold, H., Mendling, J., Onnasch, L., & Plattfaut, R. (2024). Interdisciplinary Directions for Researching the Effects of Robotic Process Automation and Large Language Models on Business Processes. *Communications of the Association for Information Systems*, 54(1), 579–604. <https://doi.org/10.17705/1CAIS.05421>
- Sarferaz, S. (2025). Implementing Conversational AI Into ERP Software. *IEEE Access*, 13, 160238–160250. <https://doi.org/10.1109/ACCESS.2025.3608477>
- Aguinis, H., Beltran, J. R., & Cope, A. (2024). How to Use Generative AI as a Human Resource Management Assistant. *Organizational Dynamics*, 53(1), 101029. <https://doi.org/10.1016/j.orgdyn.2024.101029>
- Acharya, D. B., Kuppan, K., & Divya, B. (2025). Agentic AI: Autonomous Intelligence for Complex Goals—A Comprehensive Survey. *IEEE Access*, 13, 18912–18936. <https://doi.org/10.1109/ACCESS.2025.3532853>
- Islam, M. S., Islam, M. I., Mozumder, A. Q., Khan, M. T. H., Das, N., & Mohammad, N. (2025). A Conceptual Framework for Sustainable AI-ERP Integration in Dark Factories: Synthesising TOE, TAM, and IS Success Models for Autonomous Industrial Environments. *Sustainability*, 17(20), 9234. <https://doi.org/10.3390/su17209234>
- Brynjolfsson, E., Chandar, B., & Chen, R. (2025). *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence* [Working Paper]. Stanford Digital Economy Lab. <https://digitaleconomy.stanford.edu/publication/canaries-in-the-coal-mine-six-facts-about-the-recent-employment-effects-of-artificial-intelligence/>

Transparency Note — Generative AI

Generative AI tools were used for source organization, prose revision, translation support, and formatting checks. The authors retained full responsibility for source selection, interpretation of findings, reference validation, and final content. AI outputs were reviewed, edited, and verified by the authors. No AI tool is credited as an author.