

[← Back to Insights](#)

STRAVORIS

Most Agentic AI Is Just Expensive RPA

Executive Summary

The enterprise world is spending aggressively on "agentic AI" in early 2026, but the evidence suggests most of what is being built does not qualify as agentic in any meaningful sense. Deloitte's 2026 Tech Trends report finds that organizations are "trying to automate existing processes—tasks designed by and for human workers—without reimagining how the work should actually be done."^[1] Gartner predicts over 40% of agentic AI projects will be canceled by the end of 2027 due to escalating costs, unclear business value, or inadequate risk controls.^[2]

The core problem is not the technology. It is organizational. Enterprises are applying AI to existing human-designed workflows, effectively creating sophisticated RPA dressed in agentic language. Deloitte calls this "agent washing"—the rebranding of existing products such as AI assistants, robotic process automation, and chatbots "without substantial agentic capabilities."^[1] The Futurum Group confirms the pattern at the vendor level: Salesforce rebranded its entire platform around "Agentforce 360," and Microsoft repackaged Copilot features as part of a larger intelligence layer—in both cases reorganizing existing capabilities rather than delivering fundamentally new agent architectures.^[4]

Despite the hype, genuine progress is occurring at the margins. The Arcade.dev 2026 State of AI Agents survey reports that 80% of respondents see measurable economic impact from agents, 57% deploy multi-step agent workflows, and 81% plan to expand into more complex use cases.^[3] But only 14% of organizations have deployable agentic solutions, and a mere 11% are using them in production.^[1] The gap between experimentation and production is the defining challenge of 2026.

The enterprises getting results share three traits: they redesigned processes rather than automating existing ones, they invested in data architectures that make enterprise information searchable at runtime (not through ETL pipelines), and they treat agents as a managed workforce with onboarding, performance monitoring, and lifecycle governance. Organizations making architectural decisions about agent platforms in Q1/Q2 2026 have a narrow window to get the foundations right before the 2027 reckoning Gartner predicts.

Evidence Base & Methodology

This brief synthesizes findings from 10 primary sources spanning analyst reports, industry surveys, vendor analyses, and enterprise case studies. Research was conducted on March 11, 2026, covering evidence from mid-2025 through early 2026.

Search Approach

- **5 web searches** covering: agentic AI failures and hype; Gartner predictions; Deloitte Tech Trends; agent washing and RPA rebranding; production case studies
- **4 full-text page reads** of highest-signal sources (Deloitte, Arcade.dev, Futurum Group, XMPRO)
- **2 seed URLs** from the idea file fetched and analyzed (Deloitte, Arcade.dev blog)

Evidence Characteristics

Dimension	Assessment
Date range of evidence	June 2025 – March 2026
Analyst sources	Gartner (2 reports), Deloitte (Tech Trends 2026), Futurum Group
Survey data	Arcade.dev 2026 survey; Gartner poll of 3,412 attendees (Jan 2025); CrewAI survey of 500 executives
Notable gaps	Gartner's full methodology behind the 40% prediction is paywalled; limited independent academic research on agent vs. RPA effectiveness; most case studies are vendor-published

1. The Agent Washing Problem

1.1 What "Agentic" Actually Means

The term "agentic AI" has become so overused that it risks losing meaning. Based on the collected evidence, a useful working definition requires three capabilities that distinguish true agents from automation:

Capability	Agentic AI	Traditional RPA / Automation
Decision-making	Autonomous, goal-oriented reasoning with context awareness	Rule-based, deterministic execution of predefined steps
Adaptability	Handles exceptions, learns from outcomes, adjusts approach	Fails or escalates when encountering unscripted scenarios
Scope	End-to-end workflows spanning multiple domains and systems	Single-task or single-system process execution

Deloitte frames the distinction through a spectrum: *augmentation* (agents enhance human capabilities), *automation* (agents handle human-defined tasks), and *true autonomy* (minimal oversight, requiring AGI-level capability). Most current deployments sit firmly in the augmentation tier, with some reaching automation—but few approach genuine autonomy.^[1]

1.2 How Agent Washing Manifests

The Futurum Group's analysis of major vendor announcements through 2025 found a consistent pattern: existing capabilities were reorganized and relabeled as "agentic" without substantive architectural changes.^[4]

Vendor	Announcement	Assessment
Salesforce	"Agentforce 360" platform rebrand	Existing core modules reorganized under agentic branding ^[4]
Microsoft	365 Copilot → "Work IQ" intelligence layer	Existing features repackaged as part of larger agent narrative ^[4]
Oracle	"User AI" experience across applications	Expanded existing embedded AI capabilities ^[4]
ServiceNow	Pre-configured agents for IT, HR, customer service	Genuine orchestration platform with real multi-agent capability ^[4]
Adobe	AI Foundry for brand-specific model training	Genuine new capability enabling enterprise model customization ^[4]

The Futurum Group notes that vendor "claims wars" around superlatives like "unique" and "the only vendor" do "little to actually move the needle, in terms of driving buyer interest."^[4] The analyst verdict: progress is real but incomplete, and hype remains significant.

1.3 The Pricing Opacity Problem

By Q3 2025, "the debate around per-user, consumption-based, and 'AI included' pricing models intensified," with customers "scrutinizing whether vendors are delivering measurable value or using AI to justify higher subscription tiers."^[4] This pricing confusion compounds the agent washing problem: organizations cannot easily determine whether they are paying premium prices for genuine agentic capabilities or for relabeled automation.

2. Why 40% of Agentic AI Projects Will Fail

2.1 Gartner's Prediction and Its Root Causes

Gartner's June 2025 prediction that over 40% of agentic AI projects will be canceled by end of 2027 cites three primary drivers: escalating costs, unclear business value, and inadequate risk controls.^[2] A January 2025 Gartner poll of 3,412 webinar attendees provides context on investment posture:

Investment Level	Percentage
Significant investments made	19%
Conservative investments made	42%
No investments	8%
Wait-and-see / unsure	31%

The adoption funnel narrows dramatically at each stage. While 30% of organizations are exploring agentic options and 38% are piloting, only 14% have deployable solutions and just 11% are in production.^[1] The strategy gap is equally stark: 42% are still developing roadmaps, and 35% have no formal agentic AI strategy at all.^[1]

2.2 Three Infrastructure Obstacles

Deloitte identifies three infrastructure-level obstacles that explain most project failures:^[1]

Legacy System Integration. Enterprise systems lack "real-time execution capability, modern APIs, modular architectures, and secure identity management." The Arcade.dev survey corroborates this: 46% of organizations cite integration with existing systems as their primary challenge.^[3]

Data Architecture Constraints. Traditional ETL and data warehouse models create "friction for agent deployment." Deloitte argues the paradigm must shift from extract-transform-load to "enterprise search and indexing—similar to how Google made the World Wide Web discoverable."^[1] Nearly half (48%) cite searchability of data and 47% cite reusability of data as primary AI automation obstacles.^[1]

Governance and Control. Traditional IT governance models "don't account for AI systems that make independent decisions."^[1] Security concerns rank as the third-highest barrier, with 40% of organizations identifying security and compliance as adoption blockers.^[3]

2.3 The Process Automation Trap

The most fundamental failure mode is attempting to automate existing human-designed processes without questioning whether those processes should exist. As Dell's John Roesse warns: "If you don't have solid processes, you should not proceed. Figure that out first."^[1] Intel's Brent Collins adds: "Don't simply pave the cow path. Instead, take advantage of this AI evolution to reimagine how agents can best collaborate."^[1]

Deloitte invokes Henry Ford's observation: "Many people are busy trying to find better ways of doing things that should not have to be done at all."^[1] The implication is that organizations applying AI to existing workflows are optimizing processes that may be fundamentally misdesigned for agent execution.

3. What Actually Works: Evidence from Production

3.1 The Production Reality

Despite widespread failure, organizations that have reached production are seeing returns. Among executives who report deploying AI agents in production, 74% achieved ROI within the first year, and 39% have deployed more than 10 agents across their enterprise.^[6] Reported ROI ranges from 5x-10x per dollar invested.^[6]

The Arcade.dev survey shows 80% of respondents report measurable economic impact, with 88% expecting ROI to continue or increase in 2026.^[3] However, this survey skews toward organizations already invested in agent infrastructure, introducing selection bias.

3.2 Case Studies: Genuine vs. Relabeled

Organization	Implementation	Result	Assessment
Walmart	End-to-end agent: signal detection → demand forecasting → inventory action	22% increase in e-commerce sales in pilot regions; reduced out-of-stock incidents ^[6]	Genuinely agentic: multi-step, cross-domain, autonomous decision chain
Ramp (Fintech)	AI agent reads policy docs, audits expenses, flags violations, generates approvals	Autonomous compliance workflow launched July 2025 ^[6]	Genuinely agentic: reasoning over unstructured docs, autonomous judgment
HPE	"End-to-end process transformation rather than solving for a single pain point"	Not disclosed ^[1]	Process-redesign approach aligned with agentic principles
Banking sector (aggregate)	KYC/AML workflow agents	200%-2,000% productivity gains reported ^[6]	Wide range suggests variable implementation quality

3.3 The Hybrid Reality

The evidence consistently points to a hybrid model rather than wholesale replacement of RPA. Blue Prism's analysis concludes that "the future is fusion, not replacement."^[5] The most successful strategies deploy RPA for high-volume, structured execution while leveraging agentic AI for complex decision-making, unstructured data, and exception handling. This is pragmatic but also reveals that most

"agentic" deployments are really augmented RPA—agents handling the exceptions that RPA cannot, within workflows still fundamentally designed for deterministic execution.

4. Silicon Workforce Management: A New Operating Model

4.1 The Deloitte Framework

Deloitte introduces the concept of a "silicon-based workforce that complements and enhances the human workforce," requiring HR-equivalent management practices:^[1]

Practice	Human Workforce Equivalent	Agent Workforce Requirement
Onboarding	Training and orientation	Dual-approach training for both agents and human supervisors on collaboration
Performance Management	Reviews and KPIs	Systems proving "what agents did, why they made specific decisions, and under whose authority"
Identity & Access	Badge and permissions	Digital identity, cryptographic receipts, immutable audit logs, zero-trust architecture
Lifecycle Management	Career development, offboarding	Ongoing training updates, redeployment to priority areas, retirement planning
Financial Operations	Compensation and benefits	FinOps frameworks tracking token-based pricing, resource tagging, autoscaling

4.2 The FinOps Imperative

A frequently overlooked failure mode is cost management. Deloitte warns that poor agent configuration can trigger "cascading actions like unpredictable resource consumption and ballooning costs."^[1] This requires specialized financial operations frameworks that differ fundamentally from traditional IT cost management, incorporating token-based pricing, real-time resource monitoring, and autoscaling controls. Organizations accustomed to fixed RPA licensing costs are unprepared for the variable-cost model of LLM-powered agents.

4.3 Orchestration Protocol Landscape

The multi-agent orchestration space is fragmenting across competing protocols, each with limitations:

Protocol	Originator	Approach	Limitation
Model Context Protocol (MCP)	Anthropic	Standardized AI-to-data connections	Complex enterprise security scenarios ^[1]
Agent-to-Agent Protocol (A2A)	Google	Cross-platform agent communication	Early-stage adoption ^[1]
Agent Communication Protocol (ACP)	Open standard	RESTful API for agent messaging	Network coordination complexity ^[1]

The lack of a dominant standard adds friction to enterprise adoption and increases the risk of vendor lock-in—a concern that maps directly to the RPA era's integration challenges.

5. Is It Agentic, or Is It Automation?

5.1 A Diagnostic Framework

Based on the collected evidence, the following diagnostic separates genuine agentic implementations from relabeled automation:

Question	If Yes → Likely Agentic	If No → Likely Automation
Did you redesign the workflow, or automate the existing one?	Process was reimagined for agent-first execution	Existing human process was preserved and automated
Can the system handle exceptions it was not explicitly programmed for?	Agent reasons over novel scenarios using context	System fails or escalates on unscripted inputs
Does the system span multiple domains or organizational boundaries?	Agent orchestrates across systems, teams, data sources	Operates within a single system or task boundary
Is enterprise data searchable by the agent at runtime?	Knowledge graph or indexed data layer available	Data accessed via batch ETL or static integrations
Do you manage this system like a workforce member?	Onboarding, performance monitoring, lifecycle governance	Deployed and left to run like traditional software

5.2 The Maturity Spectrum

Cross-referencing Deloitte's autonomy spectrum with the Arcade.dev survey data reveals where most organizations actually sit:

Maturity Level	Description	Estimated Share of Market
Level 0: Exploring	Evaluating vendors, building strategy	~30% ^[1]
Level 1: Piloting	PoCs with limited scope, typically single-task	~38% ^[1]
Level 2: Deployable	Production-ready but limited rollout	~14% ^[1]
Level 3: In Production	Multi-step workflows, measurable ROI	~11% ^[1]
Level 4: Cross-Functional	Agents spanning teams and domains	~16% of those deployed ^[3]

Inference: approximately 2% of organizations have achieved genuinely cross-functional agentic AI (16% of the 11% in production). The remaining 98% are either exploring, piloting single-task automation, or deploying what is functionally enhanced RPA under an agentic label. This is a low-confidence estimate based on combining two different survey populations.

Key Assumptions & Uncertainties

- **Survey bias:** The Arcade.dev and CrewAI surveys sample organizations already invested in agent tooling, likely overstating adoption maturity and ROI. The 80% "measurable impact" figure should be interpreted cautiously.
- **Definition variance:** There is no industry-standard definition of "agentic AI." Sources use the term to describe everything from enhanced chatbots to autonomous multi-agent systems. This definitional ambiguity is itself a core finding.
- **Gartner's 40% prediction:** The full methodology is paywalled. The prediction may count project cancellations that are normal portfolio management (pivoting, consolidating) rather than genuine failures.
- **Vendor-published case studies:** Most production success stories (Walmart, Ramp, banking KYC) come from vendor marketing or press releases, not independent audits. Reported ROI figures (5x-10x, 200%-2,000%) should be treated as upper-bound claims.
- **RPA comparison:** The "agentic AI is just expensive RPA" framing, while directionally supported, oversimplifies. Some implementations genuinely exceed RPA capabilities even if most do not.
- **Timing dependency:** The landscape is evolving rapidly. Multi-agent orchestration protocols (MCP, A2A, ACP) are maturing quarter-over-quarter, and what fails today may succeed with better tooling in 12 months.

Strategic Implications & Actionable Insights

- 1. Audit your "agentic" projects for agent washing.** Apply the diagnostic framework in Section 5 to every initiative labeled "agentic AI." Reclassify those that are automating existing processes without redesign as what they are: AI-enhanced automation. This is not a failure—it is honest scoping that prevents budget overruns and misaligned expectations.^{[1][4]}
- 2. Invest in data architecture before agent platforms.** With 48% citing data searchability and 47% citing data reusability as primary obstacles, the bottleneck is not model capability—it is data access. Prioritize knowledge graphs and runtime-indexable data stores over agent orchestration tooling.^{[1][3]}
- 3. Redesign processes before automating them.** Every enterprise leader quoted in the Deloitte report—from Intel to Dell to HPE—emphasizes process redesign over process automation. Value-stream mapping should precede any agent deployment. If the workflow was designed for humans, bolting an agent onto it creates "workslop"—agents that actually add work to a process.^[1]
- 4. Build FinOps controls from day one.** The shift from fixed RPA licensing to variable LLM token costs is a budget risk most organizations have not planned for. Implement resource tagging, real-time cost monitoring, and autoscaling limits before scaling any agent deployment.^[1]
- 5. Adopt the hybrid model deliberately, not by default.** The evidence supports RPA + agentic AI coexistence. But "hybrid" should be a conscious architecture decision—RPA for structured, high-volume execution; agents for reasoning-intensive exception handling—not a rationalization for failing to build genuine agent capability.^[5]
- 6. Favor strategic partnerships over internal builds.** Deloitte data shows strategic partnerships for pilots are twice as likely to reach full deployment versus internal builds, and employee usage rates nearly double for externally built tools.^[1]
- 7. Establish agent governance now, not later.** With 40% of organizations citing security and compliance as blockers, and orchestration protocols still fragmenting, organizations need governance frameworks that address non-deterministic decision-making, audit trails, and ephemeral authentication before scaling.^{[1][3]}

References

1. Deloitte. "The Agentic Reality Check: Preparing for a Silicon-Based Workforce." *Tech Trends 2026*. [deloitte.com](https://www.deloitte.com). Accessed March 11, 2026.
2. Gartner. "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027." Press release, June 25, 2025. [gartner.com](https://www.gartner.com). Accessed March 11, 2026.
3. Arcade.dev. "5 Takeaways from the 2026 State of AI Agents." [arcade.dev](https://www.arcade.dev). Accessed March 11, 2026.
4. Futurum Group. "Was 2025 the Year of Agentic AI, or Just More Agentic Hype?" [futurumgroup.com](https://www.futurumgroup.com). Accessed March 11, 2026.
5. SS&C Blue Prism. "Agentic AI vs RPA: Comparing AI Agents and RPA Bots." [blueprism.com](https://www.blueprism.com). Accessed March 11, 2026.
6. Various. Enterprise AI agent case studies and ROI data compiled from: Google Cloud Blog, Warmly.ai, Ampcome enterprise use case database. Accessed March 11, 2026.
7. Gartner. "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026." Press release, August 26, 2025. [gartner.com](https://www.gartner.com). Accessed March 11, 2026.
8. Kore.ai. "AI Agents in 2026: From Hype to Enterprise Reality." [kore.ai](https://www.kore.ai). Accessed March 11, 2026.
9. Level Up Coding. "5 Real Projects Where Agentic AI Failed Badly in 2026." [levelup.gitconnected.com](https://www.levelup.gitconnected.com). Accessed March 11, 2026.
10. VentureBeat. "The Era of Agentic AI Demands a Data Constitution, Not Better Prompts." [venturebeat.com](https://www.venturebeat.com). Accessed March 11, 2026.

STRAVORIS

INNOVATE. INTEGRATE. ELEVATE.