

From Copilots to Operating Layers: Is the AI Market Moving Beyond Chat?

Executive assessment

The evidence supports a **qualified version** of the hypothesis.

The enterprise AI market is undergoing a material shift in product positioning, architecture, hiring demand, governance requirements, and workflow design. In 2023, leading products were commonly presented as copilots, conversational assistants, content generators, and productivity enhancers. By 2025–2026, the dominant enterprise vocabulary increasingly includes agents, reasoning, orchestration, enterprise context, knowledge graphs, control planes, digital labor, workflow outcomes, guardrails, and human-agent teams. Microsoft and Salesforce provide particularly clear longitudinal examples: Microsoft's 2023 Copilot launch emphasized drafting, summarizing, prompting, and working "alongside" users, while Salesforce's 2023 Einstein GPT emphasized generated CRM content; Salesforce now explicitly describes its evolution as moving "from assistants to agents" that reason, plan, and execute. ¹

However, the evidence does **not** show that chatbots are disappearing or that enterprises have already completed this transition. Chat remains the primary human interface for many systems, common production use cases still include conversational information access and content generation, and scaled agent deployment remains limited. McKinsey found that 88% of surveyed organizations used AI in at least one function in 2025, but nearly two-thirds had not begun scaling AI enterprise-wide; only 23% reported scaling an agentic system somewhere, and no individual business function had agent scaling above 10% of respondents. ²

The more accurate conclusion is therefore:

The market is not moving away from conversation as an interface. It is moving beyond conversation as the complete product.

Chat is increasingly the visible entry point to a deeper operating layer consisting of company data, semantic context, workflow state, permissions, tools, models, evaluations, policy controls, human approvals, and execution systems. The transition is strongest in **enterprise strategy and vendor road maps**, moderate in **investment and hiring**, and still incomplete in **scaled operational deployment**. ³

The advertisements that prompted this investigation are therefore useful as a market-sensing observation, but not sufficient proof by themselves. Their significance is that they independently reproduce language already appearing across major enterprise platforms, analyst reports, labor-market datasets, and investment research.

Proposition	Assessment	Evidence strength
Enterprise marketing has shifted from generation and assistance toward reasoning, agents, context, and execution	Supported	Strong
Organizational knowledge and governance are becoming central product layers	Supported	Strong
Hiring is moving from standalone prompting toward integration, operations, and workflow execution	Supported	Moderate to strong
Venture capital has moved decisively from chat interfaces to workflow and knowledge systems	Partially supported	Mixed
The creator economy is becoming a market for executable knowledge systems	Plausible but not established	Weak to moderate
Chatbots are no longer commercially central	Not supported	Strong counterevidence
Autonomous agents are already reliable at enterprise scale	Not supported	Strong counterevidence

The messaging shift from assistance to agency

In early 2023, generative AI was commonly sold through three immediate capabilities: **talk to your information, generate content, and increase individual productivity**. Microsoft introduced Microsoft 365 Copilot as an LLM integrated into Word, Excel, PowerPoint, Outlook, and Teams, emphasizing drafting, rewriting, summarizing, creativity, and productivity. Its cross-application Business Chat still operated through natural-language requests issued by a user. Salesforce introduced Einstein GPT in almost identical terms: personalized emails, customer-service responses, marketing content, code generation, real-time data, and natural-language prompts. ⁴

“Copilot” was an effective category label because it placed responsibility and initiative with the human. The user supplied intent, the model produced an artifact, and the user decided what happened next. The central unit of interaction was a prompt-response exchange, even when enterprise data was retrieved behind the interface. Microsoft also described Copilot as working “alongside” the user, while Salesforce initially called its subsequent product a conversational CRM assistant. ⁵

By 2025–2026, the language had changed substantially. The recurring verbs are now **reason, plan, coordinate, act, govern, monitor, and deliver outcomes**. Salesforce’s current AI history says that it moved from assistants to agents and that its Atlas Reasoning Engine enables AI to reason, plan, and execute across Customer 360. Microsoft’s Work Trend Index describes companies built around human-agent teams, digital labor, and workers managing or training agents. Gartner now distinguishes “assistive intelligence,” such as copilots and advisers, from platforms that commit to workflow outcomes. ⁶

This shift is visible across vendors that compete in otherwise different categories:

Palantir presents its Ontology as a unified representation of enterprise data, logic, actions, security, and decision-making. Its architecture models the “nouns” and “verbs” of operations in a form usable by humans and agents, rather than treating AI as a detached chat window over documents. ⁷

Glean now describes itself as “Work AI” connecting organizational knowledge, systems, and context. Its platform language moves from finding answers toward taking action, automating work, and converting fragmented enterprise information into trusted context. ⁸

WRITER markets a library of more than 100 enterprise agents that can research, analyze, reason, plan, and act, explicitly describing them as more than chat interfaces. ⁹

ServiceNow emphasizes agent orchestration, agent teams, control towers, guardrails, roles, tools, business objectives, and workflows that close the loop across existing systems. ¹⁰

IBM calls watsonx Orchestrate an “agentic control plane,” emphasizing orchestration, policy enforcement, lifecycle control, governance, visibility, and coordination across agents, tools, data, and workflows. ¹¹

UiPath now positions orchestration above individual bots and agents, coordinating AI agents, deterministic robots, people, and long-running business processes. Its message is that automating isolated tasks is insufficient when the work spanning those tasks is still coordinated manually. ¹²

SAP argues that agents require a trusted business-data layer and a knowledge graph that encodes relationships between processes and data. The stated purpose is to ground agent decisions and actions in operational context across finance, procurement, service, and supply chains. ¹³

The labels vary—AI operating system, control plane, agent platform, enterprise AI platform, reasoning engine, orchestration layer—but they are converging on a common architectural claim: **the valuable system is no longer only the model that produces language. It is the surrounding structure that allows probabilistic models to interact with a real organization.**

That convergence is meaningful, although marketing terminology should not be mistaken for technical standardization. “Reasoning engine,” “ontology,” “agent,” and “AI operating system” do not have uniform definitions across vendors. Some products may contain genuinely new planning, tool-use, and state-management capabilities; others may be conventional workflow automation relabeled with agentic language. Gartner has explicitly warned about “agentwashing” and places agentic AI at the Peak of Inflated Expectations. ¹⁴

The enterprise architecture emerging beneath the language

Across the current vendor landscape, a recognizable enterprise stack is forming.

The first layer is **organizational context**: permissions-aware search, enterprise data connectors, taxonomies, metadata, entity relationships, knowledge graphs, business ontologies, and retrieval systems. These systems attempt to answer not merely “What text resembles this query?” but “What does this information mean inside this company, who may use it, and how does it relate to a business object or

process?” Palantir’s Ontology, SAP’s Knowledge Graph, and Glean’s enterprise-context layer are three different implementations of this general direction. ¹⁵

The second layer is **planning and reasoning**. Here, a model decomposes an objective, selects tools or information, evaluates intermediate results, and determines subsequent steps. This is the capability vendors are trying to distinguish from one-shot generation. Salesforce’s Atlas messaging, WRITER’s agent library, and SAP’s description of context-grounded multi-step agents all emphasize this layer. ¹⁶

The third layer is **orchestration and state**. Long-running enterprise work rarely consists of one prompt and one answer. It involves dependencies, waiting periods, exceptions, multiple systems, identity changes, approvals, escalation paths, and recovery from failure. ServiceNow, IBM, and UiPath therefore emphasize coordinating agents, applications, people, robots, and process state rather than maximizing model autonomy in isolation. ¹⁷

The fourth layer is **governed execution**. An enterprise system must determine which actions an agent may take, which require human authorization, which data may be exposed, how decisions are logged, and how outputs are evaluated. ServiceNow advertises central agent management and guardrails; IBM emphasizes policy enforcement and lifecycle control; UiPath combines probabilistic agents with deterministic automation and people. ¹⁷

This hybrid architecture is important. The emerging enterprise model is not simply “replace every workflow with autonomous AI.” It is more often:

probabilistic interpretation → bounded planning → deterministic tools → policy checks → human approval where required → logged execution

That architecture preserves the flexibility of language models without assuming they are reliable enough to own every consequential decision. McKinsey found that AI high performers are more likely to redesign workflows and institute human validation, while Deloitte found that governance is lagging behind agent adoption. ¹⁸

Industry adoption data confirms both the strategic shift and its incomplete state. McKinsey’s 2025 survey found that AI use had reached 88% of organizations in at least one function, up from 78% the prior year. Knowledge management joined IT, marketing, and sales among the most common areas of use, and 62% of respondents said their organizations were at least experimenting with agents. Yet approximately two-thirds had not begun enterprise-wide scaling, and only 39% reported enterprise-level EBIT impact. ²

Deloitte’s 2026 enterprise survey reports that employee access to AI rose 50% during 2025 and that the number of companies with at least 40% of AI projects in production was expected to double within six months. But only 34% of surveyed organizations were using AI to reimagine products, services, core processes, or business models; 37% were still applying AI with little or no process change. ¹⁹

Accenture’s July 2026 research found similarly strong commitment but limited realized value: 82% of surveyed C-suite leaders were increasing AI investment, 49% of companies were piloting or deploying agents, and 55% expected agentic initiatives to produce board-reportable outcomes within a year. Only 23%, however, reported widespread and sustained AI value, down from 32% earlier in the year. ²⁰

Microsoft's 2025 Work Trend Index found that 82% of leaders considered 2025 pivotal for rethinking strategy and operations, 45% identified expanding capacity with digital labor as a near-term priority, and 78% planned to hire for new AI roles. Its 2026 report adds a readiness warning: only 19% of AI users occupied its "Frontier" category, where employee ability and organizational readiness reinforced each other, while many capable employees were constrained by weak governance, management practices, or organizational systems. ²¹

These findings support a distinction between **tool adoption** and **organizational transformation**. Chat assistants can be distributed quickly because they leave the surrounding organization mostly unchanged. Workflow agents require data access, process redesign, authority boundaries, evaluation, ownership, and exception handling. The growing attention to organizational intelligence and orchestration is therefore partly a response to the discovery that model access alone does not produce enterprise transformation. ²²

Knowledge capture and governance are becoming enabling markets

The knowledge-capture component of the hypothesis is strongly supported, although it encompasses several distinct problems.

One problem is **fragmentation**. Organizational knowledge is distributed across documents, email, messaging platforms, tickets, databases, meetings, individual memory, and undocumented routines. Enterprise AI cannot reliably act on that knowledge unless it can retrieve the correct information, interpret its organizational significance, respect permissions, and connect it to an operational object or workflow. Glean's context layer, SAP's semantic bridge, and Palantir's ontology are direct responses to this condition. ¹⁵

A second problem is **tacit knowledge**: pattern recognition, judgment, situational awareness, and practical understanding that experts use without always being able to articulate it. Deloitte defines tacit knowledge as unwritten, intuitive pattern recognition in context and notes that experts often struggle to explain how they know what they know. Its 2025 human-capital research proposes using AI to help harvest that knowledge and make it available to less-experienced employees. ²³

A third problem is **institutional memory loss**. An APQC survey of 1,000 professionals reported that 58% of C-suite respondents were very concerned about knowledge loss associated with retirement, 92% of organizations did not consistently capture knowledge from soon-to-retire employees, and 85% had not operationalized AI to automate knowledge-management processes. Because the survey was sponsored by a knowledge-management vendor, the exact figures should be treated with appropriate caution, but the underlying issue is consistent with broader concern about retirement, turnover, and shrinking experience pipelines. ²⁴

AI changes the economics of knowledge capture because it can assist with transcription, classification, summarization, question generation, relationship extraction, draft documentation, and repeated interviews. Microsoft, for example, has introduced knowledge-management agents that can turn cases, conversations, and email into draft knowledge articles subject to supervisory review, while Deloitte has described manufacturing applications that convert technician knowledge into procedures and training materials. ²⁵

Yet capturing words is not the same as capturing expertise. An expert's interview may omit exceptions, environmental signals, implicit causal models, or the conditions under which a rule should not be applied. The strongest systems will therefore need iterative elicitation, provenance, confidence indicators, validation by other experts, observed workflow evidence, and feedback from real use. AI can reduce the cost of externalizing knowledge, but it does not automatically guarantee that the resulting representation preserves what made the expertise valuable. This conclusion follows from Deloitte's description of tacit knowledge and APQC's findings about accuracy, privacy, and compliance concerns. ²⁶

Governance is growing for the same structural reason: an AI system that only drafts text creates limited direct operational risk, while a system that can modify records, contact customers, approve transactions, change configurations, or coordinate other agents creates a much larger control problem.

Governance demand now comes from at least three directions. The first is **internal operational risk**: organizations need access control, evaluation, monitoring, audit trails, cost controls, incident response, and human intervention. The second is **customer and procurement trust**: enterprise buyers increasingly ask where data travels, how models are evaluated, and who remains accountable. The third is **formal regulation**. NIST's Generative AI Profile extends its AI Risk Management Framework with actions addressing distinctive generative-AI risks, while the European Union's AI Act entered its broad enforcement phase on August 2, 2026, with the AI Office and member-state authorities assuming implementation and enforcement responsibilities. ²⁷

The governance gap is substantial. Deloitte reports that only one in five surveyed companies has a mature governance model for autonomous agents. Gartner predicts that more than 40% of agentic AI projects will be canceled by the end of 2027 because of escalating cost, unclear value, or inadequate risk controls, and separately predicts that 40% of enterprises will demote or decommission autonomous agents after governance failures become visible in production. These are forecasts rather than measured outcomes, but they indicate how central control and accountability have become in enterprise planning. ²⁸

Gartner's 2026 Hype Cycle is especially revealing because governance, agent security, FinOps, context graphs, agent-development life cycles, orchestration, and agent-management platforms appear as distinct supporting categories around agentic AI. Only 17% of surveyed organizations had deployed agents, although more than 60% expected to do so within two years. Gartner's interpretation is that agent ambitions are advancing faster than the mechanisms required to operate them safely, economically, and consistently. ²⁹

This supports one of the central observations behind the proposed article: **structure is not peripheral to enterprise AI. Structure is what converts model capability into organizational capability.**

Labor demand is shifting toward execution, while the creator-economy evidence is less conclusive

The labor-market evidence is among the strongest indicators that the market is moving beyond standalone prompting.

Lightcast's analysis for the 2026 Stanford AI Index found AI skills in 2.5% of all U.S. job postings, up 55% in one year. Mentions of its "Agentic AI" skill cluster increased by more than 280%, from 0.06% of postings in

2024 to 0.23% in 2025, representing roughly 90,000 U.S. postings. The same dataset found declines in mentions of ChatGPT, conversational AI, and chatbots while deployment-oriented skills such as scalability, cloud infrastructure, and workflow management gained importance. ³⁰

A separate Lightcast analysis found that unique postings requesting generative-AI skills rose from 55 in January 2021 to nearly 10,000 in May 2025, with postings for “Generative AI Engineer” increasing sevenfold. LinkedIn has likewise reported AI engineer among the fastest-growing roles in numerous countries, growth in Head of AI positions, and rapid expansion in AI consulting and AI-literacy skills. ³¹

By contrast, dedicated prompt-engineer jobs appear to have remained rare. A 2025 academic study examined 20,662 LinkedIn postings and found 72 prompt-engineer positions—less than 0.5% of its sample. The study represents one platform and one collection window, so it cannot establish the full size of the occupation, but it undermines the early assumption that “prompt engineer” would become a large, durable standalone profession. ³²

Prompting is not vanishing. It is being absorbed into broader jobs. The emerging implementation role must typically do more than phrase instructions. It may need to map a business process, select models, connect enterprise systems, design retrieval, define tool permissions, construct evaluation datasets, create exception paths, manage human approvals, observe production behavior, control cost and latency, and train users. Those responsibilities resemble systems integration, product management, automation engineering, knowledge architecture, risk management, and organizational change more than a narrow prompt-writing occupation. This interpretation is consistent with Lightcast’s finding that hiring demand is shifting toward infrastructure, workflow management, and system-scale execution. ³³

Titles such as AI operator, AI workflow designer, automation specialist, AI implementation consultant, agent engineer, AI operations lead, and knowledge architect remain inconsistently defined. Skills data is currently more reliable than title counts because similar work may appear under consulting, operations, product, data, automation, software, or transformation roles. Microsoft’s finding that 78% of leaders plan to hire for new AI roles should therefore not be interpreted as evidence that one standardized occupation will dominate. ³⁴

The creator-economy portion of the hypothesis has weaker evidence.

There is clear evidence that creators are seeking more direct, durable relationships with audiences rather than depending entirely on algorithmic reach. Patreon’s 2025 research, based on conversations and research involving thousands of creators and fans, described declining dependable reach on large platforms and a shift toward creator-controlled relationships. Creator advertising spending in the United States was projected by the Interactive Advertising Bureau to reach \$37 billion in 2025, an increase of 26% year over year. ³⁵

There is also visible commercial language around packaging expertise. Kajabi describes itself as an operating system for human expertise, while coaching, course, membership, newsletter, community, and consulting platforms encourage people to turn experience into repeatable intellectual assets. These are real market signals, but they come primarily from vendors whose business models benefit from portraying expertise as a monetizable asset. ³⁶

The stronger creator-economy conclusion is therefore limited: **creators are increasingly trying to own audience relationships and package expertise into repeatable formats.** The evidence is not yet strong enough to conclude that a broad market has formed around AI-generated “knowledge operating systems” for individual experts. AI can already help interview a creator, organize ideas, search archives, produce derivative materials, and support audience interaction, but the market remains distributed across coaching, education, subscriptions, communities, publishing, and general AI tools rather than one well-defined category.

The enterprise and creator trends are adjacent rather than identical. Enterprises need institutional knowledge to remain usable across employees, systems, and time. Individual experts need personal knowledge to remain coherent across content, services, and audience interactions. Both involve externalizing judgment, but their governance, scale, economics, and risk requirements differ substantially.

Capital allocation and the major research reports

Venture-capital evidence supports a shift toward infrastructure, vertical applications, and agents, but it does not show a simple transfer of money from chat interfaces into organizational knowledge systems.

CB Insights reports that AI companies raised a record \$225.8 billion in 2025, including \$83.2 billion during the fourth quarter. Its 2025 research repeatedly highlights large infrastructure rounds, agent companies, and agent-supporting infrastructure; agents and supporting infrastructure represented 21% of the companies in its 2025 AI 100 selection. ³⁷

PitchBook’s second-quarter 2025 AI analysis found that horizontal AI platforms continued to lead in deal value, while vertical applications led in deal count. That distinction matters: the largest capital amounts remain concentrated in compute-intensive foundation models and broad platforms, while a larger number of application companies are being formed around domain-specific workflows. PitchBook’s 2026 work now frames AI investment as a multi-layer “super-cycle” associated with digital labor, rather than as a single chatbot category. ³⁸

Stanford’s 2026 AI Index reports that worldwide corporate AI investment reached approximately \$581.7 billion in 2025 and private investment reached \$344.7 billion, with generative AI accounting for nearly half of private AI funding. These totals are heavily influenced by very large model, infrastructure, and platform investments, so they should not be read as direct evidence that knowledge-management or governance startups are receiving most of the capital. ³⁹

The venture picture is therefore better described as **stack expansion.** Capital remains enormous at the model and infrastructure layers, while investors and enterprise buyers are increasingly looking for applications that can own a workflow, integrate proprietary context, or produce measurable outcomes. Menlo Ventures estimated that infrastructure captured roughly half of enterprise generative-AI spending in 2025, while a16z’s survey of 100 enterprise CIOs found AI spending moving from experimental innovation budgets into recurring IT and business budgets and procurement beginning to resemble conventional enterprise-software purchasing. ⁴⁰

The requested industry reports converge around several findings:

McKinsey finds broad adoption, substantial experimentation with agents, and limited enterprise-wide scaling. Its strongest differentiator between ordinary users and high performers is not model access but workflow redesign, growth orientation, operating-model changes, and human validation. ²

Gartner predicts rapid agent adoption and major workflow effects, but simultaneously warns about agentwashing, inflated expectations, weak data foundations, project cancellations, cost-to-value gaps, and governance failures. Gartner's analysis is supportive of the strategic shift but skeptical that the market's current language accurately reflects production maturity. ⁴¹

Deloitte reports rapidly expanding access and production expectations but finds that only 34% of organizations are genuinely reimagining their businesses and only one in five has mature autonomous-agent governance. It also identifies infrastructure, data, risk, and talent as readiness constraints. ⁴²

Accenture reports continued investment and agent pilots, but only 23% of organizations realizing widespread and sustained value. Its analysis argues that the differentiating work is redesigning infrastructure, roles, work, and operating models. ²⁰

Microsoft's Work Trend Index presents the strongest "digital labor" thesis: human-agent teams, agent management, new AI roles, and organizational structures built around on-demand intelligence. Its 2026 findings nevertheless show that employee capability frequently exceeds organizational readiness. ²¹

Stanford's AI Index shows record investment and widespread organizational AI use, while also showing that agent deployment remains in the single digits across nearly all business functions and that benchmark agents still fail a meaningful share of structured computer-use tasks. ⁴³

CB Insights identifies record funding, billion-dollar infrastructure rounds, verticalized agents, and supporting agent infrastructure as major investment themes. This supports market formation around agents but also demonstrates continued concentration of capital in foundational layers. ⁴⁴

PitchBook finds horizontal platforms dominating capital value while vertical applications dominate deal count. Its evidence suggests that infrastructure and domain-specific workflow companies are both expanding, not that one has replaced the other. ⁴⁵

a16z finds enterprise optimism about reasoning models but limited production adoption, growing preference for packaged applications, recurring budget treatment, multi-model procurement, and stricter cost and security evaluation. This is consistent with a maturing software market rather than an experimental prompting market. ⁴⁶

Sequoia distinguishes the exceptional consumer product-market fit of general assistants from the fragmented needs of professional work. It also warns that AI capital expenditure has not yet produced commensurate end-user revenue and argues that the next phase must focus on execution and enterprise value realization. ⁴⁷

Taken together, these reports do not describe a market abandoning chat. They describe a market encountering the limits of **ungrounded, isolated, user-initiated generation** and redirecting attention toward the organizational changes required for AI to participate in consequential work.

Counterevidence and final judgment

The strongest counterargument is that chat remains the most successful distribution mechanism in AI.

General-purpose assistants continue to attract enormous audiences. a16z's 2026 consumer analysis estimated that ChatGPT had approximately 900 million weekly users and retained a large engagement lead over competing assistants. Its earlier consumer research found limited cross-shopping among general assistants and strong concentration around a small number of products. ⁴⁸

Enterprise use also remains substantially conversational. McKinsey's 2025 respondents most commonly reported using AI to capture, process, and deliver information through conversational interfaces; provide content support for marketing; and automate customer service. These are not merely obsolete 2023 behaviors. They continue to represent real, widespread commercial use. ²

A second counterargument is that "reasoning" and "agentic" language may be running ahead of technical reliability. Gartner reports that only 17% of surveyed organizations have deployed agents and places agentic AI at the Peak of Inflated Expectations. Stanford's benchmark reporting shows major progress in agents operating computer environments but also indicates that leading systems still fail roughly one-third of structured operating-system tasks. ⁴⁹

A third counterargument is economic. Agent systems may consume more tokens, require more model calls, introduce more failure modes, and demand expensive integration, evaluation, observability, and governance. Gartner forecasts that more than 40% of current agentic projects will be canceled because of cost, unclear value, or control problems. Sequoia has likewise questioned whether end-user revenue can justify the capital expenditure supporting the broader AI buildout. ⁵⁰

A fourth is organizational inertia. The technologies associated with structured organizational intelligence require companies to clarify processes that may never have been formally defined, reconcile inconsistent data, establish decision rights, and decide who is accountable for machine-mediated actions. Deloitte found that employee education remains more common than actual role or workflow redesign, while Microsoft found that many capable workers lack supportive organizational conditions. ⁵¹

These counterarguments do not invalidate the observed shift. They define its current boundary.

The evidence supports the following final judgment:

There is a strategic market shift from AI as a conversational productivity feature toward AI as an organizational execution capability. This is visible in vendor language, product architecture, labor demand, governance frameworks, enterprise surveys, and venture narratives. ⁵²

The shift is additive rather than substitutive. Chat remains the dominant interaction pattern and a major commercial category. In many emerging systems, chat will survive as the command surface through which users invoke structured context, agents, workflows, and tools. ⁵³

The scarce capability is moving below the interface. As models become more broadly available, differentiation increasingly comes from organizational context, proprietary data, semantic structure, process integration, permissions, evaluations, governance, and the capacity to execute reliably. 54

Knowledge is becoming operational, but not automatically executable. Companies are attempting to convert documents, interactions, business objects, expert judgment, and process history into context that machines can use. Preserving meaning, exceptions, provenance, and authority through that conversion remains an unresolved design problem. 55

Governance is becoming part of the product architecture rather than an external policy document. Once AI can act, control must be embedded in access rights, tool definitions, approval gates, monitoring, evaluation, and lifecycle management. 56

The strongest defensible thesis for an article is therefore not:

“The chatbot era is over.”

It is:

The chatbot was the market’s introduction to AI. The next contest is over who can give AI enough context, structure, authority, and restraint to participate in real work.

Your advertisements were not proof of that transformation. They were compressed expressions of it: one foregrounded business ontology and reasoning, another organizational structure, another tacit knowledge, another packaged expertise, and another the human labor needed to operate AI workflows. The broader research shows that this convergence is real—but also that the operating layer being advertised is still being built.

navlistRecent enterprise-agent developmentsturn13news41turn13news40

-
- 1 4 5 **Introducing Copilot for Microsoft 365 | Microsoft 365 Blog**
<https://www.microsoft.com/en-us/microsoft-365/blog/2023/03/16/introducing-microsoft-365-copilot-a-whole-new-way-to-work/>
- 2 3 18 22 53 **The State of AI: Global Survey 2025 | McKinsey**
<https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>
- 6 52 **Artificial Intelligence (AI) at Salesforce | Salesforce**
<https://www.salesforce.com/artificial-intelligence/>
- 7 15 54 **AIP architecture overview • Palantir**
<https://palantir.com/docs/foundry/architecture-center/aip-architecture/>
- 8 **Glean – Work AI that Works | Agents, Assistant & Search**
<https://www.glean.com/>
- 9 16 **WRITER Launches AI HQ to Revolutionize Agentic Work in the Enterprise**
<https://writer.com/blog/writer-ai-hq-press-release/>

10 17 56 **AI Agents - ServiceNow**

<https://www.servicenow.com/products/ai-agents.html>

11 **IBM watsonx Orchestrate**

<https://www.ibm.com/products/watsonx-orchestrate>

12 **Business Orchestration Software for Enterprise Processes | UiPath**

<https://www.uipath.com/platform/agentic-automation/business-orchestration>

13 **How SAP Uniquely Delivers AI Agents with Joule | SAP News Center**

<https://news.sap.com/2025/02/joule-sap-uniquely-delivers-ai-agents/>

14 **Gartner Predicts 40% of Enterprise Apps Will Feature Task- ...**

https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025?utm_source=chatgpt.com

19 28 42 51 **The State of AI in the Enterprise - 2026 AI report | Deloitte US**

<https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/content/state-of-ai-in-the-enterprise.html>

20 **Accenture Pulse of Change**

<https://www.accenture.com/us-en/insights/pulse-of-change>

21 **The 2025 Annual Work Trend Index: The Frontier Firm is born - The Official Microsoft Blog**

<https://blogs.microsoft.com/blog/2025/04/23/the-2025-annual-work-trend-index-the-frontier-firm-is-born/>

23 26 55 **Closing the skills gap | Deloitte Insights**

<https://www.deloitte.com/us/en/insights/focus/human-capital-trends/2025/closing-the-experience-gap-through-talent-development.html>

24 **APQC Study Warns of Looming “Great Retirement” Crisis, Highlights Role of AI and Knowledge Management to Mitigate Risk | APQC**

<https://www.apqc.org/about-apqc/news-press-release/apqc-study-warns-looming-great-retirement-crisis-highlights-role-ai>

25 **Update knowledge base with Customer ...**

https://learn.microsoft.com/en-us/dynamics365/release-plan/2025wave1/service/dynamics365-contact-center/use-customer-knowledge-management-agent-update-knowledge-base?utm_source=chatgpt.com

27 **AI Risk Management Framework | NIST**

<https://www.nist.gov/itl/ai-risk-management-framework>

29 49 **2026 Hype Cycle for Agentic AI | Gartner**

<https://www.gartner.com/en/articles/hype-cycle-for-agentic-ai>

30 33 **The Stanford AI Index 2026**

<https://lightcast.io/resources/research/stanford-ai-index-2026>

31 **The Generative AI Job Market: 2025 Data Insights**

https://lightcast.io/resources/blog/the-generative-ai-job-market-2025-data-insights?utm_source=chatgpt.com

32 **Prompt Engineer: Analyzing Skill Requirements in the AI ...**

https://arxiv.org/html/2506.00058v1?utm_source=chatgpt.com

34 **Executive Summary: 2025 Work Trend Index Annual Report**

<https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/final/en-us/microsoft-product-and-services/ai/pdf/executive-summary-work-trend-index-annual-report.pdf>

35 **Patreon | News**

<https://news.patreon.com/articles/state-of-create>

36 **The Operating System for Human Expertise**

https://www.kajabi.com/product?utm_source=chatgpt.com

37 44 **State of AI 2025 Report - CB Insights Research**

https://www.cbinsights.com/research/report/ai-trends-2025/?utm_source=chatgpt.com

38 45 **AI VC Trends**

https://pitchbook.com/news/reports/q2-2025-ai-vc-trends?utm_source=chatgpt.com

39 **Inside the AI Index: 12 Takeaways from the 2026 Report**

https://hai.stanford.edu/news/inside-the-ai-index-12-takeaways-from-the-2026-report?utm_source=chatgpt.com

40 **2025: The State of Generative AI in the Enterprise**

https://menlovc.com/perspective/2025-the-state-of-generative-ai-in-the-enterprise/?utm_source=chatgpt.com

41 50 **Gartner: Over 40% of Agentic AI Projects Will Be Canceled ...**

https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027?utm_source=chatgpt.com

43 **Economy**

https://hai.stanford.edu/assets/files/ai_index_report_2026_chapter_4_economy.pdf?utm_source=chatgpt.com

46 **How 100 Enterprise CIOs Are Building and Buying Gen AI in 2025 | Andreessen Horowitz**

<https://a16z.com/ai-enterprise-2025/>

47 **AI in 2025: Building Blocks Firmly in Place | Sequoia Capital**

<https://sequoiacap.com/article/ai-in-2025/>

48 **The Top 100 Gen AI Consumer Apps — 6th Edition**

https://a16z.com/100-gen-ai-apps-6/?utm_source=chatgpt.com