

# The State of Enterprise Agentic AI in 2026

Agentic Reality Check:  
Hype or Not?

**Looking past the hype: What Agentic AI really looks like in large, complex enterprises.**

Beyond the Silicon Valley narrative, what does agentic AI truly look like inside a \$5B+ revenue organization? This research explores the gap between "Agent-Washing" and the realities of deploying agents in complex, regulated, and legacy-heavy environments.

# Executive Summary

**The enterprise AI conversation is saturated with superlatives.** “Agentic AI” has become the defining term of 2025–2026, promising a future where autonomous digital workers independently pursue goals, orchestrate complex workflows, and transform business operations. Venture capital is pouring in. Vendor marketing has reached a fever pitch. And yet, when we ask 740 senior executives from companies generating \$1B–\$20B+ in annual revenue what agentic AI really looks like inside their organizations, a very different picture emerges.

This research was designed to cut through that noise to establish, with rigorous survey data, where the enterprise market truly stands on its journey toward agentic AI. What we found confirms some suspicions, challenges others, and surfaces a set of structural challenges that the vendor community and enterprise leaders alike must confront honestly.

## Five Core Findings

### 01 The Adoption Reality: agentic AI is still rare despite the narrative

51% of respondents report having AI agents in live production environments. On its face, a stronger adoption signal than many anticipated, but the reality of these deployments tells a different story. Only 24% have deployed anything that qualifies as a true agent (autonomous task execution and tool use or multi-agent collaboration), and **only 10% have deployed true agentic AI capabilities**. The majority are operating assistants and workflow-guided AI — sophisticated tools, but not the autonomous agents the market is promising. This misunderstanding of agentic AI is made worse by the second key finding: pervasive agent-washing.

### 02 The Agent-Washing Crisis: a systemic problem with real consequences

84% of enterprise leaders encounter products marketed as “agents” that are, in reality, sophisticated chatbots. This is not a minor nuisance; 88% report it has negatively affected their trust in AI broadly, with 29% saying this has made it materially harder to secure budget for legitimate agentic projects. Market credibility is eroding from within.

### 03 The Trust & Governance Gap: the defining barrier

The top three barriers enterprises cite are all variants of the same fundamental problem: they cannot trust the outputs. Reliability/hallucinations (43%), security and privacy (42%), and accuracy (40%) dominate. 86% of executives cite one of these three barriers as a top blocker to implementation. Enterprises are not being held back by a lack of willingness to invest; 71% say they are investing “a lot.” They are being held back by agents that behave unpredictably.

### 04 The Knowledge Access Gap: limiting AI value

The question of what agents know is as critical as how they behave. Only 19% of respondents use basic/manual knowledge management — most have progressed to some form of RAG architectures. But most of these organizations are still using naïve RAG and lack the advanced agentic RAG knowledge fabric needed to handle agentic complexities and tackle high-value use cases. The result is that the knowledge pipeline is leaking at multiple points: 38% struggle with data that doesn’t update, 38% fear giving agents access to sensitive or regulated data, and 31% are hamstrung by data silos. Not all RAG is created equal, and agents are only as good as what they know.

### 05 What Enterprises Are Building Toward: the decade of agents

Andrej Karpathy was right: this is the decade of agents, not the year. The foundations for transformative agentic AI are being laid right now: investment is high, production deployments are growing, and the capability roadmap is clear. But the structural challenges around trust, governance, and knowledge are deep enough that widespread, high-value, autonomous agent deployment is a multi-year journey. Enterprises that build the right foundations today will be positioned to capture disproportionate value when the technology matures.

The sections that follow examine each of these findings in depth, drawing on granular survey data across industries, company sizes, and geographies. For enterprise leaders navigating investment decisions, and for technology vendors designing the next generation of agentic platforms, this report offers the most grounded, data-driven picture of enterprise agentic AI available today.

# The Adoption Reality

## The “Agentic AI” Adoption Misconception

Agentic AI is everywhere, or soon to be. At least, that’s the narrative that has been circulating corporate circles throughout much of 2025 and early 2026. In our survey of 740 senior business and technology leaders at some of the world’s largest enterprises, that perception has certainly taken hold. When asked about the maturity of their AI adoption, an eye-popping 51% report having AI agents in live production. On the surface, this figure looks genuinely significant, and it could be one of the reasons for some of the misleading headlines around agentic AI adoption today.

**Figure 1 – Self-reported Agentic AI Adoption Maturity**




Figure 1 – Over half of respondents report having agentic AI live in production, but there’s more to the story.

But dig a little deeper, and it becomes clear that the data is more a reflection of confusion in the market than true adoption. Only 13% of the leaders surveyed reported that they could easily tell the difference between a real AI agent and an AI assistant or even a simple chatbot. When probed further, the reality of AI deployment tells a very different story.

## The Deployment Sophistication Gap

Our research defines a spectrum of six deployment types, from simple chatbots to fully collaborative multi-agent systems:

**Figure 2 – Deployment Sophistication: What “Agentic AI” Really Looks Like**




Figure 2 – Deployment Sophistication: The Spectrum from Chatbot to Multi-Agent

The critical finding here is that 71% of respondents, a great majority of the market, describe their most advanced deployment as assistive AI or below. These are knowledge-retrieval systems, not autonomous actors. They cannot independently pursue goals, make decisions, or take actions in the world. Calling them “agents” is, at best, aspirational marketing.

## The Scaling/Mature Paradox

76% of organizations describe themselves as “scaling” or “mature” in their agentic AI journey. Yet when you examine what those same organizations have actually deployed, less than half of them (46%) are running true agents or multi-agent systems. This is not dishonesty. It reflects genuine confusion in the market about what “agentic” means. When the vocabulary is this imprecise, self-reported maturity becomes an unreliable signal.

Agent capability, or the ability for an AI system to independently decide how to pursue a goal, select tools, take actions, and adapt to results, is reported by only 29% of respondents. Multi-agent deployments, where multiple specialized agents collaborate, sit at just 11%. And these results include organizations reporting to be in some stage of investigation or pilot. When removing this group, only 24% demonstrate any true AI agent deployments. Only 10% of the enterprises demonstrate multi-agent systems: true agentic AI. These are the systems that represent the transformative potential the market is promising. They exist, they are growing, but they remain a minority.

**76%**

do not demonstrate any AI agent deployments

**14%**

demonstrate true AI agents, deployed in production

**10%**

demonstrate fully agentic AI deployments in production

## The Autonomous Budget Signal

One of the most telling indicators of genuine agentic intent is the proportion of the AI budget allocated to fully autonomous agents, those that make decisions and execute tasks without human intervention. The distribution is revealing:

**40%**

allocate only 0-24% of budget to fully autonomous agents

**33%**

allocate 25-49%: careful, monitored experimentation

**27%**

allocate 50%+ to fully autonomous execution

The majority of enterprise AI investment is still oriented toward human-assisted or human-validated workflows. Genuine autonomy, where the agent executes without human checkpoints, is a minority budget allocation. This is rational given current trustworthiness concerns, but it also means the productivity and scale benefits associated with full autonomy remain largely unrealized.

## Geographic Reality: USA Leads, Europe Follows Carefully

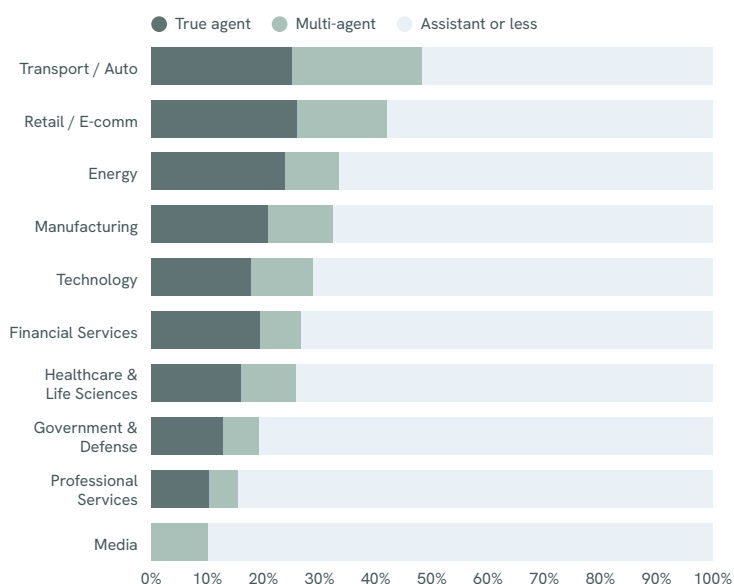
The geography of adoption reflects differing risk appetites, regulatory environments, and investment cultures. The United States is moving faster and more ambitiously; Europe is moving more deliberately.

Section 01: The Adoption Reality

The sharpest divergence is in multi-agent deployments: 12% of US organizations vs. 8% in Europe. The investment gap is also material. 72% of US respondents invest “a lot” in AI, compared to 60% in Europe. Europe’s more cautious posture reflects several converging factors: stricter regulatory frameworks (GDPR and the EU AI Act), more conservative investment cultures, and, critically, higher rates of agent-washing fatigue that make it harder to build the internal case for real agentic investment.

One particularly striking European data point: only 5% of European leaders say they can “easily tell the difference” between a real agent and an agent-washed chatbot, versus 15% in the US. This lower baseline of vendor literacy contributes to a more skeptical and slower-moving market.

Figure 3 – The Industry Dimension: Who Is Leading and Why




Industry context matters enormously. The most surprising leader in our data is Transportation & Automotive. Nearly half of all respondents in that sector are running true or multi-agent systems, and uniquely, multi-agent (23%) almost matches true agent (25%), suggesting a sector that has moved quickly to coordinated, collaborative AI rather than stopping at single-agent deployments. The operational imperative is clear: when the cost of supply chain failure is measured in millions per hour, the ROI case for intelligent process automation is self-evident.

Technology being fifth is perhaps the most counterintuitive finding in the entire dataset. Despite the highest overall investment and production rates, tech companies are actually mid-table on genuine agentic sophistication. Only 29% have moved beyond workflow AI. The sector appears to be scaling volume (many deployments) rather than depth (true autonomy).

Government & Defense and Professional Services bring up the rear, which is consistent with the expectations. They are conservative environments with high scrutiny and significant barriers around trust and governance.

Figure 4 – Business Functions Currently Running AI Agents or Agentic AI



IT Operations and Cybersecurity stand as the dominant deployment functions at 78%, more than 12 percentage points ahead of Customer Service. This reflects a pragmatic deployment logic: IT operations generate high volumes of structured, repetitive actions where agent errors are recoverable and measurable. Security monitoring and alert triage represent similar profiles. These are ideal proving grounds for agentic AI.

The relatively lower penetration of strategic functions — industry-specific operations at 39%, Legal/Risk at 30%, Engineering/R&D at 29% — reflects the sophistication required for these use cases. Deploying agents in legal risk or R&D environments requires deeper domain knowledge, higher accuracy standards, and more sophisticated governance than IT operations automation. These will be the next frontier as trust is established in lower-stakes deployments and knowledge layer is built out.


# The Agent-Washing Crisis

## How Market Hype Is Poisoning the Agentic AI Well


Perhaps no finding in this research is more consequential for the long-term health of the agentic AI market than the scale and impact of agent-washing. The term describes a practice that has become pervasive: vendors rebranding existing tools (chatbots, AI assistant search interfaces, workflow automation platforms) with the “agentic AI” label, without the underlying capability to match.

Our data reveals that this is not a fringe phenomenon. It is the dominant market experience.

**Figure 5 – Among “Scaling” or “Mature” enterprises, actual deployment type**



**Figure 6 – How often do vendors market chatbots as “agents” during evaluations?**



**84%**

encounter agent-washed products “sometimes,” “often,” or “always”

**38%**

encounter it “often” or “always.” Virtually every evaluation

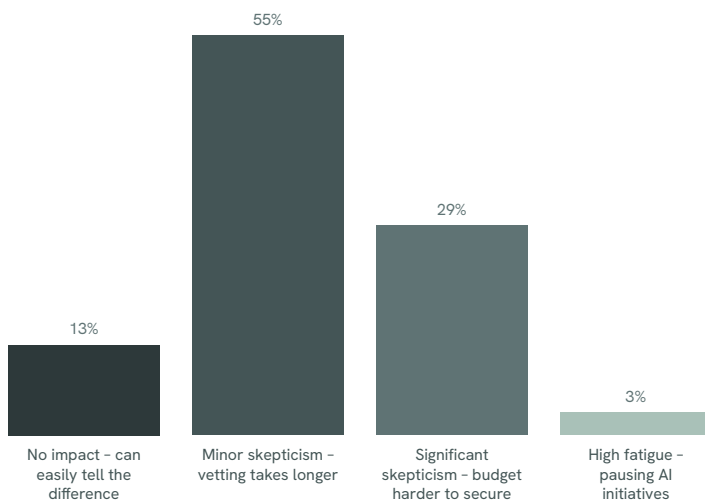
**88%**

report some negative impact on their trust in AI broadly

## Why Agent-Washing Is More Than a Marketing Problem

The conventional wisdom treats agent-washing as a vendor credibility issue, a problem of misleading marketing that sophisticated buyers can navigate. Our data suggests the consequences run deeper and are more systemic. More than 87% of leaders report that agent-washing has had a material impact on the trust of vendors and AI initiatives.

**Figure 7 – Impact of Agent-Washing on Enterprise Trust in AI**



### Budget Destruction at Scale

32% of senior executives say agent-washing has made it materially harder or even impossible to secure budget for legitimate agentic AI projects. In practical terms, this means the noise of fraudulent claims is actively reducing investment in real capability. A technology with genuine transformative potential is being slowed by its own marketing ecosystem.

This trust destruction is cumulative. Each time a technology leader evaluates a product marketed as an “autonomous agent” and finds a glorified chatbot, they return to their organization with a more skeptical lens. Over enough iterations, this skepticism becomes organizational orthodoxy, or a generalized distrust of “AI agents” that makes every future evaluation harder.

There is a secondary consequence that is perhaps even more damaging: enterprise leaders begin to doubt their own ability to evaluate AI claims. When 95% of European leaders and 85% of US leaders cannot reliably distinguish a real agent from an agent-washed product, the evaluation process itself breaks down. Procurement decisions get made on brand familiarity and sales relationships rather than genuine capability assessment, or don’t get made at all.

## What Constitutes a True Agent? The Enterprise Definition

One root cause of agent-washing is definitional ambiguity. When the market lacks a shared vocabulary for what “agentic” means, vendors are free to claim the label with whatever justification suits them. Our survey asked enterprise leaders directly: **what capabilities must a system have for you to consider it genuinely agentic?**

The capabilities most commonly cited as minimum requirements for true agentic status:	What is notably absent from the “minimum viable” definition:
<ul style="list-style-type: none"> <li>• <b>Goal pursuit</b> – ability to work toward an objective autonomously: 55%</li> <li>• <b>Self-correction</b> – reasoning through errors and adapting: 50%</li> <li>• <b>Multi-step planning and reasoning</b>: 49%</li> <li>• <b>Workflow execution</b> with branching decisions: 49%</li> <li>• <b>Tool use</b> – taking actions in external systems: 42%</li> <li>• <b>Company-specific knowledge</b>: 43%</li> </ul>	<ul style="list-style-type: none"> <li>• Memory and self-improvement ranked lowest at 27% – enterprises understand this is aspirational, not foundational</li> <li>• Natural language understanding alone (36%) is explicitly insufficient – the market understands that “it talks to me” does not make something an agent</li> </ul>

## The Social Proof Trap

There is a more subtle driver of adoption inflation that our data illuminates: the pressure to appear on the leading edge. When peer organizations, industry media, and vendor case studies create the impression that everyone is deploying agents, there is a powerful social incentive to claim similar progress, even when the underlying deployment is more modest.

This helps explain why 28% of organizations that self-identify as “scaling” or “mature” agentic AI adopters have, upon closer examination, only deployed assistants or chatbots. These are not intentionally misleading responses. They reflect genuine uncertainty about where on the capability spectrum their deployments fall, combined with a natural inclination to frame one’s organization favorably.

The implication for enterprise leaders is important: treat competitive benchmarking on agentic AI with significant skepticism. The peer comparisons being made in board presentations and analyst reports are likely systematically overstating genuine capability deployment.


# The Trust & Governance Gap

## Why Trustworthiness is the Defining Barrier

If there is a single word that defines the enterprise agentic AI challenge in 2026, it is trust. Enterprises lack confidence that an AI system will behave reliably, accurately, and safely enough to be given meaningful autonomy to act.

Our data is unambiguous on this point. When enterprise leaders rank the barriers preventing them from deploying agentic AI at scale, the top three are all variants of the same fundamental problem:

Figure 8 - Top Barriers to Enterprise Agentic AI Implementation

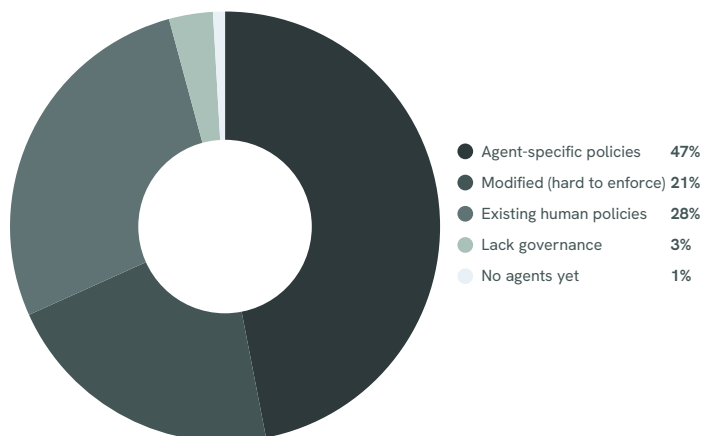


Reliability/hallucinations (43%), security & privacy (42%), and accuracy (40%) are not independent concerns. They are three dimensions of a unified trust deficit. Enterprises are not asking “can we afford this technology?” They are asking “can we trust this technology to behave correctly?” That is a fundamentally different question, and it has a fundamentally different solution path.

## The Governance Infrastructure Gap

Behind the trust question lies a governance infrastructure that is still being built. The frameworks, tools, and organizational practices required to control, monitor, audit, and constrain AI agents operating at scale simply do not yet exist at the maturity level that large enterprises require.

Figure 9 - Enterprise Governance Readiness: Framework Status



### The Governance Math Does Not Add Up

Of the organizations that have true agentic deployments, 53% do not have agent-specific governance policies. They are either applying human-centric frameworks that were not designed for AI agents (28%), attempting to build modified policies that they acknowledge are difficult to enforce (21%), or frankly lacking adequate governance altogether (4%). Only 47% have purpose-built policies for AI agents taking autonomous action.

This gap between deployment velocity and governance maturity is one of the most consequential structural risks in the current market. Organizations are granting AI systems the ability to take real-world actions — executing trades, updating CRMs, routing customer inquiries, managing supply chain decisions — with governance frameworks that were designed for human actors operating within human-scale processes.

## What Governance Capabilities Actually Exist

The good news is that the market is actively building governance capability. The bad news is that the most safety-critical capabilities are also the least deployed:

Most commonly deployed:	Less commonly deployed, higher effort, higher value:
<ul style="list-style-type: none"> <li>• <b>Comprehensive Observability</b> - monitor performance and alert on failures: 53%</li> <li>• <b>Multi-Agent Conflict Resolution</b> - manage contradictory agent actions: 52%</li> <li>• <b>Data Traceability</b> - trace agent-used data to source: 51%</li> <li>• <b>Real-time Kill Switch</b> - ability to terminate all agent sessions: 44%</li> </ul>	<ul style="list-style-type: none"> <li>• <b>Usage Limits</b> - token, budget, tool caps: 31%</li> <li>• <b>Dynamic Entitlement Management</b> - agent-specific identity/access: 34%</li> <li>• <b>Cross-System Audit Trails</b> - logs across siloed systems: 40%</li> <li>• <b>Human-in-Loop Verification</b> - human approves before action: 41%</li> <li>• <b>Agent Guardrails</b> - prevent undesired inputs/outputs: 42%</li> </ul>

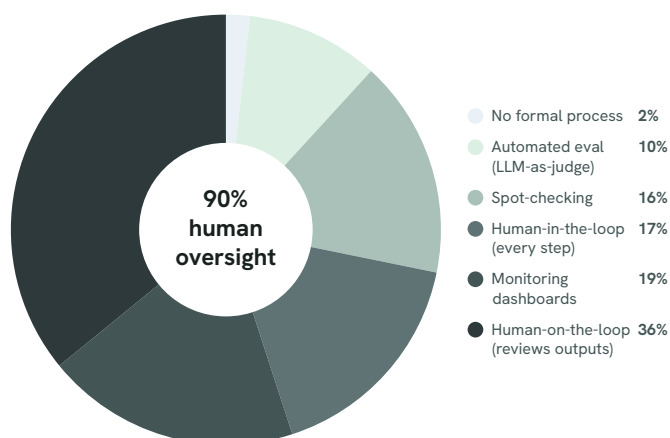
### The Pattern

The easiest capabilities to implement (kill switches, monitoring) are most deployed. The capabilities that require deep integration with existing enterprise systems (dynamic entitlements, structured human checkpoints, agent guardrails) lag significantly.

## The Validation Problem: Who Is Watching the Agents?

Even where governance frameworks exist, the practical question of how organizations actually validate agent outputs reveals a market still heavily dependent on human oversight.

Figure 10 - How Enterprises Currently Validate AI Agent Outputs



Fully 88% of validation approaches involve human review of some kind, either directly (human-in-the-loop at 17%, human-on-the-loop at 36%) or indirectly through spot-checking (16%) and monitoring dashboards (19%). Only 10% of enterprises have moved to automated evaluation frameworks (LLM-as-judge), and 2% have no formal validation process at all.

This is the right posture given current maturity levels. Deploying AI agents without any human oversight would be genuinely reckless. But it also means the productivity promise of autonomous AI agents remains largely unrealized. An agent that requires human review of every output is, in practice, an assistant rather than an autonomous actor.

### The Skills Gap Hidden Inside the Trust Problem

The fifth-highest barrier in our data is a skills gap (35%). Internal teams lack the AI expertise to validate whether agent outputs are correct. This creates a dangerous situation wherein organizations are deploying agents, governance frameworks require human validation, but the humans responsible for validation lack the technical grounding to do so reliably. The quality of human oversight is inversely correlated with the complexity of agent behavior.

## Trust and Governance by Industry and Size

Trust and governance challenges are not uniform. Healthcare and Government & Defense face the highest stakes. Patient safety and national security create zero-tolerance environments for agent misbehavior. Healthcare's governance gap (10% lack adequate governance) is concerning given these stakes. Government & Defense, by contrast, has high governance coverage but very low production deployment, suggesting that governance requirements are themselves a deployment bottleneck.

On company size, a counterintuitive pattern emerges: the largest organizations (100K+ employees) have the lowest rates of agent-specific governance (33%), relying instead on monitoring dashboards. This likely reflects the organizational complexity of implementing agent-specific policies across a 100,000+ person enterprise. Size creates governance challenges that smaller, more agile organizations can sidestep.

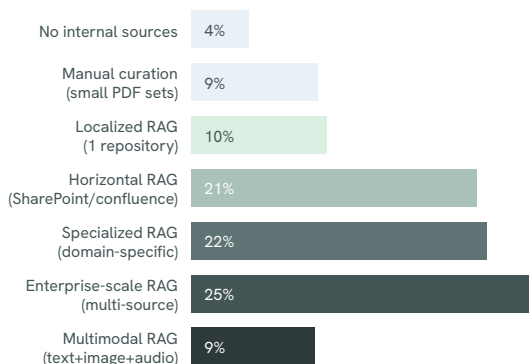
# The Knowledge Access Gap

## What Agents Know — and What They Don't


There is a popular misconception that what makes an AI agent capable is primarily a question of reasoning sophistication, tool access, or model intelligence. These matter. But for enterprise deployment, where agents must navigate the specific, idiosyncratic, constantly-evolving knowledge environment of a large organization, dynamically and securely connecting agents to siloed, multimodal, and domain specific knowledge with context is even more important.

An agent that can reason brilliantly from incomplete information will still fail in production if that information does not include, for example, the current pricing policy, the specific regulatory requirement for this jurisdiction, or the proprietary methodology that distinguishes your organization's approach from a generic one.

**Figure 11 – Knowledge Readiness: How Enterprises Manage AI Knowledge**

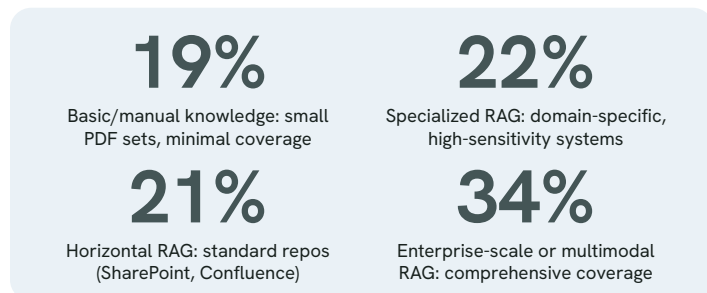


**Figure 12 – Knowledge Readiness: Where Enterprises Get Stuck**



## The Knowledge Management Spectrum

The evolution from basic knowledge management to enterprise-scale knowledge fabric represents one of the most significant technology investments an organization can make in service of agentic AI. Our data shows the market distributed across this spectrum:



The most significant gap in the market is between organizations that have achieved enterprise-scale knowledge fabric (35%) and those still operating on limited, siloed, or manually-curated knowledge bases. An agent connected to enterprise-scale RAG has access to the full breadth of organizational knowledge, documents, databases, and operational systems, with appropriate security controls. An agent connected to a small set of PDFs uploaded by hand is confined to a very narrow slice of organizational intelligence.

This becomes clear when comparing organizations that have true agentic deployments in production. Here we see 49% that use enterprise-scale or multimodal RAG. This is 14.5 percentage points higher than the average across all respondents. Truly agentic organizations are nearly twice as likely to have the most sophisticated knowledge architectures. The reason for this divide is obvious: without advanced RAG knowledge architectures, the reliability and value of agents is limited.

## The Low-Value Use Case Trap

Without enterprise-scale knowledge access, agents are constrained to use cases that can be adequately served by generic knowledge or a small curated corpus. These tend to be the low-value, task-automation use cases: scheduling, basic Q&A, form processing. The high-value use cases, like clinical decision support, regulatory compliance monitoring, complex financial analysis, strategic market intelligence, all require deep, current, domain-specific knowledge. Organizations investing in agentic AI without a parallel investment in the knowledge layer are funding the wrong half of the stack.

## The Five Knowledge Obstacles

Our survey identified a hierarchy of knowledge obstacles that prevent agents from being fully informed. Understanding this hierarchy is essential for prioritizing knowledge layer investments:

# Obstacle	Enterprises Affected
1 <b>Dynamic data:</b> agents work with static data; real-time information doesn't auto-update	38%
2 <b>Data sensitivity:</b> fear of giving agents access to confidential, PII, or regulated data	38%
3 <b>Data silos:</b> too many disparate and disconnected systems	31%
4 <b>Connectivity:</b> source systems that agents can't connect to	31%
5 <b>Data permissions:</b> no consistent way to handle data access rights for agents	30%
6 <b>Volume overload:</b> too much data overloads context windows or causes hallucination	22%

## The European Knowledge Disadvantage

The knowledge readiness gap is most pronounced in Europe. Dynamic data obstacles affect 49% of European respondents versus 35% in the US, a 14.6 percentage-point gap. Connectivity issues affect 37% in Europe versus 29% in the US. These are not primarily technology gaps; they reflect the reality that European enterprise data architectures tend to be older, more siloed, and more conservative in their approach to data integration.

GDPR creates an additional structural layer. The regulation's requirements around consent, purpose limitation, and data minimization are in tension with the way agentic AI systems naturally operate, pulling from broad data sources, retaining context, and potentially combining information in ways that create new privacy implications. European data sensitivity concerns reflect both the regulatory reality and a cultural orientation toward data protection that will shape European agentic AI deployment patterns for years to come.

## The RAG Architecture Evolution

For organizations that have progressed beyond manual knowledge curation, Retrieval-Augmented Generation (RAG) has become the de facto standard for connecting agents to enterprise knowledge. Our data shows a clear progression from localized to enterprise-scale RAG, with each tier unlocking a qualitatively different range of use cases:

### Naïve RAG

**Basic RAG tiers** (19% of market): while these may be decent if not limiting for assistive AI use cases, they are not practical for agentic AI.

- Manual curation: small PDF sets, high maintenance, narrow coverage
- Localized RAG: single repository, useful for specific use cases but not scalable

**Mid-tier RAG** (43% of market): severely limiting agentic AI to low-value use cases, despite being broadly adopted.

- Horizontal RAG: standard enterprise repos (SharePoint, Confluence); broad but shallow
- Specialized RAG: domain-specific, high-sensitivity systems, deep coverage in a vertical

### Advanced RAG

**Enterprise-grade RAG** (25% of market): appropriate for broad use cases and enterprise-wide deployments.

- Enterprise-scale RAG: secure multi-source, domain-specific silos, full organizational coverage

### The Future: Multimodal RAG

9% have deployed multimodal RAG, extending beyond text to images, audio, and video. For industries like healthcare (medical imaging), manufacturing (equipment photos), and financial services (charts and visual reporting), multimodal knowledge access is transformative. It remains nascent today but represents the direction of travel.

# What Enterprises Are Building Toward


## Investment Priorities, Capability Gaps, and the Decade Ahead

The research to this point has been largely diagnostic, describing where the enterprise market stands today. This section turns forward. What are organizations trying to achieve with agentic AI, what capabilities do they most need to unlock that potential, and how should the next several years of this technology evolution be understood?

### Strategic Investment Priorities: What Drives the Business Case

Understanding why organizations are investing in agentic AI is as important as understanding where they are deploying it. Our survey asked leaders to rank six strategic outcomes from most to least important to their agentic AI investment.

Figure 13 – Strategic Investment Priorities for Agentic AI



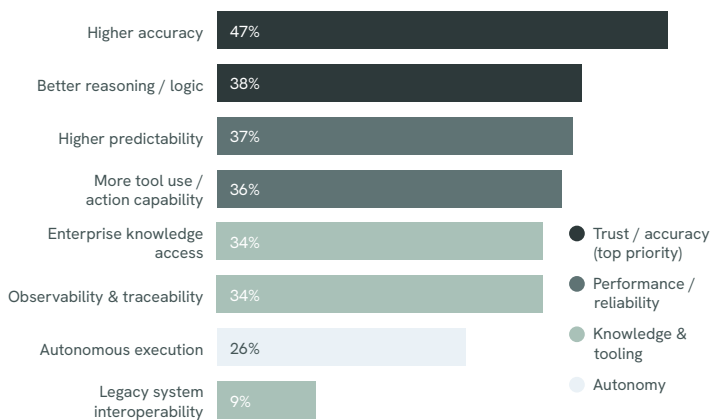
Ranked by enterprise leaders, 1=most important, 6=least important

The ranking is instructive. Productivity and acceleration (doing more with the same resources) edges out pure efficiency/cost-cutting as the primary value driver. This is a meaningful distinction. Efficiency framing is about reducing costs; productivity framing is about expanding capacity. Organizations are more attracted to the idea of agentic AI enabling them to do things they couldn't do at all, rather than just doing current things more cheaply.

Compliance sits at the bottom of the priority ranking (average rank 3.84). This may reflect where organizations are in their adoption journey. Compliance-driven deployment tends to come after value-driven deployment is established. It also suggests an underappreciation of regulatory risk in agentic AI deployment that may create problems later.

## The Capability Gaps That Must Close

Figure 14 – Capabilities Most Needed to Unlock High-Value Enterprise Use Cases (select up to 3)



When looking at the gaps that organizations need to close in order to unlock higher value use cases, accuracy (47%) tops the list, consistent with the evidence of trust and knowledge barriers throughout this research. That finding is even more pronounced in enterprises with true agentic deployments, where 55% cite higher accuracy as the top gap. Organizations know what they want to build; they need to be able to trust what agents produce before they can deploy them in high-stakes environments.

Better reasoning/logic (38%) and higher predictability (37%) follow closely. These are related to accuracy but distinct: reasoning quality determines what agents can tackle; predictability determines whether outputs are reliable enough to act on without extensive review.

The lower-ranked needs are also revealing. Autonomous execution (26%) and legacy system interoperability (9%) are notably deprioritized. The implication is that the market has largely accepted that full autonomy is a future state, not an immediate requirement, and that the value of connecting to legacy systems, while real, is less urgent than improving core trustworthiness.

### The Decade of Agents: A Realistic Timeline

Andrej Karpathy's observation, that "2025-2035 is the decade of agents," provides a useful frame for interpreting this research. Viewed through a one-year lens, the market looks frustratingly immature: most enterprises are not running true agents, governance frameworks are incomplete, knowledge layers are not fully built, and trust barriers are high.

Viewed through a ten-year lens, the market looks like exactly what you'd expect at the early inflection point of a transformative technology. The core capabilities exist and are demonstrably valuable. The adoption curve has crossed the production threshold. Investment is committed and growing. The obstacles are structural and addressable. They are the kind of challenges that take three to five years of focused engineering and organizational adaptation to solve, not indefinite blocking problems.

## The Compounding Advantage

The organizations that are building trust infrastructure, knowledge architecture, and governance frameworks today, even when it slows their initial deployment velocity, will have compounding advantages as the technology matures. When models become more capable, when governance standards crystallize, when knowledge architecture tools mature, the organizations that have the underlying readiness will be able to move fast. Those that cut corners today will face legacy cleanup problems that constrain their ability to accelerate.

The path from today's constrained deployment reality to tomorrow's transformative vision runs through four domains:



### Trust Infrastructure

Building the evaluation, monitoring, and validation systems that make agent behavior inspectable and improvable. Usage limits, access controls, and behavioral guardrails are foundational.



### Knowledge Architecture

Progressing from manual curation to enterprise-scale RAG, and eventually to multimodal, real-time knowledge pipelines. This is the investment that determines the ceiling of agent value.



### Governance Maturity

Moving from human-centric oversight to agent-native governance: dynamic entitlements, kill switches, cross-system audit trails, and multi-agent coordination protocols. Governance is not a constraint on value; it is an enabler of autonomy.



### Organizational Capability

Building the internal expertise to design, evaluate, validate, and continuously improve agentic systems. The skills gap is real and will not be solved by technology alone. It requires deliberate investment in human capability alongside AI capability.

# Strategic Recommendations

## What Enterprise Leaders Should Do Now

The findings of this research do not call for a pause on agentic AI investment. They call for investment that is grounded in reality. One that acknowledges the current constraints honestly, builds the foundational infrastructure systematically, and sets expectations that align with the decade-scale transformation underway.

### For C-Suite and Senior Business Leaders

- 01 Demand genuine capability assessment, not vendor claims.**  
Require vendors to demonstrate the specific capabilities that define true agency: goal pursuit, tool use, autonomous action, and self-correction. Request live demonstrations against your specific use cases, not polished demos. The 84% agent-washing rate means the probability that a marketed "agent" is actually an agent is low. Assume it isn't until proven otherwise.
- 02 Invest in a knowledge fabric as a prerequisite, not an afterthought.**  
Enterprise-scale and multimodal RAG is not a nice-to-have feature. It is the foundation that determines the ceiling of agent value. Organizations that invest in knowledge architecture in parallel with agent deployment will have access to a fundamentally different quality of use cases than those that do not. Budget accordingly.
- 03 Build governance architecture before you need it.**  
Agent-specific governance frameworks, dynamic entitlement management, and cross-system audit trails take time to design and implement. Organizations that wait until they have scaled deployments to address governance will face a retrofit problem that is exponentially harder. Start with agent-native governance while deployments are small.
- 04 Calibrate expectations to the decade, not the year.**  
Set internal success metrics that reflect the developmental stage of the technology. A 2026 KPI of "deploy 10 production agents in IT operations" is more appropriate than "achieve autonomous workflow coverage across all departments." Celebrate real milestones on the right timeframe.

### For AI Leaders and Technology Practitioners

- 01 Establish an internal agentic AI taxonomy.**  
Define precisely what your organization means by "agent," "assistant," "workflow AI," and "multi-agent." Shared vocabulary is prerequisite to shared understanding of where you are and what you're building toward.
- 02 Prioritize trust infrastructure.**  
Before expanding agent scope or autonomy, invest in the evaluation frameworks, testing pipelines, and monitoring systems that make agent behavior observable and improvable. You cannot improve what you cannot measure.
- 03 Map your knowledge architecture gap.**  
Audit the delta between the knowledge your agents can currently access and the knowledge required for your target use cases. That gap is your roadmap.
- 04 Build human oversight into the design, not as a retrofit.**  
Design human checkpoints deliberately, knowing which decisions require human review and which can be automated as trust is established.
- 05 Plan for multi-agent coordination.**  
Even if you are not deploying multi-agent systems today, the architecture decisions you make now will either enable or constrain future multi-agent deployment. Build with composition in mind.

# Conclusion

## Hype or Not? The Honest Answer

So is all this talk about agentic AI in the enterprise hype? The honest answer is: yes and no.

The hype is real. Agent-washing is pervasive. Self-reported maturity levels significantly overstate genuine capability deployment. The transformative promises being made in vendor presentations, conference keynotes, and analyst forecasts are substantially ahead of operational reality in large enterprises today. Organizations claiming to have deployed AI agents at scale typically mean something considerably more modest than that implies.

But the underlying potential is also real, and the investment commitment is genuine. 71% of the enterprises in our survey are investing significantly in agentic AI. 24% have production deployments. The use cases that are working, like IT automation, customer service triage, supply chain optimization, are delivering measurable value. The capability roadmap, from current-generation assistants to tomorrow's multi-agent collaborative systems, is technically credible.

The gap between today's reality and tomorrow's promise is real, but it is not permanent. It is the gap that separates a technology at the beginning of its deployment decade from a technology at the height of its maturity. Bridging that gap requires three things: honest reckoning with current constraints, systematic investment in the foundational infrastructure (trust, knowledge, governance, organizational capability) that mature deployment requires, and the patience to build for a decade rather than a quarter.

The enterprises and vendors that approach agentic AI with that combination of honesty and ambition are the ones who will be standing at the far end of this decade having built something genuinely transformative. The ones who chase hype, cut governance corners, or mistake early assistants for mature agents risk building technical and organizational debt that makes the real journey harder, not easier.

The decade of agents has begun. The question is not whether it will transform enterprise operations. It will. The question is which organizations will be ready when it does.

## Research Methodology

This research was conducted in March 2026. The survey targeted senior business leaders with direct involvement in AI strategy, implementation, oversight, or operational management within their organizations.

<b>Sample Size:</b>	740 qualified respondents
<b>Company Size:</b>	5,000+ employees (all respondents)
<b>Revenue:</b>	\$1B-\$20B+ annual revenue
<b>Geography:</b>	North America (64%), Europe (36%)
<b>Respondent Roles:</b>	C-suite/senior leadership, AI strategy, operations, digital transformation heads
<b>Qualification:</b>	All respondents have direct involvement with AI implementation, oversight, or strategic planning

Respondents were screened to ensure direct involvement with AI implementation, oversight, or strategic planning. All responses were collected anonymously. Statistical significance was not calculated for subgroup comparisons, which should be interpreted as directional rather than definitive given subgroup sizes.



## About ChapsVision

ChapsVision is a global leader in agentic AI and data solutions. Trusted by over 2,000 customers across 40 countries, ChapsVision delivers industry-specific software solutions powered by agentic AI to meet the complex needs of data-intensive sectors. Its innovative and scalable Data & AI platform is powered by proven technology enablers that accelerate data acquisition, preparation, and analysis.

Driven by significant R&D investments in large-scale data and AI and supported by a targeted strategy of acquisitions and international expansion, ChapsVision has rapidly established itself as the trusted partner in agentic AI for business and government organizations around the world.

**More information: [www.chapsvision.com](http://www.chapsvision.com)**