



State of the Agentic AI Market Report

June 2025



Table of Contents

3 Summary

7 Methodology

9 Introduction: What is Agentic AI?

13 Study Findings

15 Horizontal Offerings Targeting Function-specific Processes

18 Industry-specific Solutions: Agentic Innovation by Industry

22 Spotlight on Agentic and Generative AI Software

24 Enterprise Needs and Concerns for Agentic AI

25 How can enterprises meet the need to accelerate agentic AI adoption?

26 How should agentic AI be priced and purchased?

27 How should enterprises manage multi-agent coordination?

27 What is the appropriate level of autonomy for agentic AI?

28 How can enterprises deal with data woes when adopting agentic AI?

29 Are there concerns around model or cloud lock-in for agentic AI?

30 How to Get Started with Agentic AI

31 The Data Requirements: Strong Data Tower & New Data Frameworks

33 Getting Your Organization Ready to Achieve ROI

34 From Promise to Performance: Operationalizing the Agentic AI Enterprise

37 Authors & Contributors



[Summary](#)

[Methodology](#)

[What is Agentic AI?](#)

[Study Findings](#)

[Spotlight on Software](#)

[Enterprise Needs & Concerns](#)

[How to Get Started](#)

Summary

Summary

Generative AI has held the spotlight in the last few years, but the emergence of autonomous AI agents changes the game all over again. If this were the industrial revolution, GenAI would be like the original steam engines, and agentic AI is akin to more efficient and pervasive power sources like internal combustion engines and electricity.

Agentic AI is gaining momentum as a new way for enterprises to get work done – enabling systems to act with more independence, context and coordination. But while the potential is real, so are the challenges. This ISG study shows that, while early use cases are delivering value, most organizations are still navigating fragmented data, unclear governance and operating models not yet built for this kind of capability. The goal of this report is to help leaders focus: to cut through the noise, understand where agentic AI is already working, and identify what’s holding it back. For those that get it right, the opportunity isn’t just about efficiency – it’s about building a more competitive and resilient enterprise.

1

Most agentic AI use cases pick up where generative AI left off

At its core, generative AI (GenAI) is most often applied to solve problems related to the extraction and summarization of information. In doing so, GenAI creates a faster and more effective means of accessing collective organizational knowledge. However, those systems are limited in their ability to use that information for any logical subsequent task. For example, we could use GenAI to collect and summarize macro-economic statistics for an effective perspective of the market and suggest appropriate investment actions to financial advisors. But such a model could not take action against investment assets.

Today, we are seeing agentic AI systems build on that extracting capability by supplying the additional steps that can take this information to an action with limited or no human intervention or interaction. Having said that, knowledge management use cases for AI remain very relevant – it’s just that they are now being repositioned as data engineering use cases that support the development of the context-aware data backbone needed for scaling agentic AI.

2

A few horizontal offerings and industry-specific value chains are leading the agentic AI market

Among more general, non-industry-specific applications of agentic AI, over half are found in IT, while marketing, sales and finance each account for only about 10%. For industries, we’re seeing specialization as well: 70% of agentic AI use cases are concentrated in just three industries: banking, financial services, and insurance (BFSI), retail, and manufacturing.

3

More autonomy is not always better, and sometimes RPA works just as well

Simple and model-based agents, which are mostly task-focused, are great at delivering value and represent an average of 43% of use cases in the market. Based on the market investments we're seeing from our study, we expect that, as enterprises mature and start to address more complex problems, this will evolve. As it evolves, goal-based agents (for example in transportation: complex optimization exercises like route scheduling to balances efficiency, safety and predictive and preventive maintenance needs against organizational revenue targets) will become a larger focus. However, leading providers need to continue to engage with their enterprise clients to ensure that the right agent type is getting matched to the right use case, including the possibility that an expensive AI agent is not needed to do what regular RPA can easily accomplish.

4

Right now, providers are focused on delivering rapid results to enterprise clients... while also investing in a more complex multi-agent future

To help accelerate adoption of agentic AI, providers are creating pre-built solutions for quick wins within a specific process. Many are specializing in certain functions or industries by investing in really mature solutions before going to market. We see many enterprises that are still struggling to get GenAI pilots into production with only a few, mature enterprises working on piloting agents, so this resonates well as the most practical approach.

However, we also found that, in parallel, most providers are already investing in more complex multi-agent solutions, envisioning a future where orchestration becomes the most challenging aspect of agentic AI workflows. In other words, the provider ecosystem is readying itself for organizations to grow in adoption of agentic AI and for demand that will be very different from solutions on the market today.

5

Agentic AI has not eliminated humans in the loop (HITL), and it likely will not for a while

Only 25% of current solutions allow agents to operate independently, while 45% position AI as advisors supporting human decision-makers. The role of human oversight remains loosely defined, creating uncertainty around how much control should be retained. In response, many providers are investing in orchestration and governance with the aim of helping enterprises strike the right balance between autonomy and oversight as well as enabling scaled up adoption when the time comes.

6

Top adoption challenges include data and organizational readiness

Thirty eight percent of organizations engage a service provider to improve data management, but despite this effort, more than half of organizations still struggle with legacy data (*source: ISG Market Lens, IT Budget and Spending Study, n=300*). For organizations with poor AI outcomes to date, the biggest adoption challenges we're seeing are a lack of AI skills and the viability of business cases. For enterprises with more success in their AI programs to date, they still lack strong data towers or frameworks and organizational readiness to reimagine processes. The challenge of organizational readiness mimics the challenges seen during the start of offshoring, when companies had to do things differently because someone else (then: an offshore worker, now: AI) was doing the job.

7 Data is key to a return on investment (ROI) from agentic AI

To accelerate adoption, leading providers are leveraging agentic AI to address complex data challenges within enterprise organizations. These solutions act as a bridge across disparate data sources, using intelligent agents for tasks like monitoring, profiling, enrichment and governance. While this helps achieve results from agentic AI faster, it also highlights that agentic AI does not consume data in the same way as humans or other data-driven technologies. Therefore, the basic data architecture widely accepted as the standard today needs to be revisited.

8 Enterprise goals for agentic AI aren't just in efficiency, but also revenue growth

The data clearly showcases that agentic AI is pushing organizations to adapt, and this goes beyond jumping onto the hype wagon. Market leaders hope that agentic AI is a path to reinvention, both for their processes and also for their products and business models. In context of how ISG has chosen to define agentic AI (i.e., **executes business processes through autonomous actions**), most organizations start implementing agents further down the organizational layers (e.g., infrastructure, back office) by simply increasing the levels of complexity of tasks that can be automated.

This initial effort progresses upward towards goal-oriented workflows in the higher layers of front office and service delivery. This represents a significant shift up the value chain for automation – from broad, functional use cases to more vertical, industry-specific applications. But more importantly, it highlights that, while organizations are starting their agentic AI initiatives by looking for cost efficiency opportunities, down the road (in two to three years' time) the focus will shift to growth opportunities (*source: ISG Market Lens, Generative AI Adoption Study, n=201*).





[Summary](#)

[Methodology](#)

[What is Agentic AI?](#)

[Study Findings](#)

[Spotlight on Software](#)

[Enterprise Needs & Concerns](#)

[How to Get Started](#)

Methodology

Methodology

This ISG State of Agentic AI report outlines and analyzes the most recent developments in the agentic AI solutions and services market. Our research approach uses a combination of primary and secondary research, including surveys with end users, enterprise interviews and discussions with providers offering agentic AI solutions and services. The report offers insights intended to help enterprises make informed decisions and improve the impact of their AI projects.

ISG (Information Services Group, Inc.) is solely responsible for the content of this report. Unless otherwise cited, all content, including illustrations, research, conclusions, assertions and recommendations in this report were developed by and are the sole property of ISG.

The research and analysis presented in this report include insights from [ISG Provider Lens™](#) service provider evaluation research, [ISG Buyers Guide™](#) software provider evaluation research, [ISG Market Lens™](#) buyer behavior research, ongoing ISG Research programs, interviews with ISG advisors and analysts, briefings with service providers, and analysis of publicly available market information from multiple sources. The data collected for this report represents information that ISG believes to be current as of May 2025 for providers that actively participated as well as for providers that did not. ISG recognizes that several active use cases and projects that were in the pipeline have taken place since that time, but those changes are not reflected in this report. For further information on our process, see the appendix.

ISG interviewed the following service and technology providers as part of our research for this study: Accenture, Cognizant, Infosys, Genpact, HCLTech, HTC, LTIMindtree, NTT DATA, Tech Mahindra, Wipro, Apexon, AccionLabs, Brillio, EXL, Happiest Minds, Harman, Hexaware, Movate, Persistent Systems, Trigent, Unisys, Virtusa, Xoriant, MathCo, Sigmoid, Tiger Analytics, Eviden, Accion Labs, Hitachi Digital Services and WNS Analytics.

The study was divided into the following steps:

1. Definition of the agentic AI solutions and services market.
2. Use of questionnaire-based surveys of service providers on the topic of agentic AI in an RFI model.
3. Interactive discussions with service providers/ vendors on capabilities and use cases.
4. Use of ISG's internal databases and advisor knowledge and experience (where applicable).
5. Detailed analysis and evaluation of services and service documentation based on the information received from providers and other sources.
6. Peer review by ISG subject matter experts, research analysts and advisors.

What is Agentic AI?

Introduction: What Is Agentic AI?


At the highest level, **agentic AI represents an evolutionary leap forward in how technology gets work done.** Agentic AI shifts enterprises from rule-bound, task-based automation to systems capable of dynamically pursuing goals and making decisions on their own. It enables software to act – not just respond – by sensing context, analyzing data, deciding among alternatives and taking action to achieve outcomes.

Yet, despite the growing excitement, **there is no universally accepted definition of what agentic AI actually is.** Across vendors, analysts and adopters, the term is being used inconsistently.

This has led to confusion in the market and has made it difficult for enterprises to discern between genuine innovation and marketing buzz – a phenomenon ISG refers to as “agent-washing.”

To cut through this ambiguity, **ISG Research defines agentic AI as “designed to execute business processes through autonomous actions,** potentially controlling multiple processes and systems through the orchestration of one or more AI or algorithmically determined rules-based models.


Phase 1 | BEFORE 2022 | Conceptual



AI Can: Learn & Repeat

- AI learns and understands
- Less about rigid programming, more about evolving knowledge
- Breaks the boundaries of specialized departments

Phase 2 | 2022 - 2025 | Generative



AI Can: Learn, Repeat, Create & Recreate

- AI creates, designs and imagines
- AI expands its horizons, synthesizing knowledge from various domains to birth innovative output

Phase 3 | 2025 - Present | Agentic



AI Can: Learn, Repeat, Create, Recreate, Act & Collaborate

- AI can make decisions and take actions with different levels of autonomy, with varied levels of human interventions or oversight
- AI can manage and calibrate itself as well as its interact with other AI, data and human entities
- Breaks the boundaries of specialized departments

These models act based on an understanding of the environment and the goals that should be achieved” (source: *Agentic AI: A Strategic Guide for Enterprise Leaders*). Critically, agentic systems differ from bots and traditional automation in that they **adapt, learn and take autonomous action** within defined guardrails. As we learn more about agentic AI through the adoption lifecycle, this definition will evolve and mature.

Agentic AI is new, and it’s arriving **unevenly across service and software providers and industries**. Many vendors are branding conventional automation features as agentic AI without meeting the essential criteria. Enterprises must navigate this landscape with clear-eyed discernment, understanding not just what’s being delivered now and what’s possible in the near-term, but also how they can achieve value for their business.

To that end, it’s helpful to look at the **essential traits that define true agentic AI**:

- **Autonomy** – Takes initiative without requiring continuous direction or step-by-step prompts.
- **Goal Orientation** – Acts to achieve defined business outcomes, not just task completion.
- **Context Awareness** – Monitors and responds in real time to changing environments and data inputs.
- **Bounded Decisioning** – Operates within constraints such as risk, policy or ethical frameworks.
- **Escalation Logic** – Knows when to pause and seek human input or oversight.
- **Evolving Behavior** – Learns from experience and refines its actions to improve over time.

Importantly, **not all agentic systems are created equal**. ISG classifies agentic AI maturity along four levels: simple, model-based, goal-based and utility-based agents. See the chart on the following page for definitions and examples of each.

These capabilities can be orchestrated as multi-agent systems, where agents of different types collaborate to fulfill broader goals, each playing a specialized role within a coordinated framework.

It is also essential to **differentiate automatic action** from autonomous action. Automatic systems operate on predefined schedules or event triggers. For example, a nightly fraud detection scan is automatic. In contrast, an autonomous system runs continuously, adapts its strategy based on emerging data and re-optimizes its fraud detection models in real time – without waiting for a prompt or script.



Why This Report Matters

This report provides a clear, research-based foundation for understanding agentic AI, separating substance from hype. It highlights where agentic AI capabilities are truly being delivered, where gaps remain and what enterprises must do to become agent-ready. We explore key traits, use cases, deployment patterns and maturity levels to help organizations make informed, future-ready sourcing and investment decisions. With the promise of significant financial savings demonstrated in early use cases, agentic AI is poised to reshape digital business and workplace environments. Getting smart about what it is – and isn’t – is the first step toward unlocking its full value.



[Summary](#)

[Methodology](#)


[What is Agentic AI?](#)

[Study Findings](#)

[Spotlight on Software](#)

[Enterprise Needs & Concerns](#)

[How to Get Started](#)



Study Findings

Study Findings

Broadly, agentic AI is still early in its adoption cycle, but we are seeing pockets of innovation. Today, providers are delivering agentic AI solutions primarily in these two ways:

1 Horizontal offerings targeting function-specific processes within an enterprise that have been historically hard to automate and are beyond the capabilities of RPA. Most providers working in agentic AI have an horizontal solution for the major functional areas, like DevOps in IT or customer experience in sales and marketing. They do this by creating a core solution and then customizing it to each organization's unique needs and workflows. Perhaps most interestingly, traditionally industry-specific offerings, such as predictive maintenance and optimization, are now being offered in a horizontal model for any industry. Unlike with industry-specific solutions where providers are specializing, horizontal offerings are likely available with your provider of choice.

2 Industry-specific solutions focused on key value chains, such as insurance underwriting or claims management, where agentic AI capabilities can reinvent the whole approach. This requires service providers to not only possess a unique understanding of the business and domain-specific engineering skills but also to have access to relevant industry data. It is therefore not surprising that, right now, for most providers, there is only room to focus on a key shortlist of industries. For providers in our study, 70% of use cases/proofs of concept (POCs) in agentic AI come from one of three industries. This demonstrates how providers are focusing narrowly, looking for unique intersections of internal capabilities and high ROI potential for their agentic AI investments.

Horizontal Offerings: Targeting Function-specific Processes

Procurement

5% Supply Chain

Operations

4% Customer Support

HR

3% Human Resource Management

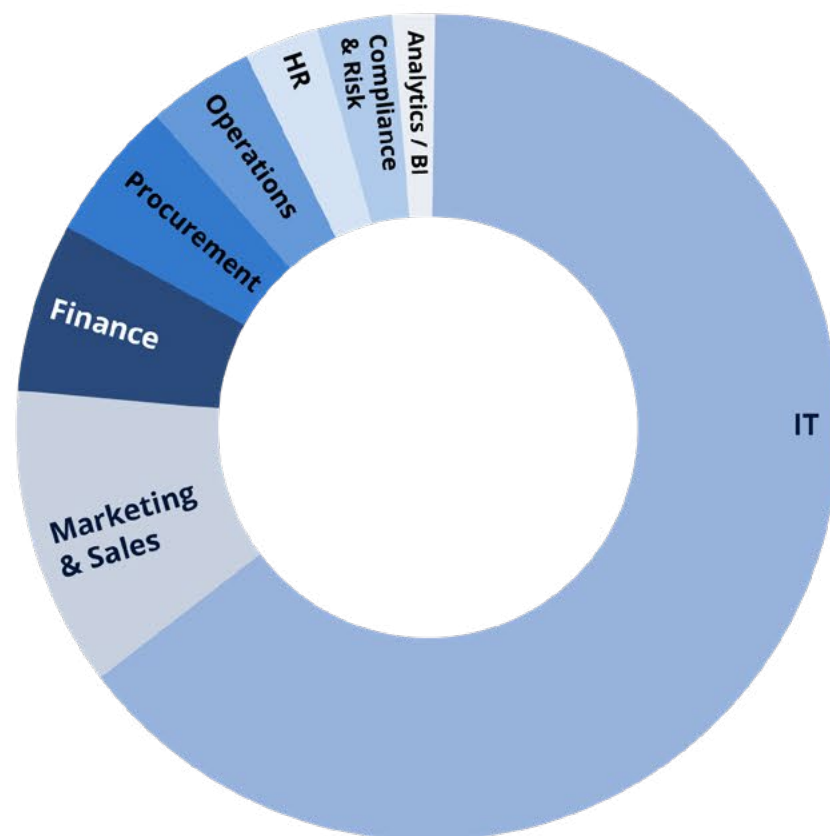
Compliance / Risk

3% Fraud Prevention

Analytics / BI

1% Predictive

1% Data Engineering



IT

48% DevOps

9% Cybersecurity

8% Infrastructure Management

Marketing & Sales

10% Customer Engagement

2% Sales

Finance

5% Accounting

2% Insights

1% Other

At a functional level, a clearer picture emerges for the types of activities that require simple agents versus goal- and utility-based agents (more autonomous agents). Goal- and utility-based agents are most common in functions such as IT, procurement, and sales and marketing. Thirty five percent of IT use cases are goal- or utility-based agents, such as highly versatile cybersecurity and infrastructure agents. Fifty six percent of procurement use cases are goal- or utility-based, such as agents used for supply chain optimization. In sales and marketing, 25% use cases are goal- or utility-based, needed because of the high degree of adaptability these agents must have to execute customer engagement interventions. On the other hand, logic-based activities found in industries like finance don't require agents to be as adaptive – 44% of finance use cases are simple agents.

IT: DevOps, Quality Assurance, Infrastructure Management and Cybersecurity make up 52% of non-industry-specific use cases

With the introduction of generative AI, a lot has changed in how software is developed, managed and tested, with providers investing heavily in code generation capabilities. Agentic AI is taking these already fairly mature capabilities and stretching them further by creating suites of agent-based solutions for DevOps (83% of IT use cases). Leading providers are enabling agents to autonomously parse complex requirement documents to extract precise deliverables, generate architecture blueprints and data models, and orchestrate end-to-end workflows with built-in traceability. Multi-agent testing frameworks deploy specialized bots to generate, execute, validate and self-heal. Integration and end-to-end test suites automate environment provisioning, script adaptation and risk-based strategy formulation. Legacy modernization can become frictionless as agents analyze existing codebases, translate them into new languages or frameworks, and continuously refactor and validate functionality against the original system.

Cybersecurity agents and infrastructure maintenance agents are starting to sound very similar to preventative maintenance in manufacturing (both at 15% of IT use cases), both in terms of autonomy and scope. Agentic AI platforms can autonomously maintain a real-time inventory of informational and operational technology (IT/OT) assets, ingesting threat intelligence and applying graph-based risk scoring to detect anomalous behavior across complex environments without human tuning. Multi-agent workflows then self-orchestrate patch discovery and deployment, incident containment, and automated remediation playbooks, continuously learning from each event to adapt detection models and response strategies. By unifying monitoring, threat hunting, compliance checks and governance into a single, self-improving system, these solutions can significantly cut remediation time and ensure consistent security.

Similarly, as cloud environments get increasingly more complex, agentic AI platforms continuously ingest telemetry from hybrid and multi-cloud environments and automatically detect and remediate anomalous behaviors without human intervention. Built-in predictive maintenance agents can analyze historical patterns to forecast component failures and schedule upkeep before outages occur, while resource-optimization agents can right-size instances, recommend reserved purchases and retire unused assets to manage costs. Together, these autonomous workflows slash unplanned downtime costs and free IT teams to focus on strategic initiatives.

Based on the breadth of these services we see in the market, we expect that organizations are going to start seeing these capabilities as foundational automation needs in their IT services solutions. Because of this, the unique value proposition from agentic AI is likely to get diluted and commoditized as customers come to expect the gained efficiencies and the associated savings in traditional IT services.

Marketing and Sales: Customer Engagement and Sales – 10% of non-industry-specific use cases

We have seen in retail, travel and hospitality industries that providers are making deeper investments in customer engagement, both in terms of customer experience and loyalty program management. However, as the benefits of agentic customer experiences become clear, leading providers are proactively building non-industry-specific solutions that allow flexibility at deployment. The same agentic AI approach to personalization can significantly benefit sales teams, both for business-to-business (B2B) and business-to-consumer (B2C) organizations, where focus can be shifted from repetitive customer inquiries to true value-adding sales activities.

Providers are seeing the possibilities of crossover between their enterprise clients in customer engagement. More mature providers are leveraging this as an opportunity for market expansion. This means that any organization that has challenges with customer engagement – regardless of their footprint or the maturity of their loyalty offering – should explore the possibilities of agent-driven solutions and the potential efficiencies and gains it may offer.

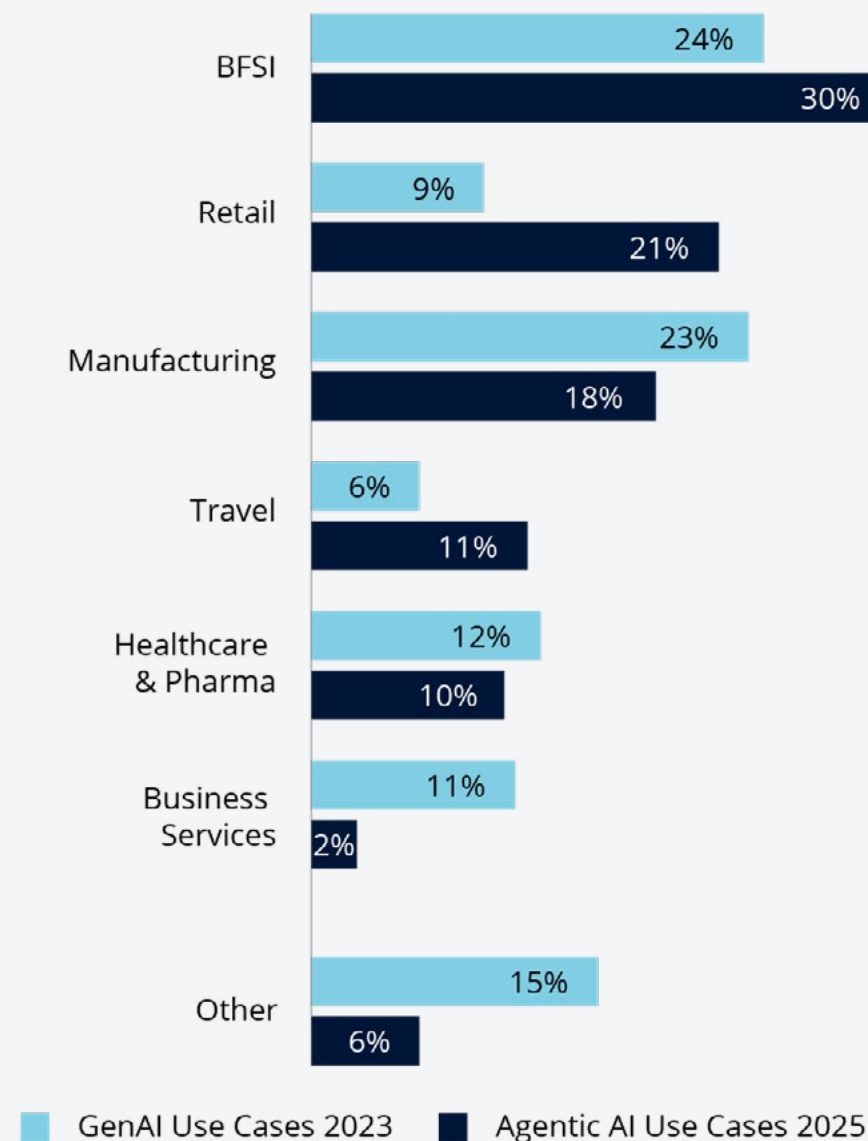
Finance: Accounting and Insights – 10% of non-industry-specific use cases

Accounting data, both structured and unstructured, contains a wealth of information and can be fairly uniform across organizations, regardless of industry. At the same time, it is an area of minimal investment for many organizations, yet it is rife with inefficiencies and manual processes that are hard to break due to domain-specific knowledge. Similarly, software providers have had almost no investment in general and administrative agents to date, but we expect to see this change rapidly in the next year.

Effective agentic AI solutions could significantly impact daily operations by autonomously ingesting and validating invoices. The AI agent could extract, enrich and tag data to purchase orders (POs), prioritizing exceptions, detecting duplicates or anomalies, and driving touchless accounts payable (AP) processing. Initial findings suggest cycle times can be reduced to one-fifth of their original length. Beyond payables, end-to-end invoice-to-cash agents could orchestrate billing validation, dispute resolution, tailored collection outreach and accounts receivable (AR) reconciliation without manual handoffs, accelerating cash conversion. In record-to-report, agents could reconcile balance sheets, auto-validate journal entries, generate narrative disclosures, flag anomalies with risk-scored recommendations and surface real-time chief financial officer (CFO) insights, cutting close cycles and enhancing reporting accuracy.

Several service providers that have a strong data foundation in finance see this as a perfect opportunity to develop a suite of agentic AI solutions for accounting across industries, effectively reimagining how the function is run. Meanwhile, software providers are embedding agentic-like AI in their existing suites of horizontal software for finance and accounting.

Industry-specific Solutions: Agentic Innovation by Industry



Adoption of agentic AI has so far been strongly focused on specific industries. We see BFSI and retail are leading the pack with 30% and 21% of use cases, followed by manufacturing (18% of use cases), travel, transportation and leisure (TTL) (11% of use cases), and healthcare and pharmaceuticals (10% of use cases). Notably, in terms of autonomy of agents, only 9% of BFSI use cases leverage goal- and utility-based agents, suggesting that the processes are more focused on efficiency gains than autonomous actions. On the other hand, in retail, manufacturing and TTL, agents need to optimize, adapt and evolve in order to deliver value, so we see more goal- and utility-based agents.

BFSI - 30% of use cases / POCs

BFSI continues to be the beacon for AI innovation. This is driven by the industry's rich structured and unstructured data, which is the first step to successful adoption of agentic AI in critical workflows. When we looked at this sector in 2023 just as generative AI had emerged, the most mature AI use cases in financial services and insurance were in customer support and personal banking. However, there was a lot of experimentation with fraud analytics and other capabilities around data extraction and summarization. The shift we are seeing now is the perfection of those original use cases with the addition of using agentic AI capabilities to perform proactive actions.

Agentic AI use cases in BFSI can be grouped under financial services and insurance. In financial services, top use cases include fraud detection (19%), customer engagement (16%) and autonomous financial services (14%). Looking specifically at fraud detection, some of the more advanced providers are not just developing solutions to detect subtle fraud patterns. The leaders here are working to help prevent fraud actors from evolving financial crime tactics, e.g., sophisticated layering and trade-based laundering, which requires advanced detection capabilities in real time beyond traditional rule-based systems.

Early results show that, with effective proactive actions, organizations can effectively drive down costs – for example, by significantly reducing the rate of false positive alerts. Autonomous financial services – such as credit scoring and decisioning, autonomous portfolio management, financial planning, proactive loan management, and intelligent tax optimization – drive higher value from advisory services across the financial services value chain. In customer engagement, we see providers develop solutions that make it easier and faster to onboard new enterprise clients, including document processing and verification and client profiling.

Data engineering use cases in financial services were some of the early use cases for generative AI, where new technology offered more effective ways for summarizing and extracting multi-modal information. Agentic AI is able to build on these initial wins by combining insights from multiple generative and traditional AI models. Doing so, it is possible to optimize not only general insights but also to create advisor copilots for employees and customers. In this way, BFSI is where we see the first practical applications of multi-agent collaboration. Numerous agents are trained to analyze different types of indicators and can then work together to achieve deeper understanding of financial and economic signals.

In insurance, use cases in claims and underwriting far outnumber other agentic AI applications. Claims processing faces significant operational challenges due to heavy reliance on manual processing, resulting in slow turnaround times, high administrative costs and increased risk of fraud. Claims involving complex and unstructured medical documents often require heavy manual review, delaying decisions and increasing the likelihood of missed critical information.

These inefficiencies reduce accuracy, delay service, and create compliance and customer satisfaction risks in an evolving regulatory landscape. Agentic AI solutions have the potential to streamline insurance processes by automating medical record summarization, claims validation and document handling. For underwriting and quote processing, agentic AI provides contextual data extraction, summarization and chat-based insights, significantly reducing manual effort and turnaround time across the insurance value chain.

Retail - 21% of use cases / POCs

Retail companies have become one of the primary consumers of agentic AI. Likely, this is because the sector has historically struggled to deliver strong customer engagement across its digital channels, especially in areas of personalized shopping, upsells and customer support. Generative AI offered some early wins, but soon inconsistent service quality from chatbots limited its capability to act autonomously. This became a challenge in terms of driving further value from AI.

With 26% of retail use cases in customer engagement, we are seeing agents use natural language processing to convert customer requests into structured query language (SQL) commands, effectively working with structured data via unstructured multi-modal query framework. AI agents can then access data quickly to answer user questions with minimal human intervention. This application of agentic AI also helps improve online product searches by optimizing relevant product suggestions. In addition, agentic AI has provided new support to loyalty programs with specialized providers delivering personalized customer loyalty solutions that autonomously create, execute and optimize marketing communications.

Data engineering represents 16% of retail use cases where agentic AI can help tackle problems of image and metadata management in online shopping. Currently, the industry spends a great deal of human effort and investment in correctly identifying product features and managing related metadata. Agentic AI can be used to extract multi-modal product features and optimize the process of capturing and updating this data in enterprise systems of record. There is an opportunity to further improve the autonomy of these agents by allowing them to identify new product trends and create new tags that have not yet been directly built into product management systems. Early results indicate that agentic AI will be able to provide accurate and consistent information, leading to improved customer experience, lower bounce rates and a decrease in lost sales.

Manufacturing - 18% of use cases / POCs

Within manufacturing, generative AI opened the door to a lot of new possibilities and predictive maintenance continues to offer room for innovation. Today, we see predictive maintenance as the top use case in this industry (42% of manufacturing use cases) with a varying degree of autonomy. More advanced providers are building a supervised set of orchestration agents that autonomously manage inventory and assist their human users. For example, agents can optimize production processes, warning of faults and scheduling maintenance by monitoring equipment health. Agents can also monitor human behavior and safety by detecting potential hazards and preventing issues in real time. These unique suites of tools often utilize digital-twin technology and really push the envelope for agentic AI. They show that agentic AI can be more cost effective than a human orchestrator and can offer built-in analytics and just-in-time decisions, which are becoming standard in the future of manufacturing.

Very much adjacent to predictive maintenance are use cases in optimization (34%) where agentic AI is helping improve operations as well as aspects of research and development. Specialized providers are driving innovation across every aspect of manufacturing with AI-driven simulation and optimization of energy usage, logistics planning and coordination. The key to these use cases is the autonomy of agents to make adjustments with minimal human intervention. In research and development (R&D), the focus is on autonomous identification of market trends for product development as well as material optimization.

Transportation, Travel & Leisure (TTL) - 11% of use cases / POCs

Similarly to retail, TTL saw a significant jump in adoption of agentic AI when compared to the generative AI study in 2023. There are further similarities: both industries originally went after customer engagement and customer support as major use cases, leaning in on virtual assistants for booking travel or having chatbots take over concierge services. Today, agentic AI is most heavily applied in customer engagement (25% of TTL use cases) as well as customer support (13%), where organizations are looking to agentic AI to solve similar customer experience challenges as with retail.

The more advanced providers are focusing on a complete solution that adjusts customer experience to suit their preferences in real time, both in their physical locations and online. This not only has the potential to improve customer service scores, but it also impacts margins by driving more customers to book directly as opposed to through a third party. Customer experiences personalized at scale have been a long-time goal for the industry, and agentic AI is bringing it closer to reality. Specialized providers are also helping the industry in its ongoing struggle to balance awarding points with associated costs. An AI agent would be able to analyze loyalty member behavior, determining when it would be financially beneficial to apply a promotional offer or reward to effectively impact customer behavior while maximizing profitability.

Two other big areas of opportunity are specifically in transportation: optimization (17% of TTL use cases) and preventive maintenance (13%). Similarly to the agentic AI approach to preventive maintenance in manufacturing, advanced providers are developing solutions for transportation that orchestrate numerous agents to work together in the management of a fleet of vehicles and drivers to promote safety and to optimize schedules and routes. These agents can autonomously determine when and how to schedule routine maintenance to optimize downtime and can monitor driver behavior or even autonomously call emergency services.

Healthcare & Pharmaceuticals - 10% of use cases / POCs

In pharmaceuticals (pharma), providers are most notably focused on addressing inefficiencies in clinical trials (32% of healthcare and pharma use cases). The market is investing in transforming labor-intensive clinical study templates and regulatory reports –traditionally requiring lengthy manual synthesis across structured and unstructured data –into agentic AI solutions. By doing so, it could cut the process down to hours, dramatically cutting lead times and content-development costs. By automatically generating and customizing clinical study report (CSR) templates that adhere to regulatory standards, these systems ensure high template quality, compliance and consistency with minimal human intervention.

Beyond report drafting, AI-driven content-authoring agents can read and interpret diverse source documents to populate study details – objectives, design and evaluations – boosting both the accuracy and relevance of output while slashing publication timelines. More specialized providers are working on a centralized “command-center” agent framework, which unifies communication among research teams, healthcare providers and commercial partners. The goal of this framework would be to resolve fragmented workflows and speed up decision making to improve patient outcomes. Providers are also investing in several minimum viable product (MVP) development agents to drive new healthcare solutions – automating prototype iterations, quality assurance checks and compliance reviews.

The healthcare industry is under significant cost pressures, and therefore the AI focus here is analysis and insights that can help drive operational efficiency quickly. Traditional analytics methods of dashboards and discovery are proving too slow, and so organizations are looking for ways to address these challenges. Agentic AI “insight-generation” agents autonomously comb through vast volumes of unstructured program and patient data

using natural language processing (NLP), sentiment analysis and deep-learning. These agents then can surface emerging key performance indicators (KPIs), at-risk patient cohorts and market trends without human intervention. The agents can ultimately synthesize and prioritize actionable findings, e.g., program strengths and weaknesses or personalized risk predictions, into easy-to-digest reports or real-time alerts, continuously learning from new data to improve accuracy. By automating complex data ingestion, analysis and decision-triggering, they free up leadership and clinicians from manual effort – dramatically cutting report-creation times, reducing errors and optimizing resource allocation. The result is faster, more data-driven program improvements, earlier preventive interventions, lower costs and more effective, personalized care.



[Summary](#)

[Methodology](#)

[What is Agentic AI?](#)

[Study Findings](#)

[Spotlight on Software](#)

[Enterprise Needs & Concerns](#)

[How to Get Started](#)

Spotlight on Agentic and Generative AI Software

Spotlight on Agentic and Generative AI Software

For software providers, the market for agentic and GenAI software is still nascent, though evolving rapidly. Last year, nearly every software provider was adding GenAI capabilities to their products in the form of assistants and copilots. This year, the big push is around adding agentic AI capabilities, although exactly what that means varies from one software provider to another. In our ISG Buyers Guide for agentic and GenAI, we evaluated more than 75 different criteria, including generative and agentic capabilities in addition to data preparation, scalable processing, tooling, large language model (LLM) operations and which investments software providers are making.

We found that, like most advances in business technology, the age of agents is being proclaimed ahead of its actual arrival and practical accessibility. There is a fair amount of “agent washing” taking place as software providers rebadge their existing automation and programmatic software elements as new technology. Nonetheless, agents, along with traditional AI and GenAI, are set to have a profoundly positive impact on productivity and performance in business for enterprises.

Our software research finds that 85% of the 35 AI software platform providers we evaluated support the use of LLMs in their products. However, only one quarter offer agent design tools that go beyond providing APIs. Less than one-third offer bias detection mechanisms, and less than 30% have adequate toxicity testing capabilities. What these findings suggest is that agentic and GenAI efforts still require many manual processes to supplement the software tools that are available today. Software providers are aware of these shortcomings, and many have new capabilities that are in various stages of pre-release, so we expect the situation to improve, but it will take time.

Perhaps one of the biggest disappointments today is the lack of integration with common business applications. The objective of agentic AI systems is to accomplish tasks within the enterprise. Typically, those tasks involve a variety of software applications, but less than 10% of software providers offer pre-built integrations with these products. The recent emergence of the agent-to-agent (A2A) protocols should help accelerate interoperability between agents and applications, but, in the meantime, agent development still requires a lot of manual coding.

Agents or agent-like systems are beginning to appear in a wide range of business software and will continue to proliferate at an accelerating pace over the next three years. Agentic AI can help any enterprise, but the readiness of an organization or enterprise to adopt the existing technology from specific software providers is likely to vary significantly.

Enterprise Needs and Concerns for Agentic AI

Enterprise Needs and Concerns for Agentic AI

How can enterprises meet the need to accelerate agentic AI adoption?

Understanding where to start with agentic AI is a challenge for many enterprises, especially as the technology is evolving so quickly. Some organizations are mid-way through multi-year roadmaps for generative AI adoption and need to pivot to avoid missing out on agentic AI. Market leaders are looking for quick ways to adopt agentic AI without getting stuck in organizational alignment problems. quick ways to adopt agentic AI without getting stuck in organizational alignment problems.

How is the market supporting this need?

Leading providers are helping enterprises succeed by providing pre-developed agentic solutions that allow enterprises to target both horizontal tasks (most notably in IT operations, software development, incident management, etc.) and function-specific tasks (such as

customer support, sales operations, wealth management and supply chain). To accelerate adoption, specialized providers are bringing their industry experience to bear to ensure that the domain and business knowledge within these solutions maximizes time to value.

In addition, more comprehensive solutions on the market today are made up of modular components where an enterprise can engage any part of a suite of agents to address the needs of its business. As an example, consider a suite of agentic AI solutions for the accounting function where an enterprise can choose to engage agents that capture data from invoices, ingest these and then separately process for duplication and anomalies across the payment data. These kinds of off-the-shelf solutions make agentic AI much more approachable and require little experimentation.

Notable market gaps:

While solutions that provide quick wins within a specific function or organizational process will be reaching maturity soon, there is a notable lack of strategic cross-functional agents that collaborate across the organization. Many organizational processes are fragmented due to misalignment between stakeholders, and here agentic AI has the capacity to help orchestrate cross-functional human collaboration. For example, a great use case for many organizations is around mergers and acquisitions (M&A) and the associated integration activities.

How should agentic AI be priced and purchased?

There is currently no consensus on how agentic AI should be priced. Enterprises and providers alike are experimenting with a range of models, each reflecting a different interpretation of what agentic AI actually is – a tool, a service or a digital labor force. This lack of alignment is creating friction in procurement and slowing down enterprise adoption.

How is the market supporting this need?

Leading providers are exploring multiple pricing approaches, including:

- **Tool-based pricing:** Treating agents like software licenses, priced per seat or per function.
- **Consumption-based pricing:** Charging based on usage metrics such as API calls, compute time or number of tasks executed.
- **FTE-equivalent pricing:** Framing agents as digital workers and pricing them similarly to outsourced human labor.
- **Value-based pricing:** Charging based on business outcomes such as cost savings, revenue uplift, or cycle time reduction.

While value-based pricing is attractive to both providers and enterprises, it remains difficult to operationalize due to the lack of standardized metrics for measuring agentic AI's impact. Procurement, IT and business stakeholders often approach pricing from different perspectives, further complicating the conversation.

The simpler the agent, the less compute it typically consumes, making it more likely to be priced at a fixed rate, such as a subscription or license. As agents become more complex and autonomous, they require more resources to run, and pricing models are shifting accordingly. Expect to see fixed plus variable structures emerge, where the variable component reflects compute usage – either as pass-through costs or tiered consumption levels.

Some providers are beginning to differentiate pricing based on the maturity of the agentic solution. For example, early-stage agents that accelerate existing tasks may be priced like traditional automation tools, while more advanced agents that operate autonomously or coordinate across workflows may be priced based on outcomes or digital labor equivalence. This mirrors the evolution of pricing in BPO, where contracts shifted from effort-based to value-based pricing as the market matured.

Notable market gaps:

The absence of a shared definition of what constitutes agentic AI and how to measure its value is a major barrier to pricing standardization. Without clear maturity models or outcome frameworks, enterprises struggle to compare offerings, and providers struggle to justify premium pricing.

There is a clear need for the development of standardized pricing frameworks that:

- Reflect both the cost to build and maintain agentic solutions and the value they deliver
- Account for different levels of agent maturity and autonomy and
- Align procurement, IT, and business stakeholders around a common understanding of what is being purchased.

Until such frameworks are established, pricing will remain a source of confusion and misalignment. Enterprises and providers must work together to co-create pricing models that are transparent, flexible and tied to measurable outcomes.

How should enterprises manage multi-agent coordination?

Looking at agentic AI solutions offered on the market today, multi-agent solutions are an area of heavy investment in the provider landscape. Additionally, as agents become more integrated into enterprise workflows, the need for effective orchestration grows. It's true that not many organizations are yet mature enough to face this challenge – for an average provider, only 17% of deployed agentic solutions are multi-agent based. But as adoption grows, this is emerging as a critical concern for many early adopters.

How is the market supporting this need?

Leading providers are already investing heavily in standardization, orchestration and multi-agent frameworks with leading providers treating this as a key area of innovation for the next 12 to 24 months. The emergence of industry-standard protocols for agent collaboration in the recent past – most notably, MCP and A2A – have removed a lot of roadblocks to broader agent-to-agent integration. Already, we are seeing advanced providers deliver several agent networks for various enterprise functions, such as sales, finance, customer service and supply chain management. As organizations mature in their use of agentic AI, questions like, “how would agents work together to dynamically decompose a single task or synchronize multi-agent memory and do so at scale?” will become more common.

Notable market gaps:

Many organizations venture into agentic AI via a PoC as opposed to wholesale workflow adoption. There's a risk that, much like with early GenAI adopters' governance challenges, organizations will struggle to adopt proper orchestration practices when scaling. This is concerning given the isolated functional focus of agentic AI adoption, which doesn't lend itself to broader cross-organizational alignment on orchestration.

What is the appropriate level of autonomy for agentic AI?

One of the most critical – and least standardized – aspects of agentic AI adoption is determining how much autonomy to give agents. Our data shows that most enterprises are still cautious: only 25% of agentic AI solutions currently allow agents to operate independently, while 45% position agents in an advisory role, with humans making the final decisions. As use cases mature and agents take on more complex tasks, the question of how to structure human oversight will become increasingly urgent.

How is the market supporting this need?

Leading providers are embedding human oversight into orchestration frameworks, not by making the orchestrators themselves autonomous, but by using them as structured control layers that enforce governance. These frameworks manage agent behavior through goal decomposition, task sequencing, performance monitoring and escalation logic. In this way, the orchestrator acts as a proxy for human intent, translating high-level objectives into bounded agent actions.

To support this, providers are implementing role-based execution and permission systems that define what agents can do, when they must escalate and how they should respond to changing conditions. While some orchestration components may be automated, the governance logic remains human-defined. This ensures that autonomy is granted deliberately, not by default, and that agents operate within clearly-defined boundaries.

Notable market gaps:

The orchestration layer is only as effective as the governance logic it enforces – and that logic is still evolving. Governance for agentic AI must go beyond traditional concerns like model bias or drift. What’s fundamentally new is the need to govern action. Each agent’s ability to act autonomously introduces operational risk, and current frameworks are not yet equipped to manage it.

This becomes even more complex in multi-agent systems, where agents interact and influence one another. Without fine-grained controls, observability and clearly-defined human roles, enterprises risk deploying agents that act unpredictably or fail to escalate when needed. Until shared frameworks mature, enterprises will rely on trial-and-error learning to find the right balance between autonomy and oversight – a risky proposition for mission-critical workflows.

How can enterprises deal with data woes when adopting agentic AI?

A strong data foundation and robust data capabilities are essential for enterprises that want to deliver value from agentic solutions. There have been heavy investments by enterprises to bridge data siloes across the organization and establish foundational governance and data quality standards. Our research shows that 50% of organizations engage a service provider to deliver improved data management (source: *ISG Market Lens, IT Budget and Spending Study, n=300*). Despite this effort, more than half of organizations still struggle with legacy data residing in old applications, compounding the overall data silo issue and creating further roadblocks to value from AI (source: *ISG Market Lens ADM Study, 2024; Americas only; n=126*).

How is the market supporting this need?

To help accelerate adoption, many leading providers are turning to agentic AI to solve the problems inherited by their enterprise clients. Essentially, agentic AI is being used as the bridge that brings together disparate data sources and managed context across domains. Agentic AI solutions enhance data management by deploying intelligent agents for tasks like data monitoring, profiling and enrichment. These agents proactively identify data quality issues, provide contextual insights through cognitive search and enforce governance by monitoring data flows to ensure compliance.

We see providers engaging in one or a combination of the two solutions below:

- Leveraging a combination of LLMs and agentic AI to generate rich, context-aware, domain-specific data. These solutions tackle both structured and unstructured data and are often supported by a centralized data hub designed to facilitate data interactions for all models and agents across the enterprise. These central hubs can also act as a place for managing common data challenges – such as hallucinations, security and regulatory compliance – uniformly across agents. The main benefit is that agentic AI data workflows can be built directly into overall agentic AI workflows, enabling a direct lineage from data to outcome.
- Treating agent capabilities as a broad accelerator which is part of an agentic AI data management solution. This approach involves creating agents designated to tackle typical data functions like data quality and data governance not as part of a specific agentic workflow, but rather as a dedicated data management solution. The advantage of this approach is that it tackles the data challenges of a typical organization, in general allowing the benefits to extend beyond AI use cases

Notable market gaps:

Many solutions discussed in this study show how agentic AI has been able to serve itself with the right data and context as a part of a larger, automated workflow. We also see how these capabilities are evolving into broader data management solutions. However, for most organizations, the traditional data tower remains the same, built on the same elements of medallion architecture, extract, transform and load (ETLs), and pipelines. This leaves a lot of room for opportunity for foundational data architecture improvement.

Are there concerns around model or cloud lock-in for agentic AI?

Just like with cloud adoption, there are many aspects of agentic AI that can be considered interchangeable. As models and technology continue to improve, one of the major concerns that enterprises have is how to keep up. 18% of enterprise organizations responded that this was one of the top three challenges for agentic AI adoption (source: *ISG Market Lens, Generative AI Adoption Study, n=201*). The journey of AI adoption is one that requires enterprises to balance investment with the risk of missing out on something new due to high switching costs.

How is the market supporting this need?

Many hyperscalers are offering cost-effective measures to help organizations get PoCs off the ground. However, as these begin to scale, more long-term sustainable architectures will be needed, and lock-in is one important consideration to keep in mind. Leading providers who are recognizing this concern are ensuring that the solutions they build leave the door open to avoid both cloud and model lock-in. This is achieved by treating models and agents as modules that are put inside containers so they can be moved from cloud to cloud and combined easily with normal open-source, cloud-agnostic architecture. This work will help to ensure lock-in risks are kept in check.

Notable market gaps:

While many providers claim their agentic AI solutions are modular and portable – promising the ability to swap out models, agents or even entire orchestration layers – the reality is that the underlying protocols enabling this flexibility are still immature. The recently introduced agent-to-agent (A2A) protocol shows promise, but it is too early to determine whether it will become the industry standard. Until such protocols are widely adopted and matured, true interoperability will remain aspirational.

This lack of standardization creates a risk of de facto lock-in, even when architectures appear cloud-agnostic on paper. Without robust, widely accepted protocols for agent communication, memory sharing and orchestration, enterprises may find themselves tied to specific provider ecosystems or orchestration frameworks. A more strategic, industry-wide approach to protocol development and governance is needed to ensure long-term portability and reduce switching costs as the agentic AI landscape evolves.

How to Get Started with Agentic AI

How to Get Started with Agentic AI

The Data Requirements: Strong Data Tower & New Data Frameworks

For adoption to succeed, most providers are designing agentic AI solutions that effectively bypass a lot of legacy data challenges. Amidst this heavy focus on rapid adoption, data is a vital component, but it is not the goal driving most AI projects. Our research shows that, when it comes to AI adoption, a lack of skills and the viability of business cases are much bigger worries than data in the minds of enterprises (source: *ISG Market Lens, Generative AI Adoption Study, n=201*).

“My PoC works great! Why is data still a roadblock for scaling agentic AI?”

Fast-moving PoCs are able to procure the necessary data via support of AI-enabled accelerators or other agentic AI tools, but the core data challenges of siloed, low-quality data remain. As organizations want to deploy more agents and scale, this data problem will come back to haunt them. We can learn from history by looking at the emergence of digital. By 2016, three to four years into the digital revolution, data quality became a major challenge for nearly a quarter of organizations as they progressed down their journey to full digital adoption. This bottleneck is likely to cause problems in AI adoption in the next one to two years, just like it created roadblocks for digital adoption in 2016.

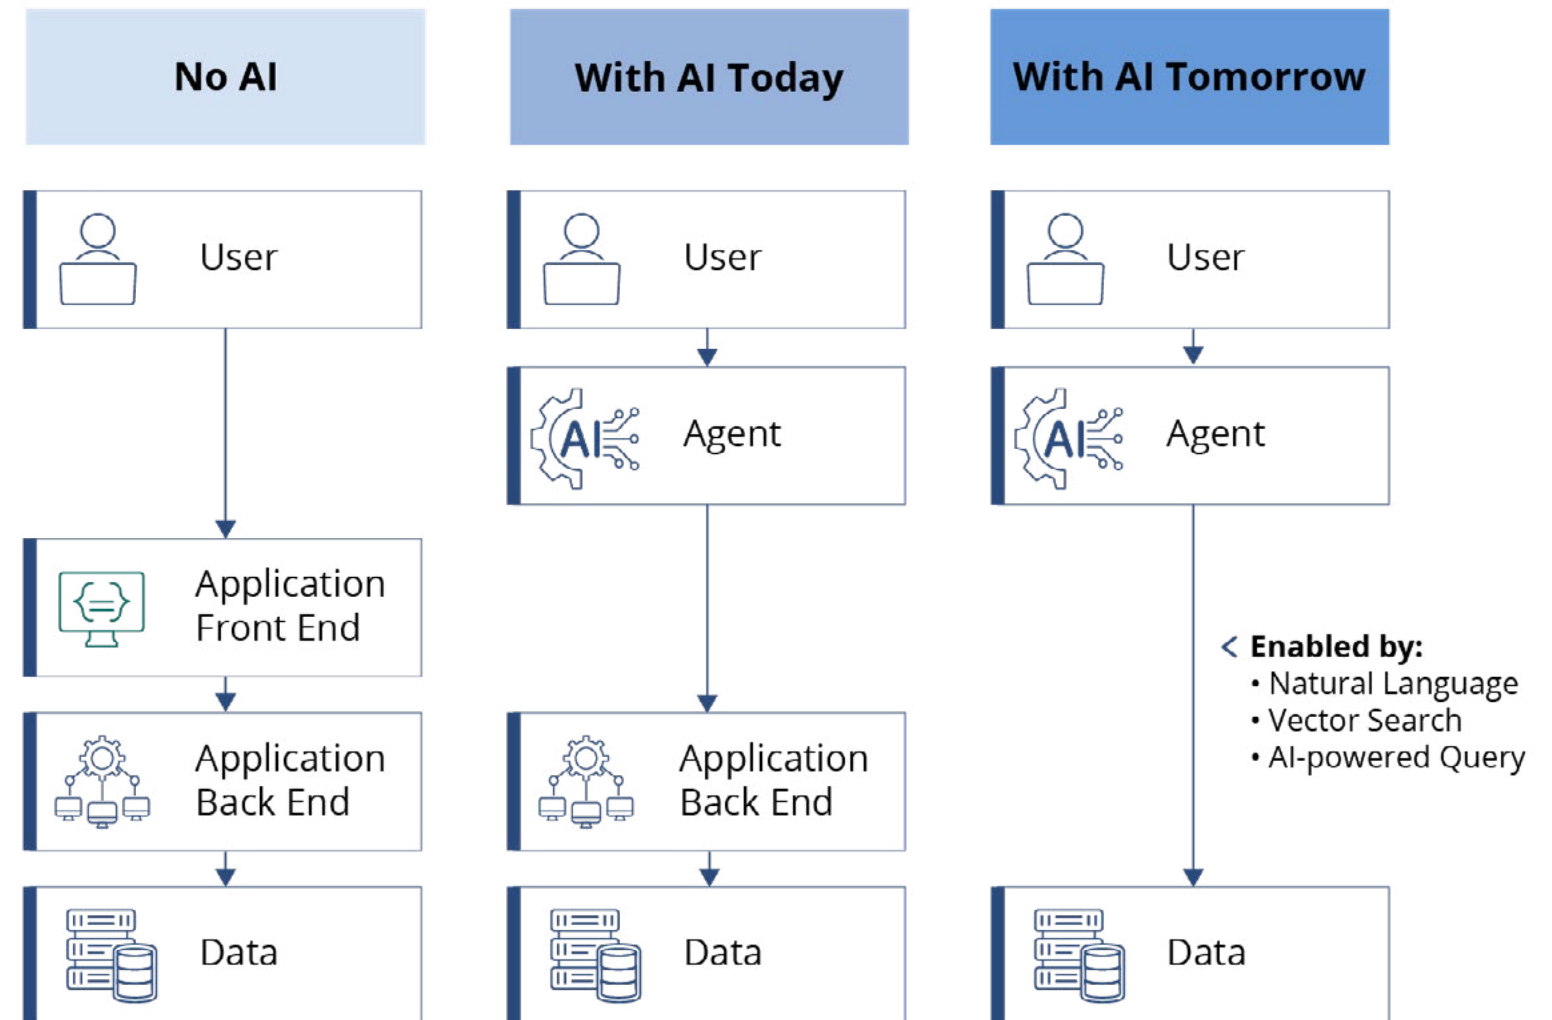
The lack of progress in extracting data from silos lies deep in the data structures of most applications. In most cases, the value of an app is that it adds context to raw data, a lot of times via black box semantic logic that only exists inside that application. During the digital boom, organizations started to leverage cloud technologies to establish data platforms that pull raw data from applications, recreate contextual logic in a central repository and deliver governed, transparent data products. Thus emerged the current standard medallion architecture framework, which has remained mostly the same with little innovation for decades.

Even today, the implementation of medallion architecture is synonymous with multi-year roadmaps. There is a growing sense that this high standard is a fruitless effort and that something needs to change.

At the same time, we are seeing organizations continue to stockpile data that they cannot use because of quality issues, and yet they still have to pay for its storage, security and compliance. In this way, data is quickly moving from being an asset to a liability.

Does data need to be done differently and why does this matter now?

Agentic AI has the potential to change everything for data.



Agents interact very differently with data in comparison to humans. Agents do not require context from applications to understand data; they can create their own from raw data and can do so at scale.

To really extract value from AI, organizations need to not only redesign critical processes and rethink how the task is done but also redesign the way data is incorporated into workflows.

The question is, is there a better way to do data? Can AI value be realized without some dramatic change in how data is managed? Is the medallion architecture finally ready to evolve and perhaps go away altogether or be redefined into something new? We can see that leading providers are leaning in by creating new data hubs – hubs driven entirely by agents for agents. But much more change is still to come.

Similarly, what role will applications play in the future? Without the need for direct application interactions, will agentic AI processes redefine the way we use, design and interact with data so significantly that applications and their purpose also significantly evolve?

One thing is clear, organizations have to stop thinking about their data approach in terms of establishing the data foundation and then spending years to get to value. It's much more urgent than that. If you are planning to invest in a data tower, but you're still using medallion architecture, it might be time to pause and reflect on whether that's the right way forward. Data continues to be the lifeblood of organizations, but it's time for data to be to put to work for you instead of you working for your data.

Getting Your Organization Ready to Achieve ROI

Agentic AI is not an IT solution – it's a new way of conducting business. Much like the early days of business process outsourcing (BPO), where enterprises had to re-architect their operations to fully realize its benefits, agentic AI demands a fundamental shift in how organizations are structured, how roles are defined and how work gets done.

From Technology Adoption to Organizational Transformation

While many enterprises begin their agentic AI journey with small, function-specific pilots, the path to ROI lies in scaling these efforts across the organization. This requires more than deploying agents – it requires becoming agent-ready. That means:

- **Redesigning roles:** As agents take on more decision-making and execution, human roles must evolve. Managers become orchestrators of outcomes, not just task supervisors. Analysts become curators of context, not just data processors.
- **Rewiring workflows:** Traditional linear processes give way to dynamic, goal-driven workflows where agents and humans collaborate in real time.
- **Reframing governance:** Oversight shifts from static controls to embedded escalation logic, real-time monitoring and adaptive feedback loops.

A New Operating Model for a New Class of Intelligence

To achieve ROI, enterprises must move beyond experimentation and build an operating model that supports agentic execution. This includes:

- **Capability Mapping:** Identify which roles, processes and decisions are best suited for agentic augmentation. Not every task needs an agent – but every function should be evaluated for agentic AI potential.
- **Skill Development:** Equip teams with new skills in agent orchestration, prompt engineering, escalation design and AI governance. These are not traditional IT skills; they are business-critical capabilities for the agentic era.
- **Change Management:** Prepare the organization for a shift in how work is done. This includes aligning incentives, updating performance metrics and fostering a culture of trust in autonomous systems.

Start Small, Scale Smart

Most successful agentic AI journeys begin with targeted use cases that deliver quick wins, such as invoice processing, customer onboarding or IT incident resolution. But the real value emerges when agents are connected, coordinated and scaled across functions.

To do this, enterprises should:

- **Invest in modular architectures** that allow agents to be reused and recombined across workflows.

- **Establish cross-functional governance** to ensure alignment, compliance and continuous learning.
- **Measure what matters** – not just efficiency gains, but improvements in decision velocity, customer experience and business agility.

From Promise to Performance: Operationalizing the Agentic AI Enterprise

As agentic AI capabilities accelerate, a gap is emerging and widening between the technology's potential and how much value enterprises are able to realize. Conversations with providers and early adopters consistently point to a clear pattern: while the technology is still evolving at pace, the critical barriers to impact are no longer technical; they are centered around a lack of organizational readiness.

To move from technology-led experimentation to enterprise-wide value realization, enterprises must overcome two foundational constraints: data that is unfit for dynamic decisioning, and operating models that were never designed for the decision velocity of an AI-driven enterprise.

Shifting the Focus: From Tech Deployment to Becoming AI-Ready

Agentic AI is not an overlay. Rather, it demands an AI-ready enterprise. That means rethinking how decisions are made, how workflows are orchestrated, and how data, people and technology interact to drive intelligent, adaptive outcomes.

Here's what mature agentic AI enterprises should look like – the North Star:

- **Decisions are distributed, dynamic and governed**, executed where they are most effective.
- **AI is embedded within value streams**, enabling workflows to adapt continuously based on goals and context.
- **Roles are redefined** with humans overseeing AI performance and managers becoming decision architects.
- **Architectures are modular and event-driven**, integrating predictive, generative and agentic forms of AI.
- **Governance is embedded** with explainability, escalation logic and trust mechanisms designed into AI execution.

In a mature agentic AI enterprise, the enterprise doesn't just use AI – it evolves through it.

The Enterprise as a Decision System

Real transformation begins when enterprises reframe themselves not as a set of siloed functions, but as a decision system, composed of interdependent layers that must evolve, both independently and together.

To move beyond PoCs and pilots, enterprises must rewire decision-making across the five functional layers that make up an organization:

1. **Infrastructure Layer** – where enablement decisions are made about platforms, security and observability.
2. **Back Office** – where transactional decisions are governed, e.g., payments, compliance and procurement.
3. **Middle Office** – where strategic alignment, risk and policy decisions take shape.
4. **Front Office** – where real-time, revenue-driving decisions engage customers.
5. **Industry-specific Service Delivery** – where execution decisions fulfill enterprise promises.

Each layer requires a tailored AI application guided by the level of autonomy, data context and governance maturity appropriate to its function.

Transformation typically begins within individual layers, where decision-making context is focused, and governance is manageable. Cross-layer coordination unlocks greater value, but it also requires greater maturity in data, architecture and organizational alignment.

To make any changes, the enterprise needs to look at adjusting four capability levers:

- **Data and intelligence:** Decision-grade, contextual data that flows in real time.
- **Technology and architecture:** Modular systems that support agentic AI interaction and explainable autonomy.
- **Process and decision architecture:** Goal-oriented execution models with embedded logic for escalation and governance.
- **People and organization:** Teams built to govern, trust and evolve with AI.

The more complex the decision system, the more foundational change is needed to sustain it. The journey begins functionally but scales systemically.

From Linear Process to Decision Network

One of the most profound shifts in the decision operating system (OS) is the move away from linear business processes to **composable decision networks**, wherein:

- Processes become recursive, context-aware workflows
- Execution becomes goal-driven, coordinated in real time by AI agents and
- Value creation becomes adaptive by design

This doesn't mean abandoning current operations. It means restructuring them to emphasize outcomes, decision velocity and system adaptability.

Automation accelerates tasks. Agentic AI redefines how the enterprise thinks and acts.

The Bottom Line: Decision Velocity as a Competitive Advantage

The future will not be won by those that pilot the most agents but by those that transform how they make and manage decisions. Agentic AI rewards those who act with purpose, adapt continuously and build systems where decisions – not processes – become the core unit of value creation.

Enterprises that rewire for decision velocity today will lead the AI-powered markets of tomorrow. Those that treat agentic AI transformation as a journey – not a one-time initiative – will build the resilience and agility required to stay ahead in an increasingly autonomous enterprise landscape.

Authors



**Loren
Absher**

Director,
Americas Lead –
AI Advisory



**Olga
Kupriyanova**

Principal Consultant,
AI & Data Engineering

Contributors (alphabetically):

Shafqat Azim

Alex Bakker

Taylor Card

Namratha Dharshan

Steve Hall

Stanton Jones

Prashant Kelker

Robert Kugel

David Menninger

Hemangi Patel

Srinivasan P N

Gowtham Sampath

ISG (Nasdaq: III) is a global AI-centered technology research and advisory firm. A trusted partner to more than 900 clients, including 75 of the world's top 100 enterprises, ISG is a long-time leader in technology and business services that is now at the forefront of leveraging AI to help organizations achieve operational excellence and faster growth. The firm, founded in 2006, is known for its proprietary market data, in-depth knowledge of provider ecosystems, and the expertise of its 1,600 professionals worldwide working together to help clients maximize the value of their technology investments.