

SUPPLY CHAIN RESEARCH

Independent · vendor-neutral research

Agentic AI in the Supply Chain

The 2026 Buyer's Guide & Vendor Landscape

A practitioner's guide to understanding, evaluating, and adopting agentic AI in the supply chain: what it is, where it applies, how the market and vendors stack up in 2026, what it costs, and how to separate real capability from hype.

Prepared by **Supply Chain Research** · June 2026 · Vendor-neutral: we take no payment from the vendors covered.

Contents

- 1 Executive summary
- 2 What agentic AI is
- 3 The market and adoption in 2026
- 4 The vendor landscape
- 5 How to evaluate agentic AI
- 6 Cost and pricing
- 7 Implementation: where programs succeed or fail
- 8 Trends shaping 2026
- 9 Segment-specific guidance
- 10 ROI and the business case
- 11 Frequently asked questions
- 12 Recommendations
- 13 Methodology and caveats
- 14 Sources

1. Executive summary

Agentic AI is the most-hyped supply chain technology of 2026, and the most misunderstood. The term describes autonomous AI agents that perceive, decide, and act toward a goal with limited human intervention, a step beyond the generative-AI copilots that assist a human and well beyond the traditional machine learning and optimization that has driven planning for years. The promise is a supply chain that senses disruption and responds on its own. The reality, today, is narrower: most production value comes from supervised copilots and simple single-task agents, not autonomous multi-agent systems.

This guide is written for supply chain, operations, and IT leaders trying to separate real capability from marketing. It is deliberately vendor-neutral: we accept no payment from the vendors covered, and we name no single best platform, because the right approach depends on your use case, data readiness, and appetite for autonomy. Our central message is to size the opportunity honestly. The pages that follow define agentic AI against the alternatives, present the market and the adoption gap side by side, profile the vendor landscape, and lay out how to evaluate, govern, and phase a program so it lands on the right side of Gartner's cancellation statistic.

\$53B

Gartner's 2030 forecast for SCM software with agentic AI, from under \$2B in 2025

>40%

of agentic AI projects Gartner expects to be cancelled by the end of 2027

~15%

of day-to-day work decisions Gartner expects to be made autonomously by 2028

What the evidence says

The projections are explosive, and they start from a tiny base. Gartner expects SCM software with agentic AI to grow from under \$2B in 2025 to \$53B by 2030, implying roughly a 90 percent annual growth rate off a very small starting point.

Adoption is far behind the hype. Across independent surveys the trajectory is consistent: roughly a third of enterprises have adopted agentic AI in some form, about 11 percent run it in production, and fewer than 10 percent have scaled it to tangible value.

Agent washing is rampant. Gartner estimates only about 130 of the thousands of vendors claiming agentic AI are real, and warns that many use cases marketed as agentic do not require agents at all.

Value today is in copilots and narrow agents. Exception management, document automation, and bounded procurement and planning workflows deliver now; end-to-end autonomous decision-making remains aspirational.

Data quality is the gating constraint. Roughly eight in ten organizations cite data limitations as the roadblock to scaling agentic AI, ahead of the model or the tooling.

2. What agentic AI is

An AI agent is software that pursues a goal across multiple steps: it perceives a situation from data, reasons about what to do, takes actions using tools and systems, and adjusts based on the result, all with limited human intervention. Agentic AI is distinct from a chatbot or copilot, which responds to a prompt and then stops, and from traditional optimization, which solves a fixed problem within fixed rules. The practical spectrum runs from assistants, to simple single-task agents, to advanced multi-agent systems that coordinate, and most of the supply chain market sits at the simpler end.

The spectrum from copilots to agents

Vendors apply the word agent inconsistently, so it helps to place capabilities on a spectrum rather than treat agentic as a binary label.

Capability	What it does	Autonomy	Maturity in SC
ML / optimization	Forecasts or optimizes within fixed rules	None (advisory)	Mature
Generative copilot	Answers, drafts, and explains on request	Low (human acts)	Emerging-to-mature
Simple single-task agent	Completes a bounded task end to end	Medium (supervised)	Early production
Advanced multi-agent	Coordinates several agents across a process	High (on-the-loop)	Mostly pilots

What makes a system truly agentic

- **Goal-directed autonomy:** it pursues an objective across multiple steps rather than answering one prompt.
- **Tool use and action:** it can call systems, trigger workflows, and change state, not just produce text.
- **Memory and context:** it retains state across steps and draws on enterprise data and history.
- **Guardrails and oversight:** it operates within defined limits, with a human in or on the loop for consequential decisions.

The enabling stack


Modern agents are built on large language models for reasoning, orchestration frameworks that chain steps and tools, and interoperability protocols, notably the Model Context Protocol (MCP) and emerging agent-to-agent (A2A) standards, that let agents connect to data and to one another. These protocols matter for buyers because they reduce lock-in and let a planning agent, an execution system, and a data platform work together rather than in isolation.

Agentic AI versus agent washing

The single most useful filter in this market is to ask whether a tool is truly agentic or simply rebranded automation. Gartner has been blunt that agent washing is widespread, estimating that only about 130 of the thousands of vendors claiming agentic AI are real, and noting that many use cases positioned as agentic do not require agents. A rules-based workflow with a chat interface is not an agent. When evaluating, look for genuine autonomy, multi-step reasoning, and the ability to act within guardrails, and discount the label itself.

3. The market and adoption in 2026


The agentic AI market is growing quickly, but sizing it is difficult because the boundary is unsettled and the projections are unusually aggressive. Even the starting numbers diverge by definition: estimates for the current market range from under \$2B for SCM-specific agentic AI to nearly \$8B for AI agents across all sectors. Treat the figures as directional, and read the explosive growth forecasts alongside the far more sober adoption data.



Source: Supply Chain Research analysis of published estimates, 2024-2025. Gartner projects the SCM-specific slice alone to reach \$53B by 2030 (see Figure 2).

Figure 1. Published estimates for the agentic AI market in 2024-2025, by definition. Even the current size is unsettled.

Research firm	Scope	Current size	Forecast / CAGR
Gartner	SCM software w/ agentic AI	< \$2B (2025)	\$53B by 2030
Grand View Research	Enterprise agentic AI	\$2.58B (2024)	46.2% CAGR
Mordor Intelligence	Agentic AI (all)	\$6.96B (2025)	42.1% CAGR
MarketsandMarkets	AI agents (all sectors)	\$7.84B (2025)	\$52.6B by 2030



Source: Gartner, April 2026 (implies roughly a 90% CAGR off a very small base). Broad agentic-AI market CAGRs across firms run about 42% to 46%.


Figure 2. Gartner's projection for SCM software with agentic AI. The curve implies roughly a 90% CAGR off a very small base.

Why the estimates diverge

The spread is a definition problem compounded by novelty. Some firms size agentic AI across all industries, some isolate enterprise deployments, and Gartner's headline figure measures SCM software that embeds agentic capabilities rather than standalone agent platforms. The growth rates, 42 to 46 percent for the broad market and roughly 90 percent for Gartner's SCM-specific slice, are extraordinary precisely because the base is so small. For planning, treat agentic AI as a fast-emerging capability being added to software you may already own, not a settled market with reliable numbers.

The adoption reality

The most important data in this guide is not the market size but the adoption gap. Independent research consistently shows a steep drop from experimentation to scaled value. McKinsey found that nearly two-thirds of enterprises have experimented with agents but fewer than 10 percent have scaled them to tangible value; MIT Sloan Management Review and BCG put adoption at about 35 percent within two years, with most efforts stuck in what they call pilot purgatory; and Deloitte found only about 11 percent running agentic AI in production. Gartner expects more than 40 percent of agentic AI projects to be cancelled by the end of 2027.



Source: McKinsey (June 2025), MIT Sloan Management Review / BCG (Nov 2025), Deloitte (Dec 2025), Gartner (June 2025). Figures are from separate surveys.

Figure 3. From pilots to production: the agentic AI reality gap, drawn from separate 2025 surveys that point the same direction.

Drivers and barriers


- **Drivers:** supply chain volatility and the demand for faster response, planner and analyst productivity pressure, the rapid maturation of large language models, and heavy vendor investment embedding agents into existing platforms.
- **Barriers:** data quality (the leading constraint), integration with legacy ERP and execution systems, trust and explainability, governance and risk controls, unclear return on investment, and agent washing that makes real capability hard to identify.

4. The vendor landscape

Agentic capability in the supply chain is arriving from four directions: the supply chain platforms building agents into planning and execution, the broad enterprise suites adding agents across their footprint, horizontal agent platforms that are powerful but not supply-chain-specific, and a wave of workflow startups automating narrow operational tasks. This section profiles the vendors that matter in 2026, grouped by role, with strengths and limitations. Because the category is young and claims vary widely, treat every capability statement as something to verify in a proof of concept.

How to read the landscape

The two questions that separate vendors are how supply-chain-specific the agents are and how mature and autonomous they really are. Supply-chain-native platforms embed agents in domain workflows; horizontal platforms offer general agent frameworks you must shape to supply chain; enterprise suites sit in between, strong where you already run their software; and workflow startups attack a single high-value task.



SCR's directional interpretation of vendor positioning in 2026, not an analyst ranking. Agentic capabilities are evolving rapidly and vendor claims vary widely; many tools marketed as agentic are not yet autonomous.

Figure 4. Agentic AI vendor positioning, SCR's directional interpretation for 2026. This is not an analyst ranking, and many tools marketed as agentic are not yet autonomous.

Supply-chain-native agent platforms

Kinaxis (Maestro Agents)

Kinaxis embeds agents in its concurrent-planning environment, with a human-in-the-loop design and guardrails, plus a Maestro Agent Studio for building custom agents and a planned agent marketplace for 2026. It supports MCP and agent-to-agent protocols, and cites customers including Jabil, Procter and Gamble, and Reckitt. IDC has highlighted the flexibility-and-guardrails approach. The strength is depth in planning; the limitation is that value is concentrated in planning rather than across the whole supply chain.

Blue Yonder (Cognitive Solutions)

Blue Yonder's Cognitive Solutions add a set of AI agents on a platform that integrates the Snowflake AI Data Cloud with a knowledge graph. Lenovo, cited by the vendor, reported a 5 percent forecast-accuracy gain, a 4 percent on-time-delivery improvement, and 10 percent higher delivery accuracy. The breadth across planning and retail execution is the draw; the outcomes are vendor-reported and should be validated.

o9 Solutions (Digital Brain)

o9 layers agentic capability onto its AI-driven digital twin of the supply chain, with clients including Kroger, Keurig Dr Pepper, and Coca-Cola, and a vendor-cited 1 to 3 percent EBITDA uplift. It is a strong fit where an enterprise wants planning and agents on one knowledge-graph platform.

Manhattan Associates (Active + Agent Foundry)

Manhattan extends its cloud-native Active platform with agentic capabilities and Manhattan Agent Foundry, letting users build agents for warehouse and order workflows. It is most relevant where Manhattan already runs the warehouse or order management.

Enterprise suites adding agents

SAP (Joule)

SAP is embedding more than 40 specialized Joule agents across its applications, with Joule Studio for custom agents reaching general availability in early 2026, and has demonstrated an order-fulfillment proof of concept with Nestle and IBM. The caveats are significant: a 2026 DSAG survey found only about 3 percent of SAP customers run Business AI in production, and Joule requires RISE or GROW contracts, with pricing in AI Units that the base allotment rarely covers. SAP is the default to evaluate for SAP-centric estates, with realistic expectations about production maturity.

Oracle and Microsoft

Oracle is adding agents across Fusion Cloud applications, and Microsoft is extending Dynamics 365 and Copilot with agents and an agent-building studio. Both are natural candidates where an organization already runs their suites, and both blur the line between supply-chain-specific and horizontal capability.

Horizontal agent platforms

Several platforms offer powerful, general agent frameworks that are not supply-chain-specific but are increasingly used to build supply chain agents: **Palantir AIP** (strong in data integration and operational decisioning), **Salesforce Agentforce**, **ServiceNow**, and **Microsoft Copilot and agents**. They offer flexibility and maturity in agent orchestration, at the cost of the domain depth that comes built into a supply-chain-native platform; expect to invest in shaping them to supply chain processes.

Workflow and emerging specialists

A distinct layer automates narrow, high-volume operational tasks rather than planning. **Kognitos** automates document-heavy processes, and one customer, Century Supply Chain, processes more than 50,000 bills of lading per month through it. Others include **SourceDay** (purchase-order and supplier management), **Leverage AI**, and **MarkIt**. These tools often deliver the fastest, most measurable value precisely because the use case is bounded.

Vendor summary

The table is an orientation aid, not a scorecard. Use it to frame where to look, then evaluate against the framework that follows, weighting genuine capability over the agentic label.

Vendor	Category	Best fit	Notes
Kinaxis	SC-native	Planning agents	Maestro Agents; human-in-the-loop; MCP/A2A
Blue Yonder	SC-native	Planning + retail	Cognitive Solutions; Snowflake + knowledge graph
o9 Solutions	SC-native	Planning digital twin	Digital Brain; vendor-cited 1-3% EBITDA
Manhattan	SC-native	Warehouse / order	Active + Agent Foundry
SAP (Joule)	Suite	SAP estates	40+ agents; ~3% in production; RISE/GROW
Oracle / Microsoft	Suite	Existing suite users	Agents across Fusion / Dynamics 365
Palantir / Agentforce	Horizontal	Custom agents	Powerful, not SC-specific; needs shaping
ServiceNow / MS Copilot	Horizontal	Cross-function	General agent frameworks
Kognitos	Workflow	Document automation	Century: 50,000+ bills of lading/month
SourceDay / Leverage	Workflow	Procurement tasks	Narrow, fast-to-value use cases

5. How to evaluate agentic AI

Agentic AI should be evaluated differently from packaged software. The question is not which product has the longest feature list, but whether a specific, bounded use case will deliver value on your data, under governance you trust. Score candidates against the same defined dimensions, and weight genuine capability and data readiness far above the agentic label.

The five evaluation dimensions

1. **Genuine capability versus agent washing.** Confirm real autonomy, multi-step reasoning, and the ability to act within guardrails. A rules engine with a chat front end is not an agent, and Gartner warns most claimants are not real.
2. **Use-case fit and value.** Start from a narrow, high-value, bounded use case (exception handling, document automation, a procurement workflow) with a measurable baseline, not a vision of full autonomy.
3. **Data readiness and integration.** Assess whether your data is clean and accessible and how the agent integrates with your ERP and execution systems. Data quality is the leading reason programs stall.
4. **Governance, trust, and explainability.** Decide where a human must stay in the loop versus on the loop, and require auditability, guardrails, and the ability to explain a decision before you let an agent act.
5. **Viability, cost, and lock-in.** Weigh the vendor's roadmap and stability, the true multi-year cost including data and integration, and whether open protocols such as MCP reduce lock-in.

Buy, build, or use embedded agents

Most organizations should start with agents embedded in software they already run (a planning platform, an ERP suite) or with a bounded workflow tool, rather than building multi-agent systems from scratch. Building is justified only where the use case is core, differentiated, and unavailable off the shelf, and where the data and engineering capability exist to support it.

A pragmatic evaluation process

6. **Pick one bounded use case** with a clear baseline and owner.
7. **Shortlist two or three options** that fit it well, mixing native, suite, and workflow where appropriate.
8. **Run a time-boxed proof of concept on your own data**, measuring outcome and effort against the baseline, with governance in place.
9. **Assess data, integration, and governance fit** as rigorously as the model output.
10. **Decide, then plan to scale in supervised stages**, not in a single leap to autonomy.

6. Cost, risk, and governance

The cost of agentic AI is highly uncertain, the risks are real, and governance is the gating factor that decides whether a program scales or is cancelled. Budget for far more than the software line: the data preparation, integration, and oversight usually cost more than the agent itself.

Cost element	Indicative figure	Notes
Embedded agents (suite)	Add-on to existing license	e.g., SAP AI Units, min 100 at ~7 euros/unit
SC-native platform agents	Custom subscription	Priced with the underlying planning platform
Workflow / point tools	Per-task or subscription	Often fastest payback on a bounded use case
Data and integration	Frequently the largest cost	Cleaning, access, and connecting systems
Run-time and governance	Ongoing	Inference, monitoring, and human oversight

The risks

Autonomy introduces risks that advisory analytics did not. An agent that acts on bad data or a flawed objective can make wrong decisions at scale and quickly. Explainability and auditability are immature, security and access expand the attack surface, and return on investment is unproven beyond narrow use cases. These risks are why Gartner expects more than 40 percent of agentic AI projects to be cancelled by the end of 2027, often because costs escalate, value is unclear, or risk controls are inadequate.

Governance: in the loop versus on the loop

The central governance choice is whether a human stays in the loop, approving each consequential action, or on the loop, monitoring an agent that acts within guardrails and intervening by exception. Most supply chain deployments today should start in the loop for anything consequential, with clear guardrails, audit trails, and the ability to explain and reverse a decision, and move to on-the-loop only as trust and evidence accumulate. Govern before you scale, not after.

Most vendor pricing is gated and usage-based; confirm via direct quotes, and model the data, integration, and oversight cost explicitly, because it typically exceeds the software.

7. Implementation: where programs succeed or fail

Agentic AI programs fail for reasons that have little to do with the model. The Gartner forecast that more than 40 percent will be cancelled by the end of 2027 reflects escalating cost, unclear value, and weak governance, not a shortage of capable technology. The programs that succeed start narrow, fix their data, and govern before they scale.

Why programs fail

- **Data quality is inadequate**, the leading constraint; roughly eight in ten organizations cite data limitations as the roadblock to scaling.
- **The business value is unclear**, because the program chased autonomy rather than a measurable, bounded outcome.
- **Governance and risk controls are weak**, so the organization cannot trust an agent to act.
- **The ambition outruns the maturity**, attempting multi-agent autonomy before proving a single supervised agent, often after buying a tool that was agent-washed.

>40% of agentic AI projects expected to be cancelled by end-2027 (Gartner)	<10% of enterprises have scaled agentic AI to tangible value (McKinsey)	~8 in 10 cite data limitations as the roadblock to scaling (McKinsey)
---	--	---

Three principles that separate success from failure

Start narrow and supervised. Prove one bounded, high-value use case with a human in the loop before expanding scope or autonomy.

Fix the data first. Clean, accessible data is the precondition; an agent acting on poor data fails faster and more expensively than a dashboard.

Govern before you scale. Put guardrails, audit, and explainability in place early, and move from in-the-loop to on-the-loop only as evidence accumulates.

A phased path

The lowest-risk pattern climbs the autonomy spectrum deliberately: begin with a copilot that assists planners and analysts, add a simple single-task agent for a bounded job such as document processing or exception triage, then allow supervised multi-step action, and only later move selected, well-understood decisions to human-on-the-loop autonomy. Each stage delivers value and builds the data foundation and organizational trust the next stage requires.

8. Trends shaping 2026

From copilots to agents

The clear trajectory is from generative-AI copilots that assist toward agents that act. Nearly every major supply chain and enterprise vendor shipped or expanded agentic capability in 2025 and 2026, and the center of gravity is shifting from answering questions to completing tasks, within guardrails.

Multi-agent orchestration

Vendors are moving from single agents toward orchestrated multi-agent systems in which specialized agents coordinate across a process. This is where much of the long-term promise lies and also where most deployments remain pilots, because coordination, governance, and data demands rise sharply.

Human-in-the-loop versus human-on-the-loop

The operating model is itself a design decision. The pragmatic 2026 posture keeps humans in the loop for consequential actions and on the loop for routine, well-bounded ones, with the boundary moving as trust and evidence grow rather than by default.

Interoperability: MCP and A2A

The Model Context Protocol and emerging agent-to-agent standards are becoming the connective tissue that lets agents reach enterprise data and work with one another across vendors. Buyers should favor platforms that support open protocols, because they reduce lock-in and make a multi-vendor agent estate workable.

Value today versus aspiration

The defining tension of 2026 is between the projections and the production reality. Value is real now in exception management, document automation, and bounded procurement and planning workflows; end-to-end autonomous supply chain decision-making remains aspirational. The vendors and buyers who acknowledge that distinction are the ones avoiding the cancellation statistic.

9. Where agentic AI applies across the supply chain

Agentic AI is not uniformly mature across the supply chain. The table summarizes where it applies and how ready it is today; the notes that follow add detail. Maturity ratings are SCR's assessment of production readiness, not vendor claims.

Function	Where agents help	Maturity today
Exception management	Detect, triage, and resolve disruptions and alerts	Higher
Document automation	Process bills of lading, invoices, customs paperwork	Higher
Procurement / sourcing	Analyze bids, draft and chase purchase orders	Medium-high
Order management / service	Answer order queries, resolve routine cases	Medium
Planning (demand / supply)	Propose forecast and supply adjustments, run scenarios	Medium
Transportation	Procure capacity, monitor and reroute shipments	Medium
End-to-end autonomy	Coordinate decisions across the whole chain	Low (aspirational)

The pattern is consistent: agents deliver soonest where the task is bounded, high-volume, and rule-rich, such as exception triage and document processing, and where a wrong action is cheap to catch and reverse. They are useful but still supervised in procurement, order service, planning, and transportation, where they propose and a human approves. And they remain aspirational for end-to-end autonomy across the chain, where coordination, data, and governance demands are highest. Sequence adoption from the top of the table downward.

10. ROI and the business case

The return on agentic AI is real in narrow applications and unproven at scale, so the business case should be built use case by use case on your own baseline. Vendor-reported outcomes are useful to size the opportunity but should be treated as vendor-sourced until validated.

5% forecast-accuracy gain reported by Lenovo on Blue Yonder (vendor-sourced)	1-3% EBITDA uplift cited by o9 for its platform (vendor-sourced)	50,000+ bills of lading per month automated at Century via Kognitos
--	--	---

The value levers

The value of agentic AI today concentrates in a few levers that compound. Productivity is the largest: agents remove manual data gathering, triage, and document handling, freeing planners, buyers, and service staff for judgment work. Faster exception resolution protects service and cost when something goes wrong. Document and workflow automation cuts cycle time and error on high-volume tasks. Where planning agents are used under supervision, modest forecast-accuracy and service gains can free working capital and improve delivery. The most reliable business case is built on your own productivity, service, and error baselines, with vendor figures used only to size the prize.

Use a simple frame: ROI equals annual benefit (labor freed, faster resolution, error and cycle-time reduction, and any planning or service gain) minus the annual all-in cost of the agent, its data and integration, and its oversight, divided by that cost. Because most value today is productivity rather than autonomy, build the case on bounded use cases and resist crediting speculative end-to-end automation.

11. Frequently asked questions

What is agentic AI?

Software in which autonomous AI agents perceive a situation, decide what to do, and act toward a goal across multiple steps with limited human intervention, going beyond copilots that only respond and beyond traditional optimization that works within fixed rules.

How is agentic AI different from a copilot or from machine learning?

A copilot assists a human who then acts; traditional machine learning forecasts or optimizes within fixed rules; an agent pursues a goal, uses tools, and takes action on its own within guardrails. They sit on a spectrum from advisory to autonomous.

What is agent washing?

The practice of rebranding existing automation or chatbots as agentic AI. Gartner estimates only about 130 of the thousands of vendors claiming agentic AI are real, so the label alone is not evidence of capability.

Is agentic AI production-ready in the supply chain?

Partly. It is in early production for bounded tasks such as exception triage and document automation; it is supervised in planning, procurement, and transportation; and it is still aspirational for end-to-end autonomy. About 11 percent of enterprises run it in production, and fewer than 10 percent have scaled it to value.

Where does agentic AI deliver value today?

In narrow, high-volume, rule-rich tasks where a wrong action is cheap to catch: exception management, document processing such as bills of lading and invoices, and bounded procurement and planning workflows under human supervision.

What does agentic AI cost?

It varies widely. Embedded agents are priced as add-ons to existing suites (for example, SAP AI Units), native platform agents are priced with the underlying platform, and workflow tools are per-task or subscription. The data, integration, and oversight cost usually exceeds the software.

What are the main risks?

Acting on bad data at scale, weak explainability and auditability, expanded security exposure, and unproven return on investment. These drive Gartner's expectation that more than 40 percent of projects will be cancelled by the end of 2027.

Human-in-the-loop or human-on-the-loop?

Keep a human in the loop, approving each action, for anything consequential, and move to on-the-loop monitoring only for routine, well-bounded decisions as trust and evidence accumulate.

What is MCP, and why does it matter?

The Model Context Protocol is an open standard that lets agents connect to enterprise data and tools, and agent-to-agent protocols let agents work together. Supporting open protocols reduces vendor lock-in and makes a multi-vendor agent estate workable.

How should we start?

Pick one bounded, high-value use case with a measurable baseline, run a time-boxed proof of concept on your own data with governance in place, prove value with a human in the loop, then scale in supervised stages rather than leaping to autonomy.

12. Recommendations

- 11. Filter for genuine capability, not the agentic label.** Require real autonomy, multi-step reasoning, and guardrails, and discount agent-washed tools; Gartner warns most claimants are not real.
- 12. Start where value is real today.** Target exception management, document automation, and bounded procurement and planning workflows before considering broader autonomy.
- 13. Use what you own first.** Evaluate agents embedded in your planning platform or ERP suite (Kinaxis, Blue Yonder, o9, SAP, Oracle, Microsoft) and bounded workflow tools, before building multi-agent systems.
- 14. Treat data and governance as the program.** Clean the data first, define in-the-loop versus on-the-loop boundaries, and put audit and explainability in place before letting an agent act.
- 15. Build the case on bounded use cases and your own baseline.** Validate vendor-reported outcomes in a proof of concept, credit productivity rather than speculative autonomy, and scale only what proves out.

13. Methodology and caveats

- **Market sizing diverges sharply and is unusually uncertain** (from under \$2B for SCM-specific agentic AI to nearly \$8B for AI agents across all sectors in 2024-2025); the explosive growth rates reflect a very small base, and all figures are directional.
- **Adoption statistics come from independent surveys** (McKinsey, MIT Sloan Management Review and BCG, Deloitte, Gartner) that use different definitions and samples; they are presented together because they point the same direction, not as a single dataset.
- **Vendor outcome figures are vendor-sourced** (Lenovo on Blue Yonder, o9 EBITDA, Kinaxis customers, the Century volume via Kognitos) and should be validated against your own baseline before use in a business case.
- **The vendor positioning in Figure 4 is SCR's directional interpretation** for 2026, not an analyst ranking. Agentic capabilities are changing rapidly, vendor claims vary widely, and Gartner does not endorse any vendor.
- **Product, pricing, and adoption details are accurate as reported** at the time of writing; this is a fast-moving category, most pricing is gated and usage-based, and roadmaps evolve quickly, so confirm directly.

14. Sources

- [1] Gartner (Apr 2026). [Gartner Forecasts SCM Software With Agentic AI Will Grow to \\$53 Billion by 2030](#). From under \$2B in 2025; 60% adoption by 2030.
- [2] Gartner (Jun 2025). [Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027](#). Includes the agent-washing estimate (~130 of thousands real).
- [3] McKinsey (Jun 2025). [Seizing the agentic AI advantage](#). Fewer than 10% have scaled; ~8 in 10 cite data limits.
- [4] Kinaxis (2025). [Kinaxis accelerates the agentic era of supply chain orchestration with Maestro](#).
- [5] Kinaxis. [AI agents for supply chain \(Maestro Agents, Agent Studio\)](#).
- [6] Supply Chain Digital. [Top 10 supply chain analytics software \(Blue Yonder agents; o9 EBITDA\)](#).
- [7] Manufacturing Digital. [How Kinaxis, o9, and Blue Yonder fix fragmented supply chains \(Lenovo outcomes\)](#).
- [8] SAP (Oct 2025). [SAP connects Business AI with new Joule agents and embedded intelligence](#).
- [9] Innobu (2026). [SAP Joule and the agentic enterprise \(production adoption and pricing\)](#). DSAG: ~3% of SAP customers run Business AI in production.
- [10] Kognitos (2026). [Top AI automation tools for supply chain operations \(Century Supply Chain\)](#). 50,000+ bills of lading per month.
- [11] Mordor Intelligence (2025). [Agentic AI Market](#). \$6.96B (2025) to \$57.42B by 2031, 42.1% CAGR.
- [12] Grand View Research (2024). [Enterprise Agentic AI Market Report](#). \$2.58B (2024), 46.2% CAGR.
- [13] MarketsandMarkets (2025). [AI Agents Market](#). \$7.84B (2025) to \$52.62B by 2030.

Additional figures drawn from: MIT Sloan Management Review and BCG (2025, agentic AI adoption and pilot purgatory); Deloitte (2025, ~11% in production); and vendor materials for Kinaxis, Blue Yonder, o9, Manhattan, SAP, Oracle, Microsoft, Palantir, Salesforce, ServiceNow, and Kognitos. Adoption statistics come from separate surveys with differing definitions. Vendor outcome claims are vendor-stated unless otherwise noted. Gartner does not endorse any vendor, product, or service depicted in its research.

Supply Chain Research is an independent, vendor-neutral research platform for supply chain and IT leaders. We accept no payment from the vendors covered. Figures should be validated against your own requirements before any purchasing decision.