

Supply Chain Research · Cornerstone Reference

What Agentic AI

Actually Means

How Is Agentic AI Different from Automation, RPA, and Traditional AI?

A system is agentic when it pursues a goal on its own, not when it runs a smart script. The test is five properties together: autonomy, planning, tool use, memory, and a feedback loop. Most supply chain systems marketed as agentic satisfy some but not all of them.

CONTENTS

- 01 The short answer
- 02 What makes a system agentic rather than automated?
- 03 How is agentic AI different from RPA?
- 04 How is it different from predictive and generative AI?
- 05 Is agentic AI the same as an AI agent?
- 06 How autonomous is supply chain agentic AI in practice?
- 07 Frequently asked questions
- 08 Method, sources, and where to go deeper

01 The short answer

Agentic AI describes a system that pursues a goal with a degree of independence: it decides what steps to take, carries them out using tools, and adjusts based on the result. It differs from robotic process automation, which follows fixed rules and breaks when the process deviates; from predictive AI, which produces a forecast or classification but does not act on it; and from generative AI, which answers a prompt but holds no goal across tasks. A useful test is whether the system exhibits five properties together: autonomy, planning, tool use, memory, and a feedback loop. Much of what is marketed as agentic in supply chain today is recommendation with human approval, which satisfies some of these but not autonomy.

5	>40%	10%
properties that together define an agentic system	of agentic AI projects Gartner expects to be canceled by end of 2027	of practitioners in one 2026 survey would trust fully autonomous decisions

Key takeaways


- **Agentic means goal-directed, not sophisticated.** A complex model that outputs a number is not agentic; a simpler system that pursues a goal and acts can be.
- **The five-property test is the practical filter.** Autonomy, planning, tool use, memory, and a feedback loop. Missing any one weakens the agentic claim.
- **RPA and agentic AI sit at opposite ends of one axis.** RPA executes fixed rules; agentic reasons over goals and self-corrects.
- **An AI agent and agentic AI are not identical terms.** The peer-reviewed distinction reserves agentic for orchestrated, multi-agent, memory-persistent systems.
- **Autonomy is a spectrum, and current deployments sit low on it.** Most production supply chain use is human-in-the-loop, not autonomous.

02 What makes a system agentic rather than automated?

Five properties recur across academic, standards, and major-vendor definitions, and a system needs all five to be meaningfully agentic. **Autonomy:** it can act toward a goal without a human directing each step. **Planning:** it decomposes a goal into steps rather than following a fixed script. **Tool use:** it calls external systems, APIs, or functions to gather information or act. **Memory:** it retains context across steps and often across sessions. And a **feedback loop:** it observes the outcome of its actions and adjusts, rather than running once and stopping.

The United States National Institute of Standards and Technology describes agentic AI as systems that function as autonomous agents capable of independently making decisions, learning from interactions, and adapting to changing environments. That framing matches the five properties and is useful precisely because it is not a vendor definition. The point of the test is discrimination: a system that scores a forecast but takes no action fails on autonomy and feedback, and a scripted bot that never plans fails on planning, however much either is described as agentic.

The same task handled four ways: only the last one pursues a goal



The distinction is not sophistication but goal directedness. RPA, predictive AI, and generative AI each perform one bounded step. An agentic system is defined by five requirements together: autonomy, planning, tool use, memory, and a feedback loop. Most systems marketed as agentic satisfy some but not all of these, which is what the definitional test in this page is for.

Figure 1. The same task handled four ways. The distinction is not how sophisticated the system is but whether it pursues a goal, which only the rightmost approach does. RPA, predictive AI, and generative AI each perform one bounded step.

Because the test has five parts, most systems on the market are partly agentic rather than fully so, and saying which parts are present is more useful than arguing about the label. A tool may plan and use other systems but escalate every consequential choice to a person, which is autonomy deliberately withheld. Another may act without asking but never revise its approach when the outcome disappoints, which is a missing feedback loop. Scoring a product property by property, across autonomy, planning, tool use, memory, and feedback, turns a marketing question into an engineering one, and it tends to reveal that the honest description of most current systems is agentic in some respects and scripted in others.

03 How is agentic AI different from RPA?

Robotic process automation executes predefined, rule-based steps. It is deterministic and fast, and it is well suited to stable, high-volume processes such as moving data between systems or generating a routine document. Its defining limitation is that it does not reason: when the process deviates from the rules it was given, it fails rather than adapts. A common analogy is that RPA follows a recipe exactly, while an agent is closer to a cook who adjusts to what is in the kitchen.

Agentic AI is built for the opposite condition. It reasons over a goal, chooses steps, and self-corrects when something does not work, which suits variable processes where the path cannot be fully specified in advance. The two are not competitors so much as different tools: RPA remains the better choice for a stable, well-defined task, and paying for agentic reasoning where a rule would do adds cost and unpredictability for no benefit.

The fair case for RPA, which the current wave of agent marketing tends to obscure, is that determinism is a feature in many supply chain processes. An operator often wants a task performed the same way every time, auditably, with no possibility of the system deciding to do something different. Where that is the requirement, a rule-based bot is not a primitive version of an agent, it is the correct answer.

04 How is it different from predictive and generative AI?

Predictive AI, the category that covers most supply chain machine learning in use today, forecasts or classifies. It answers how much demand to expect, or whether a shipment is likely to be late, and outputs a number or a label. It does not decide what to do about that output, and connecting a prediction to an action is a separate step that a human or another system performs. Predictive AI is not a weaker form of agentic AI; it does a different job, and most planning value in production today comes from it rather than from agents.

Generative AI produces content in response to a prompt, whether text, code, or a summary. It is well suited to drafting and explanation, but on its own it does not carry a goal from one task to the next, does not call external systems, and does not act. The accurate way to state the relationship is that agentic AI extends generative AI with planning, tool use, memory, and action. A language model is frequently the reasoning core inside an agent, which is why the two are conflated, but the model alone is not the agent. The agent is the model plus the scaffolding that lets it plan, remember, and act.

05 Is agentic AI the same as an AI agent?

Not quite, and the distinction is worth holding because vendors use the terms interchangeably. A peer-reviewed taxonomy published in 2025 separates the two. An AI agent, in that framing, is a single, largely task-specific system: a language model equipped with tools that performs a defined job. Agentic AI refers to the broader pattern of multiple agents coordinating, decomposing tasks dynamically, and maintaining persistent state across a workflow. In that reading, agentic is an architectural step beyond a single agent, not a synonym for it.

For a buyer the terminology matters less than the question it forces. Ask whether a product is one tool-using agent performing a task, or an orchestrated set of agents managing a process, because the second is substantially harder to build, govern, and trust, and the word agentic is applied to both. A single agent recommending a replenishment order and a multi-agent system autonomously running replenishment across a network are markedly different commitments wearing the same label.

06 How autonomous is supply chain agentic AI in practice?

Less autonomous than the marketing implies, and deliberately so. The systems in production today are predominantly human-in-the-loop: the software analyzes a situation, proposes an action, and a planner approves it. A representative deployment runs exception analysis overnight and stages a recommended scenario for a planner to review the next morning, which is a genuine capability and also plainly not autonomous goal-pursuit. It maps to the first four properties in the test while stopping short of the fifth acting without approval.

Independent signals support this reading. In one 2026 practitioner survey conducted for a supply chain software vendor, a majority preferred a model in which AI recommends and a human decides, and only around ten percent said they would trust fully autonomous supply chain decisions, even as confidence in AI overall rose. And Gartner has predicted that more than forty percent of agentic AI projects will be canceled by the end of 2027, citing escalating cost, unclear value, and immature risk controls, alongside a judgment that only a small fraction of self-described agentic vendors are doing anything that meets the definition, a pattern it calls agent washing.

The fair reading of that maturity picture is not that agentic AI is vaporware. It is that the category is early, that the credible deployments are assistive rather than autonomous, and that the gap between the two is exactly where a buyer should apply the five-property test. Verifying which properties a product actually exhibits, rather than accepting the label, is the whole of the discipline here.

In supply chain specifically, the deployed reality sits well short of full autonomy, and the reason is accountability rather than capability. Most operational deployments keep a person in the loop for consequential decisions, so the system proposes and a human approves, or the system acts inside narrow, pre-authorized bounds and refers anything outside them. This is a reasonable design where a wrong autonomous action carries real cost, such as committing inventory or dispatching freight, and it means that a large share of what is marketed as agentic in this domain is, by deliberate choice, a recommendation engine with a human decision at the end. That is not a criticism of the systems. It is a description of where the responsible boundary currently falls.

07 Frequently asked questions

Is agentic AI just RPA with a language model on top?

No. Adding a language model to a rule-based bot does not make it agentic unless it also plans, uses tools, retains memory, and adjusts based on outcomes. A model bolted onto a fixed script still follows the script. The agentic properties are about goal-directed behavior, not about which model sits inside.

Do we already use agentic AI if we run machine-learning forecasting?

Almost certainly not. Forecasting is predictive AI: it outputs a number and stops. That is valuable and is where most supply chain AI value sits today, but it is a different capability. An agent would take the forecast and act toward a goal, adjusting as conditions change.

What is the difference between agentic AI and generative AI?

Generative AI produces content in response to a prompt and holds no goal between tasks. Agentic AI adds planning, tool use, memory, and action, usually with a language model as its reasoning core. The model alone is not the agent; the agent is the model plus the scaffolding that lets it plan and act.

Does agentic mean the system runs without any human involvement?

It can in principle, but in supply chain today it rarely does. Most production systems are human-in-the-loop, recommending an action for a person to approve. Full autonomy is a point on a spectrum, and current deployments sit low on it by design, because trust and governance are not yet in place for higher autonomy.

How do we spot agent washing?

Apply the five-property test and ask what the system decides and acts on without a human. If it only produces a recommendation, or only generates text, or only follows fixed rules, it is not agentic regardless of the label. Ask specifically which external systems it calls and whether it adjusts based on outcomes.

Is a single AI agent the same as agentic AI?

Not in the precise usage. A single tool-using agent performs a task; agentic AI, in the peer-reviewed distinction, implies multiple agents coordinating with persistent state across a process. The second is harder to build and govern. The label is applied to both, so ask which one a product actually is.

Which supply chain tasks are actually run by agents today?

The most mature are assistive: exception management, replenishment recommendations, and procurement exception handling, generally with human approval. Fully autonomous re-routing or end-to-end orchestration is mostly piloted rather than in production. The honest current state is recommendation with a human in the loop.

What has to be true in our data before an agent can act?

The same foundations any automated decision needs: accurate, current, accessible data and a system the agent can act through. An agent acting on poor inventory data will act confidently and wrongly, and the autonomy makes the error faster rather than smaller. Data readiness is a precondition, not a later concern.

What is human-in-the-loop, and do agentic systems need it?

Human-in-the-loop means a person approves or can intervene in the system's decisions rather than letting it act unsupervised. Whether an agentic system needs it depends on the cost of a wrong action: where a mistaken decision commits inventory, money, or freight, keeping a person in the loop is a reasonable design rather than a failure of autonomy. Most current supply chain deployments keep a human in the loop for consequential choices for exactly this reason.

08 Method, sources, and where to go deeper

Method

- The five-property definition and the distinction between an AI agent and agentic AI follow a peer-reviewed taxonomy and the definitional framing published by the United States National Institute of Standards and Technology, rather than vendor definitions.
- Maturity and adoption signals are drawn from a named analyst prediction and a disclosed-methodology practitioner survey, each identified below and each treated as an interested source.
- Supply Chain Research is independent and vendor-neutral. We accept no payment from the vendors or categories covered, and this page names no products.

Caveats

- This is a fast-moving field and the specific figures date quickly. The Gartner cancellation prediction was published in mid-2025 and the practitioner survey in early 2026; both should be re-checked before relying on the numbers, and both are interested sources.
- The practitioner survey was conducted for a supply chain software vendor. Its methodology is disclosed, but a vendor-sponsored survey should be read as directional rather than definitive, and it is cited here for the autonomy-preference finding, not as a market measurement.
- Figure 1 is a conceptual distinction, not measured data. The boundaries between these categories are contested at the edges, and the five-property test is a practical filter rather than a formal standard.

Where to go deeper

This page defines the terms; two SCR resources go further. The SCR agentic AI guide covers the vendor landscape, capability areas, and adoption considerations for buyers evaluating these systems, which is the market view this definitional page deliberately does not attempt. The SCR editorial on AI-washing examines how AI capability is overstated in vendor marketing and how to test claims, which is the argument behind the agent-washing point raised in section 06. Readers scoping across categories should start with the SCR supply chain software category map.

Sources

- [1] Sapkota, Roumeliotis, and Karkee. [AI Agents vs. Agentic AI: a conceptual taxonomy, applications and challenges](#). Published in [Information Fusion](#). Peer reviewed; primary definitional source.
- [2] Sapkota, Roumeliotis, and Karkee. [AI Agents vs. Agentic AI, preprint](#). Open-access version of the peer-reviewed taxonomy.
- [3] National Institute of Standards and Technology. [Agentic AI](#). Standards body; neutral definitional framing.
- [4] arXiv. [AI Agents and Agentic AI: navigating a plethora of concepts for future manufacturing](#). Academic; supply chain and manufacturing relevance, including the autonomy-as-spectrum framing.
- [5] Gartner. [Gartner predicts over 40 percent of agentic AI projects will be canceled by end of 2027](#). Interested source: analyst that sells research. Origin of the agent-washing framing.
- [6] RELEX Solutions, via PR Newswire. [AI moves into core supply chain decisions as volatility persists, State of Supply Chain 2026](#). Interested source: vendor survey, disclosed methodology, n=514.

[7] Association for Supply Chain Management. [SCOR Digital Standard overview, used for process placement of automation and decision support.](#)

Supply Chain Research is an independent, vendor-neutral research platform for supply chain and technology leaders. We accept no payment from the vendors or categories covered. This page is analysis, not procurement advice, and its conclusions should be validated against your own circumstances before any decision.