

● SUPPLY CHAIN

Agentic AI in S&OP: who do you fire when the agent is wrong?

Every vendor is selling autonomous planning. Gartner predicts machines will resolve 60% of disruptions by 2031 — and that 40% of agentic projects will be cancelled by 2027. Both are true. Here's how to hold them in the same head.

MAY 22, 2026 · 11 MIN READ

There is a question that ends most agentic-AI demos when you ask it out loud: *who do you fire when the agent is wrong?*

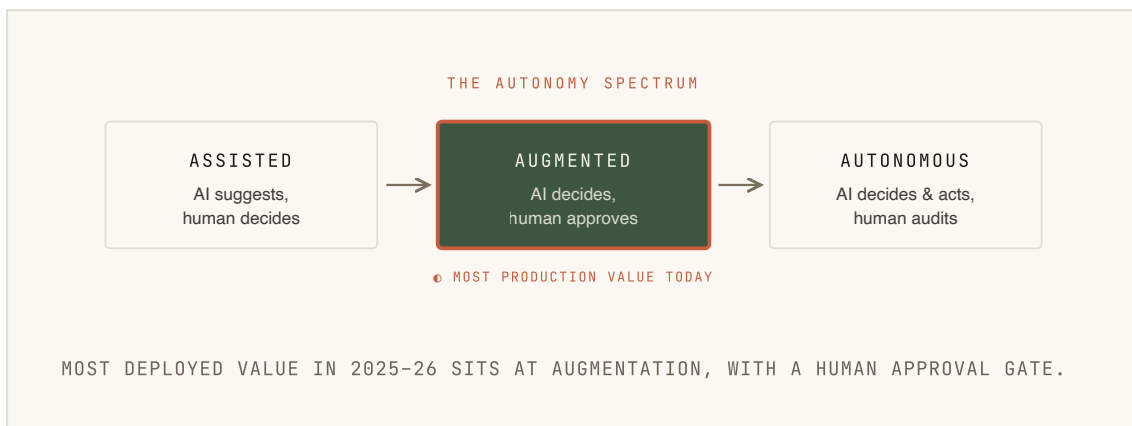
The vendor will say the agent operates within guardrails, that there is a human in the loop, that the decision is auditable. All true, and all beside the point. Because the entire promise of agentic AI in planning is that it removes the human from the loop for a growing share of decisions. The moment it does that at scale, accountability has to live somewhere — and in a cross-functional process like S&OP, accountability was already the thing that didn't work.

This is the gap between the agentic-AI promise and the agentic-AI reality. It is not a technology gap. It is the same governance gap that has been quietly killing S&OP processes for thirty years, now running at machine speed.

What "agentic" actually means this time

It is worth being precise, because the term is doing a lot of marketing work. The substantive difference between agentic AI and the machine learning that planning systems have used for a decade is three things: multi-step reasoning (decomposing a problem and calling tools to solve the parts), orchestration across systems of record, and memory — the agent records why it made a decision and what happened next.

This is a real step-change in *interface and orchestration*. It is not a step-change in optimisation accuracy. The optimiser is the same; the thing wrapped around it can now act on more steps without a human clicking through each one.



Gartner frames the destination as a spectrum: task automation, then decision augmentation, then full supply chain autonomy. *Gartner expects supply chain autonomy to become mainstream in the next five to ten years.* The word "mainstream" and the phrase "five to ten years" are both carrying weight there. Most deployed value today sits firmly in the middle band — augmentation, with a human approval gate.

The predictions, read carefully

Two Gartner headlines define the optimistic case, and both are quoted constantly without their qualifiers.

The first: *by 2030, half of supply chain management solutions will use intelligent agents to autonomously execute decisions.* The second: *Gartner predicts 60% of supply chain disruptions will be resolved without human intervention by 2031.*

Read precisely, these are predictions about *vendor product availability* and the *upper bound of automatable decisions* — not about what the average enterprise will actually run. And the same analyst house publishes the counterweight: *Gartner says autonomous planning has passed the Peak of Inflated Expectations on its Hype Cycle*, and that trust remains the binding constraint. The community of chief supply chain officers, in Gartner's own surveys, expects to trust autonomous decisions over only a small percentage of their cost of goods sold.

The most important number, then, is not the 60% or the 50%. It is this: *Gartner predicts over 40% of agentic AI projects will be canceled by the end of 2027, due to escalating costs, unclear business value or inadequate risk controls.*

40%+

of agentic AI projects predicted to be cancelled by end of 2027 — escalating cost, unclear value, inadequate risk controls

GARTNER, JUNE 2025

Hold both in your head. The same firm forecasting that machines will resolve most disruptions by 2031 is forecasting that nearly half of the projects trying to get there will be dead by 2027. These are not contradictory. The first is the destination; the second is the casualty rate on the road.

What is actually in production

Here the honesty test gets harder, because the gap between "generally available" and "running autonomously in production" is wide.

What is genuinely deployed clusters at the *task* and *execution* layers, not the cross-functional decision layer:

- **Kinaxis** early adopters use Maestro for decision *support* with human approval — the pharma example reduced inventory-risk identification from 40 clicks to 4, which is augmentation, not autonomy.
- **o9** reported more than 30 go-lives in a single quarter in 2025 — but these are planning-solution go-lives, not autonomous agents running unattended.
- **Blue Yonder** reported large volumes of warehouse tasks optimised — at the execution layer, not in S&OP.

The pattern is unambiguous: autonomy is being achieved in the warehouse, in transportation, in last-mile — bounded environments with clear feedback and low political content. It is not being achieved in the cross-functional S&OP decision, because that is where the organisation's unresolved conflicts live.

The S&OP problem agents can't solve

S&OP is not a technology process. It is a general-management process that supply chain has been handed responsibility for running without the authority to make it work. I have written about this before; the relevant point here is what happens when you point an autonomous agent at it.

The supply chain analyst Lora Cecere has made the structural argument repeatedly: *organisations try to work together in S&OP but are unclear on the rules, with leaders carrying different definitions of supply chain excellence*. The failure modes are ownership, incentives, and decision rights — not algorithms.

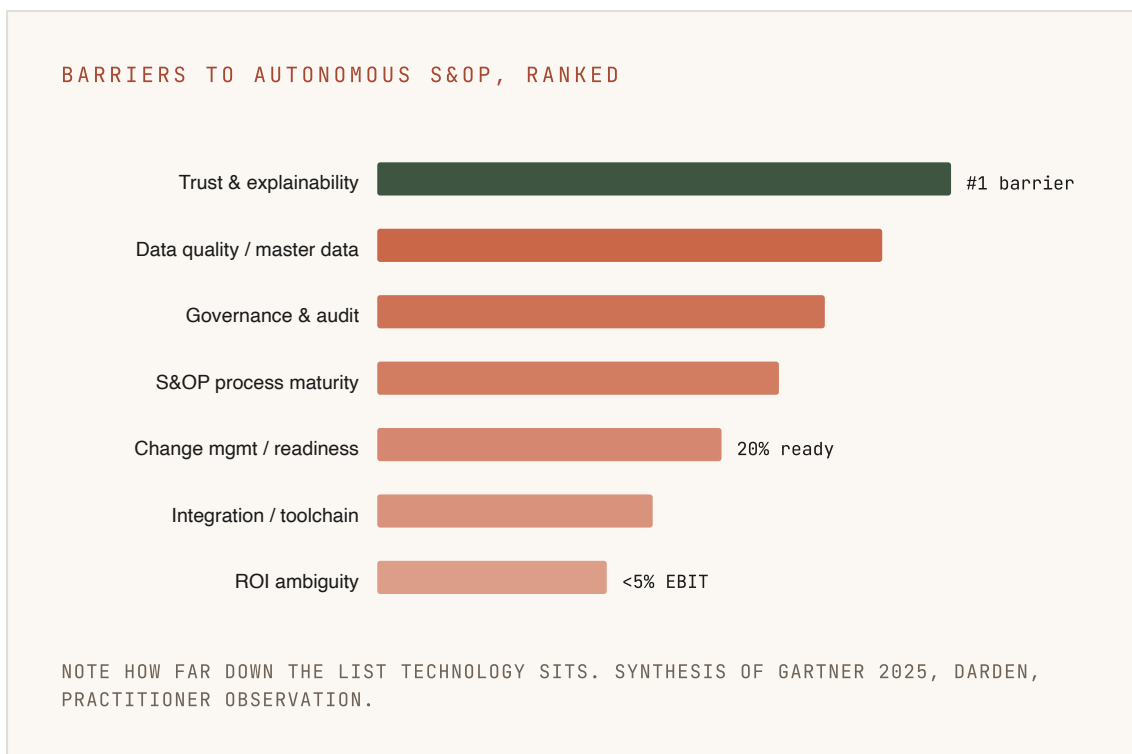
An agent deployed into that environment does not resolve the conflict. It automates one side of it. If commercial sandbags the forecast and finance pads the budget, an agent trained on that data learns to reproduce the sandbagging and the padding at speed. You have not removed the politics. You have encoded it and made it faster.

The University of Virginia's Darden school frames the path out of this as *moving from pilot purgatory to autonomous supply chains by deliberately matching autonomy levels to risk profiles*. Translated: there is no general-purpose autonomous S&OP. There is a portfolio of specific decisions, each with its own risk profile, that can be automated progressively as trust is earned. The exception triage can go first. The consensus demand number cannot — not until the process underneath it has an owner.

*An agent deployed into a broken S&OP doesn't resolve the conflict.
It automates one side of it.*

The barriers, ranked

Synthesising Gartner's 2025 surveys, the Darden analysis, and direct practitioner observation, the barriers to autonomous S&OP rank in roughly this order — and notice how far down the list "technology" sits.



1. **Trust and explainability** — the number-one barrier in Gartner's data.
2. **Data quality and master-data governance** — agents trained on test data fail in production.
3. **Governance and audit infrastructure** — who approves, who audits, who rolls back?
4. **S&OP process maturity** — pre-existing dysfunction gets automated.
5. **Change management and workforce readiness** — *Gartner found only 27% of executives have a comprehensive AI strategy and just 20% believe their workforce is truly AI-ready.*

6. **Integration and toolchain fragmentation.**

7. **ROI ambiguity** — most organisations attributing any EBIT impact to AI say it is under 5% of EBIT.

The risk nobody is pricing: the bullwhip at machine speed

The classic bullwhip effect propagates demand distortions upstream because each tier overreacts to the tier below it. The propagation takes weeks, which is — perversely — a safety feature. Weeks give a human time to notice that something is wrong and intervene.

Agentic systems acting at machine speed remove that buffer. A miscalibrated agent at one tier can trigger automated reactions at three more tiers before any human sees a dashboard. The distortion that used to take a quarter to ripple through a supply network can now do it in a day. We have spent decades designing S&OP to *dampen* the bullwhip. Pointing fast, unsupervised agents at the same network risks amplifying it instead — and the failure would be both faster and harder to trace.

What to actually do in 2026

For a planning leader, the defensible posture is neither refusal nor moonshot.

Buy the capability from your incumbent vendor rather than as a standalone bet; *Gartner notes AI is in the Trough of Disillusionment through 2026 and will most often be sold by the incumbent provider*

rather than bought as a new moonshot. Deploy at task level first — exception triage, parameter recommendation, post-cycle analysis — with human approval gates. Invest in data governance and audit infrastructure *before* you let any agent execute a decision. And resist the specific temptation to drop agents into the S&OP consensus before the process has clear ownership and aligned incentives, because an agent that votes without an accountability path will accelerate the loudest stakeholder's bias, not remove it.

Then watch the cancellation rate as your own leading indicator. If your pilot passes eighteen months without measurable economic impact, it is in the 40% cohort, and the honest move is to say so.

The uncomfortable conclusion

The autonomous supply chain is technically closer than the sceptics think and organisationally further than the vendors admit. The constraint was never the algorithm. It is that S&OP is a process about who gets to be right when functions disagree — and you cannot automate your way out of a disagreement you have not resolved.

Agentic AI will deliver real value at the edges of planning, fast. It will deliver autonomous S&OP roughly when organisations fix the ownership and incentive problems that already prevent S&OP from working — which is to say, not because of the technology, and not on the technology's timeline.

Who do you fire when the agent is wrong? Until you can answer that, you are not ready to let it decide.

Sources

- Gartner. (2025). *Gartner Predicts Half of Supply Chain Management Solutions Will Include Agentic AI Capabilities by 2030*. gartner.com
- Gartner. (2026). *Gartner Predicts 60% of Supply Chain Disruptions Will Be Resolved Without Human Intervention by 2031*. gartner.com
- Gartner. (2025). *Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027*. gartner.com
- Gartner. (2025). *Autonomous Planning Has Passed the Peak of Inflated Expectations*. gartner.com
- Gartner. (2025). *How CSCOs Can Maximize Growth With Autonomous Supply Chains*. gartner.com
- Gartner. (2026). *2026 Supply Chain Surge: Growth, Control Costs, and AI Truth*. gartner.com
- Kinaxis. (2025). *Maestro Agents*. kinaxis.com
- o9 Solutions. (2025). *o9 Reports Strong Q3 2025: 30+ Go-Lives*. o9solutions.com
- Blue Yonder / Business Wire. (2026). *Blue Yonder Transforms Retail Supply Chains With New AI-Driven Innovations*. businesswire.com
- Cecere, L. *In Search of Supply Chain Excellence*. Supply Chain Movement; LinkedIn.
- Laseter, T., & DuVall, T. (2025). *From Pilot Purgatory to Autonomous Supply Chains*. UVA Darden Ideas to Action.
- McKinsey & Company. (2025). *The State of AI in 2025*. mckinsey.com