ABS TechEDGE

Your Newsletter on AI, Analytics & Innovation @ ABS | November 2026

The Agent that doesn't wait to be asked

How AI is shifting from a tool you query to a system that acts first



For most of the past two years, using AI at work meant starting a conversation. You typed a question, wrote a prompt, or asked for a summary, and something answered back. However capable the tool, a person still had to go first.

FROM RESPONDING TO INITIATING

The industry term for this shift is **"Agentic AI"**, and it is worth defining plainly: it means AI systems that pursue a goal across several steps on their own, deciding when to check something, when to flag something, and sometimes when to act, without a person prompting each step along the way.

The scale of this shift is not subtle. Gartner expects 40% of enterprise applications to embed task-specific AI agents by the end of 2026, up from under 5% just two years ago. That is a faster adoption curve than cloud computing saw in its own breakout years. The interesting part is not that AI got smarter. It is that AI stopped waiting to be asked.

ABS TechEDGE
Your Newsletter on AI, Analytics & Innovation @ ABS | November 2026

WHAT THIS ALREADY LOOKS LIKE AT ABS

This is not a future-tense conversation for us. It is already how several parts of the business run:

Sales insight agents review activity and surface executive-level takeaways automatically, before anyone requests a report.

Trigger agents watch incoming emails and start a workflow the moment a relevant one arrives.

Notification agents monitor processes in motion and speak up mid-process, not after the fact.

Vendor performance agents track supplier metrics against consistent criteria, giving procurement clearer, more current data before a sourcing decision is made.

Process agents are taking over full steps of previously manual workflows, not just individual tasks within them.

None of this is a chatbot answering a question. It is software that noticed something, decided it mattered, and acted, or spoke up, without being asked first.



THE CAPABILITY LEAP, AND THE CATCH

This is a genuine leap from where this newsletter has taken you before. In August 2025, we covered AI that responds when asked: computer vision that flags a defect when you point a camera at it, a chatbot that answers when a customer types a question. What is live at ABS today is a different category: AI that initiates.

But the data on this shift comes with an honest caution attached. Industry research shows 79% of enterprises say they have adopted AI agents, yet only 11% actually run them in production. And Gartner expects more than 40% of agentic AI projects to be canceled by 2027. The reason is rarely that the technology fails. It is that organizations build the autonomy faster than they build the structure around it: who is watching, who gets escalated to, and who owns the outcome when an agent's decision turns out to be wrong.

ABS TechEDGE
Your Newsletter on AI, Analytics & Innovation @ ABS | November 2026

WHAT GOOD AGENT DESIGN REQUIRES

The organizations getting this right are asking three questions before scaling any agent:

- What is this agent allowed to decide alone, and what must it escalate?
- Can we see what it decided and why, in language a person can act on?
- Who owns the outcome if it gets this wrong?

None of those are technology questions. They are management questions, applied to software that is now doing management-adjacent work.

THE ABS PERSPECTIVE

The foundation for this was laid earlier than it might seem. The data discipline covered in January's issue is exactly what makes an agent's decisions trustworthy rather than confident-sounding guesses. The judgment and checkpoint thinking from last month is exactly what determines where a person still needs to be in the loop, and where an agent has earned the room to act on its own.

As agents take on more of the analysis behind ABS's decisions, not just the flagging, the real differentiator will not be which agents ABS has. Plenty of companies will have similar tools within a year. It will be whether ABS knows precisely what each agent is trusted to decide, and who is accountable when it does.

KEY TAKEAWAY

The capability question is settled. Agentic AI works, and it is already working here. The open question is no longer whether AI can act on its own. It is who owns what it decides.

The tools that changed how we work over the past year taught us to ask better questions. The tools taking shape now are starting to answer them before we ask. The advantage will not go to whoever deploys the most agents. It will go to whoever builds the clearest structure around them.



Stay Focused.
Stay Decisive.

atlantic-bearing.com