


AUTOMATE, THEN INNOVATE: HOW AI IS FLIPPING THE SCRIPT

To achieve full automation for agentic systems, enterprises need a partner with an engineering culture built for zero-error tolerance, not a "beta product" mindset

BY HOLDEN SPAHT, MANAGING PARTNER

APRIL 2026



AI IS FLIPPING THE CONVENTIONAL SCRIPT

Automate, then innovate.

For decades, that's been the conventional order of operations on how enterprises integrate new foundational technologies. Automation means using technology to do what the enterprise *already does* but faster, cheaper, better. It's about *exploiting* what you already have and know. Innovation, meanwhile, means discovering and doing new things that the enterprise couldn't previously do or in some cases even imagine. It's about *exploring* new possibilities and unknown white spaces.

This logic stands on the assumption that automation is easier than innovation. Automate the known and learn as you go about the practical challenges of execution. Then explore the newly possible, combining imagination and creativity with the practical knowledge and experience gained along the way.

But AI seems to be flipping that script. This is partly because generative AI models excel at exploring uncharted spaces by exposing connections between ideas that people don't typically see. It's central to how today's generative AI models are built and structured. Ask ChatGPT what you haven't considered, where the blind spots are in your arguments. Ask Claude to surface subtleties you missed. You often get surprisingly insightful and creative results. And this is how many people use AI today, to good effect.

It's also true that many seemingly straightforward and even simple AI automations are hard to execute successfully in practice, in part because the technology is immature and evolving so quickly. Automation of complex workflows is yet another step harder. And when organizations look at the promise of full automation — taking humans out of a decision-and-action loop — they frequently find that the practical demands and roadblocks are extraordinarily high. The coming wave of agent-based automation will likely reveal the sharp edges of this dilemma, by exposing just how hard it is to implement and just how much risk is carried along the way.

This flip of the automate-innovate script isn't really about AI capability itself. It's about where and when an organization will choose to deploy beta agent products and risk cascading failures in mission-critical systems, as opposed to using AI to surface interesting new ideas and possibilities.

THE HIGH STAKES OF BEING “ALMOST” RIGHT

Consider a scenario where an enterprise is migrating benefits systems from one provider to another — perhaps changing health insurance companies. Using AI to research and explore how to put together a benefits package that employees might prefer is a form of innovation. Automating the actual migration is the harder problem for AI. Autonomous agents must precisely identify which and how many employees and dependents are covered. And they must perfectly transfer employee coverage elections and contributions from one system to another. All while safeguarding protected health information along with other sensitive details.

Bad things happen when something goes wrong in the migration process. Coverage gaps. Payroll breaks. Employees lose health insurance because a field was misread or an instruction misinterpreted. Compliance violations turn into lawsuits. The CFO wakes up at 2 a.m. to a crisis that could spread across the entire employee base. Confidence and trust are in deep trouble.

Consider another example, where the transport industry is grappling with these risks in real time. Using AI to explore route optimization or predict traffic patterns is a fantastic form of innovation (and useful even when imperfect). Using AI for Level 5 vehicle autonomy (full self-driving without human control) is much harder and has essentially zero tolerance for error. But we have to try, because the benefits are huge: Road crashes kill an estimated 1.19 million people and cause a multiple of that in serious injuries and disabilities each year. A recent Waymo study found a 96% reduction in any-injury-reported vehicle-to-vehicle intersection crash events and a 91% reduction in airbag deployment events. Based on this data, full automation right now would vastly reduce human suffering along with the massive economic costs of road accidents. But so far at least, society won't allow it, because the one accident caused by a robot car seems less acceptable than perhaps a thousand caused by humans.

WHY FULL AUTOMATION IN THE ENTERPRISE DEMANDS ZERO TOLERANCE

In enterprise contexts, the stakes differ. But the consequences of (and tolerance for) agent automation error are often equally unforgiving. Yet as with autonomous vehicles, enterprises have strong incentives to move toward full automation because the potential benefits are so great.

The stakes depend on the scope, so a little more granular reasoning helps to clarify how this can work.

- There's narrow automation that has become part of the baseline: for example, rules-based task automation where AI speeds things up a bit, like auto-routing support tickets based on keywords. Or simple work process improvements where AI assists with tasks but humans remain in control, such as flagging duplicate vendor records for a human to review. AI likely saves meaningful time and money here.
- Then there's human-in-the-loop innovation, where AI augments human capability to do new things. If a CFO uses an AI agent that asks her questions, points out unobserved connections, or “advises” or “augments” what the CFO is doing, the agent is acting as an adjunct to human decision making. That's incredibly valuable innovation, but it's not full automation.
- Full automation is where the agent takes on the entire workflow, makes decisions and executes — removing the human in the loop. This is where huge efficiency gains are to be had.

The scope for full automation is massive — after all, organizations are full of “mundane” processes that are rarely the best and highest use of human effort. Automating them would be the modern equivalent of simple first-generation industrial machines. It would relieve humans and animals from having to pull a plow or chop wood with an axe.

But unlike simple industrial tools, in automated business processes, there's no do-over. If a plow digs too deeply or a sawmill cuts some wood unevenly, the error is contained, visible and fixable. But an error in major enterprise systems isn't contained: It's likely to cascade through interconnected systems unpredictably. The autonomous agent can't afford to be wrong. The more enterprises rely on agents to fully automate intricate systems and decisions, the lower the tolerance for error anywhere in the process. This demands not just accuracy, but full data integrity across multiple workflows.

WHY LIVED EXPERIENCE TRUMPS RAPID PROTOTYPING

Companies need software partners who can deliver that level of performance and assurance. And I believe three concrete advantages separate those partners from everyone else:

1 Information density and “ground truth”

Extensive data across many enterprises, engaged with over a long period of time, is a major advantage in charting a path to ground truth. The data reveals patterns, and patterns become anomaly identifiers. The agents know when they encounter an exception, distinguishing routine scenarios from edge cases that require (rare but essential) human intervention. A company that's seen thousands of benefits migrations understands where the edge cases are likely to show up. A startup running its first migrations doesn't.

The advantage isn't that patterns guarantee accuracy. It's that pattern matching constrains the risk of full automation. You automate more confidently when you know the terrain.

2 Institutionalized risk mitigation

The deeper advantage for enterprise software companies is cultural and structural. Software systems of record have an operational and cultural DNA calibrated to near-zero error tolerance. These organizations think about production systems differently than startups do. They've spent decades building systems where downtime costs millions, where data loss is unacceptable and where compliance violations trigger lawsuits. This lived experience shapes culture and culture shapes product — from how engineering works to how you test, deploy and decide what to ship.

This isn't a negative judgment on startup culture. It's a structural incompatibility. When a CFO hands over a mission-critical benefits migration to an autonomous agent, she's making a bet about organizational discipline. Does this partner's default mode assume things could go catastrophically wrong, and build accordingly? Or does it assume you ship betas and learn from what breaks in the wild? The colloquial word here is “trust”, but this is the difference that gives that word meaning. A customer who's willing to experiment with full automation that takes humans out of the loop is much more likely to do so with a partner that has an engineering and product culture where things can't afford to go wrong, where risk tolerance is low, where the “beta product mindset” doesn't infect what gets shipped to the customer.

3 Tacit operational knowledge

Long-standing relationships with customers, particularly buyers, create a deeper understanding of the operating requirements for agents that may be granted full autonomy. The tacit knowledge of how a business actually runs — not how it's supposed to run, but how it really operates — becomes invaluable when autonomous systems must navigate the unwritten rules and informal workarounds that keep enterprise processes functioning.

The script has flipped. Innovation is now easier than automation. There's value in both, sequenced appropriately. And when it comes to automation, I believe the winning providers won't be the fastest to prototype; they'll be the organizations that have been right, repeatedly, in some of the highest-stakes environments.

This document is provided for informational purposes only and reflects current beliefs; it should not be construed as investment advice, a recommendation, or an offer to buy or sell any securities. Forward-looking statements and projections are subject to risks and uncertainties that could cause actual results to differ materially.