




How to resolve the roadblocks to enterprise AI adoption

It's not about the technology. It's about the people.

*An Intellyx Analyst Guide for BMC Helix
by Jason Bloomberg, Managing Director
October 2026*

Enterprises define and implement AI strategies to generate business value. However, many organizations are running into roadblocks on the path to strategic AI adoption.

Human-centric limitations, rather than technological deficiencies, are the primary roadblocks preventing enterprises from successfully scaling AI initiatives and creating digital capacity to deliver innovation for the business.

“Human in the loop” governance is fundamentally flawed because human constraints like reviewer fatigue and automation bias fail to match the accelerated decision-making speed of AI agents. Successful agentic AI requires a new operating model that puts humans in control, not simply in the loop.

BMC Helix ServiceOps delivers IT service and operations for agentic AI, supporting a new operating model that addresses organizational bottlenecks by allowing AI agents to handle repetitive tasks while reserving judgment-based responsibilities for human operators.

The challenge of digital capacity

Remember digital transformation? It wasn't that long ago that every enterprise struggled to implement strategies that leveraged digital technologies to break down organizational silos so they could better meet the needs of their customers.

Then ChatGPT and the rest of the AI wave came along, and everything changed.

Digital transformation initiatives suddenly went on the back burner. AI strategies rapidly became the C-suite *cause célèbre*, sucking up technology budgets to drive AI adoption across the entire enterprise. Today, despite the possibility of a looming correction, enterprise investment in AI continues to explode.



Digital capacity is the maximum business value the organization can expect to achieve by optimally leveraging all the technology options open to them (including automation and AI) to transform the organization to better meet customer needs into the future.

Despite all the hype around AI and all the money poured into it, however, many enterprises' AI initiatives have stalled. Extracting business value from the technology while adequately managing its risks has proven far more difficult than predicted. In fact, according to [a study by Upwork](#), almost half of enterprises have been unable to improve productivity with their AI efforts.

Perhaps it's time to revisit those sidelined digital transformation strategies, only now including AI among the list of digital technologies such strategies can leverage.

In other words, how can enterprises optimize their *digital capacity*: the maximum business value the organization can expect to achieve by optimally leveraging all the technology options open to them (including automation and AI) to transform the organization to better meet customer needs into the future?

Digital transformation was always about transforming IT from a cost center to a driver of business value. And business value is precisely what today's AI initiatives are struggling to achieve.

AI technologies are everywhere. They're plentiful and available. So why are organizations finding it so difficult to achieve meaningful business results from deploying AI?

To answer this question, organizations must remember the hard lessons of failed digital transformation efforts of the past: the primary roadblocks to business value are not technology-centric. They are *human-centric*.



Organizations must remember the hard lessons of failed digital transformation efforts of the past: the primary roadblocks to business value are not technology-centric. They are human-centric.

Until organizations resolve the human limitations to success with AI, they won't be able to maximize their digital capacity to achieve their strategic goals.

The agentic orchestration squeeze

The first step to resolving the human-centric bottlenecks that prevent AI initiatives from delivering strategic business value is, of course, to understand the problem.

The challenge that most enterprises are running into today is that AI agents behave differently than previous generations of technology. They potentially deliver exceptional business value – but also suffer from a corresponding increase in risk.

Both sides of this equation – value as well as risk – derive from the fact that agents transform how technology drives automation.

As generations of automation technologies have come to market, enterprises have been able to automate increasingly human-centric processes. The rise of agentic AI takes this progression to the next level, because agents can behave *non-deterministically*.

Instead of fixed, predictable outcomes, agents can learn as they go and achieve the goals set out for them in sometimes surprising ways. This *agentic orchestration*, where agents are able to decide for themselves which actions to take and in what order to achieve their goals, is both extraordinarily powerful and potentially dangerous.

As agents become an increasingly powerful driver of automation and digital capacity generally, they also become an increasingly problematic source of risk. The more AI drives automation across the organization, the less governable it becomes.



As agents become an increasingly powerful driver of automation and digital capacity generally, they also become an increasingly problematic source of risk. The more AI drives automation across the organization, the less governable it becomes.

Raising the bar on agentic AI governance

Agentic AI governance presents significant challenges for the team responsible for managing AI agents in production. Should it prevent agents from behaving nondeterministically – learning as they go and making decisions about what actions to take, even though sometimes nondeterministic agents go off the rails?

Some teams are taking that step. But locking agentic AI down to bring it under control jeopardizes the very capabilities that agents promise in the first place.

The current state of the art for addressing this problem is *human in the loop* or HIL. The idea with HIL is that to prevent agentic automation from going off the rails, we need a human to review and approve agentic decisions and actions.

Most of today's agentic AI governance approaches leverage HIL. Nobody wants agents to operate entirely autonomously, where no human is there to keep agents in line should they misbehave. The simplistic answer is to put a human approval step in each agentic orchestration.

Just one problem: HIL is fundamentally flawed. Here's why.



'Human in the loop' is fundamentally flawed.

The challenge with 'human in the loop': it's the humans

As agents mature and agentic orchestration improves, the rapid advance of automation technology squeezes the role that humans play.

At the same time, people always succumb to human constraints. We have short attention spans. We can only focus on one thing at a time. Rushing decisions leads to poor decisions. Meeting in groups to reach consensus takes time and effort, and more often than not, leads to interpersonal conflict. The list goes on and on.

Meanwhile, agents are making increasingly important decisions on accelerated timelines. The result is an inevitable impedance mismatch, as even the best laid governance plans can go awry in the face of human limitations.

To illustrate this point, here are nine human limitations that impact HIL participation in agentic orchestration. Which of these constraints does your organization suffer from?

1. Feedback Quality and Bias

How do people decide which agents to build and deploy in the first place? How do you know you're deploying the right ones?

Organizations often rely heavily on feedback loops to determine whether agents are valuable. Such feedback loops include humans, even though human feedback is often inconsistent, biased, or incomplete. Organizations may optimize agents for popularity rather than business outcomes.

2. Automation Bias

People watch newly deployed agents closely, as their behaviors are unclear. Over time, however, people come to trust the agents. As agents become more reliable, however, humans tend to trust them excessively.

Errors may be rare, but when they do occur, they can quickly run amok, because humans have become passive, reluctant approvers rather than active reviewers.

3. Reviewer Fatigue

HIL makes sense when there are only a few loops. But what happens when there are hundreds or thousands of agentic orchestrations in production?

Human reviewers can easily become overwhelmed as the number of agents grows. Governance exists on paper but becomes ineffective in practice.

4. Accountability Ambiguity

If a person is responsible for a task and something goes wrong, it's easy to figure out who's responsible. But what if an agent misbehaves? Is a human responsible? If so, who?

Agentic accountability is complex, and only becomes more convoluted when agents mature, the number of agents in production grows, and they increasingly interact with each other. Organizations struggle to assign responsibility when incidents occur. After all, nobody wants to end up holding the bag.

5. Tacit Knowledge Loss

Agentic orchestration can replace human activities, especially for people in junior roles. Over time, there will be fewer junior people to become senior people, which means that eventually, there will be fewer senior people who have a deep understanding of their respective areas of expertise. When something goes wrong, therefore, there will be nobody who knows how to fix the problem.

Furthermore, many processes – especially in IT – rely on informal expertise that is difficult to encode. Organizations eventually become dependent on systems they no longer fully understand, potentially leading to catastrophic failures.

6. Escalation and Exception Handling

AI agents will do their best to accomplish the tasks set out for them, even if they don't have the necessary capabilities to complete those tasks. Furthermore, agents may not even recognize when a situation falls outside their competence.

As a result, the agents may complete their tasks poorly or achieve the wrong outcomes altogether – with nobody the wiser until the errors build up over time. Even with HIL, humans enter the process too late after the agent has already caused extensive disruption.

7. Measuring the Wrong Things

Which tasks should agentic orchestrations automate, and which tasks should organizations reserve for humans? More importantly, how do you tell the difference?

The ostensible answer is to measure resulting outcomes to determine the best course of action. However, many valuable human contributions are difficult to measure.

As a result, organizations may prioritize automating highly measurable work while neglecting strategically important but less quantifiable activities. In the end, agents end up working on the wrong things.

8. Trust Calibration

AI teams usually implement overly simplistic AI governance constraints: establish a set of policies and then ensure that agents conform to those policies.

The reality of agentic orchestration is more complex, as different agents may exhibit different levels of autonomy. A one-size-fits-all governance model can either over-constrain low-risk agents or under-control high-risk ones.

9. Workforce and Organizational Impact

Do your human personnel like working with agents? Do they trust and appreciate the value agents bring to the table, or do they feel that agents (and AI generally) are more of a threat than a contributor of value to the organization?

Such workforce dynamics may limit successful outcomes and thus impede the digital capacity of any organization. Technical success may not translate into organizational success if people perceive agents as threats rather than useful tools.

Despite these pitfalls, it's easy to look at the power of today's AI and decide that such technology can replace humans, even for complex tasks that require insight and decision-making.

However, as the challenges with HIL attest, this human replacement strategy is doomed to fail. Instead, we must always remember that as with any technology we bring into our organizations, agentic AI offers tools in human hands that humans must decide how to wield.

Any approach to agentic AI governance that emphasizes HIL misses this point. Humans can't simply be in the loop. They must remain in control of agentic AI governance so they can bring their judgment, oversight, trust, and accountability to agentic orchestration.

The burden, therefore, shifts to the human side of any enterprise. Organizations must redesign roles, implement reskilling and adaptation across the workforce, and actively capture and build organizational knowledge.

Organizations need clear rules defining who is responsible for decisions, when humans should intervene, and how they should monitor and manage autonomous agents.

How BMC Helix helps organizations address the human challenge of agentic orchestration

Given the challenges of agentic AI governance are essentially human, no software vendor can solve such problems simply by offering better technology.

The team at BMC Helix realizes this fundamental truth and places human considerations at the forefront of their ServiceOps strategy.

BMC Helix's customers are creating new operating models for the digital workforce by providing centralized lifecycle governance for autonomous AI agents.

This operating model governs how organizations deploy, monitor, update, and eventually retire agents, empowering a human-centric workforce rather than an agentic one.

BMC Helix believes that organizations must balance agentic autonomy with the appropriate oversight – not simplistic HIL, but retaining high-value responsibilities that require judgment, collaboration, and creativity for human operators while empowering AI agents to orchestrate repetitive operational tasks.

Three BMC Helix agents in particular provide these human-centric capabilities:

- **Change Risk Advisor** – analyzes contextual risks associated with IT changes and recommends appropriate actions for human operators
- **Knowledge Curator** – keeps organizational knowledge up to date, thus supporting a reliable foundation for both human and agentic decision making
- **Insight Finder** – makes advanced analytics accessible to non-experts via dashboards and natural language, thus supporting a more productive collaborative relationship between humans and the agentic tools they use to achieve their intended business outcomes.

Via these three agents (and inevitably more to come), BMC Helix supports a human-centric model that provides the structure, visibility, and intelligence to support human-centric agentic AI governance.

As a result, BMC Helix enables organizations to scale their deployment of agentic AI in production to deliver the business outcomes each enterprise desires.

The Intellyx take

This paper began with an analysis of digital capacity. How can organizations best leverage technology capabilities to deliver on business outcomes?

Now that AI – in particular, agentic AI – is the central technology capability that enterprises are implementing, how can organizations maintain and extend their digital capacity?



Don't let the word digital trip you up. Despite this word's technological context, digital capacity is a human capacity, not a technological one.

Don't let the word *digital* trip you up. Despite this word's technological context, digital capacity is a *human* capacity, not a technological one.

Keeping this fundamental principle in mind will enable any organization to implement agentic AI successfully. Remember, we're at the very beginning of the agentic AI journey. Organizations are only now learning their way.

It's easy to become enamored of the technology. It's harder to ensure that your people have the knowledge, skills, and the will to leverage AI to drive business value.

The secret to digital capacity: don't think of agentic AI as humans serving the agents; it's always the other way around.

About the Author



Jason Bloomberg is founder and managing director of enterprise IT industry analysis firm Intellyx BV. He is a leading IT industry analyst, author, keynote speaker, and globally recognized expert on multiple disruptive trends in enterprise technology and digital transformation.

Mr. Bloomberg advises, writes, and speaks on a diverse set of topics, including artificial intelligence, cloud native computing, cybersecurity, and legacy modernization. He currently writes for SiliconANGLE, The New Stack, CIO, and other publications, and he is the author or coauthor of five books.

About Intellyx



Intellyx is the change agent analyst firm focused on customer-driven, technology-empowered enterprise transformation. Rather than categorizing vendors, our thought leadership distills insights across the rapidly evolving enterprise IT landscape; from agentic AI to quantum-safe security, cloud native development to modern mainframes. Our advisory helps you and your customers see through the hype and get beyond the fear of technology disruption to take action and realize value through change. Read and learn more at www.intellyx.com or follow us on [LinkedIn](#).

About BMC Helix



[BMC Helix](#) helps the world's most forward-thinking IT organizations turn AI into action – unlocking human potential to multiply productivity so teams can focus on the work that matters most. With an industry-leading, AI-powered, open platform, BMC Helix delivers a dynamic fleet of AI agents that augment work across enterprise IT service and operations management to anticipate needs, automate solutions, and accelerate outcomes for ServiceOps.