

BEYOND ADOPTION: THE RISE OF THE AI-NATIVE ORGANIZATION

What it actually takes for enterprises to compete and win in the age of AI

kyndryl

WIRED Consulting.

CONTENTS

BEYOND ADOPTION: THE RISE OF THE AI-NATIVE ORGANIZATION

The AI-Native Era 3

You’re Going To Have Lots More Managers—Of Agents 5

‘Perpetual Beta’ Is Your New Operating Model 10

Moats Are Getting Rebuilt 15

The Future Of The Organization 20

With thanks to:

Thomas Dohmke, former CEO, GitHub

Erik Brynjolfsson, professor at the Stanford Institute for Human-Centered AI, director of the Stanford Digital Economy Lab and co-founder of Workhelix

Andrew McAfee, co-director of the MIT Initiative on the Digital Economy, co-author of *The Second Machine Age* and co-founder of Workhelix

Jaime Teevan, chief scientist, Microsoft

Nate B Jones, prominent AI influencer

Jake Saper, general partner, Emergence Capital

Ismail Amla, global consult lead, Kyndryl

[The AI-Native Era]

The enterprise AI conversation has shifted from “show me the tech” to “show me the money”—and the money has started to show up. One recent paper from Stanford University states that around a third of organizations using AI in at least one business function have seen improved profitability. This reflects how AI strategy is maturing. Businesses have found that simply providing chatbots or copilots is a poor path to ROI, and that a better one is to pilot and scale specific use cases.

But this approach is not yet delivering a transformation. In that report, an equal percentage of organizations say AI has so far had no effect on profitability. A small percentage say it has actually worsened.

So what sets apart the businesses that are truly pulling ahead? A study of AI implementation at high-growth startups underscores a key insight: Businesses that integrate AI deeply, including redesigning production around its capabilities, can see up to 1.9 times revenues compared to businesses that simply bolt AI onto existing ways of working.

In other words, AI transformation begins with the right framing. Rather than asking the old question, “How should we use AI?”, leaders should be asking: “What would our workflows, operating model, or organization look like if we were designing from scratch today?”

Ismail Amla, global consult lead at technology services provider Kyndryl, terms this “first principles” thinking. “If you’re going to apply AI to your supply chain, for example, you don’t start with saying, ‘Which bits can I make faster?’ You start with the outcome. Then you ask, ‘If I could do anything in the world, does this supply chain even need to exist?’”

The gap between those mindsets is the difference between merely adopting AI and becoming truly AI-native. The latter is about rebuilding with AI at the core—a permanent shift that begins with a fundamental rethinking of work itself at the level of workflows, functions, and capabilities within the organization.

When it comes to unlocking new value, this is a pattern that has been replicated across every previous general-purpose technology wave. Electricity didn't transform manufacturing until factory owners reinvented the factory floor. The internet only changed retail once brands moved beyond putting their catalogs online and reimagined what a store could be.

In that sense, AI is no different. Current patterns of work and workflows are designed around people—but we are no longer the only type of intelligence in the room, and the new one doesn't think the way we do. "Right now, our metaphor for AI treats it as another human," says Jaime Teevan, Microsoft's chief scientist and technical fellow. "The really valuable new workflows are going to come from thinking about what becomes possible precisely because AI isn't human—what is different from what people already do?" She points to how humans can have a productive conversation with three or four colleagues, whereas AI can synthesize verbal information across hundreds or thousands of people simultaneously, giving a far richer picture of what an organization collectively knows than any individual or meeting ever could. Crucially, when AI's unique capabilities are channeled into improving a task, the firm-level gains must not be limited by legacy processes elsewhere in the organization.



System-level redesign might seem like a daunting prospect, but the stakes couldn't be higher. BCG's 2026 AI Radar survey of nearly 2,400 executives found that 90 percent of CEOs believe AI will redefine what success looks like in their industry by 2028.

Here, we explore three of the less obvious, but nonetheless critical, structural shifts they must confront as they go on the AI-native journey...



You're Going To Have Lots More Managers—Of Agents

Software development is the canary in the coal mine, and has lessons for us all.

1

For a preview of what AI-native work looks like, start with software engineering. For decades, coding was regarded as a difficult, specialist skill that took years to acquire and commanded a premium as a result. Then, in a matter of months, AI upended that—it can write functioning code at high velocity and the output is immediately testable. Initially, using AI for software work mostly involved autocomplete tools, but as AI became capable of self-correction and handling longer tasks without intervention, **software development became increasingly about delegating work to agents.**

One survey of more than 900 developers found that 55 percent regularly use agents. At the extreme end of that spectrum are users such as Boris Cherny, the creator of Anthropic's own coding tool Claude Code, who recently revealed that he hadn't written a line of code by hand in eight months.



Software engineering is the area where generative AI has the tightest product-market fit, and thus no sector has gone further, faster in delegating meaningful work to AI, or reinventing operations around the speed that shift brings. In particular, software engineering businesses have had to figure out what the new job of being a manager of agents actually entails. It's a skill set that may soon be relevant to an increasingly wide range of people.

Thomas Dohmke has been in the eye of that particular storm. As CEO of GitHub, he oversaw the rise of GitHub Copilot to a full agentic coding assistant. He then founded Entire, an open-source developer platform for managing AI coding agents. The rise of agents has changed what a developer's daily routine

looks like. The old flow state of “typing” (AKA writing code) has been replaced by what Dohmke describes as managing “a flow of ideas”, with agents deep in the trenches, handling everything from building new features and debugging errors to translating legacy syntax. “When OpenAI launched [the agentic coding platform] Codex and Claude Code came out, we realized we were done with AI assisting us in the old workflow,” he says. “Now we are reinventing the developer workflow with AI.”

The developer’s job today is increasingly about orchestration: directing the work, continuously, across multiple fronts, with machines setting the pace. On any given day, Dohmke runs multiple agent sessions in parallel, deciding which task to delegate next, crafting prompts and reviewing what comes back. Dohmke compares the experience to suddenly managing seven to ten direct reports who finish tasks in minutes and immediately return to the well for more. “Everyone’s realized now’s the point where the agent is the main author of the code and I’m the conductor of an orchestra of as many agents as I can handle.”

Erik Brynjolfsson, professor at the Stanford Institute for Human-Centered AI, and director of the Stanford Digital Economy Lab, sees this shift eventually playing out across every knowledge-intensive function. “There’s a lot of focus on coders and how their lives are changing, but this change will affect everyone in the organization,” he says. In late 2025 and early 2026, a number of tools ostensibly built for software developers unlocked opportunities for non-developer teams. Releases such as OpenClaw were intended as “coding agents” but users quickly realized the exact features that made them great at software development—the ability to run continuously in the background across a workspace, navigate file systems, and interact with applications—also generalized to other business tasks.

“Everyone’s realized now’s the point where the agent is the main author of the code and I’m the conductor of an orchestra of as many agents as I can handle.”

Thomas Dohmke, former CEO, Github

Brynjolfsson saw the broader potential of the technology and, along with Andy McAfee, co-founded a company called Workhelix to help large enterprises with their AI transformation. “People who weren’t coders, people who weren’t doing this stuff before, now have tools not just for coding, but for creating marketing plans, for sales,” says Brynjolfsson. “It’s changing how businesses work.” Organizations are now using agents to

manage deals, vet ideas, and compile compliance reports; some executives now have a whole staff of personal agents assisting with their daily work.

Of course, this shift comes with new trade-offs. Most work, Brynjolfsson explains, splits into three parts: framing the question, executing, and evaluating. Machines—increasingly—can handle the middle part. Humans own the two ends. This exposes an issue that every industry will recognize as AI-generated output ramps up: Creation is no longer the constraint. “We have unlimited generation,” Dohmke says. “Now the bottleneck is in deciding and reviewing.”

This turns agent management itself into a new core competency. Strong agent managers have a knack for reducing the time required to adequately review and iterate an agent’s work, making more things worth delegating to an agent in the first place. Dohmke argues that the key skill in that respect is writing instructions clearly enough that agents can run for hours unsupervised. So how do you do that? A standard prompt framework can certainly be helpful—defining the goal, the scope, the deliverables at each phase, and the final outputs—but studies show that there is no substitute for deep subject matter expertise. “You have to have the technical understanding,” says Dohmke. And once the agent is done, an effective agent manager needs to be able “to quickly surmise where things are and then provide feedback to the agent”. Another reason why subject matter expertise is invaluable.

“People who weren’t coders, people who weren’t doing this stuff before, now have tools not just for coding, but for creating marketing plans, for sales.”

Erik Brynjolfsson, professor at the Stanford Institute for Human-Centered AI, and director of the Stanford Digital Economy Lab

There is clearly, then, a very human aspect to agentic labor. With that comes a need to consider the wellbeing of the human in the middle of it all. Leaders will need to recognize that this is a new form of work that, perhaps counterintuitively, involves significant cognitive load. When work speeds up and more tasks than ever involve decisions and judgment, the intensity can be unrelenting. “Now you have a team of 10 agents that all constantly nag you for more work,” Dohmke says, “and it feels so magical that you’re literally getting hooked.” But the magic can wear off. A recent study published in Harvard Business Review of nearly 1,500 employees found that a meaningful share reported symptoms of acute cognitive fatigue linked to heavy AI use—especially when managing multiple AI systems simultaneously. Workers described mental

fog, headaches, and slower decision-making. One senior engineering manager quoted in the study said: “It was like I had a dozen browser tabs open in my head, all fighting for attention.”

There will be a question for leaders around how they address that and design work patterns accordingly. Some software companies have pivoted to four-day weeks for this very reason. But traditional tactics also have a role to play. “As in any healthy company, you have to have breaks,” says Dohmke. “And you have to be very conscious about ending your day. Then you have to get away from the computer.”

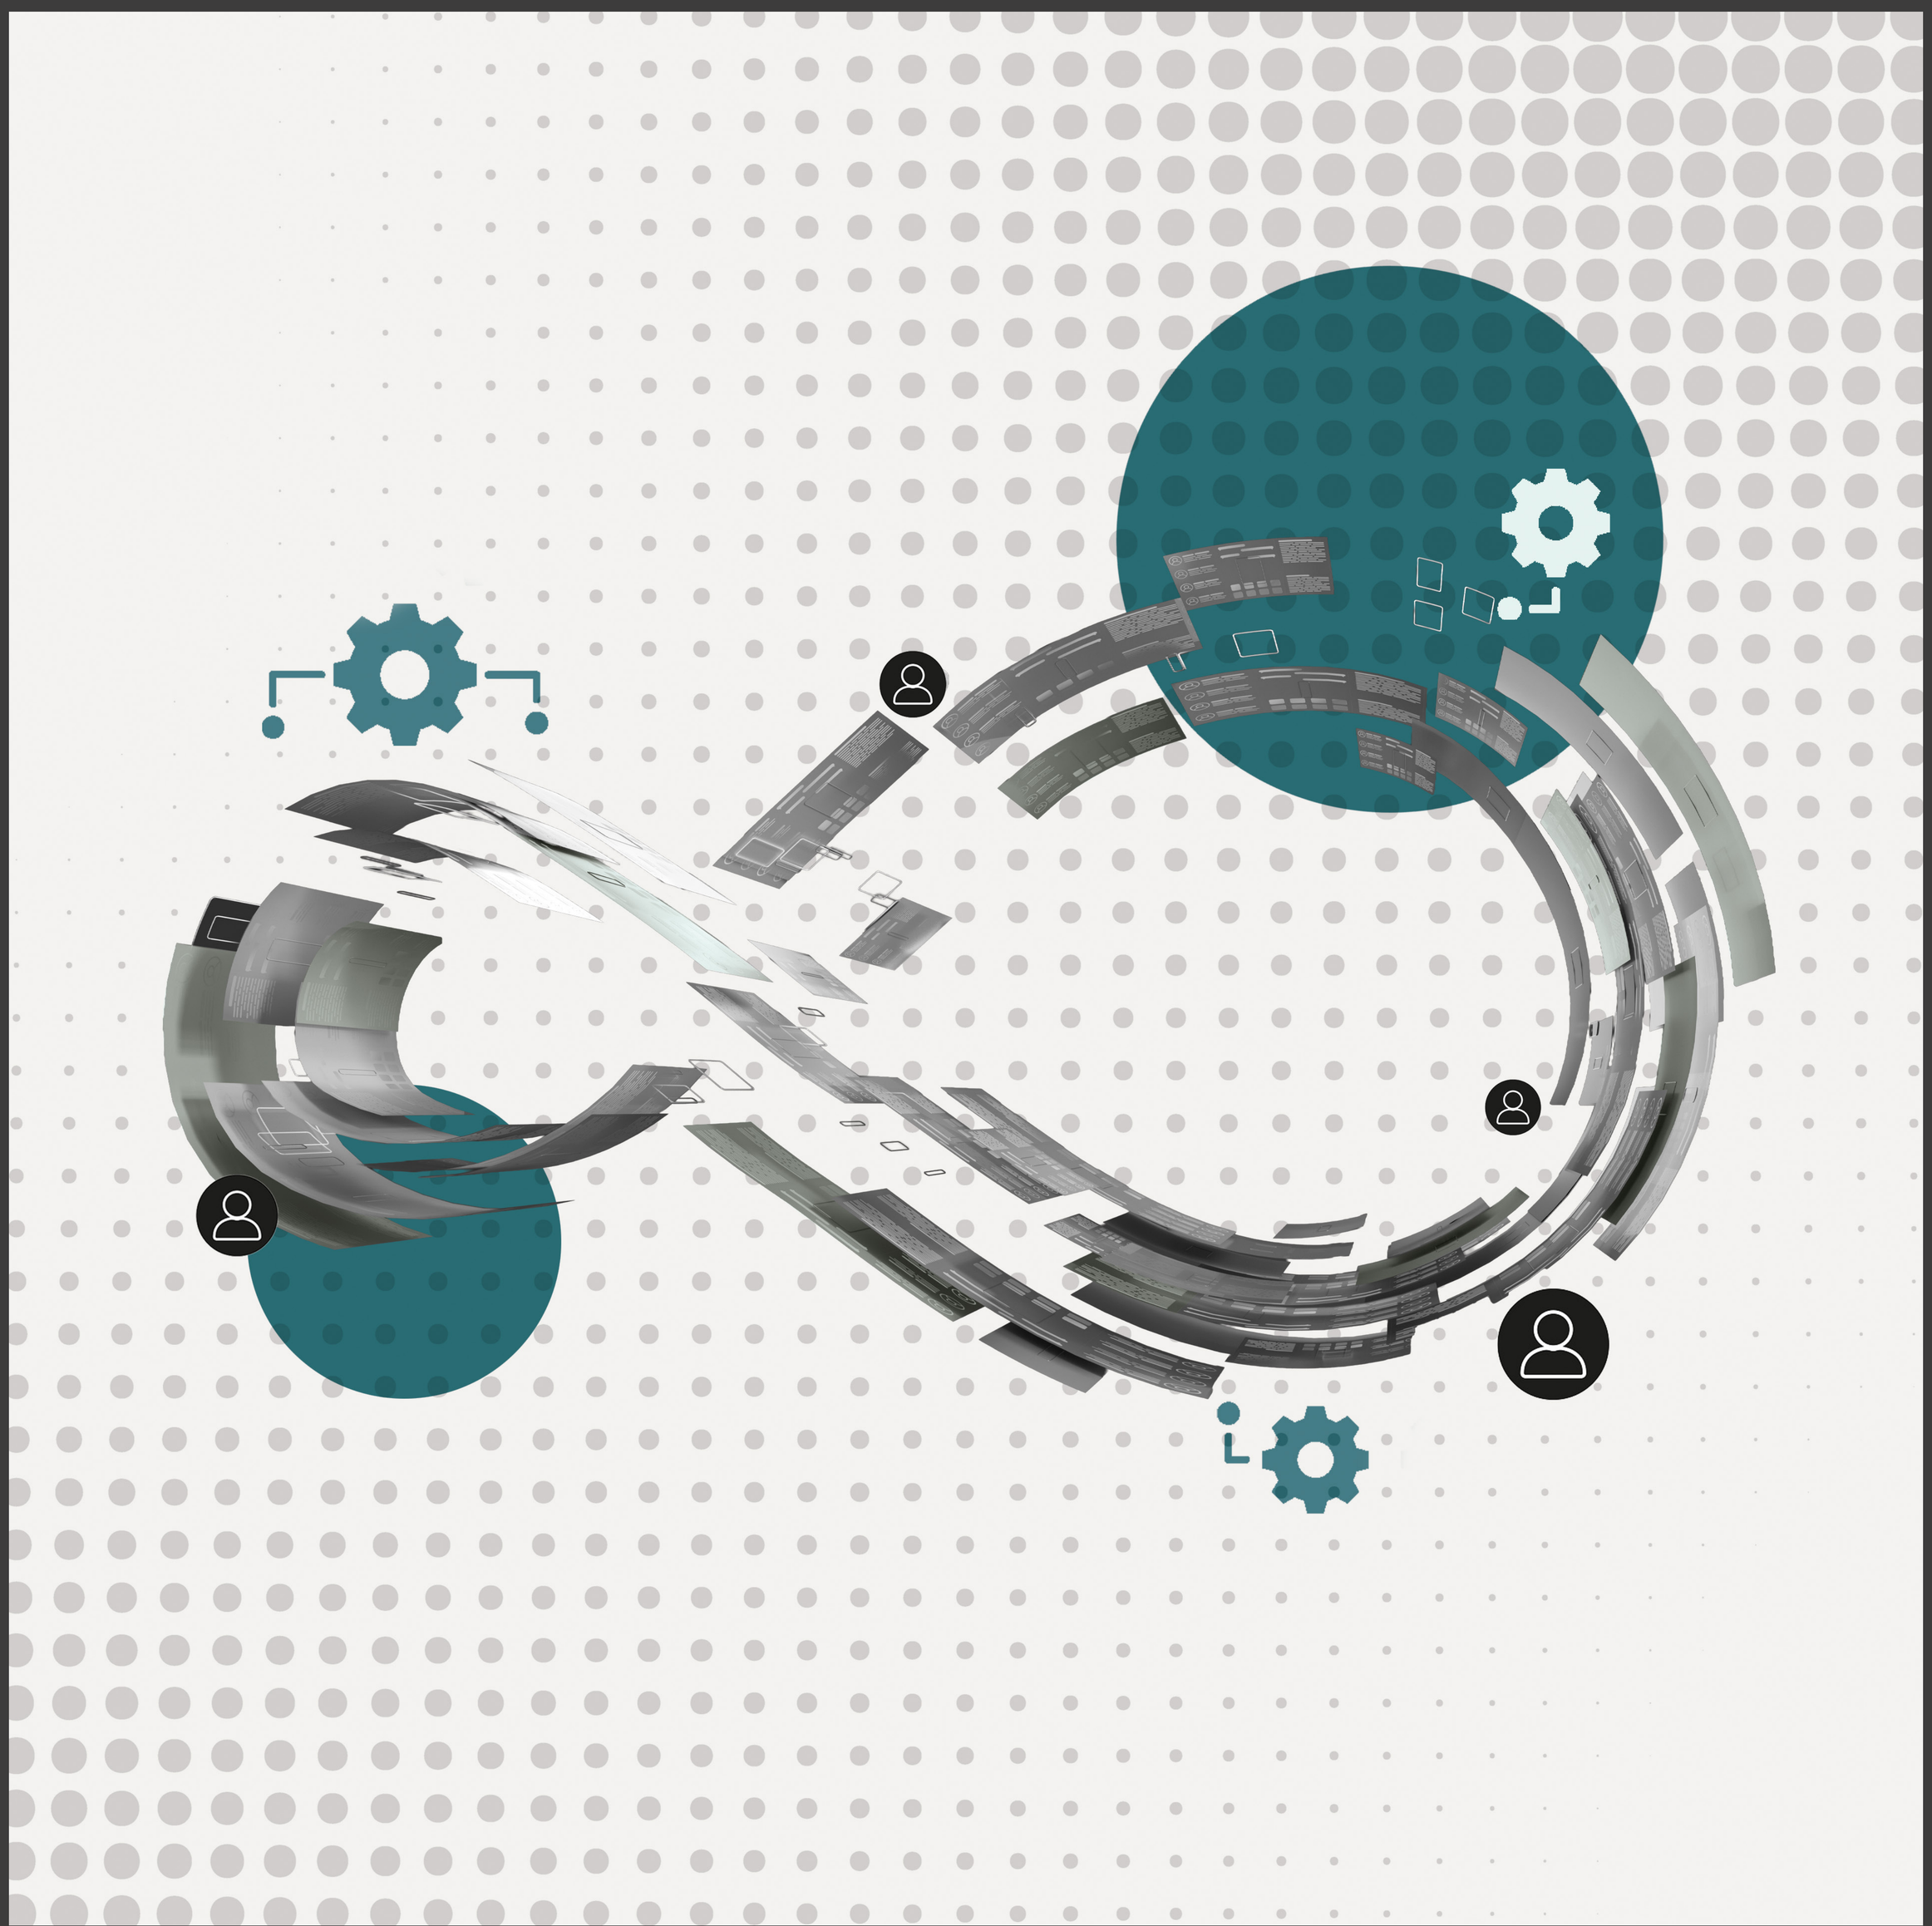
PUTTING IT INTO ACTION

1 **Decide where agents can add maximum value:** LLM-based agents are probabilistic systems. If you need a process to run deterministically, it can be better to automate this with a rule-based system or piece of code. Of course, a coding agent can help with this.

2 **Create delegation templates:** Assigning work to an AI agent is only worthwhile if it doesn’t need regular intervention or extensive iteration. This requires clear prompts and ironclad guardrails.

3 **Recognize the cognitive load:** Managing a fleet of agents can be demanding. Redesign work patterns around human capacity, and take steps to help reduce the risk of employee burnout.





`Perpetual Beta' Is Your New Operating Model

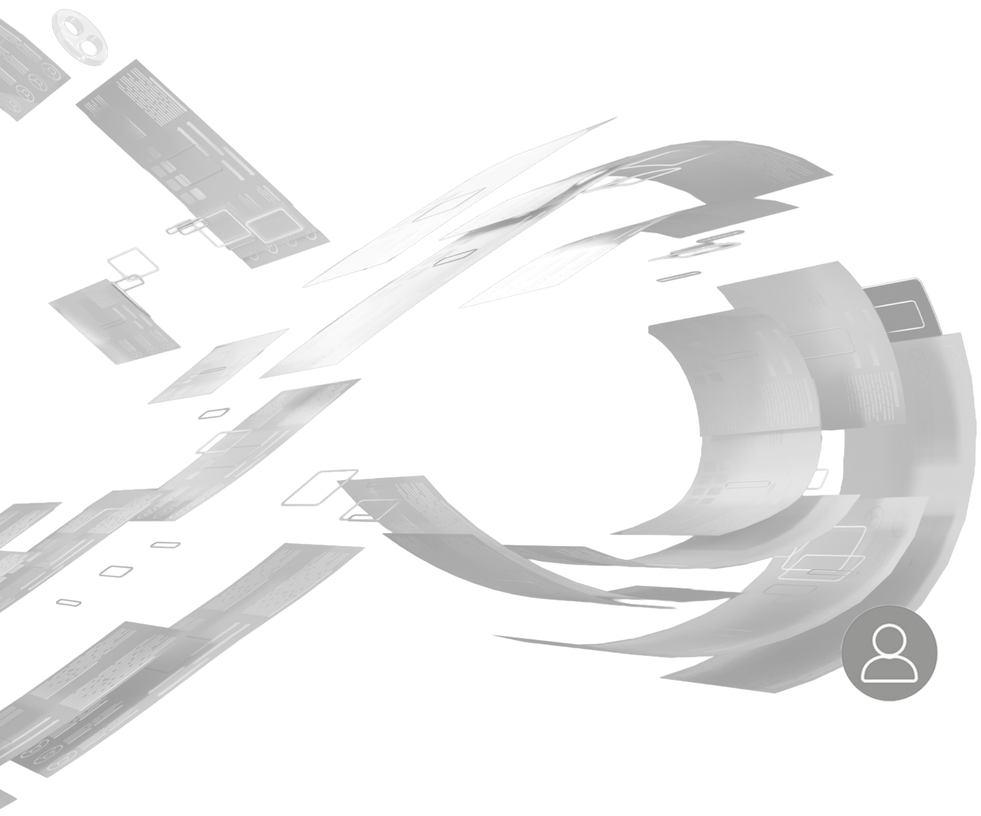
Your people—and your organization—will need to be continuously learning and adjusting.

2

Not only is AI a technology that keeps evolving, but the velocity of change only compounds. The length of tasks that AI can complete doubles roughly every seven months. Step changes in capability can arrive with apparent suddenness.

Any enterprise still running a traditional transformation playbook that treats it as a one-off project is building for an earlier era. The companies that thrive now will design continuous adaptation into the core of their operating model.

Andrew McAfee, co-director of the MIT Initiative on the Digital Economy and co-author of *The Second Machine Age*, has a mantra for the leadership teams he works with: "Plan a little and iterate a lot." The instinct in most large organizations, he argues, is the reverse. They over-index on upfront planning, analysis, anticipating the contingencies and trying to map everything out in advance. "That's an absolute ironclad guarantee of failure in a world where the technology, the building blocks, the things that you're working with are changing this quickly," McAfee says.



On the technical side, continuous adaptation requires architecture built for interchangeability. Kyndryl's Ismail Amla cautions that whatever technology you deploy today will most probably be superseded within a very short timeframe. As firms modernize and build, he advises model agnosticism, open-source foundations, and composable frameworks rather than monolithic platforms, so that customers avoid becoming bound to any single vendor or strategy. "This is the principle that underpins our whole Kyndryl Agentic AI Framework—it is by design a framework and not a platform or a piece of software, because we want the control to stay with the user," he explains.

This flexibility is not just about what's in the stack—it extends to the whole philosophy around how IT systems are designed. Take governance, which is traditionally a slow-moving domain involving roadmaps, workshops, and reporting. That isn't fit for purpose in a world where system components may need to be swapped rapidly, and agents can be spun up on-demand. In its place, organizations are exploring concepts such as "policy as code": creating machine-readable versions of rulebooks that are automatically enforced at runtime. As policies evolve, the corresponding controls can be updated so agents operate according to the latest requirements.

For Amla, however, the technical layer is only a small part of what it takes to operate in perpetual beta. "10 percent of value will be in how you use the models and the data, 20 percent in how you integrate the technology, and 70 percent in how you manage change for your people," he says. Leaders need to continuously reappraise how humans and agents work together. Employees will have to keep themselves perpetually up to speed with the latest developments and practices, and take that knowledge across the business. That means building a culture of continuous learning is paramount. After all, the LLM era has only just begun and everyone is figuring it out as it unfolds—organizations that learn and adapt quickly better position themselves to succeed.

Encouragement alone will not be enough to get this culture in place. Firms may want to explicitly hire and promote people who are good learners and have a talent for sharing what they've learned with others. They may also wish to intentionally design how knowledge flows around the organization, and make using and learning about AI an expected behavior. McAfee suggests going as far as making it an OKR: an objective measured by key results. "Without that push, you're going to go much more slowly and it's going to be very uneven," he says. So leaders need to be blunt: "AI falls into that category of a thing that is going to change the world, and without being continually up-to-date with it you are going to be in very serious trouble."

But there's a trap here. Some firms are tying bonuses and promotions to raw AI use. But usage is not the same as value; it just as easily leads to so-called "tokenmaxxing", where employees learn to game activity rather than generate results.

"10 percent of value will be in how you use the models and the data, 20 percent in how you integrate the technology, and 70 percent in how you manage change for your people."

Ismail Amla, global consult lead, Kyndryl

The metrics need to reward learning and sharing (Did you participate in an AI challenge? Did you pass an AI fluency test? Did you create a custom GPT or a skill that others are using?), and encourage good judgment about when and how to use AI, not simply that it was used at all.

AI itself introduces a novel dimension to the conversation around continuous learning. "How do you create organizations that are learning systems?" asks Microsoft's Jaime Teevan. When work is produced, the output and its impact are fresh signals. So too is every agent interaction, workflow rewire, and prompt that works or fails. This can be fed back into the system so agents learn, do better work, and drive products and services forward.

For that to happen, it's critical that agents are set up for success. They need a detailed schema of how the business operates, with easy access to all the engines that drive work around the system, from databases to strategy documentation to product files. "A document, for instance, is in some ways an artifact inferred from a bunch of knowledge and processes," says Teevan. "It's an outcome. Increasingly, in order to be a learning system as an organization, you need to think about the underlying pieces that actually produce that knowledge."

Companies are therefore placing increasing emphasis on knowledge management—not so much creating a static lake of information that can quickly get out of date, but building connections between places where knowledge resides to turn it into context that AI can use.

Teevan believes that companies that capture and compound those data points to speed up feedback, iteration, and decision-making will find that "innovation" ceases to be a risky special project and becomes inherent to the functioning of the business.

"How do you create organizations that are learning systems?"

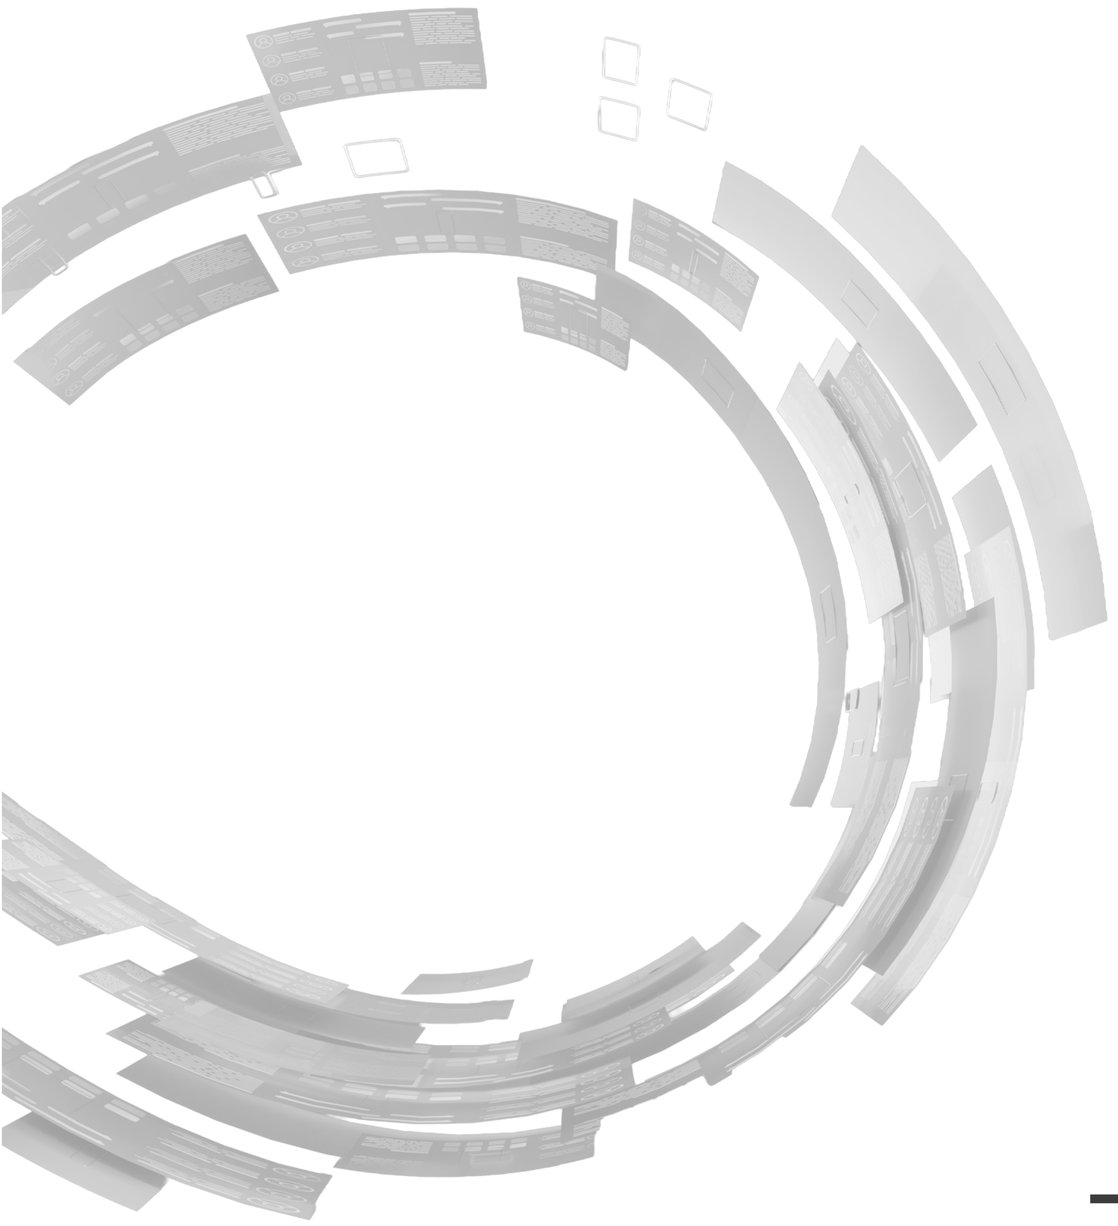
Jaime Teevan, chief scientist, Microsoft

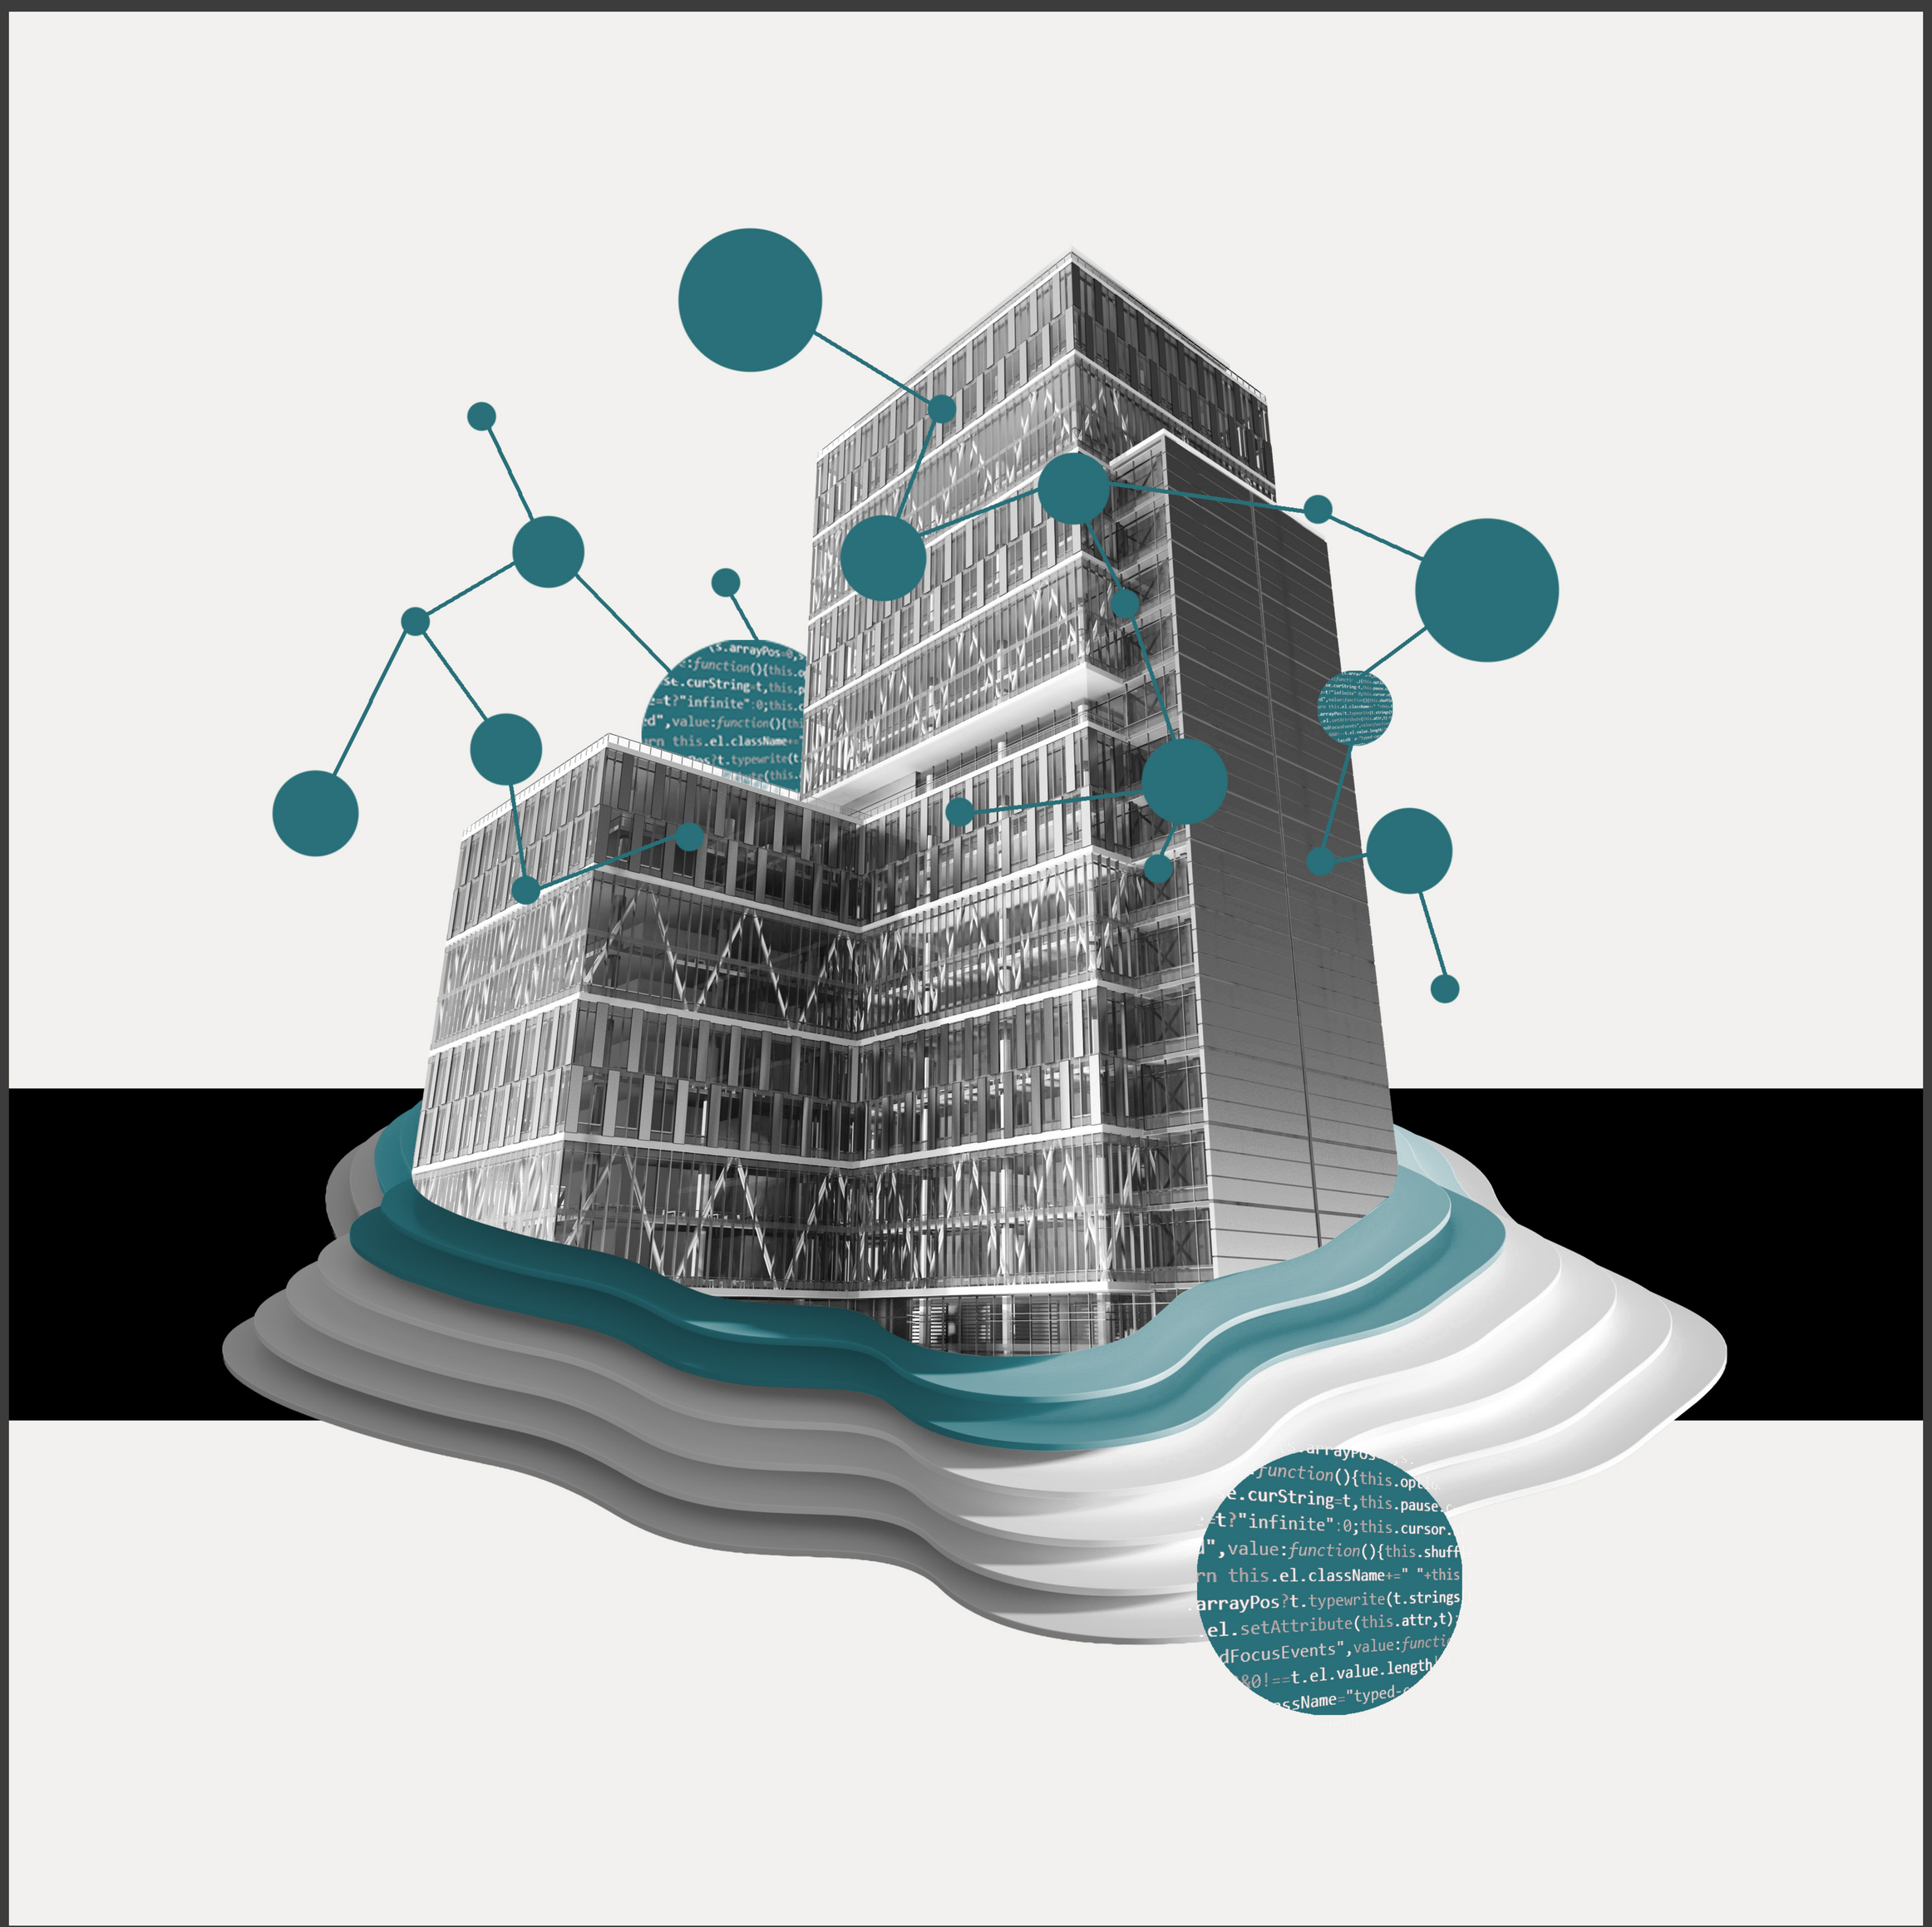
PUTTING IT INTO ACTION

1 **Build for switchouts:** Whatever technology is deployed today will likely be superseded within months. Choose frameworks over platforms, open-source over proprietary, and avoid vendor lock-in for flexibility and agility.

2 **Make AI an OKR:** But don't just measure activity. Set metrics that reward learning and encourage good decisions about when and how to use AI.

3 **Become a learning organization:** Every interaction generates signal. Build systems that feed those data points back into how agents operate, and products or services evolve.





Moats Are Getting Rebuilt

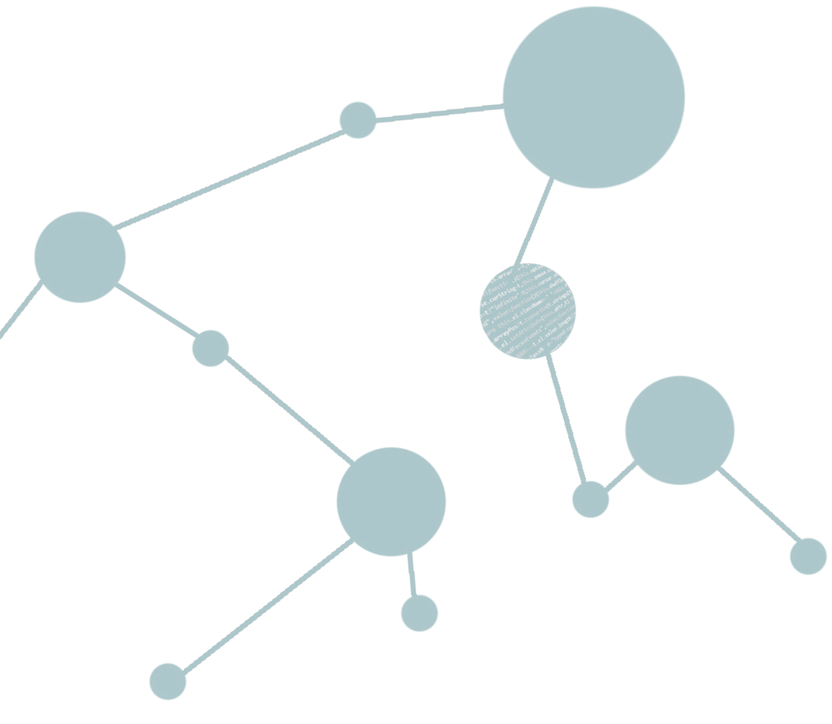
As intelligence becomes a commodity, it’s no longer an edge.

3

Internal expertise is a key area of competitive advantage. But foundation models effectively provide intelligence on tap—and any business can access them. You might expect that to level the playing field between firms but the opposite is happening. The gap between leading and lagging companies is actually widening, because the advantage is not simply having the tech but knowing what to do with it.

Erik Brynjolfsson has long been tracking this pattern. As far back as 2008, he and Andrew McAfee observed that "IT has sharpened differences among companies instead of reducing them. This reflects the fact that while companies have always varied widely in their ability to select, adopt, and exploit innovations, technology has accelerated and amplified these differences." Nearly two decades later, he sees the same dynamic playing out with AI specifically. "Right now we're seeing just such a difference in the way [organizations are] using the technology. And I think that's going to continue to grow. The AI-native firms are pulling away."

For now, then, argues Brynjolfsson, the real edge lies in "being AI native": reinventing workflows and redesigning how work gets done. This is not straightforward, which is why it confers advantage. Working out what needs to change within the enterprise, and how to achieve that, is itself a considerable undertaking. Then there's putting the vision into practice—and for a legacy firm a lot will need to change before that is possible. At a technical level, does your system have the right APIs and permissions to expose everything to an LLM in the right way? Have you built the semantic layer that will allow an agent to understand how data in one function relate to data in another? And then at a human level: How will you improve



decision-making practices so they don't restrict the very speed you want AI to unlock? In an organization built around functional silos, who runs the agents that operate across functions to optimize whole business processes? There is much to wrestle with.

But as more firms put in the hard yards of figuring out the answers, additional differentiators will count for more. This includes traditional advantages such as distribution, economies of scale, and brand. Proprietary data also becomes especially important. If everyone has access to the same AI, it's the context you provide it with that makes all the difference.

Brynjolfsson highlights one form of proprietary data as becoming especially powerful: tacit knowledge. This is the "know-how" that is essential to doing your job and which is difficult to write down or turn into rules. AI now makes it possible to automate its capture. "With machine learning, you get a bunch of inputs and outputs and it figures out the relationships even if you haven't been able to articulate them," he says.

For enterprises, systematizing tacit knowledge collection is challenging but potentially powerful: It means AI stands to capture the expertise and intuition that distinguishes top talent, and then scale that via agents.

"Right now we're seeing just such a difference in the way [organizations are] using the technology. And I think that's going to continue to grow. The AI-native firms are pulling away."

Erik Brynjolfsson, professor at the Stanford Institute for Human-Centered AI, and director of the Stanford Digital Economy Lab

The other approach to thinking about competitive advantage is through the lens of the business model: Does AI let you rethink what you're offering in the first place? Startups are leading the charge in business model experimentation—but the principles are equally applicable to incumbents who are serious about reimagining themselves for AI-native advantage.

Jake Saper, general partner at VC firm Emergence Capital, believes the dominant business model of the AI age will be "AI-native services companies". His framing is simple: Sell the fish, not the pole. The customer buys a result and doesn't care whether tasks are performed by AI, humans, or both. In accountancy, say, an AI-native firm would not sell accountancy software but would use AI—and all the new processes, efficiencies, and capabilities it brings—to close a client's books

for them with greater speed, cost competitiveness, and convenience than traditional players. He points to a range of companies across health, insurance, law, and more that are already embracing this ethos. “The arc of technology history will bend towards outcomes,” he predicts. “AI, for the first time, allows technology not to just be a tool, an implement that a human wields to get a job done. AI can be the thing that gets that job done.”

He cautions that embracing this business model is difficult because you effectively need to run a consultancy alongside your product company. “And they have to be very tightly integrated so that the learnings from one go right back to the other. It’s really hard to do.”

But say you pull it off: What are the implications for enterprise leaders? Saper says that if your customers are paying for an outcome, then the relationship fundamentally alters. Pricing shifts from per seat or per hour to “by result”. And the competitive question flips from “Is our product better than our rival’s?” to “Can we guarantee the result more reliably than the market?”

The rationale of switching to an outcome-based model is particularly compelling for SaaS companies. If you sell software, every improvement in foundation models threatens to make your product redundant. With selling outcomes, those same advances improve your service and lower your costs.

And Saper’s most radical prediction? At the limit, these companies start to resemble insurance companies. If you’re delivering outcomes and guaranteeing them financially, your core value is underwriting risk. After all, if AI makes knowledge work a commodity, what’s left?

“Doing this inside a large, complicated enterprise is genuinely difficult, but the scale of the challenge is dwarfed by the size of the prize.”

Ismail Amla, global consult lead, Kyndryl

Becoming AI-native is not easy. Getting the technology right is hard enough. Harder still is everything else: the willingness to rethink and rewire how work gets done, who performs it, and what the enterprise is ultimately for. It’s a redesign that will need to be continuous, open-ended, and as much cultural as technical.

“Doing this inside a large, complicated enterprise is genuinely difficult,” says Ismail Amla. “But the scale of the challenge is

dwarfed by the size of the prize. We work with leaders daily on this exact challenge; and candidly, the most important thing is that you start: modernizing your foundations, focusing on core workflows where AI-native redesign stands to deliver meaningful value, and recognizing that change is now a constant condition for competitive advantage."

Becoming AI-native demands leaders who can do two things at once: run the business as it is today while rebuilding it for a future that won't wait. The window for doing that from a position of strength, with existing customer relationships, distribution, and institutional knowledge intact, is open now. It may not stay open for long.



PUTTING IT INTO ACTION

1 **Take the pain:** Going AI-native is hard. But if you can do it, you'll certainly position yourself to be ahead of the competition.

2 **Capitalize on tacit knowledge:** Capture the expertise your best people carry but have never written down. Then invest in encoding it into your AI systems.

3 **Reappraise your business model:** Is there an entirely new kind of service or product you can offer as a result of AI? What was prohibitively hard before that now isn't? What might you do that you simply never imagined before?



The Future Of The Organization

Predictions on what's coming next.

4

AI changes what's possible inside organizations, giving individuals and small teams far more leverage. Here are four predictions about what's coming over the horizon from Nate Jones:

1 **Entire processes may disappear:** The back-office choreography of drafting, routing, redlining, and approving exists because humans need to coordinate with humans. Agents don't need all of those handoffs, so some processes may disappear altogether. "Most leaders haven't realized that yet," says Jones.

2 **Smaller organizational units will become more common:** In traditional companies, the building blocks might be departments or project teams. In AI-native firms, we may see more "super scouts"—one-person teams exploring ideas—and agent-equipped teams of about five: large enough to bring together the required skills and a diversity of perspectives, but small enough to keep coordination costs low. "The leader's job," says Jones, "then becomes to ask: 'Do you have the right people on those teams of five? What missions are they on? And are they equipped to go after those missions?'"

3 **AI could surface problems directly to leadership:** Today, issues generally reach the C-suite through layers of directors and VPs. But AI could help surface important problems and ideas from across the organization directly to leadership, giving them a faster way to understand what is happening on the ground. "It could mimic the ability of a great leader like Nvidia's Jensen Huang to get into the weeds with his engineers and find the problems that you really need to think about," says Jones. "That's a really exciting thought."

4 **AI agents will join board-level conversations:** Over the next few years, AI agents may begin participating in board conversations, stress-testing assumptions and bringing alternative perspectives. "You'll be talking with the CEO and his team, but there will be an agent there whose job it is to broaden the conversation."

Nate Jones advises Fortune 500 businesses on AI. He has held multiple roles at Amazon and led product across venture- and private-equity-backed companies.

Inspired by the insights shared here? Explore expert perspectives and practical resources to help your organization navigate the next era of AI.

Continue your AI-native journey:
https://www.kyndryl.com/us/en/about-us/events/wired-rethink-ai

WIRED Consulting.

WIRED is where tomorrow is realized. It is the essential source of fresh thinking and deep expertise on the technological, scientific and societal trends that are changing our world. Consulting is a division of WIRED that brings the unique WIRED network, insights and brand to commercial organizations - helping them to build internal knowledge, develop strategy and create thought-leading content that positions them at the cutting edge.

To learn more, visit: **consulting.wired.co.uk**

kyndryl®

Kyndryl runs and reimagines the mission-critical technology systems that drive advantage for the world's leading businesses. We are at the heart of progress; with proven expertise and a continuous flow of AI-powered insight, we enable smarter decisions, faster innovation, and a lasting competitive edge.

For more information, visit: **www.kyndryl.com.**