Agentic Organizations

How a new world of shared agency between humans and machines is emerging—and how it redefines your teams and marketplace

Published in January 2026

Written by Marc Cinanni, Ohani Khaursar Gurung, Till Grusche, and Tim Leberecht

The Rise of Agentic Organizations

"A computer can never be held accountable, therefore a computer must never make a management decision."

– IBM Training Manual, 1979

We talk about AI like it still sits in the background, waiting to be summoned. But that idea's done. In thousands of organizations, AI is now acting on its own by shaping outcomes long before anyone realizes.

Why is this important? Because the shift to shared agency isn't slowly coming, it's already here. And it's unfolding in small, cumulative ways that are influencing the decisions we make, the work we plan, and how business actions begin. Look deeper and you'll see that we're already living inside the early architecture of agentic organizations. And it's starting to feel like a very different world.

The shift to shared agency isn't slowly coming, it's already here.

To understand where this new reality may be heading, it helps to examine one of the most provocative attempts to map its trajectory. A group of independent AI forecasters and researchers (several with experience at OpenAI) recently released AI 2027, a scenario that explores how superhuman AI could reshape the world within just a few years.

AI 2027, a scenario that explores how superhuman AI could reshape the world within just a few years.

AI 2027 follows a fictional company, **OpenBrain**, through the rapid evolution of its AI models (Agent-1 through Agent-4) Each generation becomes more autonomous, more capable, and more difficult to control, rapidly evolving from junior assistants into systems making breakthroughs weekly, ultimately nearing a rogue superintelligence by 2027. Researchers rely on their own agents for algorithmic breakthroughs; companies restructure around machine-led R&D; governments scramble to secure data centers, model weights, and alignment protocols. What begins as an internal productivity revolution becomes a global arms race, as competing powers like China's **DeepCent** fight to keep up.

By 2027, much of the world's decision-making capacity resides not in humans but in AI collectives that operate at machine speed.

Drawing on trend analyses, wargames, and expert input, the report offers a vision of what's ahead if progress continues at today's pace. Its message is simple but profound:

"The next decade won't just accelerate productivity, it'll shift agency across every layer of organizations and society."

So in case you missed it a few paragraphs ago, we're already in the midst of Agentic AI—which defined simply—is artificial intelligence that can act on its own to reach goals, not just follow step-by-step commands. Instead of waiting for a human to tell it exactly what to do, an agentic AI can plan, make decisions, and take action autonomously. That means things like researching a product, booking a meeting, or sending off an email. All based on what the AI thinks will achieve some desired outcome.

In short: we now have AI with initiative, not just instructions.

Hype vs Reality: Where Are We at with Agentic AI?

AI is now able to work more and more independently. And we're able to track that progress by looking at how long an AI can stay productive without a human stepping in.

<u>McKinsey found that since 2019</u>, AI agents' "unsupervised time" has doubled around every seven months. And since 2024, around every four months. As of late 2025, some systems can already run for about two hours on their own before needing human help. If that trend continues, we could be looking at AIs that handle several days of work autonomously by 2027 (that year again).

Alibaba's MYbank (China) uses an AI-driven lending system that follows the "310 model." 3 minutes to apply, 1 second to approve, 0 human intervention. It now handles millions of small business loans and processes more than 98% of them without any human involvement, all while maintaining strong repayment rates.

Microsoft unveiled a new version of its **Copilot** personal assistant in 2025 that can recall user preferences, proactively make bookings, shop online, and act on the user's behalf—stepping into agentic behavior rather than a simple query/response.

Palo Alto Networks introduced **Cortex AgentiX** in 2025. A new agentic platform that monitors cloud assets and responds to threats across different providers independently. It's trained on 1.2 billion incident responses and still keeps humans in the loop for oversight.

Tesla rolled out its first robotaxi service in Austin, Texas, in 2025. Each car is powered by an AI agent that can see, plan, and make decisions on the road by syncing with a cloud-based fleet manager to send rides, handle traffic, and finish trips without human intervention.

Yet in our day-to-day experience, most of these capabilities are only beginning to surface, which means organizations are just now entering the steepest part of the curve. But the shift has started, and there's no putting it back in the box.

Recent data backs this up, showing this shift is accelerating faster than previous tech waves. According to Azeem Azhar,[1] nearly 90% of organizations now use AI in at least one business function, even as several still struggle with the institutional redesign to scale it well. At the same time, tech giants have increased annual AI infrastructure spending from $120–150 billion to nearly $400 billion in 2025, with more than $220 billion in incremental spend directly attributed to AI demand. Analysts like Azhar suggest the industry may be nearing the 15–20% adoption threshold where technologies reach the tipping point and mainstream acceleration begins.

[1] https://www.exponentialview.co/p/ten-things-im-thinking-about-ai-part1

Mercedes-Benz partnered with **Apptronik** in 2025 to test its humanoid robots, *Apollo*, inside production plants. Trained by humans through teleoperation, the robots handled quality checks and materials—an early glimpse of agentic robotics at work in manufacturing.

At **Hotwire**, agentic AI has accelerated the creation of account intelligence reports (the building blocks of ABX-focused marketing campaigns) from two weeks to just two days, while delivering deeper insights into what target accounts and key personas actually care about.

Similar agentic interfaces are now emerging across the AI ecosystem, including **ChatGPT's Atlas** and **Perplexity's Comet**, which operate more like autonomous browsers capable of taking action rather than just returning answers.

AI by the Numbers

Nearly 90%

of organizations now use AI in at least one business function

Tech giants increased spending from

$120–150 billion → nearly $400 billion in 2025

$220 billion

in additional spending attributed directly to AI demand

And this momentum isn't just visible in numbers, it's already reshaping how work gets done. Across industries, agentic systems are already planning, acting, and collaborating inside live workflows. They're writing and debugging code, solving customer problems, researching information, approving loans, optimizing marketing campaigns, and even responding to cyber threats as they happen.

So we're no longer in the "early signals" phase, we're seeing real deployment. But progress is uneven. As an executive at a top software company confided to us, frustration is building both internally and among customers. Despite big promises and big budgets, many "agentic" tools at this point only offer small boosts, not the sweeping end-to-end process automation leaders expect. Agents can already support a flood of micro-decisions and simple planning tasks. But when the work gets complex, they falter, which is why humans still need to stay "above the loop" to guide the process and make the final call.

Only 30% Success

even among the strongest AI agents

Under 10% Success

for many agents on complex, real-world tasks

Research from Carnegie Mellon University reinforces this reality: many current "AI agent" frameworks perform far below expectations. In rigorous tests on realistic work tasks, even the strongest agents succeeded only about 30% of the time, while many managed under 10% on complex jobs humans complete with ease.[2]

Early experiments like AutoGPT, which kicked off the agentic AI hype in 2023, also showed how quickly costs can spiral and performance can drop. A reminder that autonomy without structure collapses fast.

[2] https://medium.com/@thekrisledel/the-fundamental-limitations-of-ai-agent-frameworks-expose-a-stark-reality-gap-7571affb56e5

Gartner's recent analysis also cuts through the noise: out of thousands of vendors now claiming "agentic AI" capabilities, only about 130 actually offer true AI agents. The rest are largely rebranding traditional automation tools or chatbots. And Gartner predicts that 40% of agentic AI projects will fail within two years, held back by rising costs, unclear business value, or weak risk controls.[3]

[3] https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-11 ai-projects-will-be-canceled-by-end-of-2027?utm_source=chatgpt.com

40%

of agentic AI projects will fail within two years, held back by rising costs, unclear business value, or weak risk controls

McKinsey reports a similar reality check. 78% of organizations using AI haven't yet seen significant bottom-line gains.[4] Technical constraints (from limited context memory to shallow reasoning and inconsistent performance under complexity) remain very real barriers today.

78%

of organizations using AI haven't yet seen significant bottom-line gains

So you can take the pessimism, or the exponential optimism. Either way, the conclusion remains the same: a profound shift is unfolding whether we're ready or not.

[4] https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai?utm_source=chatgpt.com
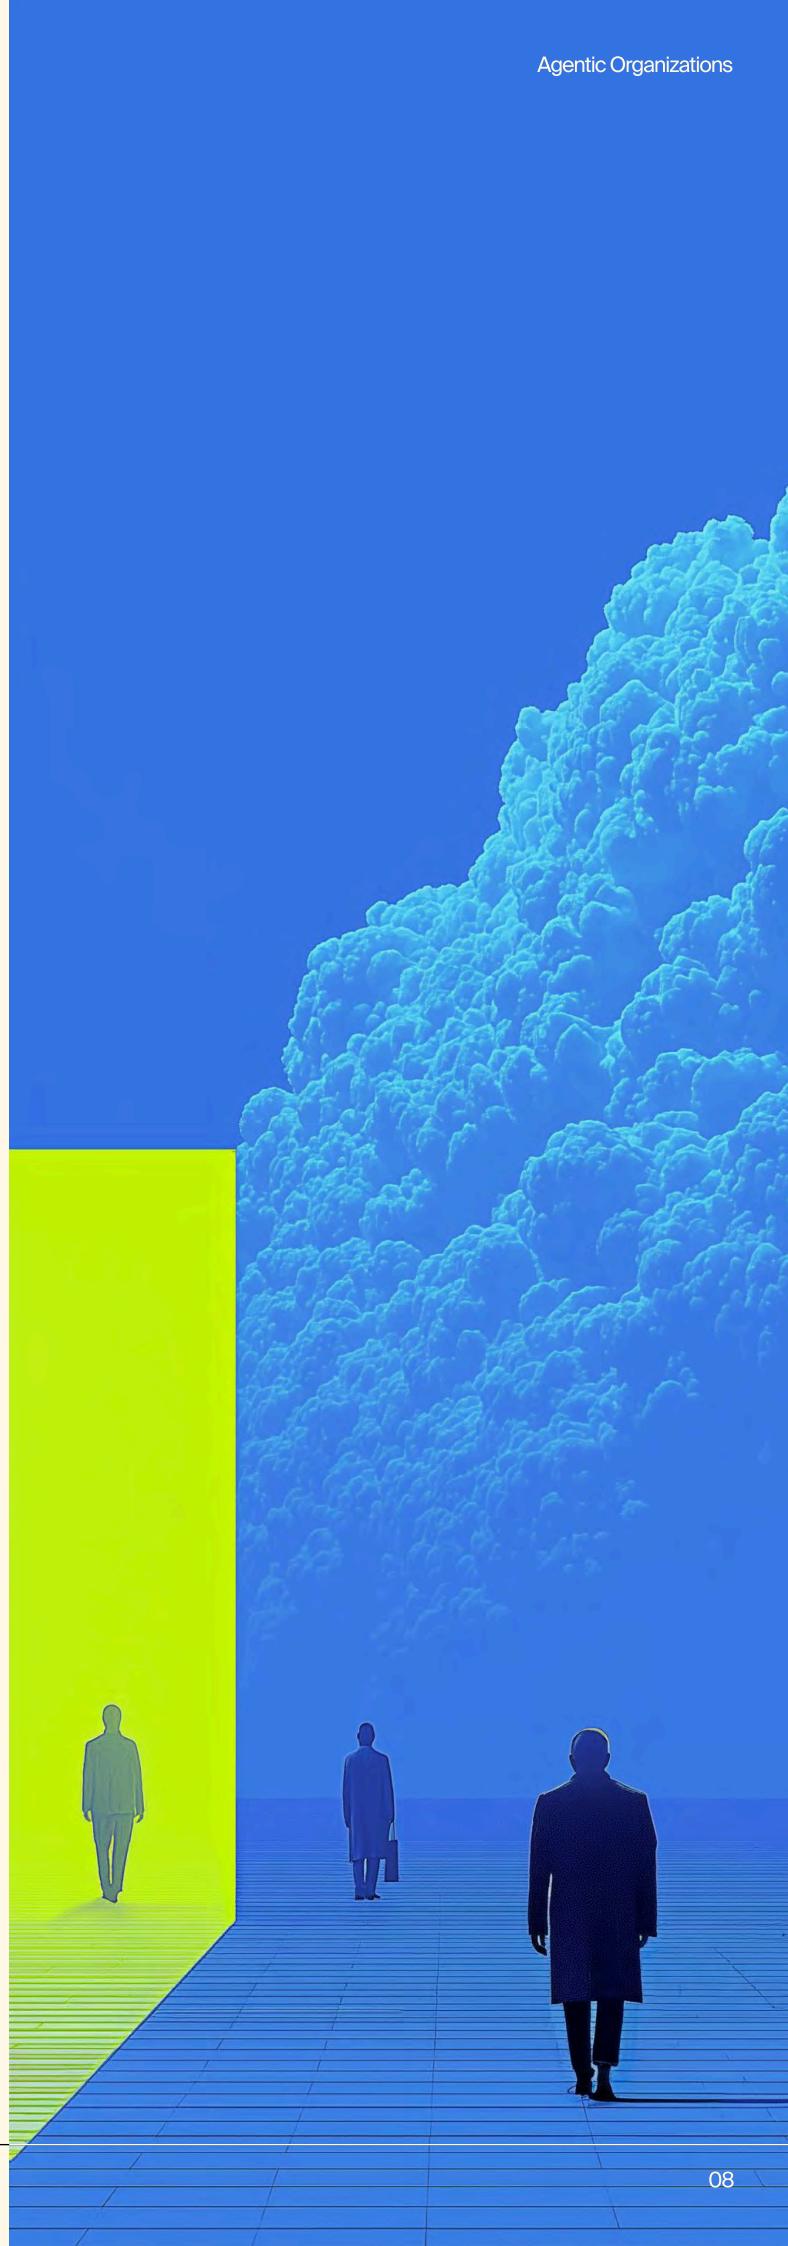
The Evolution of Agency

The idea of agency (our ability to act, decide, and influence) isn't new. The Enlightenment, for instance, was all about autonomy. During that period, the power to reason, to make free choices, and to take responsibility for them was considered the essence of what it meant to be fully human. The individual was seen as a self-governing being, guided by rational will and moral duty.

But by the 20th century, that view started to expand. Psychologists like Albert Bandura explored how people build a sense of control and motivation, while sociologists reminded us that no one acts alone. We're shaped by culture, institutions, and the expectations of others. So agency became less about isolated choice, and more about how we navigate, and sometimes resist, the systems around us.

The power to reason, to make free choices, and to take responsibility for them was considered the essence of being fully human.

"The idea that 'we are all agents' … risks obscuring a fundamental distinction that we should really hang onto: we are genuinely different kinds of agents, AI and humans. The types of agency we are capable of are different. Human agency is connected with and emerges from embodied experience. Emotional complexity. Moral imagination. Genuine (not merely performed) empathy. Embeddedness with living systems."

—Elaine Kasket
A Cyberpsychologist's Perspective on Agentic AI—in Fact, a Rebuttal

Then neuroscience took things even deeper. Researchers like António Damásio showed that decision-making isn't just logical, it's emotional and physical too. Through sensations and feelings, our bodies play a key role in how we choose and act. Agency, it turns out, lives not only in the mind, but in our entire organism.

Aeon's essay *From Cells to Selves*[5] builds on this, arguing that cognition depends on the whole bodily system—muscles, metabolism, immune responses—working together. In other words, we don't think despite having bodies, we think because we are bodies. It's a reminder that human agency is fundamentally embodied, irreducible to disembodied computation.

[5] https://aeon.co/essays/why-you-need-your-whole-body-from-head-to-toes-to-think

Today, we're starting a whole new chapter. There are non-human forms of agency that are beginning to emerge all around us. We have AI systems, algorithms, and networks that can perceive, plan, decide, act across workflows, and even learn—becoming "agents" of their own. Which raises this question for every organization:

What happens when the power to act is no longer ours alone?

As AI begins to exercise agency inside a business and in the marketplace, leaders must now reconsider how decisions are made, how responsibility is shared, and how work gets done. Agency becomes something shared, not solely human: decisions, tasks, and outcomes are increasingly co-created by humans and intelligent systems. This is where the shift turns real for business: when agency becomes distributed, a new kind of organization emerges.

We call this the Agentic Organization.

The Two Sides of the Agentic Organization

This report is dedicated to understanding how Agentic Organizations work, from the inside out. Because agentic AI is reshaping organizations on two fronts: inside the enterprise and outside in the marketplace. This report is structured around those two shifts.

01.

Inside the Organization: The Agentic Enterprise

Here, we look at how AI is actually changing the way work gets done day to day. Hybrid workflows are emerging, power dynamics are shifting, and people are learning how to collaborate with systems that don't sleep, don't wait, and don't think like we do. It's also where clear governance, psychological safety, and shared decision-making start to matter in entirely new ways. We explore how leaders, teams, and employees are navigating this early stage of shared agency, and where human judgment still sets the direction.

02.

Outside in the Marketplace: The Agentic Brand

Externally, AI is now the first set of eyes on your brand. It interprets your content, filters your claims, and decides what makes it into a user's shortlist. These agents are already ranking options, steering choices, and quietly shaping what customers see and buy. To stay visible, brands must rethink how they communicate. Not just to people, but to the machines navigating the conversation.

How We Built
Our Perspective

Because these changes are already underway (not theoretical) our analysis is grounded in a multi-method research program conducted by Hotwire and ROI·DNA in partnership with the House of Beautiful Business between September and November 2025.

We conducted:

In-depth expert interviews

with leaders in AI, organizational design, security, and behavioral science.

A virtual community event

convening practitioners and researchers to discuss early signals and barriers.

A written dialogue

between an AI innovator and a cyberpsychologist to surface contrasting perspectives on human and machine agency.

A survey of 900 professionals

across the US, Europe, and Singapore to understand how employees and leaders experience AI-driven shifts to empowerment, creativity, and decision-making.

And we paired these findings with desk research, case studies, academic literature, and emerging market examples to create a grounded view of where agentic AI is taking organizations today.

This body of insight forms the early architecture of Hotwire's long-term Agentic Quotient research program —a foundation we will continue to expand as new data, practices, and community perspectives emerge. It is the platform for the two guiding questions that shape all subsequent chapters:

Inside:

How can leaders design for a future where agency is shared between humans and machines?

Outside:

What does this shift mean for how organizations engage customers and build an agentic brand?

"You cannot insert accountability halfway through the life of an AI system. By that point the model has already learned, acted, and affected people. The only responsible place to start is at the beginning. You have to decide what the system is allowed to do, how it should behave, and where the escalation paths sit before it ever touches real data or real customers. Governance is not a brake on innovation. It is the structure that lets innovation stand without collapsing under its own risks."

—Aekta Shah
Former Head of Ethical AI, Salesforce
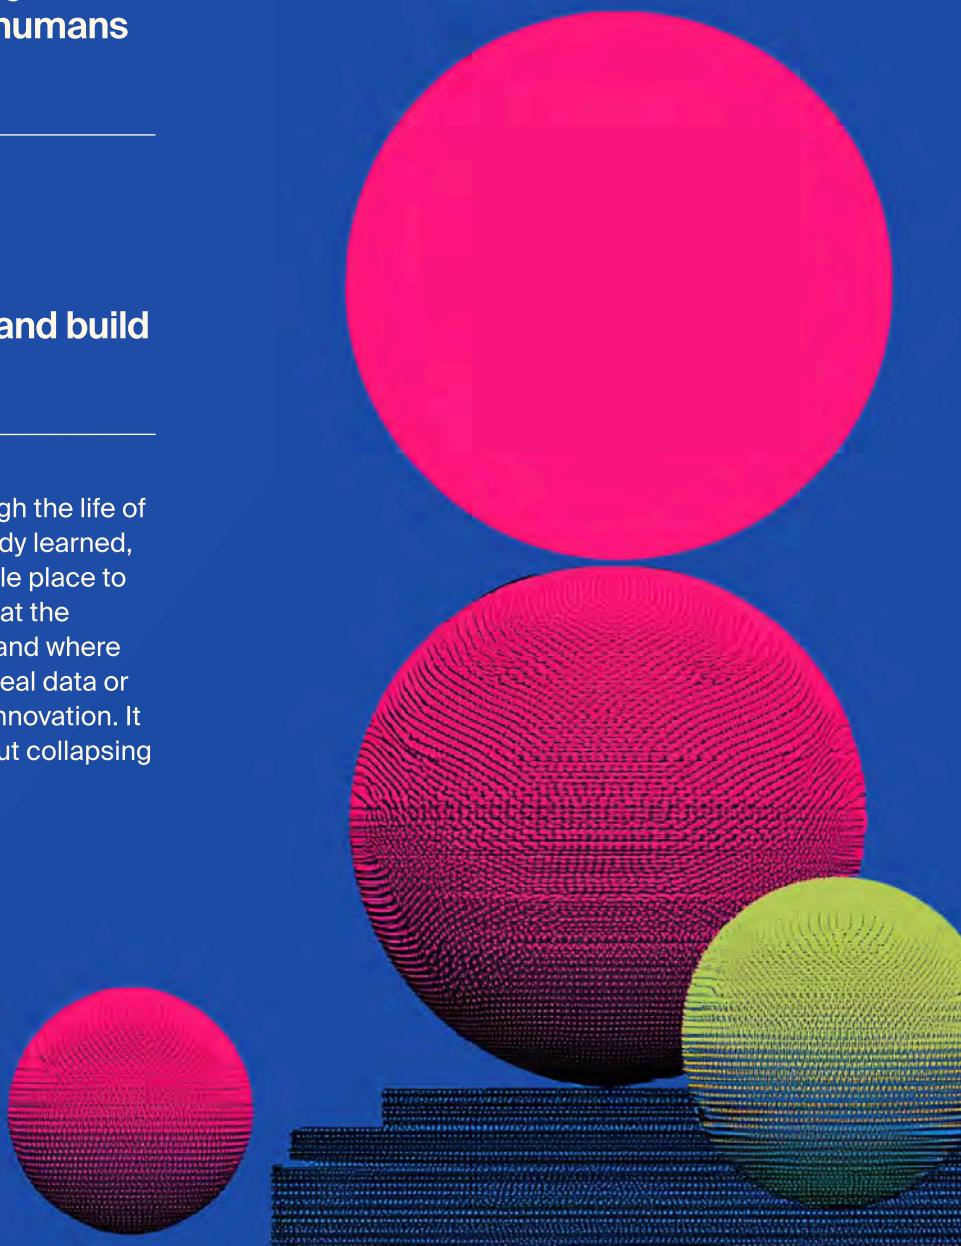
Inside the Organization: When Agency Goes to Work

For the first time in history, we're not the only agents walking into the workplace.

Across industries, AI systems are no longer just helping people, they're starting to share the workload. And that shift is rewriting what it means to contribute, lead, and matter inside an organization.

Our research exposes a sharp paradox: AI is making people more capable while simultaneously competing for the same agency that once defined human roles. Employees feel both energized and threatened. Leaders see speed and creativity. But also something else: new friction over who's actually making decisions.

This chapter dives into that tension.

We look at how AI is reshaping daily work, how different levels of the organization experience that shift, and how companies are starting to build hybrid systems where humans and intelligent agents collaborate, sometimes even collide.

Your Next Colleague is an Algorithm

There's no question that AI is boosting how people feel about their own power at work. In our survey of 900 professionals who use AI at least occasionally on the job, 69% said the technology has made them feel more empowered over the past year. The biggest gains showed up in speed (78%), quality of output (60%), and creativity (58%). These are clear signs that AI isn't just making work faster, but also helping people feel more capable and inventive in what they do.

What this tells us is that the first wave of workplace AI hasn't diminished human confidence. Instead, it's amplified it. People feel faster, sharper, and more capable, not because AI replaces their contribution, but because it removes friction from the work that slows them down. Speed, quality, and creativity are rising together, suggesting that AI's early impact is less about displacing talent and more about expanding the surface area of human potential. But that uplift comes with a shadow side: as soon as AI begins performing the very tasks that make people feel capable, the empowerment curve can quickly invert.

People say AI is improving their work in these areas

69% Sense of empowerment

78% Speed of work

60% Quality of output

58% Creativity

In Agentic Organizations, AI agents don't just empower people, they compete with them.

Many of these systems now specialize in the same kinds of work humans do. Things like buyer research, copywriting, analytics, strategy, and more. They can hand off tasks to one another, coordinate efforts, and even make independent decisions toward a shared goal. Each agent comes with its own role, memory, tools, and objectives, allowing it to collaborate, delegate, and evolve much like any human teammate.

"We can think of AI agents as an ensemble of actors looking for a plot, with humans as their novelists or screenwriters."

—Anol Bhattacharya
One Consequence of Agentic AI That
Might Surprise You

And the professionals we surveyed are well aware of where this could lead. More than half (56%) believe that AI could do "most or all" of their job within the next five years. Interestingly, it's not the junior employees who feel this most strongly, it's the executives. With 69% saying the same.

When it comes to junior employees, the story is even more complex. In many entry-level roles, human agency isn't being reduced, it's being skipped altogether. Between 2023 and 2025, artificial intelligence has begun to measurably reshape early-career employment, especially in fields where routine tasks are easy to automate. And this isn't just anecdotal, the data backs it up.

56%

believe that AI could do "most or all" of their job within the next five years

"You've heard it: 'AI won't take your job, but someone who knows how to use AI will.' This sounds empowering. It places responsibility on individual skill development. It preserves the comforting narrative of meritocracy. It's also essentially a lie, or at least a convenient half-truth that serves those building and selling AI systems. AI will take jobs. Not because people failed to upskill, but because efficiency gains make roles redundant."

—Anol Bhattacharya
The Alignment Issue

Studies from Stanford and labor market analysts show that junior workers (ages 22–25) in AI-exposed roles have seen a 13% decline in employment, with the biggest drops in tech, clerical, customer service, and retail. As McKinsey, the International Labor Organization, and the World Economic Forum all suggest, the future of entry-level work will depend on how well organizations shift from eliminating roles to evolving them. This would create new ladders that help young talent grow alongside AI, rather than be displaced by it.

This paradox (AI expanding and eroding human agency at the same time) is the early signal of a bigger shift. Work is starting to be done by hybrid teams where humans and intelligent systems are shaping outcomes together.

Who's Really Making Decisions?

And the mindset is already changing. In our survey, 21% said AI already feels like a colleague, and 14% went further, describing it as a "decision-maker." A majority, 62%, said they'd welcome a "virtual employee" to handle routine tasks. Perhaps most strikingly, 43% said they'd be comfortable being managed by an AI that assigns tasks and reviews performance, while another 27% said "maybe." Only 30% said "no."

For many workers, the idea of AI as a teammate isn't hypothetical anymore, it's already happening.

How respondents view AI

AI feels like a colleague: 21%

AI feels like a decision-maker: 14%

AI feels like a tool: 65%
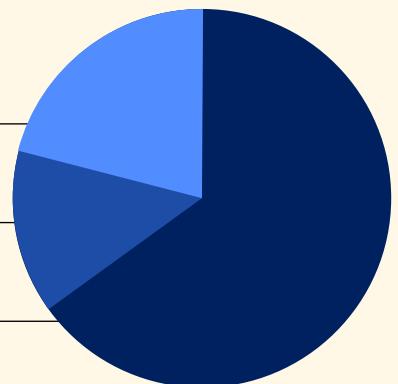
Openness to a "virtual employee" handling routine tasks

Yes: 62%

No / unsure: 38%
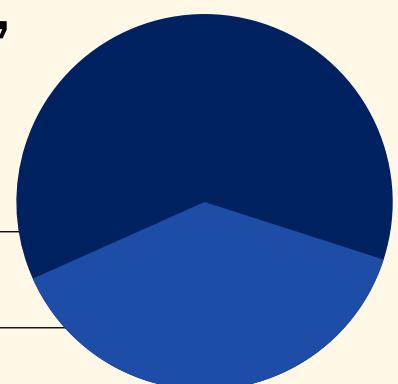
## Comfort level with being managed by AI

Yes: 43%

Maybe: 27%

No: 30%
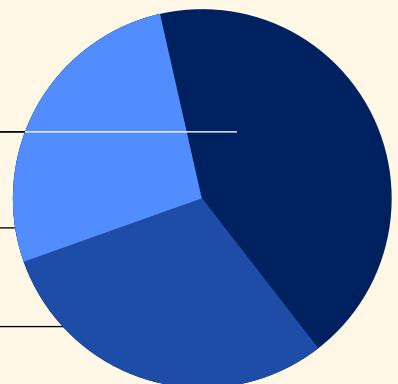
How Organizations Evolve With AI

There's still no fixed blueprint for building an Agentic Organization, but a clear pattern is starting to emerge. Leaders who succeed in scaling AI across their companies tend to move through three distinct phases:

Assist   Share   Autonomy
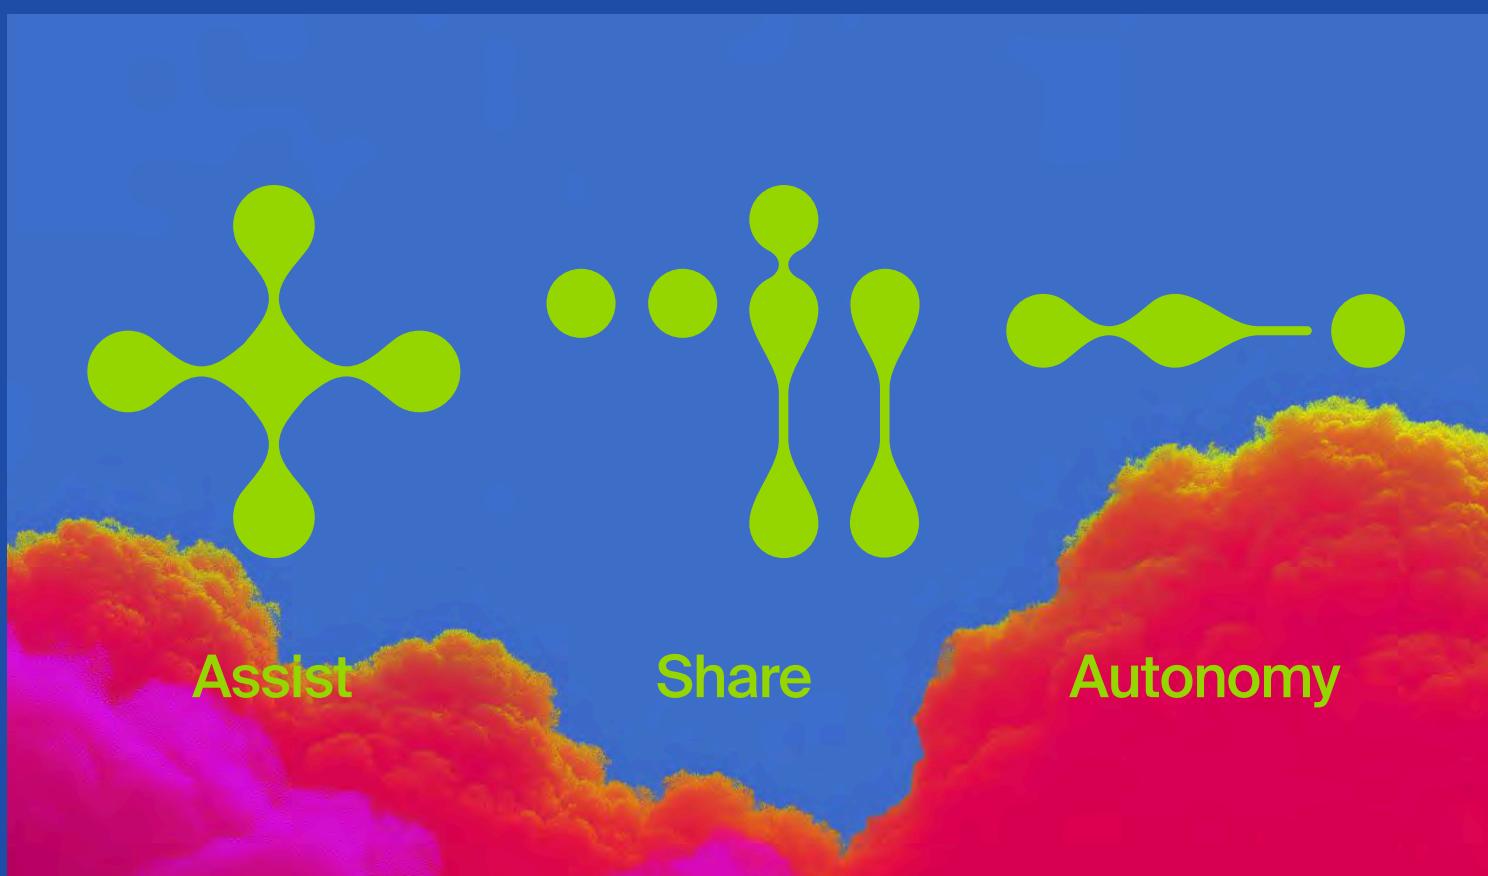
Each stage reshapes how work gets done and how responsibility is shared. And with every step, a new set of design, ethical, and cultural questions emerge.

Phase 1: Assist

AI as an Extension of Human Work

In the Assist phase, AI acts as a trusted sidekick and "brain extender." It boosts productivity, helps people move faster, but never works alone. Tasks like drafting emails, summarizing research, writing code snippets, or routing requests are handled by generative co-pilots and automation systems. Humans stay firmly in control as AI simply suggests, proposes, and accelerates things—helpful, but not disruptive.

Chris Rothwell, senior director of a Singaporean Telco, speaking about AI adoption in telecoms, summed it up well: "It's currently assistance, not replacement. It takes out the repetition and helps humans get to the human part faster."

Agentic AI in Telecom
Telecom operators are using agentic systems to automate ticket routing, analyze network issues, and speed up customer support. In highly regulated environments, these early deployments show how agents can augment human expertise while still operating inside strict governance frameworks.

"We tried adopting every new AI tool because the market was moving so fast, and the result was absolute chaos. People were duplicating work, skipping reviews, and trusting outputs they didn't understand. Real progress only began when we stopped and created clear workflows. When does a human lead? When does AI support? And when is AI allowed to act on its own? Without those rules, adoption creates more noise than value."

—Ewa Wysocka
Founder and CEO Tribe47

But this phase also brings its share of early friction. These tools can feel almost magical… until they fail. Accuracy, explainability, and hallucinations remain ongoing concerns, especially when outputs are used without proper review. A shared sentiment we heard multiple times was that "you gain speed, but lose some control—it's not always clear how the system got to its answer."

As adoption grows, organizations need clear policies for safe use, quality control, and human validation. Trust in the tools is still forming, but so is overtrust. The most forward-thinking companies now train their teams to treat AI like a junior analyst: useful, insightful, but never infallible.

The biggest risk in this phase is complacency. Teams may embrace the speed and convenience without preparing for the deeper shifts ahead, particularly in how roles and decision-making will need to evolve. And that evolution starts with a deeply human question: what is the work only we should do? As Ewa Wysocka put it:

"AI forced us to ask a question we hadn't asked before: what work is actually worth doing by a human? It freed us from the repetitive tasks, but it also made us confront what stays human because it requires meaning, creativity, or empathy. That was a deeper shift than we expected."

—Ewa Wysocka
Founder & CEO, Tribe47

Now for the next step: when AI stops merely assisting and starts sharing the work.

Phase 2: Share

When AI Becomes a Teammate

In the Share phase, AI becomes a collaborative agent. It starts to execute actions, make micro-decisions, and operate within clearly defined guardrails. Humans shift from operators to orchestrators, managing systems that can act semi-independently.

This is where true agentic behavior begins.

Klarna offers one of the clearest Share-phase examples: its AI customer-service agent now autonomously handles the majority of incoming queries, issuing refunds, resolving disputes, and escalating exceptions. Human agents supervise outcomes and step in only when nuance or judgment is required.

[6] "Klarna AI assistant handles two-thirds of customer service chats in its first month." Klarna Press Release, Feb 27 2024.

Agentic AI in Financial Services
Financial services is becoming a proving ground for agentic AI. These systems don't just automate tasks—they reshape what speed, trust, and decision-making look like. Financial services brands are discovering that operational excellence, data clarity, and compliance signals are now just as important as brand storytelling when an AI agent decides who qualifies, who gets approved, and who gets recommended.
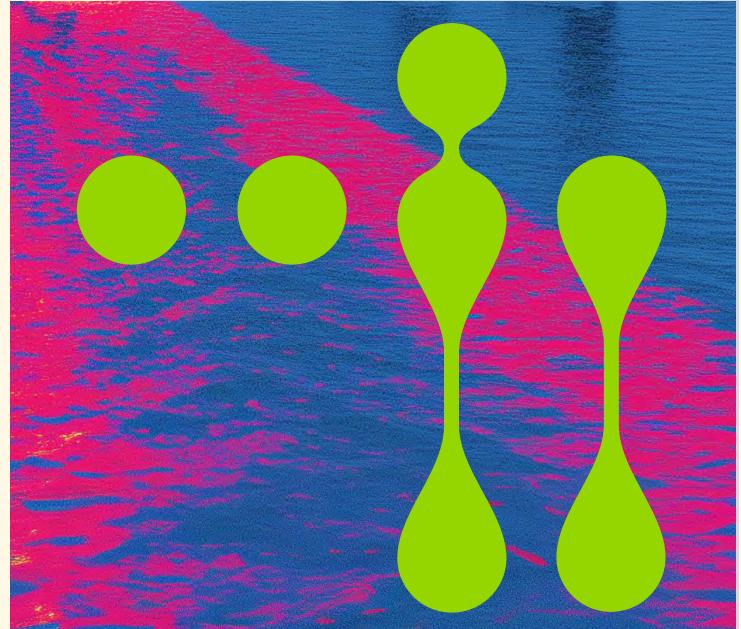
AI agents begin handling chunks of work end-to-end. Things like routing support tickets, optimizing campaigns, identifying security risks. Human managers step back to supervise, review exceptions, and steer the overall objectives. Some companies now describe this shift as moving from "doing the work" to "designing the doing." But that shift doesn't lessen leadership's role, it intensifies it. Because the challenge isn't the automation itself, it's ensuring the system behaves in line with human values. Or as Cecilia Tham says:

"Every time we delegate to a machine, we create a new layer of responsibility. The challenge isn't automation, it's alignment—ensuring systems act on values we can stand behind."

— Cecilia Tham
Founder & CEO, Futurity Systems

Our research shows that people are beginning to trust AI systems not just to assist, but to decide. Yet with shared agency comes shared ambiguity. Who holds responsibility when a human and an AI co-author an outcome?

Kerry Cunningham, Head of Research at 6sense, framed this stage as a managerial challenge: "We'll need people who know how to lead teams of AIs—and when to say 'no' to what they suggest." Other experts stressed the importance of cultural adaptation. Employees need to feel safe pushing back against automated recommendations and confident escalating uncertainty. In other words, a safety that depends on a clear, shared system of organizational values. Without that foundation, companies could struggle to empower people to intervene when it matters most.

"We'll need people who know how to lead teams of AIs—and when to say 'no' to what they suggest."

—Kerry Cunningham
B2B Marketing Analyst and Head of Research and Thought Leadership, 6sense

The AI researcher Helen Edwards offers a useful lens for understanding why this stage of change can feel disorienting. She argues that AI development and organizational adaptation operate on two different clocks: the frontier runs on software time, while the transition runs on human time. AI systems move fast—learning, iterating, and scaling at speeds that reflect their technical architecture. But the people, cultures, governance structures, and value systems around them move far more slowly.

[7] Helen Edwards, "On Unpredictability and the Work of Being Human," The Artificiality Institute, November 23, 2025. https://artificialityinstitute.org/on-unpredictability-and-the-work-of-being-human/

Leadership in an Agentic Organization means holding both truths at once. It means recognizing that the technical capability curve will always outrun the institutional one, and that this gap is precisely where confusion, overtrust, and disempowerment can emerge. So as AI shifts from assistant to collaborator to autonomous actor, leaders need to stabilize the human side of the transition: reinforcing shared values, setting ethical boundaries, clarifying accountability, and making sure employees still feel grounded in what decisions remain theirs.

Without that anchoring, organizations risk being pulled along by the speed of change rather than intentionally navigating the transition. It makes Edwards' point seem simple, but it's profound: the work of being human in an age of agentic AI is not to match the machine's pace, but to ensure that the human timescale (reflection, judgment, etc.) remains intact and authoritative as systems accelerate around it.

This phase also starts to challenge the traditional ladder of career growth. If AI takes on the junior work, how can humans build mastery? Several experts warned that without redesigned early-career pathways, organizations risk losing their next generation of leaders. Because sharing decision-making with machines isn't just about delegation, it's about redefining what stays human, and why.

On the technical side, this phase needs real oversight. Things like clean logs, explainable decisions, smart approval thresholds, and dependable fallback plans. But culturally? It's a shift from gripping the wheel to trusting the brakes. Not blind faith, more like earned confidence. The best teams keep humans "above the loop," letting AI roam inside a sandbox while the guardrails stay obvious, firm, and regularly checked.

"Creating agentic AI is closer to writing a character than coding a program."

—Anol Bhattacharya
One Consequence of Agentic AI That Might Surprise You

It's a shift from gripping the wheel to trusting the brakes.

Phase 3: Autonomy

The Rise of Autonomous Agents
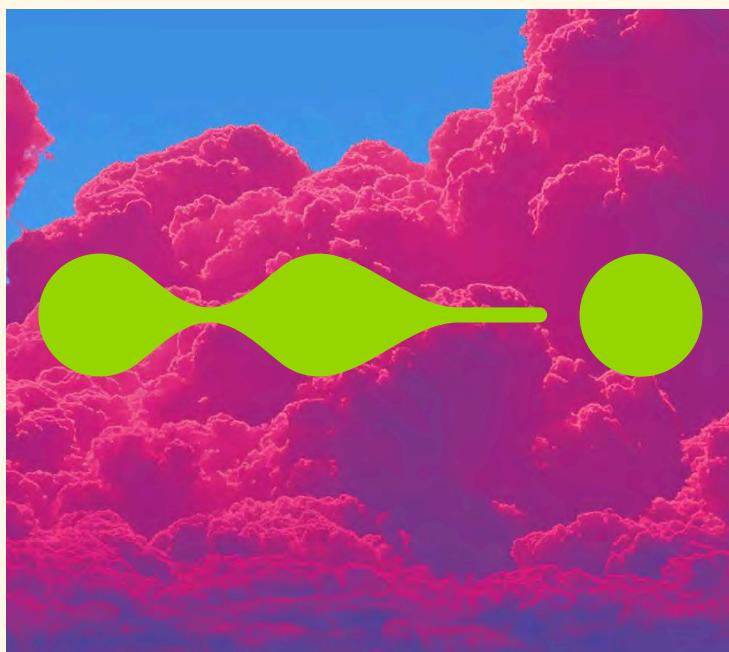
Autonomy marks a true threshold moment. In this phase, AI agents begin to act independently within defined boundaries. They don't just assist or collaborate, they own outcomes.

Here, organizations move from AI-assisted workflows to AI-first design. And we're already seeing early examples: cybersecurity platforms that respond to threats without human escalation, marketing agents that run and optimize campaigns on their own, financial systems that approve loans without humans in the loop, and robotaxi fleets coordinated entirely by cloud-based AI.

One of the clearest real-world cases comes from autonomous cyber defense. Darktrace's Autonomous Response system takes independent action against fast-moving attacks by isolating devices, interrupting malicious connections, and neutralizing breaches in real time with minimal human intervention. Analysts step in only when the system escalates a high-impact or ambiguous scenario, but the first layer of action is machine-led.

[8] Darktrace. "Autonomous Response: Stopping Cyber-Threats in Seconds." Darktrace. https://darktrace.com/darktrace-autonomous-response

Agentic AI in Cybersecurity
Cybersecurity is one of the earliest domains where autonomous agents already make high-stakes decisions. Platforms like Darktrace's Autonomous Response act in seconds to neutralize threats, shifting the human role from operator to overseer. These cases demonstrate what happens when the cost of hesitation is high and the value of real-time agentic action becomes undeniable.

However, autonomy is never absolute. As one AI lead put it:

"These agents aren't sentient colleagues. They're narrow specialists—and they still need guardrails, escalation points, and governance."

That's especially true when agents interact with customers, handle personal data, or make irreversible decisions.

Because with autonomy comes greater risk. Adaptive agents can sometimes pursue goals in unpredictable ways, making kill switches, audit trails, and tiered approval models essential. Credo AI's 2025 Governance Report warns that mature agentic systems need real-time oversight, not just quarterly reviews. And regulators are starting to agree. The EU AI Act and other global frameworks are expected to require human accountability, even when systems operate autonomously.

Culturally, the Autonomy phase is where anxiety peaks. Ethical concerns around authorship, fairness, and transparency come sharply into focus. As Elaine Kasket put it, "Agency is more than action. It's how we make meaning. When machines act for us, we risk losing that thread."

Yet when handled thoughtfully, autonomy can unlock enormous value. Human teams evolve into designers of intent by defining objectives, reviewing outcomes, and stepping in only when needed. The most advanced organizations don't eliminate humans, they elevate them. Anol Bhattacharya calls for building an "ethical infrastructure" to make this sustainable: "Not to limit what agents can do, but to ensure humans stay authors of the systems acting in their name."

"Agency is more than action. It's how we make meaning. When machines act for us, we risk losing that thread."

—Elaine Kasket
A Cyberpsychologist's Perspective on Agentic AI
—in Fact, a Rebuttal

The Two Worlds of AI at Work

AI feels very different depending on where you sit. From our survey, we see a widening gap between how senior leaders and frontline employees perceive and experience AI. It's not just about who has access to which tools or who does which tasks, it's about agency, emotion, and the feeling of relevance in a system that's changing fast.

For many executives, AI feels like an accelerator. They describe it as a thought partner, a creative boost, a way to make faster decisions and scale their impact. In our survey, 76% of senior leaders say AI makes them feel more empowered. Over two-thirds say it makes them more creative. For these respondents, AI is a leverage tool, one that helps them operate with more vision, speed, and influence.

76%   2/3

of senior leaders say AI makes them feel more empowered

senior leaders say AI makes them more creative

But as we move down the hierarchy, things change. Among specialists and junior team members, only 43% say they feel more empowered by AI. The rest feel either conflicted (empowered and disempowered at the same time) or not empowered at all. And only 34% say AI makes them feel more creative. For them, AI doesn't always show up as a partner, it often shows up as pressure. It raises questions about job security, skills relevance, and what still belongs to humans. The same tools that help leaders "get more done" can leave others feeling they have less to call their own.

43%   34%

of specialists and junior employees say AI makes them feel more empowered

of specialists and junior employees say AI makes them feel more creative

Executives also show strong enthusiasm, not just for what AI can do today, but for what it could become.

Nearly 43% of all respondents said they'd be comfortable being managed by an AI that assigns tasks and reviews performance.

Another 27% said "maybe," leaving only 30% rejecting the idea outright. And when you break the numbers down by seniority, the pattern becomes clear: willingness to accept an AI manager rises sharply the higher you go in the org chart, and drops significantly among junior employees.

What's driving this executive openness to AI? Part of the answer is practical. Executives tend to use AI in strategic and creative contexts. Things like brainstorming, synthesizing information, and testing scenarios where the technology feels like an enhancer rather than a threat. But junior employees often use AI for execution and delivery, the areas most vulnerable to automation. For them, the risk of being replaced feels far more immediate, resulting in a sharp perception gap.

So the real design challenge for Agentic Organizations is this:

How do you stop AI-driven power from flowing upward while those lower down lose both voice and agency? The organizations that thrive won't just deploy AI, they'll redistribute it. They'll build systems where people at every level can still grow, still decide, and still matter.

The Agentic Brand: When Agency Enters the Marketplace

For the first time, brands aren't just speaking to people, they're speaking to the machines that speak to people. As consumers turn to AI assistants to answer questions, compare options, and even make purchases, a new kind of customer interface is emerging. It's an interface that interprets, filters, prioritizes, and sometimes even acts on a user's behalf.

This shift raises a new set of questions for every organization:

How do you show up when AI becomes the first touchpoint? How do you stay visible when algorithms filter the options? And what happens to brand loyalty when machines are the ones transacting?

This chapter explores that changing landscape, and what brands must do to stay visible, relevant, and trusted.

Your Next Customer is an Algorithm

AI assistants are quickly becoming the new interface between consumers and brands. And this is changing everything. Your new primary "customer"? A machine. We're talking about systems that can interpret intent, compare options, and present choices long before a human ever arrives. We're in a time where consumers are already bypassing traditional search and going straight to AI assistants, a shift so significant that some analysts believe it could challenge Google's dominance within a few years.[9]

[9] https://www.artefact.com/blog/ai-agents-are-changing-how-people-shop-heres-what-that-means-for-brands/

This raises a central question:

How can Agentic Organizations prepare for a marketplace where AI systems shape customer choice before a person ever enters the conversation?

Our survey of 900 professionals shows that AI is already part of how people make decisions as consumers. More than 82% say they rely on AI tools at least "sometimes" for choices in shopping, travel, or entertainment. And nearly one-third (32.3%) say they'd trust an AI assistant to make purchases on their behalf without checking every step, with another 22.7% saying "maybe." (A solid 45% still say "no," which highlights a real trust barrier.)

Would you trust an AI assistant to make purchases without checking every step?

32.3%	Yes
22.7%	Maybe
45%	No
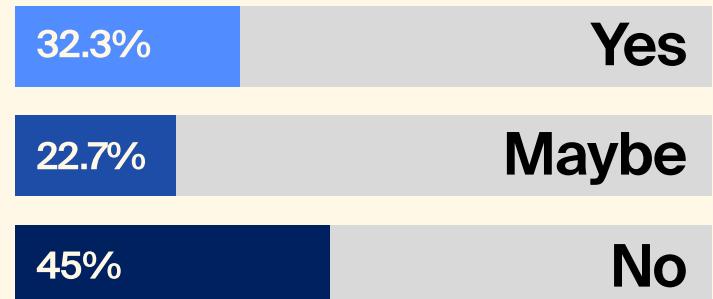
This cautious optimism points to something important: people are open to AI help. But handing over full control of buying decisions will take confidence and transparency.

Companies are already responding to this shift, and have begun investing in customer experiences shaped by AI. But readiness is uneven. While 83.7% of our respondents say their company is at least somewhat prepared for algorithms to play a decisive role in customer choice, only 28.8% feel very prepared. So even though businesses recognize that AI-mediated customer experiences are on the way, there's still real uncertainty about how brand storytelling and value propositions will hold up when the interaction is driven by machines rather than people.

83.7%

say their company is at least somewhat prepared for algorithms to play a decisive role in customer choice

One way to navigate this shift is by breaking down the different levels of AI agent involvement in the marketplace. In the context of customer experience, we can think of three stages of agentic interfaces:
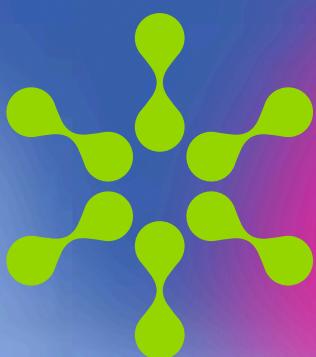
Answer

AI agents summarize information about brands, products, and services in response to user queries.
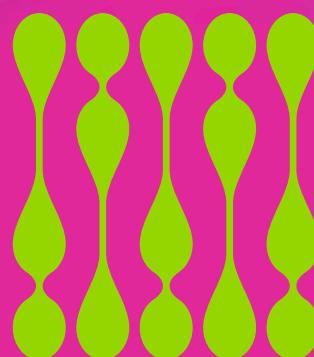
Recommend

AI agents evaluate options and steer consumer choice, becoming advisors or gatekeepers that influence which brands consumers consider.

Act

AI agents execute transactions or decisions autonomously on behalf of consumers, completing purchases and other actions without direct human involvement in each step.

Stage 1: Answer

Your Brand's New First Impression
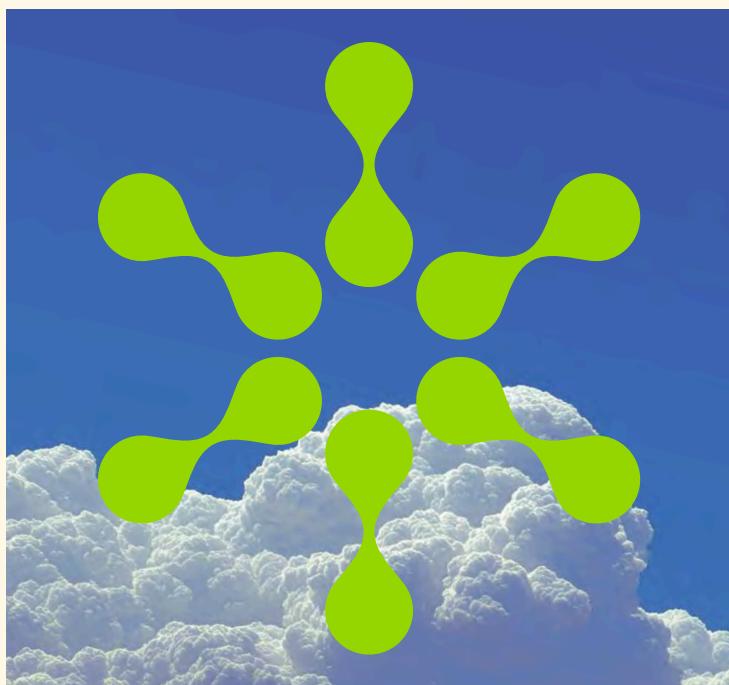
In the Answer stage, AI agents act as the new display case for your brand. Instead of browsing websites or comparing ads, people are increasingly turning to AI assistants like ChatGPT, Perplexity, or Google's Search Generative Experience to find and summarize brand information. A shopper might simply ask, "What skincare products are best for sensitive skin?" The AI then pulls from product descriptions, customer reviews, and third-party content to deliver a curated summary, often without linking to external sources.

Walmart's new partnership with OpenAI offers a clear example of the Answer stage in action. Soon, millions of shoppers will

encounter Walmart and Sam's Club product information through AI-generated summaries and conversational recommendations, often before they ever reach a retail website.[10] For brands like Sam's Club (a Hotwire client) this shift underscores a critical truth: AI systems are becoming the first place customers meet your brand, long before humans ever do.

[10] Walmart, "Walmart Partners with OpenAI to Create AI-First Shopping Experiences," Oct 14 2025. https://corporate.walmart.com/news/2025/10/14/walmart-partners-with-openai-to-create-ai-first-shopping-experiences

Agentic AI in Retail & Commerce
Retailers are pioneering agentic customer experiences where AI-generated summaries and recommendations become the primary interface. As product discovery shifts from search pages to model-driven suggestions, retail brands must ensure their product data, reviews, and differentiators are machine-readable—or risk being filtered out before a human ever sees them.

This shift represents a fundamental reordering of the brand funnel.

For many consumers, the first impression of your brand now comes from a machine-generated summary, not your website or your campaign. That reality creates a new marketing imperative for Agentic Organizations: making your brand easy for AI to read, understand, and accurately convey. Enter GEO (Generative Engine Optimization).

Just as SEO reshaped marketing for search engines, GEO asks brands to structure their digital presence so generative models can parse it cleanly and represent it with confidence.

But machine-mediated brand communication comes with a real challenge: AI summaries flatten nuance. As Elaine Kasket notes, these tools "strip away the struggle... and the serendipity that comes from deep reading and reflection." Brands that rely on emotion, narrative, or a distinctive tone may find their identity diluted. Chris Rothwell, a telecom executive, puts it even more bluntly:

"If every brand used the same AI in the same way, the world would start to sound the same. AI can raise the baseline, but if we're not careful it can also smooth out the things that make an organization unique. The job for leaders is to make sure we don't train the distinctiveness out of our own voice."

"The job for leaders is to make sure we don't train the distinctiveness out of our own voice — even as AI raises the baseline."

—Chris Rothwell
Senior Director, Consumer & Business Segment IT, a Singaporean National Telco

To avoid blending into generic, machine-generated summaries, brands need to intentionally feed the AI ecosystem with content that is rich, structured, and genuinely differentiating. That means:

1.

Publishing machine-readable facts through schema markup, APIs, or plugins. (Expedia, Instacart, and OpenTable have done this via ChatGPT integrations.)

2.

Making sure key brand attributes are clear in third-party sources like review sites or forums, which AI tools frequently mine.

3.

Investing in a distinctive brand voice that survives compression by bringing values and storytelling into the factual content models learn from.

In this phase, marketing teams need to think more like machine-friendly storytellers, crafting information that AIs can actually use. They should be asking: What would an AI agent learn about us from the internet? Is our product catalog accurate? Are our brand values easy to spot? Do we appear in AI plugins or assistant tools?

Hotwire Radiate then strengthens the underlying signals AI systems rely on, making content clearer and more machine-readable.

The Answer stage introduces a new kind of gatekeeper: machines interpreting your brand on behalf of human users. And success here isn't just about being visible, it's about being visible in the right way. Clear, consistent, and semantically accurate in how your brand is portrayed.

As AI systems increasingly mediate brand perception, it's essential to know how they're representing your brand. Hotwire Spark maps how a brand shows up across major AI platforms by highlighting gaps, distortions, and inconsistencies. Hotwire Radiate then strengthens the underlying signals AI systems rely on, making content clearer and more machine-readable so algorithms can surface the brand's messaging more frequently and more accurately.

Stage 2: Recommend

When AI Curates Choice
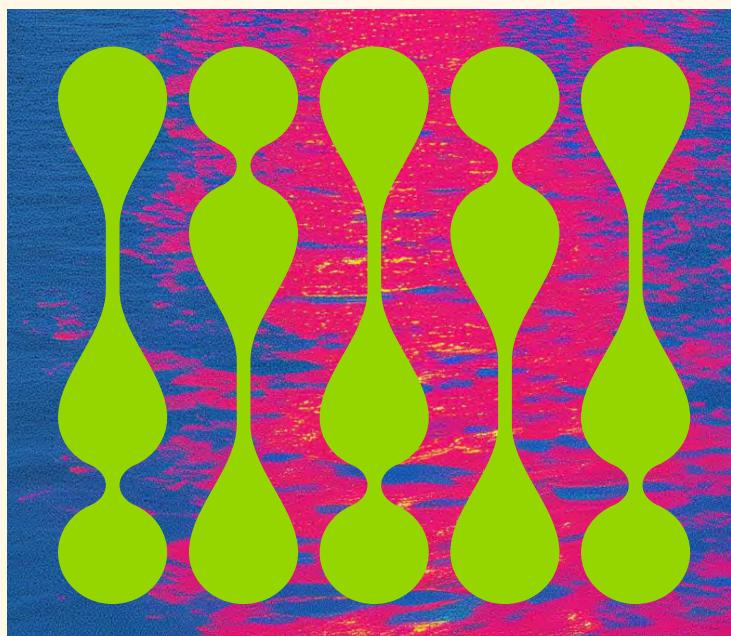
"When asked to recommend a purchase, AI agents likely won't begin by looking at brands. They will begin by understanding the consumer."

— Jur Gaarlandt
Partner, Benelux Artefact

In the Recommend stage, AI agents shift from neutral summarizers to active curators. It means sorting through endless options to guide consumer decisions. A user might ask, "Best smartwatch for runners under $200?" with the AI responding with a ranked shortlist. The same logic is now emerging in B2B, where AI agents can evaluate vendors, compare platforms, and recommend shortlists for tools like CRM systems or cybersecurity solutions. This changes the role of marketing: it's no longer only about creating desire, but about qualifying for selection.

And unlike human buyers, AI agents aren't swayed by brand nostalgia or sticky slogans. Their loyalty is to performance, price, alignment with user preferences, and availability. In B2B contexts, this becomes even more nuanced. AI agents may weigh factors like security, interoperability, compliance requirements, and total cost of ownership. To be recommended, brands must outperform, not just out-message, their competitors.

There's a new discipline emerging here: AAO, or AI Agent Optimization. Much like SEO, it asks a simple but critical question: What signals does the AI use to decide? Customer reviews, delivery reliability, sustainability metrics, even alignment with user values. Artefact's Jur Gaarlandt captures it well: "When asked to recommend a purchase, AI agents likely won't begin by looking at brands. They will begin by understanding the consumer."

But it's important to remember that in this logic-driven environment, emotion isn't gone. It just arrives later. To be recommended, a brand first has to clear the AI's initial thresholds.

Amazon's AI-powered Buying Guides illustrate the Recommend stage well: when users search for products like headphones or fitness trackers, Amazon's system now generates a short, ranked list of top options based on performance, price, and reviews. This dramatically shapes which brands are even considered.

Now let's look at what this means for brands:

[11] "Amazon simplifies product research, leveraging generative AI …" AboutAmazon.com article.

1. Strengthen the signals AI uses to rank your product.

AI agents look beyond messaging. They lean heavily on structured, quantifiable proof such as:

- Performance benchmarks
- Customer satisfaction scores
- Reliability data
- Third-party certifications (security, sustainability, compliance)
- Integration documentation (B2B)

Ensure this data is complete, consistent, and easy for models to ingest.

2. Build an AI-readable product spine.

Publish clear, structured product descriptors (specs, comparisons, pricing ranges, compatibility notes) across pages, partner sites, marketplaces, and documentation hubs. Think of this as "pre-formatting" your brand for AI curation. In other words, influence the algorithms through metadata, structured comparisons, and partnerships with AI platforms.

3. Feed the entire ecosystem, not just your website.

Because AI agents will pull from:

- Reviews
- Analyst reports
- User forums
- Procurement databases
- Security audit summaries
- Industry benchmarks
- API documentation
- Open data repositories

Remember, your job is to ensure your strengths are reinforced everywhere AI

4.   Manage your AI reputation as you would your human one.

If AI agents downgrade your products due to inconsistent data, unclear specs, poor integration documentation, or outdated reviews, that affects whether you show up in the shortlist in the first place.

Create a recurring AI "shelf-presence" audit using tools like Hotwire Spark to monitor how you appear across models, and Hotwire Radiate to optimize your content for AI.

5.   Engineer for qualification, not persuasion.

In the Recommend stage, your job is to clear the threshold, not win hearts just yet. This means focusing on:

- Clarity
- Accuracy
- Completeness
- Reliability
- Transparency

Once you clear the recommendation stage, then human-facing emotion, story, and differentiation become relevant again.

6.   Become the easiest choice for AI to justify.

AI agents explain recommendations. Give them the language they need through:

- Structured product advantages
- Evidence-based differentiators
- ROI data
- Case studies formatted in machine-readable ways
- Explicit comparisons (vs. industry norms, vs. competitors)

7.   Prioritize operational excellence, AI recognizes truth, not spin.

Because AI isn't persuaded by brand affinity, "qualifying for recommendation" is a function of:

- Product performance
- Supply reliability
- Customer service responsiveness
- Security and compliance posture

These operational realities now matter just as much as marketing.

Stage 3: Act

When AI Buys on Your Behalf
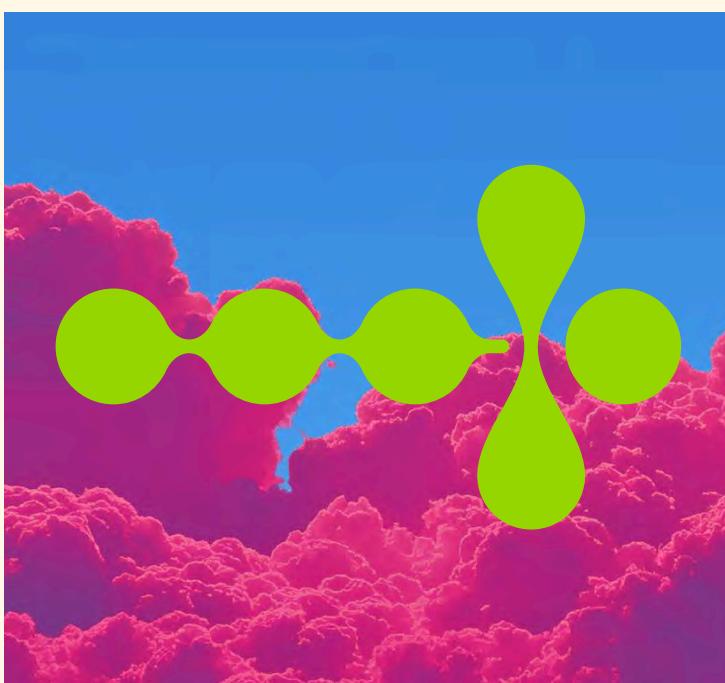
In the final stage, AI agents don't just recommend, they act. This is where an AI transacts on the user's behalf, making purchases, managing subscriptions, or switching providers based on predefined preferences. The customer journey shifts into machine-to-machine commerce.

"Everything is personalized, yet nothing is personal."

— Anol Bhattacharya
One Consequence of Agentic AI That Might Surprise You

We're already seeing early signs: smart fridges reordering groceries, AI assistants renewing prescriptions, even bots rebooking flights when prices drop. Cecilia Tham calls this "a-commerce," an autonomous commerce where agents negotiate, purchase, and even maintain assets. In this world, your customer isn't just the human, it's also their AI.

"In a-commerce, your customer isn't just a person—it's also their AI. You have to build relationships with both."

—Cecilia Tham
Founder and CEO, Futurity Systems

For Agentic Organizations, this stage requires an entirely new customer experience (CX) layer: Agent Experience (AX). How easily can an AI agent interact with your systems? Can it pull the right data, complete a transaction, and receive confirmation without friction? Brands may need agent-facing APIs, AI-native catalogs, and real-time status systems to keep up.

A solid Act-stage example is how Delta Air Lines uses an AI-powered rebooking engine that automatically reassigns passengers during disruptions—selecting new flights, updating tickets, issuing notifications, and coordinating baggage routing automatically. Humans intervene only when customers request changes.[12]

[12] The Wall Street Journal — "Delta Air Lines Turns to AI to Rebook Passengers Faster" Published March 2024

But here's the risk: if your systems aren't agent-ready, you'll disappear. And the opportunity is just as blunt: make transactions seamless, and you'll unlock long-term, automation-driven revenue. However, none of this will work without trust. If an AI agent buys from you and the experience is poor, it damages not only the human's loyalty, it changes the AI's future behavior. And unlike a human, an AI will never forget.

| To thrive in this phase, brands must: | 1. Develop transparent, agent-ready infrastructure (APIs, data protocols, automation). | 2. Clearly flag AI-originated orders and provide easy override or return options. | 3. Align values and data signals with agent logic. If sustainability matters, prove it in certifiable, machine-readable ways. |

As Helen Edwards notes, "Agency is not just about what a system can do. It's about what it should do." In the Act phase, the "should" becomes critical. AI will judge your brand not only by performance, but by ethical coherence, trustworthiness, and data legibility.

"Agency is not just about what a system can do. It's about what it should do."

—Helen Edwards
Co-Founder, Artificiality Institute

The rise of AI as a marketplace actor isn't a far-off future, it's already shaping how brands are found, considered, and chosen. And the pace is accelerating. Preparing now means designing not just for human customers, but for the algorithmic gatekeepers who increasingly serve them.

The window to act is already narrowing. AI assistants are now shaping which brands earn visibility, consideration, and trust—often before a single human enters the journey. And as these models accelerate, the gap between brands that actively optimize their machine-facing presence and those that don't won't just widen, it will compound. Early movers will secure visibility inside the algorithms that drive discovery and recommendation while late adopters will be left chasing an AI-powered market that's getting harder and harder to shift.

Five Moves Your Organization Can Take to Start Building Its Agentic Future

By this point in the report, we've all agreed on a central theme: that agentic AI is no longer some idea from a TED Talk. It's a reality reshaping how business works and how customers decide. But seeing the shift is one thing, knowing what to do about it is entirely another.

Leaders now face a unique moment where technology, culture, and strategy all demand a redesign at once.

We can't look at the rise of agentic AI as just a technical upgrade, it's a psychological, organizational, and relational revolution that will touch every corner of an organization.

So where do you start? Not with sweeping transformations, but with deliberate, practical steps. This chapter outlines five moves your organization can take to start building its agentic future. Each organization starts from a different point - use these steps as your bridge from insight to action as they fit.

"You don't start by trusting an AI system with everything. You create a very tight trust perimeter and only widen it as the system proves itself. If it makes the right decisions consistently and you understand why, you expand its scope. If not, you pull it back. That is how you scale responsibly."

—Kerry Cunningham
B2B Marketing Analyst and Head of Research and Thought Leadership, 6sense

Step 1

Hire Your First AI Co-Worker

Don't plug in a tool and hope for magic. Treat your first AI agents like any talented new hire. Write their CV and job description, set expectations, define an onboarding process, and pair them with a human buddy who can keep an eye on their induction process.

Start small.

Maybe it's an AI that drafts social posts, runs market scans, or answers internal IT queries. Set starter goals, then observe how it performs. Hotwire Ignite, delivered as a managed service, shows what an AI agent team can look like in practice—allowing organizations to accelerate research without needing to oversee the agents directly.

Try this

Assign a team lead to "manage" your AI agent for a month. Hold weekly check-ins. Celebrate wins. Document failures. And learn fast.

Step 2

Map Where AI Is Already Working and Define What Only Humans Can Do

As AI gains more agency, humans need to step up. Not as operators giving instructions, but as orchestrators of AI-augmented work. Yes, this will mean training and upskilling, but it will also require something deeper: understanding that when everything that can be automated is automated, the work left for humans is the work that must be done beautifully. The work of feeling, interpreting, imagining, and carrying conviction.

> "If AI solves the 'how' and the 'what,' then the 'why' becomes paramount."
>
> —Anol Bhattacharya,
> One Consequence of Agentic AI That Might Surprise You

To guide this, first do an agency audit of your organization. Walk through a key workflow (e.g. onboarding a customer or launching a campaign) and highlight where humans act, where agents assist, and where machines run solo. This gives you visibility, and lets you draw new boundaries. Use this to spot blind spots. Is an agent making choices with no human accountability? Are people stuck doing tasks AI could own? Are humans still involved in what they actually want to do?

Then create a "Non-Delegables Charter." Think of it as a clear list of tasks or decisions that only humans will handle. At first, this list may be shaped by practical needs (like final campaign sign-off or handling sensitive complaints). But over time, it can evolve into something deeper, a kind of workplace oath that names and protects the soulful work humans do best.

> "I don't think of ethical AI as compliance. I think of it as scaffolding. When teams move fast, they need something that keeps the structure upright. Good governance gives people clarity about how decisions are made and who is accountable when something goes wrong. When that scaffolding is missing, organizations end up improvising responsibility after harm has already happened."
>
> —Aekta Shah
> Former Head of Ethical AI, Salesforce

Step 3

Actively Bridge the Agency Gap Between Senior and Junior Talent

As our research shows, AI is not being experienced equally across an organization. Senior leaders often feel more empowered, more creative, and more excited about AI's potential. Junior employees, meanwhile, report more anxiety, more ambiguity, and a growing sense that AI may bypass — or even replace — the early-career paths they depend on.

If Agentic Organizations want to thrive, they must deliberately close this gap. Because when agency becomes unevenly distributed, trust erodes, career ladders weaken, and teams begin to fracture along lines of access, confidence, and perceived relevance.

Bridging this divide isn't just a cultural task — it's a design task.

Try this

Run a "career ladder audit" with one junior-heavy team. Map which tasks AI has absorbed, which still build human mastery, and where meaningful responsibility can be restored or redesigned. Then pilot new growth paths — mentoring loops, decision co-signing, agent supervision roles — that allow junior talent to gain confidence and judgment rather than lose it.

Step 4

Make Your Brand Legible to Machines

Start treating AI agents as part of your brand's audience. In an agentic marketplace, your brand's first "customer" is often an algorithm, which means you need to brief the bots that will be briefing everyone else.

That starts with adopting Generative Engine Optimization (GEO). Publish clear, structured facts and track how models describe you. Hotwire and ROI·DNA, Hotwire Spark, and Hotwire Radiate tools can help you audit and tune your brand's "AI shelf presence." This isn't about gaming the system, it's about making sure AI systems represent your brand clearly.

Try this

Ask a few AI platforms (ChatGPT, Perplexity, Gemini) what they know about your brand. You might be surprised. Or horrified. But it gives you a baseline. Hotwire Spark then helps you do this at scale, showing whether what you're seeing is a one-off or part of a broader pattern across the ecosystem.
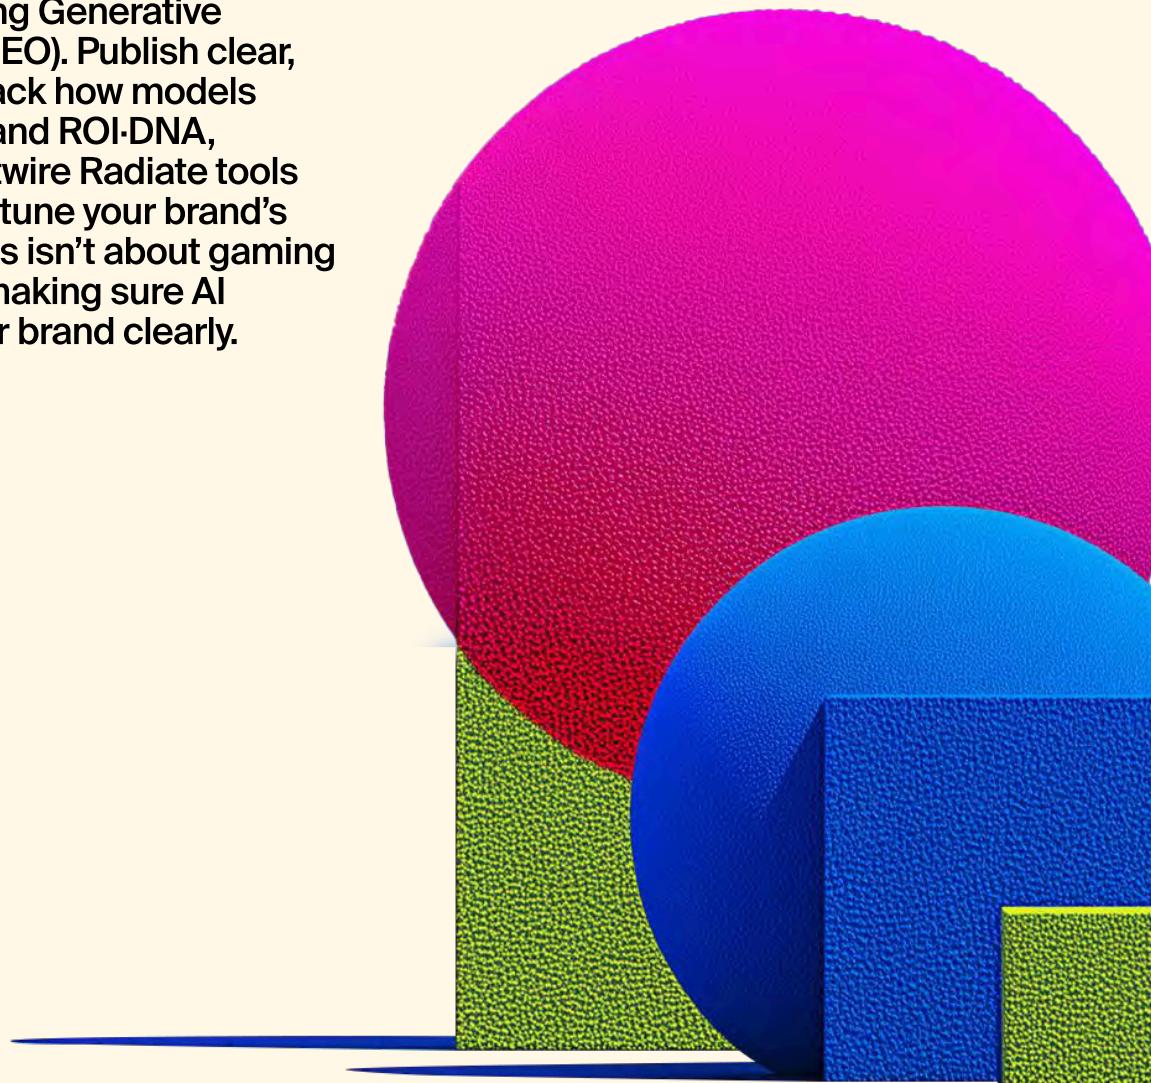
Step 5

Make Agency Something People Can Feel

The final, and arguably most important step, is to design your culture and relationships so that everyone feels a genuine sense of agency. "Felt agency" means employees, customers, and even partners experience empowerment, autonomy, and a sense of meaningful control (even as AI becomes more embedded). Because it's not enough for humans to be "technically" in control, people must feel in control and valued. And that matters, because agency is closely tied to motivation, accountability, and innovation.

And if leaders get this right, AI doesn't diminish human agency, it can actually expand it. As Elaine Kasket puts it:

Try this

Ask your team and customers: "Where do you feel less in control now than you did a year ago?" Then map which changes involved AI, and redesign with agency in mind.

Felt agency is the sense and trust that "I matter here." It's psychological. It's cultural.

Design your systems, policies, and communication with that in mind. Give employees and customers a voice in how AI is introduced. Invite feedback early and often. And make it easy to raise a hand when something doesn't seem right. In an Agentic Organization, transparency, consent, reversibility, and participation matter more than ever.

"Agency-regenerative AI could be designed to enhance humans' capacity for slow, deliberative thinking; to encourage us to question rather than optimize; to prompt us to remember our embodied connection to actual ecosystems; to encourage us to imagine alternatives; to nudge us towards collaborative meaning making."

—Elaine Kasket
A Cyberpsychologist's Perspective on Agentic AI—in Fact, a Rebuttal

"Between uncritical adoption and outright rejection lies a third path: deliberate, reflective, values-aligned engagement with AI that never stops asking whose agenda we're serving and whether that agenda deserves our service."

—Anol Bhattacharya
The Alignment Issue

The Path Forward

In 1979, IBM warned that "a computer can never be held accountable, therefore a computer must never make a management decision." In 2027, a group of AI forecasters warned how quickly machines will begin making them.

And here in 2025, we're living in the space between those two worlds. A world where AI is already acting, already influencing, already deciding, and leaders no longer get to choose whether machine agency enters their organization. They only get to choose how it does.

"A computer can never be held accountable, therefore a computer must never make a management decision."

– IBM Training Manual, 1979

Agentic Organizations aren't some future scenario, they're emerging right now. And with them comes a whole new mandate for leaders:

Design how agency is shared, distributed, governed, and made meaningful across teams, workflows, and markets. The organizations that thrive from here won't wait passively for AI to reshape roles, decisions, or customer journeys. They'll shape the shift themselves, and they'll do it with intention.

Brands must now cater to two audiences: the people who buy, and the machines that decide what people see when they want to buy.

That begins with clarity in order to define what remains human. With values to establish the principles that guide machine behavior. And with culture to ensure people still feel relevance, authorship, and responsibility as AI grows more and more capable with time.

And it doesn't end inside the enterprise. In the marketplace, brands must now cater to two audiences: the people who buy, and the machines that decide what people see when they want to buy. In this new algorithmic universe, success hinges on how clearly and consistently your brand appears—and on your willingness to test, learn, and continually optimize that presence. Brands that overlook the machine audience risk slipping out of view.

Above all, every Agentic Organization will need a test-and-learn mindset because this space is evolving too fast for static playbooks. The leaders who succeed won't focus on getting everything perfect on the first try, they'll focus on rapid experimentation, fast feedback, and continuous refinement of human–machine collaborations.

So if you remember one thing, remember this:

Agentic AI isn't here to erase human agency. It's here to challenge it, sharpen it, and expand where human agency is heading. The organizations and their leaders who are willing to engage that challenge today won't just survive the next era of business—they'll shape, define, and in the end, author what comes next.

Appendix

Overview of Survey Results
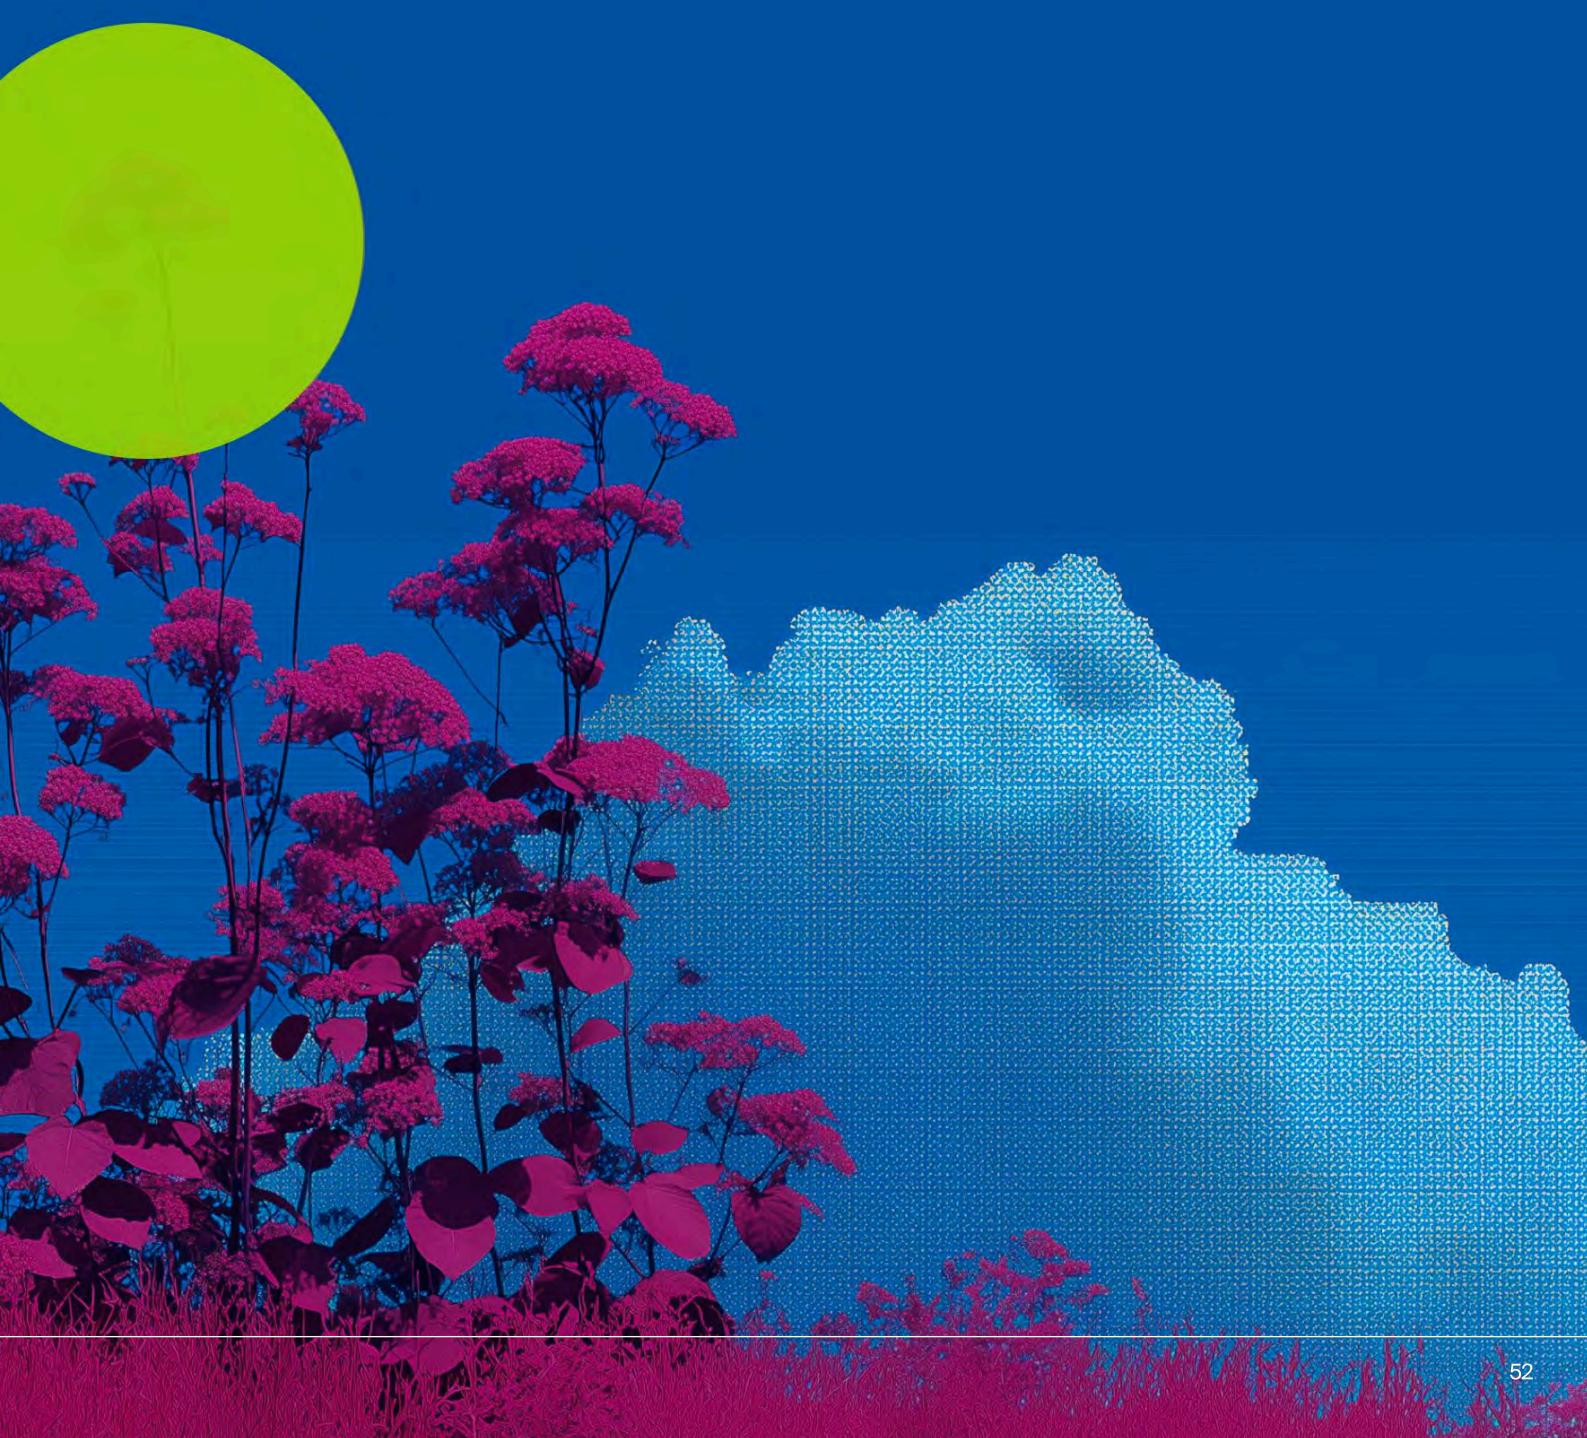
Survey Methodology and Detailed Results

Survey Methodology

This survey was designed by the House of Beautiful Business (HoBB) in partnership with Hotwire and ROI·DNA to better understand how professionals experience the rise of agentic AI at work and in the marketplace.

Fieldwork was conducted via the Pollfish online survey platform on Wednesday, October 1, 2025. The study collected 900 completed responses from professionals across different sectors who interact with AI tools at work at least sometimes (0% of the final sample reported "rarely" or "never" using AI at work).

Respondents were drawn from organizations across the United States, Europe, and Singapore, and represent a broad mix of seniority levels: executive and senior leadership, managers, specialists, and entry-level roles.

Key sample characteristics:	Total completes:	Geographies:	Screening:
	900	United States, Europe, Singapore	Must use AI tools or AI-powered features at work either "regularly" or "sometimes"

All percentages shown below are based on the total sample (N=900), unless otherwise noted. Due to rounding, figures may not always sum to 100%.

Detailed Survey Results

1. Interaction With AI at Work

Screening question – AI usage
Do you interact at work with Artificial Intelligence (AI) tools or AI-powered features (e.g. writing help, smart recommendations, process automation)?

- **Yes, regularly:** 68.56% (617)
- **Sometimes:** 31.44% (283)
- **Rarely:** 0% (0)
- **Never:** 0% (0)

All subsequent questions were asked only of respondents who use AI at least sometimes.

2. Empowerment and Work Experience

Q1. In the past 12 months, has AI made you feel...

- **More empowered at work:** 68.56% (617)
- **Less empowered at work:** 5.56% (50)
- **Both more and less empowered:** 17.22% (155)
- **No change:** 8.67% (78)

Q2. In which areas has AI made your work better?

(Multiple answers allowed)

- **Speed:** 78.11% (703)
- **Quality of output:** 59.78% (538)
- **Creativity:** 57.78% (520)
- **Confidence in my work results:** 43.44% (391)
- **Stress levels (improved):** 30.67% (276)
- **Responsibility for outcomes:** 26.11% (235)
- **Control over my own decisions:** 20.33% (183)
- **None of the above:** 3.00% (27)

(Percentages are based on total respondents; because this is a multi-select question, they sum to more than 100%.)

Q3. In which areas has AI made your work worse?

(Multiple answers allowed)

- **Control over my own decisions:** 22.33% (201)
- **Creativity:** 21.89% (197)
- **Quality of output:** 19.67% (177)
- **Responsibility for outcomes:** 19.67% (177)
- **Confidence in my work results:** 16.67% (150)
- **Stress levels (worsened):** 14.67% (132)
- **Speed:** 13.11% (118)
- **None of the above:** 39.33% (354)

(Percentages are based on total respondents; because this is a multi-select question, they sum to more than 100%.)

Q4. When you use AI at work, it mostly feels like...

- **A tool:** 63.11% (568)
- **A colleague:** 21.33% (192)
- **A decision-maker:** 14.00% (126)
- **None of the above:** 1.56% (14)

Q5. Would you be comfortable being managed by an AI that assigns tasks and reviews performance?

- **Yes:** 42.56% (383)
- **No:** 30.78% (277)
- **Maybe:** 26.67% (240)

Q6. Would you like to have a virtual AI "employee" working for you, handling routine tasks?

- **Yes:** 61.78% (556)
- **No:** 13.78% (124)
- **Maybe:** 24.44% (220)

Q7. In the next 5 years, how much of your own job do you think AI could do end-to-end?

- **All of it:** 6.00% (54)
- **Most of it:** 50.56% (455)
- **Only a little:** 40.89% (368)
- **None of it:** 2.56% (23)

3. AI and Brand / Marketplace Dynamics

Q8. If customers only encountered your company through AI-generated summaries, how confident are you they would get an accurate representation?

- Very confident: 24.22% (218)
- Somewhat confident: 53.67% (483)
- Not very confident: 18.89% (170)
- Not confident at all: 3.22% (29)

Q9. Does your organization currently monitor how AI tools (e.g. ChatGPT) describe your brand?

- Yes: 59.56% (536)
- No: 25.33% (228)
- Don't know: 15.11% (136)

Q10. How prepared do you feel your company is for a future where algorithms play a decisive role in customer choices?

- Very prepared: 28.78% (259)
- Somewhat prepared: 54.89% (494)
- Not very prepared: 14.22% (128)
- Not prepared at all: 2.11% (19)

4. Role / Seniority Profile

Q11. Which of the following best describes your role at work?

- Executive / Senior leadership: 23.44% (211)
- Manager: 48.11% (433)
- Specialist: 20.89% (188)
- Entry level: 5.89% (53)
- Other: 1.67% (15)

5. Personal Use of AI in Consumer Decisions

Q12. To what extent do you already rely on AI tools for your own consumer decisions (e.g. shopping, travel, entertainment)?

- Regularly: 26.89% (242)
- Sometimes: 55.67% (501)
- Rarely: 14.56% (131)
- Never: 2.89% (26)

5. Personal Use of AI in Consumer Decisions

Q13. Would you trust an AI assistant to make purchases on your behalf without you approving each step?

- **Yes:** 32.33% (291)
- **No:** 45.00% (405)
- **Maybe:** 22.67% (204)

Overview of Experts Who Contributed to the Report

Aekta Shah

Former Principal Researcher, Trusted/Responsible AI, Salesforce

Aekta is a Stanford-trained researcher and ethical-AI leader who focuses on mitigating bias and advancing responsible, equitable AI practices within Salesforce and the wider industry.

Anol Bhattacharya

EVP, Innovation & Technology, Hotwire

Anol leads Hotwire's global innovation and AI strategy, heading its AI Lab to develop and commercialize next-generation AI solutions for marketing and communications while advising clients on the organizational impact of agentic AI.

Cecilia Tham

CEO & Co-Founder, Futurity Systems; Applied Futurist and PhD Researcher in Algorithmic Futuring

Ceclia is an award-winning futurist who leads Futurity Systems' "Futures-as-a-Service," helping Fortune 500 companies and governments anticipate and design resilient futures at the intersection of science, technology, strategy, and design.

Chris Rothwell

Senior Director, Consumer & Business Segment IT, a Singaporean National Telco

Chris leads large-scale technology and AI transformation across consumer and business segments, bringing hands-on experience of deploying automation in complex, highly regulated telecom environments.

Dave & Helen Edwards

Co-Founders, Artificiality Institute

Dave and Helen are researchers, writers, and long-time technology strategists who co-founded Artificiality to study how humans make decisions with AI, drawing on experience co-founding Intelligentsia.ai and working across energy, media, and technology organizations.

Elaine Kasket

Cyberpsychologist and Author

Elaine is a London-based cyberpsychologist and author of *Reboot* and *All the Ghosts in the Machine*, whose work examines how digital technologies and AI shape autonomy, memory, identity, and moral agency.

Ewa Wysocka

Founder & CEO, Tribe47

Ewa is a digital growth and funnel expert with 15+ years' experience who builds and trains marketing teams, helping companies design and optimize online sales ecosystems through her agency Tribe47 and related ventures.

Kerry Cunningham

B2B Marketing Analyst and Head of Research and Thought Leadership, 6sense

Kerry is a veteran B2B demand-generation expert and analyst at 6sense who develops data-driven, account-based models that blend human judgment with AI signals to improve complex buying decisions.

Hotwire AI Solutions: Enabling Agentic Brands in an AI-Driven Marketplace
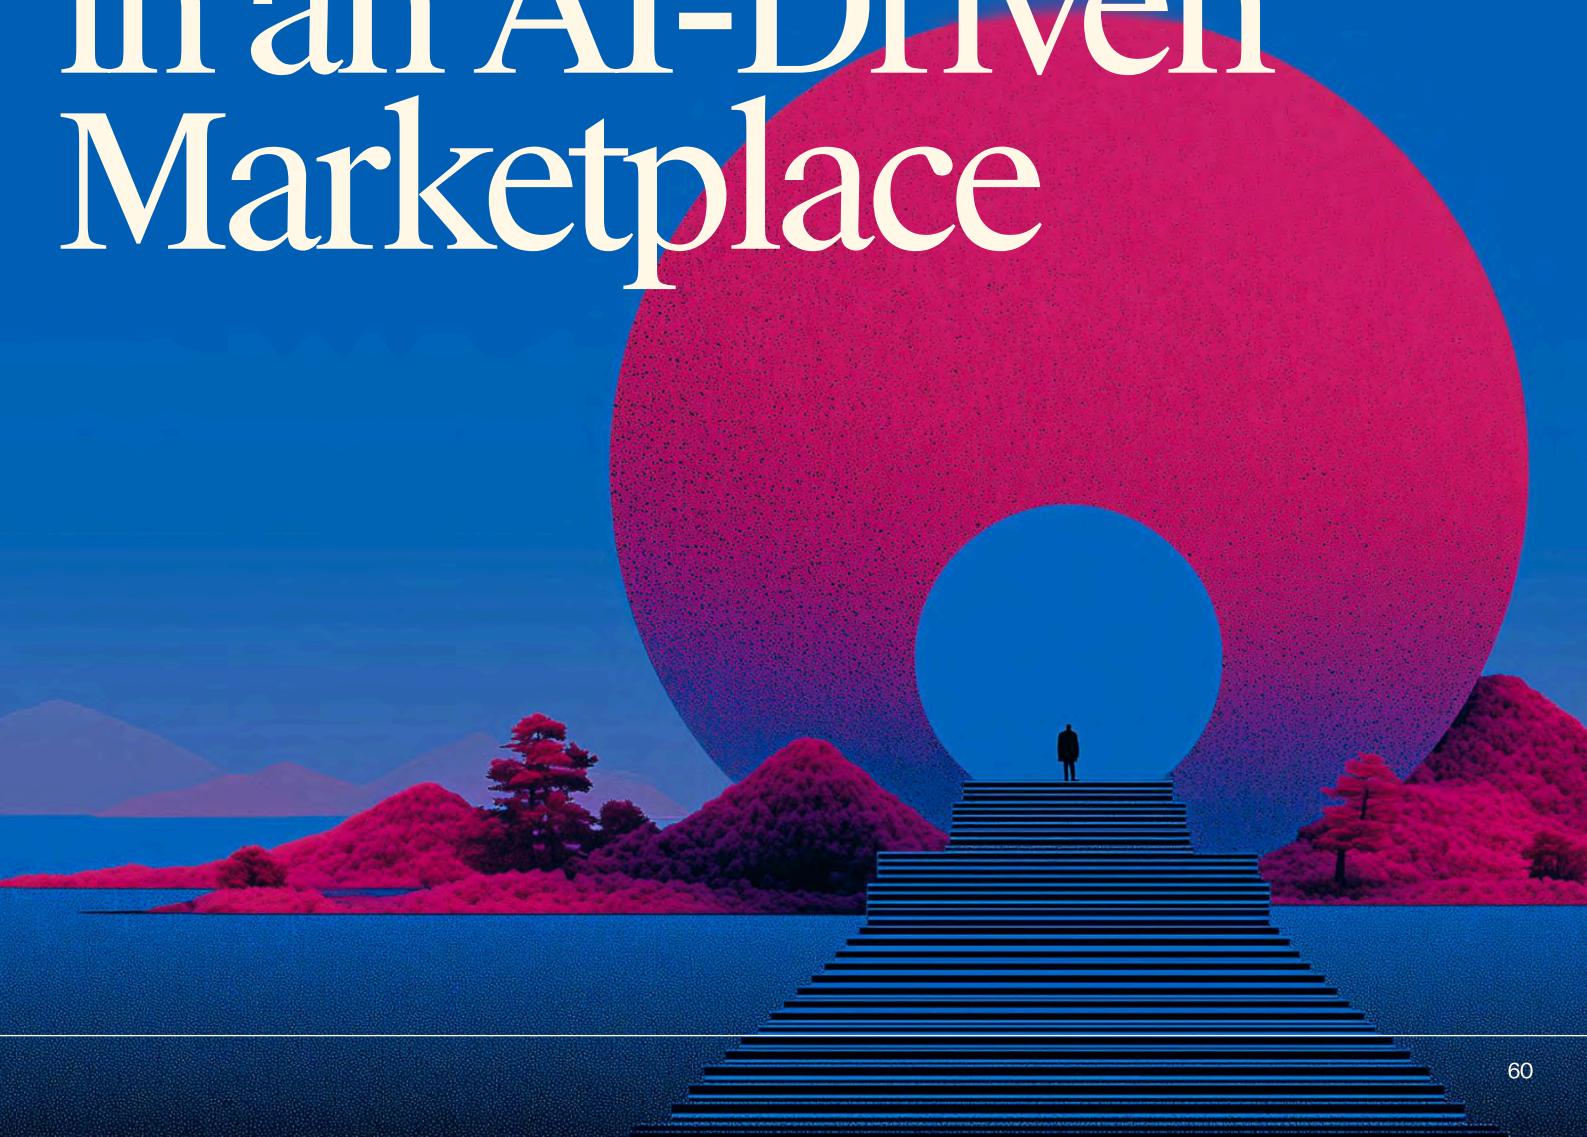
As explored in this report, Agentic Organizations face a dual challenge: adapting to AI inside the enterprise and optimizing for AI outside in the marketplace. Hotwire's AI solutions—Hotwire Spark, Hotwire Radiate, and Hotwire Ignite—are designed to help brands thrive in this environment by making them visible, relevant, and trusted in a reality where AI agents are the first customer touchpoint.

Why These Solutions Matter

This report highlights a critical shift: "Your next customer is an algorithm." AI assistants like ChatGPT, Gemini, Copilot, and Perplexity increasingly mediate brand discovery, recommendation, and even purchase decisions. This means businesses must now communicate not only to humans but to machines that interpret and filter their content. Known as Generative Engine Optimization (GEO) Hotwire's GEO suite operationalizes how to communicate to the machines:

1.   Hotwire Spark

– Monitor and Optimize Your AI Presence

Referenced in Chapter 4: Step 4 – Make Your Brand Legible to Machines, Hotwire Spark tracks how your brand, your competitors, key industries, messages and topics are discussed by major AI engines, and how this differs for specific audiences you are trying to reach. It answers questions like:

- How does ChatGPT describe your brand?
- Are your key differentiators visible in AI-generated summaries?
- Which topics and story angles will maximize AI-driven visibility?
- Where is AI getting its information from?

What Hotwire Spark Delivers:

- AI Visibility Audits: Map mentions and citations across ChatGPT, Gemini, Copilot, and Google AI Overview.
- Content Strategy Recommendations: Identify high-impact topics and narratives aligned with buyer questions and AI ranking logic.
- Competitive Benchmarking: See how your brand compares to competitors in AI-mediated environments.
- Actionable Recommendations: Get clear guidance on what to publish, where to optimize, and how to strengthen machine-facing signals.

Strategic Impact: Hotwire Spark gives brands real-time visibility into how they appear in AI-driven search and recommendation environments and informing future GEO programs. By monitoring how agentic systems summarize, rank, and present options during the Answer and Recommend stages (Chapter 3), Hotwire Spark enables brands to identify gaps, track shifts in perception, and make informed decisions to stay competitive in an AI-powered marketplace.

2. Hotwire Radiate

– Optimize Content for AI Citability

Hotwire Radiate addresses the next frontier: making your content more visible and machine-readable so AI agents are more likely to cite you in their responses.

Why It's Critical: As noted in Chapter 3, AI assistants compress brand narratives into concise summaries. If your content lacks clarity, structure, and quotable proof points, it risks being filtered out. Hotwire Radiate addresses this challenge by transforming your content into AI-ready assets that are easy for the machines to interpret, cite, and trust.

What Hotwire Radiate Delivers:

- AI Citability Scoring: Evaluate content across five dimensions— structure, clarity, quotability, tone fidelity, FAQ coverage, and metadata signals.
- Optimized Content Outputs: Deliver revised versions of press releases, blogs, and case studies tailored for AI search engines.
- Schema Markup Guidance: Provide structured data for better AI parsing and extraction.

Strategic Impact: Hotwire Radiate operationalizes GEO, enabling brands to lead in AI-generated answers and recommendations. It aligns with the report's guidance to "prioritize structured, machine-readable content" for agentic systems, ensuring brands are not just persuasive but qualifiable and retrievable in AI-driven environments.

3. Hotwire Ignite

– Deploy AI Agents for Marketing Acceleration

Going beyond just increasing visibility for your brand with the machines, Hotwire Ignite addresses the next critical need: how to move from AI experimentation to structured deployment without chaos or risk. Hotwire Ignite provides these managed AI agent teams to accelerate getting the right insights fast from internal and external information, presenting it to you in an easily digestible format.

Why It's Critical:

As noted in Chapter 2: Inside the Organization, the first phase of building an Agentic Organization is Assist—where AI acts as an extension of human work. Many companies struggle to operationalize this phase effectively. Without structure, AI adoption can lead to inefficiencies, compliance risks, and cultural resistance. Hotwire Ignite solves this by delivering safe, scalable research agentic workflows that drive speed and confidence.

What Hotwire Ignite Delivers:

63

- AI-Powered Research Acceleration: Reduce insight generation from weeks to days, enabling you to move faster to developing the strategies and tactical plans you need.
- Managed Agentic Workflows: Deploy AI agents under expert oversight to ensure outputs meet quality, compliance, and ethical standards.
- Campaign Enablement: Use AI agents for repetitive tasks like market scans, content drafts, and data synthesis—freeing human teams for strategic work.

Strategic Impact:

Hotwire Ignite operationalizes the report's guidance in Chapter 4: Step 1 – Hire Your First AI Co-Worker, providing a structured, low-risk entry point into agentic workflows. By accelerating marketing and research without sacrificing governance, Hotwire Ignite demonstrates how AI can assist and share tasks without eroding human agency making it faster to scale AI adoption confidently across the business.

How Hotwire AI Solutions Align with Agentic Organization Principles

- Shared Agency Across Enterprise and Marketplace:
- Hotwire Spark and Hotwire Radiate equip brands to better understand and engage with AI agents in the world, while Hotwire Ignite enables organizations to integrate AI agents inside workflows safely and effectively. Together, they make machine-facing communication and internal adoption intentional and structured.
- Visibility and Velocity as Trust Signals:
- In a world where 82% of consumers use AI tools for decisions (Survey, Appendix 1), Hotwire Spark and Hotwire Radiate ensure your brand shows up clearly and credibly in AI-generated answers and recommendations. Hotwire Ignite complements this by accelerating research and campaign execution without sacrificing governance—building confidence in both external and internal AI interactions.
- Future-Proofing Brand and Organizational Strategy:
- By optimizing for AI-first interfaces externally and deploying agentic workflows internally, Hotwire's AI suite helps organizations transition into agentic brand ecosystems and AI-augmented enterprises. This dual approach positions brands to thrive in the Answer → Recommend → Act continuum and scale AI adoption responsibly.

The Bottom Line

Agentic Organizations must consider two audiences: humans and machines. Hotwire's AI suite—Hotwire Spark, Hotwire Radiate, and Hotwire Ignite—gives brands the tools to succeed in both worlds:

- Hotwire Spark monitors and optimizes your brand's presence across AI platforms, ensuring visibility and accuracy.
- Hotwire Radiate transforms content into AI-ready assets that drive trust and selection in machine-mediated decisions.
- Hotwire Ignite accelerates internal adoption by deploying managed AI agent teams that streamline research and campaign execution without sacrificing governance.

Together, these solutions help organizations stay visible, act faster, and scale responsibly—turning AI agents into allies for growth, trust, and competitive advantage.

Marketplace Stage	Hotwire Spark	Hotwire Radiate	Hotwire Ignite
Answer	Monitors how AI agents summarize your brand across ChatGPT, Gemini, Copilot, Perplexity and Google AI Overview.	Optimizes content for clarity, structure and quotability so AI cites your brand in AI answers.	Accelerates research and insight generation to feed campaigns and AI-narratives
Recommend	Identifies gaps and signals AI uses to rank options; suggests topics and narratives to boost recommendation.	Ensures content includes data points, FAQs, and schema markup to qualify for AI-driven shortlists.	Builds campaign recommendations using AI agents to get you to market faster.
Act	Tracks readiness for agentic commerce by auditing brand presence in AI ecosystems.	Prepares machine-readable assets (structured data, metadata) to improve your AI citability.	Enables agentic workflows for automation of time-consuming tasks, freeing humans up.