

Why CEOs Need to Prepare for AI-Only Rivals

By [Vladimir Lukic](#), [Nina Kataeva](#), [Nicolas de Bellefonds](#), [Ulrich Pidun](#), Gabe Bouslov, and Adam Job

ARTICLE OCTOBER 23, 2025 15 MIN READ

In April 2025, Microsoft unveiled the concept of the “frontier firm”: a new type of AI-first organization in which AI agents execute and deliver core outcomes while humans oversee and orchestrate. Microsoft expects every organization to start turning into such a firm within the next five years. But the trajectory of organizational evolution isn’t likely to stop there. What lies beyond the frontier? A fully autonomous, *AI-only* firm—a network of hyper-specialized AI agents, coordinated by a central agent, operating entirely without human employees. (See “The AI-Only Firm of the Future.”)

— THE AI-ONLY FIRM OF THE FUTURE

An AI-only firm would operate entirely without human employees. One way it might organize would be as a network of hyper-specialized AI agents within a hierarchical structure. Specialized AI agents would oversee departments such as strategy, marketing, IT, and operations. Each of these departments would be coordinated by a “lead agent” who assigns tasks and manages resources (for example, by scaling up or down the number of agents in its teams).

At the head of the company would be an AI CEO, setting the overall strategy and being responsible for the firm’s success. Given that the topic is still nascent and conceptual, we should expect the structure of these firms may eventually evolve beyond imitations of human-designed firms. For now, though, studies of AI-only teams find that productivity is highest in designs that emphasize specialization and coordination.

Imagine you are an insurance executive waking up to a jarring headline: a fully autonomous, AI-only private insurance company has just launched. This competitor has no employees whatsoever—not a single underwriter or claims adjuster. The firm touts 50% lower premiums, 24/7/365 service, and settlements in minutes. Within hours of the announcement, your company’s stock tanks.

We don’t know whether AI-only firms will become a reality in five years or 15. However, the scale of their potential impact—on competition, as well as on society more broadly, given the implications for large-scale job displacement and the erosion of human control over the economy—is sufficient to warrant serious consideration from CEOs today. This article focuses on how incumbents can reimagine themselves as AI-first players to compete with AI-only firms, which would represent an entirely new category of companies.


A Wave of Investment Is Easing Constraints

A true AI-only firm—whether in insurance or other industries—does not yet exist. Yet the vast capital flows into the space suggest that investors and large technology companies see significant potential for market disruption beyond AI-first organizations.

Venture capital funding for AI companies reached roughly \$116 billion in H1 2025, already surpassing all funding in 2024 (\$101.5 billion).¹ Notably, agentic and autonomous AI startups are at the forefront of this charge. According to our analysis, more than 100 startups reporting “autonomous capabilities” have been founded since 2022, funded by 120 different lead investors to the tune of \$41 million, on average.

Roughly half of these maverick startups position themselves as offering agentic “replacements” for specific roles, and more than 10% are targeting automation of entire teams or departments. (See the exhibit.)

More Than 100 Autonomous AI Startups Are Targeting Role Automation



Sources: Internet research; Crunchbase; BCG Henderson Institute analysis.

¹Analysis based on a representative sample of 110 companies reporting autonomous capabilities; does not include all companies that meet this criteria.

²Financial analysis sample includes 78 companies with available financial data across ~200 funding rounds from January 2022 through April 2025.

³Analysis based on the scope of automation reported by companies in the sample.

With these investments, rapid progress is being made to address the technological barriers which, for now, make such firms infeasible. For example, when our report *Unlocking the AI-First Organization* was published in May 2025, flagship models could work autonomously for just over an hour. By September 2025, this task length limit has grown to more than 30 hours, far exceeding the human workday.

Still, hurdles remain across reasoning capabilities, multi-agent coordination and orchestration, and external interactivity, among other factors. These must be overcome to bridge the substantial evolutionary gap between human-supervised deployments of AI agents and fully autonomous firms. But with advancements being made on all fronts, it is unlikely that these challenges will remain insurmountable. (See “Hurdles for AI-Only Firms.”)

— HURDLES FOR AI-ONLY FIRMS

Significant *technological hurdles* remain before AI-only firms can become a reality.

Reasoning Capabilities. Many business processes—from lead generation to analysis—are multi-step chains of sub-decisions, each requiring a different set of expertise and reasoning capabilities. For example, an outbound sales agent has to identify prospects, prioritize them based on alignment to target profiles, decide at what time to reach out, how to personalize the message, and know when to give up on the lead.

On reasoning abilities, the trajectory for AI is clearly positive: leading AI models like OpenAI's GPT-5 and Google's Gemini 2.5 now show reasoning abilities equivalent to humans with an IQ of about 130. They are also able to plan and independently execute on longer, more complex tasks; Anthropic's recently released Sonnet 4.5 is able to focus for more than 30 hours on such multi-step tasks, far outlasting the seven hours achieved by the predecessor Opus 4. What's more, the "task length" limit is doubling at least every three to seven months, and this rate appears to be accelerating.

Coordination and Orchestration. All firms depend on a myriad of specialized teams collaborating and communicating seamlessly across departments and functions. Thus, at minimum, an AI-only firm's lead agents would need to be able to effectively coordinate tasks and information flows between and across their teams; to resolve conflicts when agents disagree with one another; and to determine when to spin agents up or down based on volume projections.

Though multi-agent coordination remains a significant challenge, progress is being made on this problem: in the last year, studies have demonstrated success using orchestrator or evaluator agents for coordinating and reviewing activities, and leveraging agent voting for conflict resolution. A significant number of the AI startups founded in the past few years claim to offer full team or department replacement using such multi-agent architectures: for example, CaseFlood, a player aiming to replace law firms' operations staff with AI agents, pitches to potential clients that it can double their conversion rates using a team of specialized agents for lead intake, qualification, and closing.

Memory and Knowledge Systems. Modern enterprises manage vast repositories of client and customer data, legal records, and internal documents. Therefore, the information infrastructure of an AI-only firm would necessitate efficient storing of vast amounts of both real-time and historical data. Systems would need to exist to allow agents to accurately save, retrieve and utilize this data for analysis, reporting, or decision making.

Historically limited by fixed "context-windows" (such as 100,000 words for GPT-4), emerging trends show promise on these fronts. Google research shows that, theoretically, infinite context windows are possible; Additionally, "knowledge graphs" may offer structured, contextual retrieval across ever-evolving corpora. In a sense, existing solutions like Microsoft's Semantic Index, capable of handling tens of millions of documents, can already provide the plug-and-play infrastructure for an AI-only firm's knowledge management.

External Interactivity. As participants in markets, AI-only firms would need to navigate complex networks of (human or AI) clients, partners, regulators, and

service providers. Truly contending with incumbents would require communicating naturally via email or phone, and assimilating into existing systems like payment processors, CRMs and other tools, as well as legal or tax filing sites.

Though traditionally confined to chat interfaces or codebases, agentic systems are becoming more integrated: whether interacting with smartphone users via voice, or by leveraging the innovative Model Context Protocol (MCP)—like a USB-C port for AI apps—which is enabling agents to easily connect to tools like Slack or even plug into enterprise databases. Recent examples show evidence that such AI market participation is underway, including quasi-autonomous recruiting agencies conducting interviews with humans.

Moreover, to compete in sectors like manufacturing, construction, or logistics, AI-only firms would need to interact with the physical world. This necessitates significant advances in robotics. While robots are already widespread in industrial and logistic automation—for example, roughly 75% of Amazon’s orders are fulfilled by a fleet of 750,000 robots—they have not yet replaced humans in tasks requiring fine motor skills, dexterity, or adaptive decision-making. Many nonrepetitive tasks still require human operators or supervisors to step in.

Equally important are improvements in AI-robot communication protocols; many current systems rely on pre-scripted instructions or limited APIs. For AI agents to dynamically and reliably direct robotic systems, they will require more robust machine-to-machine interfaces, possibly involving high-level goal setting, real-time feedback loops, and multimodal sensor integration.

Overall, these technological barriers suggest a phased evolution toward the AI-only firm: AI-only firms are most likely to emerge first in digital-native industries, such as software development, algorithmic trading, or digital marketing, with minimal physical-world interfaces. Moreover, in domains with kinetic requirements, hybrid setups may emerge where AI agents handle decision-making and coordination, while humans or semi-autonomous robots execute the physical components.

Beyond technical barriers, certain *regulatory and societal hurdles* will also need to be overcome. No jurisdiction currently recognizes AI as legal persons, barring them from leading a company, signing binding documents, or assuming fiduciary responsibilities. Several countries are making headway on addressing these hurdles: for example, the UAE permits AI board observers as part of its 2031 AI Strategy; a recently introduced bill in the US Congress could qualify medical AI systems to prescribe drugs. Still, in the foreseeable future, we expect that AI-only firms would be created by human founders, who define their overall mission and assume fiduciary and signatory responsibilities.

Societal hurdles will likely prove just as challenging as regulatory or technical barriers. Much like autonomous vehicles, AI-only firms may face public resistance to ceding control to artificial intelligence, especially in trust-sensitive sectors like healthcare and finance. Concerns about job displacement and loss of control over the economy could manifest as consumer boycotts or political opposition that delays adoption regardless of technical capabilities.

Competing Against an AI-Only Firm

AI-only firms would upend many foundational competitive assumptions, with an edge over traditional firms on several dimensions.

Costs. By replacing human workers with AI agents, AI-only firms could eliminate wages and associated overhead costs—such as benefits, leases, or HR infrastructure. Of course, these firms would face a different set of expenses, which would primarily consist of energy, infrastructure, and compute costs (as one proxy, A16Z’s AI startups allocate more than 80% of raised capital to compute).²

Critically, however—and unlike labor costs, which are usually inflationary over time—these costs are declining rapidly. The cost to use AI models is falling at a rate of about ten times per year, driven by 30%–40% annual improvements in AI hardware efficiency and costs, “AI price wars” due to challengers like DeepSeek, and increasing model commoditization.³ Already, there is potential for dramatic arbitrage in some areas: at \$9/hour, semiautonomous agentic coder “Devin” is far cheaper than the \$30–\$75 hourly rate a human programmer commands.⁴ If the deflationary tendencies of the cost base of AI-only firms persist, this would unlock a structural cost advantage over traditional and AI-first firms.

But being composed entirely of AI agents does not just mean lower costs. Rather, it would represent a new operational foundation that could allow AI-only firms to deliver both superior customer experience and achieve a new type of adaptability.

Customer Experience. AI-only firms could be ultra-responsive to customer demands. Fully compute-based, and with no human workforce to manage, AI-only firms would operate as 24/7/365 execution machines. Moreover, processing that would take humans hours could happen in seconds

—with some AI-first insurers already reporting more than 50-fold reductions in claims processing times.⁵

Furthermore, these agents could learn fast, leading to continuous improvements in quality: every client interaction, decision, and outcome would become data points in an AI-only firm's decision engine. With perfect recall of these data points (across all agents, regardless of which customer they have previously interacted with), agents in customer service or sales could be iteratively upskilled, leading to fewer mistakes and lower variance in service quality.

With time, the customer experience could become hyper-personalized, with microproducts tailored to individual needs on demand. At the ultimate stage, offers could even preempt customer demands. Walmart's shopping agent Sparky represents an early version of this vision, proactively addressing the perennial question of "what's for dinner?" by creating personalized meal plans and automatically adding ingredients to the shopping cart.

Adaptability. Today, most businesses struggle with adapting to a highly volatile environment, due to slow, rigid planning cycles, departmental silos, and the classic principal-agent problem (coupled with humans' natural resistance to change). AI-only firms could overcome these issues: a fully digital workforce would mean that firms can copy, deploy, or remove agents in response to demand, scaling up or down as needed. Moreover, when the CEO of such a firm would make a decision or envision a structural or behavioral change, it could be implemented instantly across the entire agent network with perfect alignment and zero resistance. Finally, agents could be continuously updated with new capabilities or performance patches.

As a result, AI-only firms could become versatile, ever-evolving organisms that continuously adapt their strategies on a whim. For example, strategy could be set weekly, not annually, and driven by always-on, parallel simulations that test pricing or marketing tactics and select statistically optimal short-term paths.⁶ Reactions to market stimuli or competitor moves could be nearly instantaneous, for example, copying innovations or countering sales initiatives with better offers.

How to Prepare for a New Age of Competition

Given the considerable advantages AI-only rivals might have, how should leaders of incumbent businesses react? Below, we outline a three-part approach that consists of first, reducing the AI-only firms' edge by becoming an AI-first company; second, capitalizing on human capabilities; and third, positioning your firm as a critical node in the hypothetical AI-dominated business ecosystems of the future.

Becoming an AI-First Company. Contending with AI-only rivals by relying purely on human strengths will be an uphill battle. Becoming an AI-first firm provides a foundation for narrowing some of the AI-only firms' advantages.

In some areas, such as coding or sales, the process-wide deployment of agents is already becoming feasible today—and may offer significant benefits. For example, in a typical outbound sales pipeline, processes like prospecting and outreach consume roughly 70% of a representative's time, while the deal-closing phase consumes about 30%.⁷ Redesigning this process such that agentic AI handles prospect research, drafting messages, and scheduling calls can significantly increase speed and reduce customer acquisition costs. Human sales agents, meanwhile, can focus on orchestrating these agents and managing the deal-closing phase, at which, for now, studies show a strong client preference for human interaction.

While few organizations today are truly ready to operate as AI-first, those furthest ahead move beyond using AI as a bolt-on to existing processes and establish it as the new core operating system. This requires rethinking how work is structured and how value is created, redesigning end-to-end workflows around autonomous AI agents, and embedding continuous learning into the organization's fabric.

Thus, three imperatives for leaders emerge: first, establish a clear AI-first north star by identifying the processes that are most advantaged when redesigned around agentic AI; second, selectively launch end-to-end transformations of the highest-impact workflows; and third, building on this momentum, scale these successes across the organization. An alternative would be a greenfield approach, whereby the firm sets up a new, dedicated business—which initially replicates the parts of the organization that best lend themselves to agentic outsourcing and is then scaled up over time.

Capitalizing on Human Capabilities. Today, humans maintain advantages over AI in areas like complex reasoning and domain-specific expertise. However, as the capabilities of AI continue to develop, it's likely that such advantages will gradually erode and fail to serve as sustainable moats. Yet, as AI commoditizes previously human-dominated domains, new constraints to maximizing value creation emerge—and human skills will be critical to addressing them.

- **Curation.** As technology expert Sangeet Paul Choudary argues, when AI commoditizes the generation of content, tasteful curation becomes a new constraint—which humans, leveraging their lived experience and cultural insights, may be uniquely positioned to address.⁸ Thus, a media brand might use AI to surface relevant content but rely on human editors and curators to craft selections that reflect brand identity and ensure emotional coherence.
- **Imagination.** As AI floods every domain with endless ideas and hypotheses, constraints shift from idea generation to evaluation and prioritization. In fact, a recent meta-analysis revealed that while AI can easily generate countless novel hypotheses, it consistently fails to distinguish between those that are merely creative and those that are feasible.⁹ Fields where this matters

most, like pharmaceuticals or deep tech, can capitalize on this by focusing resources on evaluation and selection aspects of upstream R&D, utilizing AI for ideation, while relying on human researchers to determine which problems to solve and which solutions merit investment.

- **Empathy.** As communication becomes commoditized, empathy and authenticity emerge as new areas of competitive differentiation. Today, AI struggles with emotional resonance and ethical nuance; and even as advancements are made on these fronts, many relationships rely on genuine human presence—especially in moments of vulnerability or crisis. Think of fields like health care, education, or customer relations. For example, a hybrid therapy service might use AI for background intake and analytics but retain human therapists for live sessions.
- **Human Touch.** As 24/7 AI-only operations skyrocket productivity and create abundances of physical goods, human involvement may become highly prized. For instance, the “human-made” label may take on a new cultural cachet, like artisan or locally produced goods today, conveying effort and intentionality and creating symbolic value.¹⁰ For example, some fashion companies of the future may rely purely on AI designers, which are informed by analyzing patterns on social media to detect what might be popular. Yet, others might rely purely on human designers, spotlighting how their brand is guided by the taste of a visionary creator—like major luxury fashion houses do today.

To be clear, these emerging constraints may not last forever. What will endure is the cycle itself: AI advances, creating new constraints as old ones are commoditized. Importantly, incumbents should not retreat to capabilities that AI cannot yet solve but instead seize the opportunities now possible *because* AI is so successful. For executives, this requires constantly asking: What new constraints have emerged from the latest AI developments in my industry? Could this have unlocked new opportunities for deploying uniquely human capabilities?

Serving as a Valuable Node in an AI-Only Ecosystem. AI-only firms should not be viewed solely through the lens of competition: incumbent firms can not only generate value by capitalizing on their unique strengths, but also by positioning themselves as valuable partners for AI-only players, supporting them in solving emerging constraints while capturing a share of the value created. Below are some of the positions incumbent firms could occupy in an AI-only ecosystem of the future.

- **Physical Supply Chain Integration.** Many AI-only firms would need to interface with kinetic infrastructure, like factories or logistics networks. Incumbent firms could act as execution layers or integration partners, connecting AI intelligence to physical-world workflows—for example, handling manufacturing and fulfillment—while AI-only firms handle design and sales.
- **Trust and Relationship-Building.** AI-only firms would have an ongoing need for the relational and social capital that is required for B2B partnerships, negotiations, and communities.

Relationship capital accumulated over decades could not be easily replicated by an AI-only firm which has an inability to “shake hands”; one study showed that 89% of executives say in-person meetings seal deals.¹¹ As such, incumbent firms could play the role of relationship brokers for AI-only firms. In consumer-facing contexts, trust is equally critical. For example, a newly spun-up, fully autonomous insurer might struggle to persuade potential clients that it will stand by its policies when claims arise.

Likewise, an AI-only entrant in the white goods market could find customers skeptical about the reliability of its warranties. In these scenarios, incumbent firms could play a guarantor role—lending their brand equity, reputation, and credibility to AI-only partners, or even underwriting their commitments to end customers. In doing so, incumbents could become the human face and trusted anchor in otherwise machine-driven ecosystems.

- **Ethical Stewardship.** As AI-only firms interact with more people, regulators, and society at large, demand for human oversight and governance services may emerge—whether by choice or by regulatory mandate. Incumbent firms could step in as ethical stewards, offering governance frameworks, oversight boards, or ethics-as-a-service to help AI-only firms maintain their societal license to operate. For example, companies might provide fairness audits for AI-only companies in finance, or they might develop platforms that offer reviews of health care AI firms’ decisions by licensed human clinicians to comply with medical liability laws and ethical care standards.

The Broader Societal Question

The companies that will thrive in the face of the emergence of AI-only firms are those that start today by tracking progress in AI development, piloting and iterating on new agentic tech before peers, and initiating strategic conversations about human capabilities and how they may be leveraged to solve emerging constraints in competition or cooperation with AI-only players.

Yet, as firms make these decisions, the AI-only concept raises a broader, more fundamental question: do we, as a society, want AI-only firms to become a reality? The societal consequences might be significant: potential large-scale job displacement, social and political destabilization driven by growing inequality or privacy erosion, and even the risk of catastrophic AI misalignment. These consequences could lead to loss of human control of large parts of the economy and the need to rethink the role of humans in both business and society at large.

At the same time, there are enormous competitive incentives to make these firms a reality—given that many of today’s most valuable companies are staking their fortunes on the broad adoption of AI agents, as well as the models and hardware powering them. Moreover, there is the ever-present

threat of geopolitical asymmetries, for instance, falling behind other nations that are more open to embracing AI-only firms.

Given that autonomous agents are likely to find significant use in corporate contexts early on, business leaders have both a critical role and a unique responsibility in fostering a debate that goes beyond business, and includes legal, sociological, and ethical considerations of AI-only firms.

Authors



Vladimir Lukic

Managing Director & Senior Partner; Global Leader, Tech and Digital Advantage
Boston



Nina Kataeva

Managing Director & Partner
Zurich



Nicolas de Bellefonds

Managing Director & Senior Partner
Paris



Ulrich Pidun

Partner & Director; BCG Henderson Institute Insight Leader
Frankfurt



Gabe Bouslov

Consultant
Chicago



Adam Job

Strategy Lab Senior Director, BCG Henderson Institute
Munich



ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we work closely with clients to embrace a transformational approach aimed at benefiting all stakeholders—empowering organizations to grow, build sustainable competitive advantage, and drive positive societal impact.

Our diverse, global teams bring deep industry and functional expertise and a range of perspectives that question the status quo and spark change. BCG delivers solutions through leading-edge management consulting, technology and design, and corporate and digital ventures. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, fueled by the goal of helping our clients thrive and enabling them to make the world a better place.

© Boston Consulting Group 2025. All rights reserved.

For information or permission to reprint, please contact BCG at permissions@bcg.com. To find the latest BCG content and register to receive e-alerts on this topic or others, please visit bcg.com. Follow Boston Consulting Group on [Facebook](https://www.facebook.com/bcg) and [X \(formerly Twitter\)](https://twitter.com/bcg).

- 1 CB Insights, *State of AI Q1'25 Report*, May 1, 2025.
- 2 Guido Appenzeller, Matt Bornstein, and Martin Casado, “Navigating the High Cost of AI Compute,” Andreessen Horowitz website, April 27, 2023.
- 3 Stanford University Human-Centered Artificial Intelligence, *Artificial Intelligence, Index Report 2025*; Axel Backlund, “Intelligence Too Cheap to Meter: Positioning for the Future Of AI,” *Forbes* website, May 19, 2025.
- 4 Mike Wheatley, “Cognition AI Launches Revamped Coding Assistant Devin 2.0 with Much Lower Starting Price,” SiliconAngle website, April 3, 2025; Jaye Hannah, “Software Engineer Salary Guide: Everything You Need to Know,” CareerFoundry website, December 19, 2023.
- 5 Nazar Kvartalnyi, “Exploring AI in Insurance: Applications, Examples, and Insights from Our Team,” Inxoft, May 1, 2024.
- 6 Dwarkesh Patel, “What Fully Automated Firms Will Look Like,” Dwarkesh Podcast blog, January 31, 2025.
- 7 Ilse Van Rensburg, “Prospecting Automation: Cracking the Code for More Sales,” Cognism website, April 11, 2025.
- 8 Sangeet Paul Choudary, “Humans as ‘Luxury Goods’ in the Age of AI,” *Platforms, AI, and the Economics of BigTech* Substack newsletter, April 27, 2025.
- 9 Amy Wenxuan Ding and Shibo Li, “Generative AI Lacks the Human Creativity to Achieve Scientific Discovery from Scratch,” *Scientific Reports* 15, 2025.
- 10 Nick Bostrom, *Deep Utopia: Life and Meaning in a Solved World*, Idrapress Publishing, 2024.
- 11 Harvard Business Review Analytic Services, *Managing Across Distance in Today's Economic Climate: The Value of Face-to-Face Communication*, March 2018.