




ARTIFICIAL INTELLIGENCE

Five Ways That AI Front-Runners Change How Work Gets Done

By Sagar Goel, Julie Bedard, Matthew Kropp, Ashley Sim, and Charikleia Kaffe

ARTICLE SEPTEMBER 24, 2026 15 MIN READ

AI front-runner firms don't necessarily have record-breaking amounts of AI usage. Instead, they use AI to reimagine how work gets done: reinventing their operating models and ditching established hierarchies and processes in pursuit of speed, growth, and competitive advantage. But what is really working—and what can leaders of traditional firms learn from the successes of AI front-runners?

A recent BCG Institute study showed what is at stake: companies that are able to make this transformation and redesign work and talent with AI, rather than just layering AI tools on top of

old approaches, outperform their peers by more than 11 percentage points in industry-adjusted shareholder returns annually, driven by revenue growth and margin expansion.

To better understand this journey in practice, we went to the source, interviewing senior leaders at more than 50 AI front-runner firms. These companies are headquartered in ten countries, range from startups to enterprises with over 50,000 employees, and represent more than a dozen industries. We talked to leaders at different types of AI front-runner firms, including 23 AI natives built on AI, 18 digital platforms with AI as the core engine, and 9 incumbent enterprises rewiring the business around AI, function by function. (See the sidebar, “Methodology,” for more details.)

– Methodology

This article draws on a qualitative study of AI front-runner companies conducted by the BCG Institute between March and August 2026.

How We Selected Companies. We define AI front-runner companies as organizations for which AI fundamentally redefines what the company offers to its customers, how it competes, and how it operates. The sample spans three archetypes:

- **AI-Native.** Companies where AI is the core value-creating engine of the product or system, whether digital or physical.
- **Digital-First.** Digital-native platforms that use AI as a core engine for optimization, ranking, personalization, coordination, or monetization.
- **Traditional Front-Runners.** Incumbent enterprises using AI to optimize and reshape legacy operations or products at scale, without changing their core business model. Large organizations rarely move uniformly. What we studied, and what our examples describe, are the specific functions and workflows that have gone furthest with AI.

Who We Spoke To. We conducted interviews across 50 companies with leaders positioned to describe how work is organized rather than how the technology is built: chief people officers, chiefs of staff, founders, and functional leaders. Our sample ranges from startups to major enterprises, with bases across the globe.

- **Geographic coverage** is 63% North America, 21% Europe, Middle East, and South America, and 16% Asia-Pacific.

- **Industry coverage** is half software and services and half diversified across other industries.
- **Size coverage** is across company scales (14% with fewer than 50 employees; 24% with more than 10,000), spanning startups to large enterprises.

How We Conducted the Interviews. Each interview was a semi-structured conversation following nine themes: the role of AI in the organization, culture, roles, skills, the employee value proposition, recruiting, performance management, upskilling, and ways of working and governance.

How We Analyzed the Data. We worked through each transcript against the nine themes in our interview guide, then tested which patterns recurred across companies regardless of size, geography, or archetype. Five held up consistently, and became the dimensions described here. We then went back through the sample and assessed each company against each dimension, counting a dimension only where a leader described something already in place, not an intention. That scoring underpins the prevalence figures cited throughout, including that 80% of companies demonstrated at least three of the five. We used AI tools to support the analysis—working through transcripts, surfacing patterns, and testing examples against source material—with every finding, quotation, and company example verified by the research team against the original transcript.

We've identified five dimensions along which AI front-runners have changed traditional ways of working. They are:

- The purpose of work.
- The work itself.
- The decision architecture.
- Coordination of work.
- Rewards and compensation.

Most of the leaders we spoke with described undertaking several of these shifts at once, and 80% demonstrated at least three of the five. For leaders of traditional companies, these dimensions

offer a view of how to manage the AI transition for their organizations and teams and a guide to the radical shifts in mindset and structure that are needed to enable the transformation.

The Purpose of Work: From Assigned Jobs to Owned Outcomes

The default for traditional companies is to define work top-down, with fixed roles and assigned tasks. In this environment, each employee is responsible for a bounded scope of work. However, in more than half of our sample of AI front-runners, leadership sets the direction, but individuals and teams own the outcomes. Employees are still accountable for the “what” but have greater accountability to choose the “how.”

In this structure, roles blur, with workers assembling around problems rather than org-chart positions. At AI coding agent platform Bolt.new, team borders are deliberately loose: if someone’s judgment or skill is useful to a colleague’s problem, they simply help—without needing to route the request through a manager or wait for it to be added to a roadmap. The expectation that you’ll pitch in wherever you’re useful holds regardless of formal role.



In more than half of our sample of AI front-runners, leadership sets the direction, but individuals and teams own the outcomes.

In fact, just under two-thirds of the leaders we spoke to noted their company had heavily reduced procedural checks and signoffs, instead putting real tools and budgets in people’s hands and making end-to-end ownership both safe and expected. At cloud deployment platform Railway, one leader described the shift as giving up a policing and rubber-stamping role, while more junior workers took on the responsibility to act, and the accountability for what happens next. Managers hold after-action reviews rather than needing upfront approvals. They trade signoffs for staged rollouts, and weekly status reviews for broad team access and autonomy. Meanwhile, engineers look beyond the code to understand the business, customers, and system one level higher.

AI coaching platform Valence also encourages and supports employees in finding their own AI solutions to problems. Employees own the problem and are given the autonomy and budget to obtain the right tools to solve it—subject to an IT security check and a finance review of spending patterns to ensure efficiency. Existing tools are made widely available, but team members are also expected to find new tools, experiment with them, and adopt them to make their work more effective.

The Work Itself: From “Humans Produce, Machines Assist” to “Machines Produce, Humans Direct and Evaluate”

In traditional companies, humans are the producers and their roles sit in fixed functional lanes.

Technology is used to speed up select tasks. But in AI front-runner firms, work shifts from executing to framing the problem and setting the objective. AI handles more of the doing (drafting, building, executing) and people shift to directing the technology—setting the brief, steering, and evaluating whether the output is good enough to ship. That changes the shape of a role. One person can now produce what used to take a team, so existing roles stretch to cover far more ground. People grow by directing AI to cover a larger scope, not just hiring beneath them. Entirely new roles have also emerged to build, manage, and govern the AI layer.

Among our interviewees, two-thirds describe purposefully training up their workforce so more people can direct AI, not just use it. They achieved this by setting AI-fluency expectations, hiring for AI competence, wiring fluency into each level, and making learning continuous and peer-led. We saw multiple ways that AI front-runner companies are doing this.



In AI front-runner firms, work shifts from executing to framing the problem and setting the objective.

The cloud content collaboration company Miro runs a multilevel AI-fluency program, expecting every employee to build and leverage multiple AI agents to deliver outcomes, rather than just helping on a task-by-task basis.

Cloud integration software company Zapier defines AI fluency across four dimensions—mindset (curiosity), strategy (when and when not to apply AI, matching problem to tool), building (hands-on use to produce quality results), and accountability (knowing what one is trying to achieve and holding output to it). As such, the company treats AI fluency less as “tool use,” wherein AI is applied to discrete tasks, and more about how a person thinks about their thinking. This is reflected in the company’s hiring, evaluation, and development approaches.

Software giant Salesforce expects managerial capability—foresight, judgment, setting a plan with checkpoints, and knowing when to intervene with AI—from all employees. Even a new joiner entrusted with a costly AI agent is expected to manage it with judgment, stepping in to redirect it rather than allowing it to drift off course and consume excess budget.

AI control tower ServiceNow is redesigning work to help employees focus on what’s important. Employees are rewarded for showing adaptability, curiosity, and the ability to apply context to create deeper insights. Leaders expect workers to shift from execution to insight and have built a rich orchestration layer with a library of agents and the tools to create their own.

AI front-runners are also establishing new roles to enable their operating model. Specifically, we identified five role archetypes time and again during our interviews.

Citizen Builders

Citizen builders are operators in the business, building their own apps and agents to do their own job better, reducing the old dependencies on engineering, data, or outside agencies.

Coding agent Cursor has employees outside engineering doing technical work themselves. An internal marketer replaced an entire agency by building an app that generated thousands of growth-marketing assets, and the finance team now models forecasts and digs into usage trends directly in the data, rather than sending requests to a data team.

Enterprise AI company DevRev has seen nontechnical employees ship production software—a sales-ops person with no technical background built a full Slack bot that reads from and writes back to the CRM.

Workflow Designers

People in this role rethink a company’s internal multistep processes from end to end: mapping how work flows across people today, deciding what AI should absorb versus what stays human, and repositioning the handoff points.

When enterprise GenAI platform Writer was trying to achieve more frequent references in AI assistant answers, this approach paid off. The redesigned workflow meant AI agents could mine

sales calls for customer questions and quickly check those answers against Writer’s own content to find gaps. Human workers could then focus on the part of the job where judgment was required—deciding what new content was published in response.

In another example, hiring platform Greenhouse is introducing an operations-oriented “innovation lead” within each department—a person who identifies which workflows to automate and evaluates tools to do so, then partners with a central enterprise apps team of IT engineers to build the solutions.

AI Guardians

AI guardians govern risk and compliance and keep the AI ecosystem in bounds.

One cloud-based company runs a dedicated AI council together with compliance and governance teams. When a manager elevates a proposed AI use case or agent with a business justification, it is registered and undergoes compliance and governance checks before the council reviews it and votes on whether it proceeds. After initial deployment, compliance teams stay involved in deciding whether it is fit to roll out more widely.

AI data infrastructure company Scale AI maintains a dedicated unit to strengthen the guardrails of the firm’s own models and agentic systems against hallucinations and data leakage by subjecting them to systematic attack.

Domain Anchors

Domain anchors encode expert knowledge into reusable form. Four in five firms described capturing the knowledge held by their strongest experts so that the whole organization can draw on it, not just the people who sit near them.

AI-powered skills intelligence and education technology company Workera hires one specialist per domain (such as a Ph.D. psychometrician) to encode their understanding into a document that an agent reads and amplifies across the company.

Insurtech company Reserv built its own platform on which employees create AI agents for their work, capturing each person’s expertise in how their agents operate, to be reused across the company. As these agents multiplied, the gains compounded. The time to build a new learning-and-development program fell from 18 months to four weeks, responses to regulatory changes dropped from about three days to about an hour, and junior employees were able to take on higher-stakes work than they would normally handle this early in their careers.

At Walmart, employees often move roles within the company every 18 to 24 months. To ensure their knowledge isn’t lost, experts document the different versions of how work actually gets

done, building a “digital twin” of their roles.

Agent Shepherds

Agent shepherds oversee how a fleet of agents operate: how they run and perform, what they can access, and when they are shut down.

Bolt.new describes a central, AI-native role responsible for data governance, role-based access control, where agents are hosted, and monitoring or decommissioning those that misbehave as agents proliferate across the organization.

Walmart is creating orchestration roles to manage its complex, interconnected AI systems—an extensive marketplace of agents spanning customer-facing and internal super-agents—and to support the teams building them.

The Decision Architecture: From Concentrated to Distributed

Under the default operating model, decisions flow up, information dispersion forces centralization, seniority confers authority, and accountability concentrates at the top. At AI front-runner firms, this linear structure is disrupted. Deterministic, rule-based calls can go to agents; reversible, lower-risk calls are pushed down toward whoever is closest; and irreversible ones tend to stay up. Decision rights are tied to the reversibility of the call, rather than seniority level, while accountability is more broadly dispersed.

About half the firms we spoke to described pushing decisions downward. Engineering intelligence platform Jellyfish lets the engineer closest to the work make the call. One example is in integration work: it’s highly repetitive, with similar changes made hundreds of times. Each change used to need sign off from a code owner outside the team. When the team built AI tools to write the integration code, change volume rose sharply, and reviews became an even bigger bottleneck. Today, the engineer who writes the change approves it, minding guardrails agreed once rather than a review each time, and changes clear within a day rather than three. Accountability follows the decision: you can ship it if you want, but you own it if it breaks.



About half the firms we spoke to described pushing decisions downward.

Enterprise AI company HappyRobot routes rule-bound decisions to agents and keeps humans for the judgment calls. In the past, a human would decide which engineer gets staffed on which customer deployment by weighing capacity and timelines. Now an agent applies deterministic logic, escalating to a human only when real judgment is needed, so human workers' time can be spent on higher-value work.

Of course, pushing decisions down works only if people have the judgment to make them. More than half the firms we spoke to described building this judgment deliberately—making reasoning visible in open channels, normalizing challenge and discussion, and giving less-experienced employees real decisions to make earlier in their tenure, with support from more experienced colleagues.

Reserv asks staff who are closest to problems to design and propose solutions and to lead iteration cycles. Senior leaders steer the process rather than prescriptively instructing their workers, who are able to develop better judgment as a result. At Cursor, workers share AI-generated proposals in public channels alongside their explanation of why they had the AI approach the work as it did. No one can hide behind AI output, and colleagues are expected to challenge weak AI-driven work rather than defer to it. Work AI platform Glean makes building judgment an explicit responsibility of managers, particularly when it comes to new grads who have yet to build professional experience and lack the pattern recognition that experienced staff develop over time.

The Coordination of Work: From Routine Alignment to Calibration and Relationship Building

Meetings, status updates, syncs, and approvals traditionally consume a large share of employees' time but are often necessary to ensure that key knowledge and context are transmitted from person to person. Now, AI can absorb many of the routine admin and alignment tasks and relocate the context to systems where it can be accessed more conveniently. In AI

front-runner firms, humans meet to make decisions, exchange ideas, and build relationships. About two-thirds of our interviewees said meetings have become more consequential, and a majority described middle managers shifting from being coordinators to become player-coaches.



About two-thirds of our interviewees said meetings have become more consequential.

DevRev spent two years building a proprietary knowledge system, Computer, that holds company-wide context centrally. Any employee can query it directly for answers instead of finding the right person or team to ask—removing the information handoffs that normally run through the hierarchy. Workera opens its weekly leadership meeting with an AI agent that reviews email and Slack to set out decisions made, decisions still pending, and recommendations grounded in a human-authored tradeoff document, freeing more time for genuine debate.

Other companies redesign meetings around decisions and problem solving rather than handoffs or status and protect human contact for what AI can't replicate, such as trust, tacit context, and creative friction. Miro has shifted routine data and status work to AI, handling dashboards and updates through natural-language queries rather than meetings, and redirected people toward customer relationship-building, which it regards as work AI cannot replace.

Rewards and Compensation: From Tenure and Activity to Impact, Judgment, and Leverage

In traditional companies, pay and advancement follow tenure, activity, and output volume; reviews are annual, top-down, and manager-to-employee. In AI front-runner companies, pay, promotion, and review reward demonstrated judgment, the leverage someone extracts from AI, and the impact of what they deliver. Feedback tends to be more continuous, visible, and peer-informed.

When evaluating success, AI front-runners focus on the quality of thinking, rather than deliverables. Some three-quarters weigh judgment over activity. So the test moves off the page and into real time: people are asked to explain their plan before work starts and to show where they overrode what the AI gave them.



When evaluating success, AI front-runners focus on the quality of thinking, rather than deliverables.

At data and AI platform Clay, work is evaluated by really assessing the thinking behind it. Short live conversations of 15 to 20 minutes are often relied on to check that someone genuinely understands and can act on what they've proposed, since AI-assisted output can look polished without the person having reasoned it through.

Glean assesses people on both the quality of their work and how they apply judgment over the AI's work, in hiring and in their roles. Sales candidates are given three or four prospects and are asked to use AI and their own judgment to decide which to pursue first and then to draft the outreach. The assessment is of their intuition regarding account strategy and the outreach email, their judgment of when to use AI, and how critically they reshaped what the AI produced rather than accepting it unchanged. The same standard applies once they are hired.

Zapier evaluates people on impact, not activity—for example, whether a worker improved a business metric, the customer relationship, or a teammate's work. It notes that AI is brilliant at activity, so someone can rack up heavy tool usage yet accomplish nothing.

Scale AI has made advancement impact-driven, with people reaching senior titles early in their careers when results justify it.

It's also important to recognize accomplishment publicly. Writer celebrates top contributors at all-hands meetings and has them explain what they built and why it mattered, showcasing those whose work and methods are adopted by others rather than those who simply record the most AI usage. In traditional firms, experts are often disincentivized to pass on their knowledge for fear of losing their status as “indispensable.” By making the act of sharing knowledge a visible path to success, AI front-runners ensure it's more widely dispersed, and not lost when star employees depart.

What Can Traditional Companies Do Today?

Traditional companies embarking on an AI transformation are in a challenging position. They need to make the five shifts that AI front-runner firms have already achieved, and do so while

carrying the existing business with them—a multiyear undertaking given that those systems actively resist the changes.

Critically, leaders cannot pursue all five dimensions at the same speed—they have to sequence and adapt them to their company specifics. We focus below on large companies in our sample because their coordination costs and legacy structure most resemble what incumbents will face in this journey. Leaders have pushed further and faster on some dimensions than others. Roles and judgment shift fastest, because they don't require dismantling anything: almost all large companies are building new AI-era roles and treating judgment, not activity, as the mark of good work. On the other hand, redefining work around outcomes is underway at only about two-thirds of large companies, and fewer than half have pushed authority down the corporate hierarchy.

What follows is a roadmap that traditional companies should keep in mind in their journey to AI front-runner talent practices: first building the foundation, then proving it in one place, and finally scaling it across the enterprise.

Build the Foundation

Define value and map the workforce impacts. Identify the business outcomes that your prioritized AI opportunities are meant to unlock. Don't adopt AI for its own sake. Then, translate those outcomes into which roles, functions, and skills are most affected.

Define AI fluency. Articulate what AI fluency means for your organization specifically: the mindset, judgment, and capability bar, not just tool usage. Then build it into hiring criteria. China-based real-estate platform Ke shifted its interview questions, de-emphasizing school pedigree in favor of asking candidates, especially new graduates, directly what they built with AI and what they learned from using it.

Equip and enable. Give people real tools and the space to experiment with AI in their actual work. Build deeper fluency and capability in prioritized areas, ideally through team-based workflows, and design learning journeys tailored to each function or job family. Walmart placed frontier tools in every frontline employee's hands early and explicitly framed AI as a means of elevating the work rather than removing jobs. The company reinforced this culturally by celebrating a veteran employee as the first to earn its AI certification, demonstrating that tenure was no barrier.

Accelerate the Journey

Pick your beachhead. Identify the one or two opportunity areas with the greatest value potential and the greatest readiness; these will become the template for the rest. For example, look for the function with the best business case, highest fluency, and most willing team. Rivian's software

division moved first and furthest, running its own AI framework about half a year ahead of the rest of the business. It now serves as the reference point as the company extends the model outward.

Bring the new model to life. Bring domain anchors and workflow designers into these pilot areas to redesign how work gets done and to redraw decision rights—what gets decided by agents, by individuals, or by senior leadership. Establish a new baseline for your meetings and syncs, pushing pure status and updates to async or collaboration tools and reserving real time together for what needs it.

Prove and reward it. Reconfigure incentives and recognition around the outcomes of these pilot areas. Make the outcomes visible—what got faster, what got better, what no longer needs a handoff—so the case for spreading the model rests on evidence rather than advocacy. Recognize the pilot team publicly and let them explain it to the next function.

Scale Across the Enterprise

Develop the role archetypes to scale enterprise-wide. Extend the roles in the pilot across the organization and establish the ones that become necessary at scale. Designate domain anchors in each area where deep expertise is valued and needs encoding. Widen access to workflow designers. Establish AI guardians and agent shepherds to govern and manage complexity as scale increases.

Institutionalize the capability. Make work reimagination a repeatable muscle rather than a series of one-off projects, through a dedicated team, a defined process, and a playbook that captures what each redesign learns. Set the ambition—for instance, the share of workflows you intend to transform—centrally, so progress does not depend on local enthusiasm alone. At ServiceNow, a central platform team was established to build core capabilities once and support all business units in adopting them, rather than each one reinventing its own approach.

Rewire talent and performance processes. Change what people are assessed on, so they are judged by the impact they create and the judgment behind it rather than by activity or output volume. Reward the behaviors—managers who build judgment in their teams, individual contributors who create assets others reuse—that make impact compound.

AI front-runner companies are increasingly expected to dominate their industries owing to their efficiencies and the speed with which they can scale and adapt to changes. But these changes are available to traditional companies, if they embrace the technology and the shifts in mindset required.

Critically, there is no end date for these changes. As technology moves, companies will need to keep adjusting—but the sequence we have described makes that adjustment deliberate rather

than accidental. The real work isn't catching up to where AI front-runners are today, but building the capacity to keep improving, just as those firms furthest ahead already do.

The authors thank the following colleagues for their valuable contributions to this article: Kateryna Gudziak, Josh Gyory, and Jia Chee Ooi.



BCG Institute is Boston Consulting Group's real-world think tank. We research the future of thinking about the future. Built on BCG's decades-long intellectual legacy and a global network of academics, scientists, and institutional partners, we bring together expertise across disciplines to advance ideas before they become conventional wisdom. Tracking the forces reshaping business, technology, economics, and geopolitics, we turn emerging patterns into clarity, helping leaders around the world turn early insight into measurable outcomes.

Authors



Sagar Goel

Managing Director & Partner
Singapore



Julie Bedard

Managing Director & Partner;
BCG Institute Fellow
Boston



Matthew Kropp

Managing Director & Senior
Partner; BCG Institute Fellow
San Francisco - Bay Area



Ashley Sim

Project Leader
Singapore



Charikleia Kaffe

Research Director
Athens



ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group bridges the gap between ambition and outcomes for the world's leading companies and organizations. We are built for this era of unprecedented change — bringing strategic clarity rooted in over 60 years of deep domain knowledge, combined with applied AI shaped by our practitioners. BCG works shoulder-to-shoulder with CEOs across industries and geographies to deliver transformative impact at scale: stronger returns, transferred capabilities, and change that sticks. For more information, visit bcg.com.

This publication is made available solely for your personal, informational, and non-commercial use.

You may download and retain a copy for such purposes. Except with the prior written permission of Boston Consulting Group (BCG), you may not use, reproduce, republish, distribute, sell, license, modify, create derivative works from, or otherwise exploit this publication or any portion of its content for any commercial purpose.

Without limiting the foregoing, the publication and its content may not be used for the commercial development, training, testing, evaluation, or improvement of artificial intelligence, machine learning, or other automated systems, or incorporated into any commercial database, dataset, product, or service, without BCG's prior written permission.

All rights not expressly granted are reserved by BCG.

© Boston Consulting Group. All rights reserved.