



THE APPLIED AI INDEX 2026

The Formula for Agentic AI Value

September 2026

By Jessica Apotheker, Vinciane Beauchene, Nicolas de Bellefonds, Dylan Bolden, Sylvain Duranton, Marc Roman Franke, Michael Grebe, Nina Kataeva, Santeri Kirvelä, Djon Kleine, Romain de Laubier, Vladimir Lukic, Amanda Luther, David Martin, Natasha Taylor, Jeff Walters, Rich Lesser, and Christoph Schweizer



Table of Contents

- 03 AI Value Is Becoming Real
- 04 Keys to Unlocking AI Value
- 07 Applied AI Maturity Is Rising
- 09 What Sets Future-Built Companies Apart
- 12 Five Moves to Accelerate AI Value Generation
- 20 A Functional View of AI Maturity
- 21 Playbook: How to Become an Agentic Future-Built Company
- 24 Conclusion
- 25 Appendix: Survey Design Methodology
- 26 About the Authors



AI Value Is Becoming Real

For leadership teams, AI value can seem frustratingly elusive. But that is starting to change. Almost 50% of companies are now generating measurable value from AI, even as the biggest share of that value remains concentrated among a relatively small group of top performers.

Now, agentic AI has arrived with the promise of even greater rewards. It dramatically increases the potential value from AI, and a clear playbook can enable companies to capture that value, regardless of their current level of AI maturity.

The playbook embraces a simple formula: strategic clarity plus applied AI can translate AI ambition into transformative impact. Strategic clarity is about getting the big choices right to strengthen competitive advantage, applied AI means embedding AI into how the business operates day-to-day, and transformative impact is the enterprise value created when those elements come together.



Keys to Unlocking AI Value

When strategic clarity meets applied AI, companies unlock five times as much AI value as when both are lacking. (See [Exhibit 1](#).) BCG's Applied AI Index, based on a survey of more than 1,300 CxOs and senior executives across 20-plus sectors, offers empirical support for that approach. (See the appendix, "[Survey Design Methodology](#).")

Strengthen Competitive Advantage with Strategic Clarity


To succeed in agentic AI transformation, companies need to get the big choices right. The most AI-mature companies in our sample—those that we call *future-built*—show exceptional strategic clarity in doing so, maintaining an average index score of 60 across functions, compared with an average of 36 for laggards.

For example, they concentrate AI into a focused, multiyear program that aligns with the company's biggest opportunities and business priorities, with clear accountability and a mechanism for tracking value to the P&L. In our survey sample, future-built companies are 3.5 times as likely as laggards to set up a dedicated AI transformation program.

Future-built companies know exactly what AI initiatives in their program are worth because 95% of them use clear KPIs or directly track the P&L value from AI. Most other companies are flying blind, with 47% of laggard companies measuring AI value directionally at best.

EXHIBIT 1

When Strategic Clarity Meets Applied AI, Companies Unlock Five Times as Much Value from AI



Source: Applied AI Index 2026 (n = 1,330).

Build the Capabilities for Applied AI at Scale

To get applied AI right, companies need to build capabilities across three pillars:

- Shift to an AI-first operating model with agentic controls.** Future-built companies expect nearly 70% of AI value to come from reshaping and inventing business processes, not just from deploying better AI tools. As these processes become increasingly agentic, companies will need to have the right controls in place to scale them safely. Only 5% have them in place today, although 42% expect agents to act autonomously by 2030.
- Reinvent the workforce around humans and agents.** AI transformation is fundamentally an organizational shift that requires deliberate strategic workforce planning to determine which roles should exist in an agentic world. Such planning is practiced by 55% of future-built companies versus 17% of laggards. It is especially important because AI is likely to reshape work rather than simply eliminate it: 89% of survey respondents expect AI to generate new work, versus just 11% who expect it to primarily replace existing jobs.

- Build the agentic platform, powered by data.** The technology advantage in AI is shifting from basic cloud and data foundations to agent-ready architecture—the data and platforms needed to embed agents in the enterprise’s operations. Future-built companies are already pulling ahead: 95% are undergoing a data transformation, and more than two-thirds are committing to a single enterprise-wide AI platform with a common architecture and control plane, instead of a single-vendor stack.

Our proprietary applied AI capabilities framework extends across more than 40 dimensions to provide a deep, longitudinal view of how maturity is evolving across sectors. (See the sidebar [“AI Leaders Emerge in Every Sector.”](#))

AI Leaders Emerge in Every Sector

In our rankings, technology, banking, and telecom sectors lead the way. (See the exhibit.) But a company's sector does not determine its state of AI maturity. All 20-plus sectors in our sample included at least one future-built organization, evidence that companies in any sector can become future-built. Three insights from our analysis are especially noteworthy:

- **The value gap across all sectors is markedly consistent.** Future-built and scaling companies generated two to three times as much value from AI as laggards, taking into account both cost reduction and revenue growth.
- **The gap between leaders and laggards can be bigger than it looks.** In some sectors, most companies qualify as leaders (either as future-built or scaling organizations). In technology, for example, 70% of companies fall into this category. In contrast, just 49% of biopharma companies, 48% in retail, and 46% in machinery are leaders. As a result,

laggards aren't uniformly behind. Instead, they include a long tail of companies that trail by varying degrees. For company leaders, the strategic question isn't "How does my sector compare with others?" but "Am I among the cohort of leaders?" This reinforces the point that sector doesn't determine overall AI maturity.

- **AI value is concentrated in a few core functions.** Across the 20-plus sectors we analyzed, 54% to 78% of AI value rests in core business functions, with tech/IT and R&D ranking as the first or second driver virtually everywhere. There are some exceptions, however. For example, biopharma companies include marketing among their top three, and banking companies include customer service. These points of differentiation send a clear signal to leadership teams: Don't run a generic AI rollout. Concentrate your investment in the two or three functions that deliver a higher share of value for your specific sector.

Technology, Banking, and Telecom Lead in Maturity

	Technology	Banking	Telecom	Biopharma	Aerospace	Insurance	Retail	Power and utilities	Machinery	Chemicals
Strategic clarity										
Strategic clarity across functions (0–100)	50	45	53	45	39	40	44	38	39	40
Applied AI										
Applied AI maturity (0–100)	59	54	54	51	50	50	49	49	48	46
Share of future-built and scaling (%)	70	60	56	49	52	53	48	46	46	38
Total AI spending (% of annual revenue)	4.1	5.7	3.7	3.3	3.6	3.3	2.7	3.1	3.2	2.6
Companies with agentic AI in production (%) ¹	76	61	62	50	52	48	49	36	55	36
Transformative impact										
Revenue increase through AI (% of revenue)	2.0	1.2	0.6	0.8	1.2	0.6	1.2	1.0	0.9	0.7
Cost reduction through AI (% of cost)	2.1	1.6	1.1	1.0	1.4	1.2	1.5	1.0	0.9	0.9
Top three value-generating functions										
1	R&D	Tech/IT	Customer service	R&D	R&D	Claims processing	Tech/IT	Tech/IT	R&D	Manufacturing
2	Device and software development	Customer service	Tech/IT	Marketing	Tech/IT	Tech/IT	Customer journey	Maintenance	Tech/IT	R&D
3	Customer service	Customer journey	Network infrastructure	Tech/IT	Manufacturing	Customer service	Supply chain	R&D	Manufacturing	Supply chain
Positive difference vs average Average Negative difference vs average										

Source: Applied AI Index 2026 (n = 1,330).

Note: Representative sample drawn from more than 20 sectors. Sectors not shown: payments and fintech; oil and gas; airlines; metals and mining; automotive and mobility; fashion and luxury; consumer products; medtech; public sector; postal and shipping. Each cell number is the sector average; the cell's color indicates that number's distance from the unweighted mean of the 20 sector values for the same metric, in standard deviations, capped at ± 2 standard deviations. Sectors are ordered by maturity (Applied AI Index score); left = higher. N/a = not applicable.

¹Share of firms bringing AI pilots to operational solutions or already realizing full agentic value.



Applied AI Maturity Is Rising

Since we conducted our [2025 analysis](#), global applied AI maturity has increased across all industries. (See [Exhibit 2](#).) The group of stagnating companies has shrunk significantly to just 4.5% of the total, less than one-third of last year's 14%. Companies are moving up to the emerging category (47% of the organizations in our analysis) and to scaling (41%), and as a result they are starting to see value from AI. Scaling organizations—those that have adopted an AI strategy and advanced capabilities and are already generating value—have adopted two to three times as many critical practices as of laggards have, but they still struggle to realize the full potential value from their AI investments. Their problem is that they have too many AI workflows in production and too few that are visible in the P&L.

Overall, 7.5% of the companies in our analysis are future-built, a 50% increase from last year's 5%. These companies, which apply the formula of strategic clarity and applied AI to capture value from AI, form an elite group with exceedingly strong performance. Overall, future-built companies generate 2.4 times as much top-line growth as laggards and 2.3 times as high a three-year TSR. Their EBITDA growth rate is three times as high, too, thanks in part to cost reductions (where they outperform laggards by a factor of 4.3) and to increased revenue (a factor of 6.6). This aligns with other [BCG research](#) that has found concentrations of AI value among a small group of top performers, as measured by industry-adjusted TSR.


Agentic AI is now serving as a catalyst. It accounts for a growing share of overall value and is a key point of differentiation for future-built companies. Agentic's share of AI value climbed from 17% in 2025 to 22% in 2026 and we project that it will reach 39% by 2030. On that path, it will overtake nonagentic AI solutions to become the largest single source of AI value. (See [Exhibit 3](#).)

In that way, agentic AI raises the bar for companies. It doesn't change the playbook for value, in the sense that laggards still need to build foundational capabilities across tech, data, and operating model. But it raises the payoff for companies that become future-built.

Our analysis shows that strategic clarity and applied AI can enable companies to unlock transformative impact, and it points to specific steps that companies can take, consistent with their current level of AI maturity. Organizations in the stagnating and emerging tiers that want to move up to scaling should focus on building the foundations and operating model for AI across the enterprise. Pursuing that strategy entails prioritizing project delivery, redesigning the operating model, developing a portfolio of high-value areas, and upskilling the workforce. Scaling companies that want to become future-built should focus on industrializing their AI initiatives and agentic readiness by embedding AI into products, tracking value, augmenting humans with agentic AI, and giving agents greater autonomy under suitable controls.

EXHIBIT 2

Overall, Almost Half of All Companies Generate Value from AI



Value achieved by AI leaders and laggards

METRICS BASED ON REPORTED FINANCIALS (%)			
	Laggards	Scaling	Future-built
Top-line CAGR	5	6	12
3-year TSR	23	42	54
EBITDA growth	5	5	14


METRICS BASED ON AI-SPECIFIC RESEARCH FINDINGS (%)			
	Laggards	Scaling	Future-built
Realized revenue increase	0.5	1.6	3.3
Realized cost reduction	0.8	1.9	3.4
Confidence in achieving AI RoI targets	50	70	80

Sources: Applied AI Index 2026 (n = 1,330); Financial metrics reported from Capital IQ and Orbis.
Note: Top-line and EBITDA metrics show the compound annual growth rate for 2022 through 2025, while TSR refers to the median 3-year TSR (2022–2025).


EXHIBIT 3

Agentic AI Value Share Will Nearly Double by 2030

RELATIVE SHARES OF AGENTIC AND NONAGENTIC AI (%)



AGENTIC VALUE REALIZATION (% OF COMPANIES)



Sources: Applied AI Index 2026 (n = 1,330).



What Sets Future-Built Companies Apart

How do future-built companies achieve superior performance with AI? Several factors show up in the numbers.

They focus their ambition on bigger bets in the areas that matter most. In deciding where to apply AI, future-built companies look at a bigger share of addressable revenue and costs. They report that 41% of their revenue and costs are addressable through AI. In contrast, laggards report that only 23% of their revenue and 26% of their costs are addressable.

Consistent with our research results in previous years, about 68% of the value from AI is generated in business units, 21% in enabling functions, and 11% in technology/IT. Although that broad pattern is true across industries, we noted some sector-

specific nuances. For example, in technology companies, AI value tends to be more concentrated in R&D and customer-facing work; but in power and utilities companies, it shifts to energy generation, distribution, and maintenance, with more in the tech stack. (See [Exhibit 4](#).) Still, the pattern holds true overall. Leading companies focus their AI efforts on both reducing costs and increasing revenue.


They move fast. The pace of implementation matters, and future-built companies have a clear edge in turning AI investments into P&L impact. Factoring in full process redesigns, they reach impact in about 11 months, compared with 14 months for scaling companies and 17 months for laggards.

EXHIBIT 4


The Value Distribution from AI Differs by Sector

DISTRIBUTION OF AI VALUE POTENTIAL ACROSS FUNCTIONS IN 2026 (%)

Technology



Power, utilities, and renewables



Source: Applied AI Index 2026 (n = 1,330).

Note: Customer service is a core function in technology and an enabling function in power, utilities, and renewables. Because of rounding, not all segments add up to the totals given.

¹Core IT operations (e.g., hardware, software, and IT services). ²Embedded IT systems in customer-facing products or revenue-generating products. ³IT used in internal operations and processes (e.g., automation).

They outspend on AI, not just on IT. A common view holds that future-built companies must have massive IT budgets. But enterprise IT budgets are roughly flat compared with last year, even as overall AI maturity has increased. What has changed is that AI spending has more than tripled—from approximately 1% of revenue in last year's analysis, to 1.7% in [BCG's AI Radar](#) (published in early 2026), to 3.3% today. (See [Exhibit 5](#).) And, critically, a growing share of that AI spending occurs outside enterprise IT budgets, such as on AI talent or governance (including AI transformation programs), or on embedding AI in customer-facing products. Among all surveyed companies, AI spending outside IT is roughly five to six times as high as AI spending within IT (2.8% of revenue, versus 0.5%). (See [Exhibit 6](#).) In other words, the amount that companies invest is important, but where they invest matters, too. For example, AI spending by an automotive company includes AI in the car (software-defined vehicle), AI in the manufacturing process, AI for the foundational IT platform to scale AI, AI governance and program management, and training to enable the workforce to use technologies ranging from machine learning to GenAI, agentic AI, robotics, and others.

At the same time, as the share of AI cost generated outside IT increases, it's important to establish and retain transparency about total AI spending to ensure suitable returns on those


investments. Companies benefit from virtual P&Ls that they create to compare AI benefits against underlying workforce and token costs. This reflects the broader evolution of the CIO's role in the AI-first era: as AI development and spending extend across the enterprise, technology leaders need to maintain visibility over AI economics and link costs directly to business value.

They manage token spending deliberately, rather than by default. Asked how they manage employees' use of paid AI tools, 50% of future-built companies say that they actively encourage adoption to maximize usage—roughly double the 25% of laggards that say the same—and a further 22% have set usage limits or controls to manage spending (versus 16% of laggards). Either approach can work, but the key point is that nearly three-quarters of future-built companies (75%) have made an explicit call on **token spending** and manage it toward an explicit return. Laggards for the most part have not: 50% either grant access without tracking it or monitor spending without setting a limit on it.

Collectively, these points show how future-built companies separate themselves. They spend more on AI, manage token spending deliberately, apply AI to more areas of the company, and implement it in a way that leads to faster value.

EXHIBIT 5


Future-Built Companies Realize Higher Value From AI, Reinvest the Gains, and Expect to Harvest Even More Value by 2030



Source: Applied AI Index 2026 (n = 1,330).
Note: AI spending is calculated as average per sector, then as all-sector average to avoid bias towards sectors with higher share of scaling and future-built companies.
¹Of total revenue base. ²Of total revenue/cost base.

EXHIBIT 6

Most AI Spending Now Sits Outside Enterprise IT



Source: Applied AI Index 2026 (n = 1,330).



Five Moves to Accelerate AI Value Generation

Future-built companies do more than just illustrate AI's value creation possibilities. They provide a roadmap for how other companies can follow their lead. In our analysis, the formula for value consists of three core elements: strategic clarity, applied AI (which has three subcomponents: agentic enterprise, agentic workforce, and agent-ready architecture), and transformative impact. (See [Exhibit 7](#).)


Strategic Clarity—Get the Big Choices Right

Strategic clarity is about getting the big choices right to strengthen competitive advantage in this new era. It requires connecting initiatives across the enterprise, seeing how technology, functions and people fit together. Strategic clarity also calls for having the foresight to see where advantage is about to shift, anticipating sector disruption before the market does; having the conviction to concentrate on a few high-value plays that entail reshaping processes end-to-end, rather than spreading effort across many small deployments; and having the discipline to track the outcomes.

Future-built companies lead in all of these areas. They bundle their AI efforts into a single program that the C-suite owns, with a multiyear timeline. Future-built companies are 3.5 times as likely as laggards to commit to a multiyear AI effort (61% versus 17%), instead of funding short, stop-and-start initiatives.

Future-built companies have more direct accountability for AI programs, as their C-suite leaders are 20 percentage points more likely than laggard leaders to hold the AI budget, providing senior sponsorship for the effort rather than delegating it downward. This aligns with BCG's AI Radar 2026 analysis, where nearly three-quarters of CEOs already act as their company's chief decision maker on AI. The CIO or CTO owns the transformation budget at more than half of all future-built companies, and increasingly acts that role as a chief integration officer who embeds AI into the business and moves pilots into the P&L. Companies that have no single AI transformation budget owner are about 20% less confident of hitting their 2027 targets than those that do. Accountability is a given among future-built companies: 100% of them have clear ownership of the AI agenda.

Five Moves to Accelerate AI Value Generation



Source: Applied AI Index 2026 (n = 1,330).

Value tracking is paramount. Future-built companies track AI value directly into the P&L, which gives them a clear sense of what's working and how well, in addition to providing a funding source for the next wave of AI investment. Among this group, 95% measure value on the P&L or a proxy. (See [Exhibit 8](#).) That rigor is missing at most companies: A recent [BCG analysis](#) found that while nearly nine out of ten CEOs already see cost or revenue benefits from AI, only 14% have clearly defined the P&L impact of every initiative.

These are the same principles that companies apply to other large-scale business transformations: building a business case, prioritizing among multiple options, linking initiatives to business objectives, and tracking performance in financial terms. As applied to AI, strategic clarity helps companies gain traction. On average, future-built organizations have 20 AI workflows implemented at scale (versus an average of just 4 for laggards). And when they set targets, they are over 50% more confident of hitting them.

For example, a global investment bank committed to running AI as a single funded, company-wide program and proving value at scale. The leadership team developed an AI strategy, identified a small number of high-value use cases (fewer than ten), and used a top-down and bottom-up assessment to define targets for end-to-end processes. Crucially, the company also set up a program office reporting to senior leadership to track execution and value straight to the P&L. Using this approach, the company aims to achieve financial upside of \$500 million by 2027 and \$1 billion by 2030.

Applied AI: Shift to an AI-First Operating Model with Agentic Controls

Agentic is where tomorrow's AI value sits, and leaders are already collecting it through advanced, agent-era capabilities, human-AI augmentation, and AI embedded in production, operations, and research ideation. Among future-built companies, 44% say that they already realize full value from agentic today, in contrast to 2% of laggards. And future-built companies expect nearly 70% of their AI value to come from reshaping and inventing business processes, not from deploying better AI tools (as most laggards do).


But scaling agentic AI workflows bears a risk. The larger the number of agentic workflows an organization brings into production, the greater the risk it incurs of agent sprawl—duplicated agents, conflicting policies or permissions, unclear ownership, or monitoring gaps, for example. But the deeper shift involves agents' capacity to combine data, tools, permissions, memory, and other agents across workflows, potentially causing risks to emerge dynamically and compound at machine speed. Therefore, beyond putting agents to work, applied AI requires adapting the enterprise risk and control operating model and its underlying technology architecture to orchestrate and control increasingly autonomous agentic capabilities across the enterprise. This demands changes to both the risk operating model and the underlying technology and control architecture.

EXHIBIT 8

Future-Built Firms Rigorously Track AI Value and Capture Three Times Higher Returns Than Laggards Do


Nearly half of future-built companies track the return on AI directly to the P&L

FORMS OF AI VALUE TRACKING USED (%)



Measuring direct P&L impact drives up to three times higher value

TOTAL AI VALUE (%)



Source: Applied AI Index 2026 (n = 1,330).

Companies that combine central risk and control guardrails with federated implementation across business and technology functions bring three times as many workflows to scale as companies that rely on fully centralized steering.

We have identified six essential controls:

- **Memory and Context.** Configure how agents use, store, and recall memory and context. Memory can compound learning across agents, but it can also propagate errors, sensitive information, or manipulated context across workflows.
- **Scope and Oversight.** Define each agent's purpose, scope, autonomy, and required human oversight. Authority should be bounded not only by the systems an agent can access, but by the actions it is permitted to take, for which purposes, and in what context.
- **Evaluation and Rollback.** Establish evaluation criteria, release gates, continuous monitoring, and rollback mechanisms. Controls should detect drift, anomalous behavior, and emerging risks after deployment, not merely determine whether an agent was safe at launch.

- **Tools and Integration.** Implement standardized tool integrations, APIs, and agent registries (such as the Model Context Protocol registry). Shared standards provide visibility into which agents are operating, what they can access, and how they interact.
- **Ownership.** Assign ownership across business, technology, and risk functions. Every material agent should have a named owner who is accountable for its use and outcomes, including having clear authority over who can change, suspend, or stop it.
- **Security.** Set runtime security controls, audit trails, and cost guardrails. Controls should constrain agent actions as they occur and preserve sufficient evidence to reconstruct what an agent did, under whose authority, and using which data and tools, as well as identifying which controls the agent triggered.

Of the full set of companies in our analysis, 42% expect agents to act autonomously by 2030—for example, making decisions without human approval— yet just 5% have relevant agent controls in place. Effective risk controls serve as an enabler, not a brake: putting the right controls in place enterprise-wide roughly triples the value that agents create. (See [Exhibit 9](#).)

For example, a large global business process outsourcing provider set out to reinvent its value offer to clients, shifting from selling labor and time to selling AI-enabled outcomes. To execute that change, the company developed an AI-first operating model. The company built the model as a ladder for increasing AI autonomy, advancing from insights (in which AI diagnoses the work) to copilot (AI assists employees) to autopilot (AI agents run tasks end-to-end). Autonomy is set at the process level, not at the level of individual tools. The process defines which steps exist, where decisions occur, and when a human must be involved. To implement the change, the company created a new operating model with a center of

excellence consisting of about 50 people and 100 local deployment teams, embedding new ways of working across more than 5,000 frontline staff in operations. The new model delivered tangible impact, with pilots improving conversion rates by up to 40%. Following this work and building on the experience and capabilities gained through it, the company has developed and is currently rolling out an agentic platform to take process automation to the next level for its clients.


Boards have a role to play in governance, too. As AI becomes more important to the business, boards should clearly articulate and periodically review the organization’s AI risk appetite, and management should maintain a comprehensive inventory of agents across the enterprise. Boards should have transparency into the most consequential agents—their purpose, authority, and key dependencies, and the identity of those who can intervene or stop them—and should review critical AI providers, single points of failure, contingency plans, and major incidents or control gaps.

EXHIBIT 9


Companies with All Six AI Controls in Place Enterprise-Wide Generate Three Times as Much Agentic AI Value as Those with Just One Control

Controlling agents does not diminish their value but triples it

VALUE FROM AGENTIC AI (% OF REVENUE)



Controlling for memory and context handling drives value



Six agent governance controls

- 1. Defining memory and context handling
- 2. Defining agent scope and oversight
- 3. Building evaluation and rollback gates
- 4. Standardizing tool and API integration
- 5. Assigning cross-functional ownership
- 6. Implementing security, audit, and cost guardrails

Source: Applied AI Index 2026 (n = 1,330).

Applied AI: Reinvent the Workforce Around Humans and Agents

The next element of agentic readiness involves reinventing the workforce. Laggards take a baseline view of agentic, focusing on tools that help existing employees execute current processes a bit more efficiently. Future-built companies redesign the organization around agentic AI, reshaping processes and workflows to capitalize on agents. A recent BCG analysis of [AI's impact on workforces](#) quantified the difference: 83% of companies with strong strategic clarity and extensive access to tools report impact from AI.

This element of readiness calls for an [organizational shift](#), not a technology upgrade. Companies must make an explicit strategic choice about where agents should take on work and where human judgment, relationships, accountability, or other forms of contribution should remain a source of differentiation. Rather than simply agentifying today's organization, companies need to work backward from strategy to determine what to delegate to agents, where humans should remain in the loop or in the lead, and what roles and skills that implies for the human workforce.

That degree of strategic workforce planning capability is a key differentiator. Strategic workforce planning happens at 55% of future-built companies and just 17% of laggards, a 38-percentage-point gap. The result? Future-built organizations already have 13% of their employees working in dedicated AI roles, a higher level than laggards expect to reach by 2030.

Looming workforce changes make CHROs the [architects of a hybrid workforce](#) of humans and agents. These leaders must operate at two speeds, addressing HR fundamentals while also redesigning roles and embedding agents in the companies ways of working.

That work is urgent because of where the impact from AI is likely to land. Most companies (79%) expect technology cost to rise through 2030 and labor cost to fall (49%). But the bigger story is that AI creates new work rather than simply eliminating old work. Fully 89% of companies expect AI to generate new work, versus 11% that expect it to replace existing jobs for the most part. Overall, companies expect a workforce reduction of roughly 10% to 15% from AI. That change will be concentrated primarily in middle layers of the organization—middle management and frontline leadership—where roughly half of the companies in our sample expect a workforce reduction of more than 10%. That is a far higher percentage of job loss than in senior, business-unit, or functional leadership, where only about one-fourth expect the number of roles to contract by more than 10%.

The implication is that leaders will not fully automate the workforce. Instead, they will make choices about where AI can generate a step-change in productivity and where human interventions should be preserved, elevated, or reinvented to take advantage of the strategic value they create.

Today, more than five times as many future-built companies as laggards (38% versus 7%) say that AI fully executes tasks independently, and 33% of future-built companies say that AI augments expert decisions and improves output quality. Already, future-built users are far more likely than others to let an agent run an entire task from start to finish with no human step in between (31% versus 1%), and they are three times as likely to use AI for substantial, multistep work as for quick one-off questions (78% versus 25%). But as their capabilities rise, agents will gain more decision authority: 66% of future-built companies expect agents to run with bounded or full autonomy by 2030, compared with only 32% of laggards.

A global cloud and technology company noted high AI adoption among members of one of its 1,000-person engineering teams, but found that the team's productivity gains (just 10%) were underwhelming. The issue wasn't the tools, however, but the behaviors. Engineers used the tools infrequently and mostly for basic tasks. To improve, the company ran tailored, six-week sprints that focused on coaching the engineers to use AI efficiently at work. Specific metrics tracked the depth of AI usage (shifting to end-to-end feature development), breadth (from coding to the full software development life cycle), and frequency (number of AI interactions per day). Through this approach, the engineering team's number of AI power users increased from 5% of engineers to 50%, and those engineers are now about 80% more productive than before the transition.

Similarly, a large US regional health insurer launched an AI transformation. It quickly identified customer service as a lighthouse project to reduce costs and build the company's data platform. About 80% of incoming calls were benefits- and coverage-related, with an average call time of 16 minutes (nearly double the industry average). Customer-service representatives had to search more than 20 documents per call, on average, and customer satisfaction levels were low.

In response, the company deployed three agentic AI solutions, all powered by a new platform that linked internal documents, structured data, and live APIs. The implementation focused on speed and business outcomes, with the first large-scale pilot live within just four months. Agentic chatbots can now handle about 60% of all incoming call types, reducing the average call handling time by 35% and freeing human employees to handle more complex requests. The company projects that this approach will reduce call center costs by more than \$60 million during the first wave of implementation, with broader AI use cases projected to hit \$200 million in run-rate savings.

Companies need to define trusted data in a way that users can understand and implement consistently across the organization. They also need to establish a shared understanding of semantics and operating context, and redesign information flows and controls for agents.

Applied AI: Build an Agentic Platform Powered by Data

The technology gap is no longer about basic data or cloud, which are now widely held. Instead, it has moved up the stack into the data, enterprise intelligence, and platform layers, which let agents understand context, take governed action, and run reliably at scale. That shift has implications for how companies manage the economics, complexity, and vendor dependencies associated with agentic.

AI foundations are becoming table stakes, though they are not yet universal. In 2025, more than 50% of future-built companies had a single enterprise-wide data model, versus about 4% of the least mature companies. In 2026, basic data structuring and access controls are common across maturity levels, ranging from 58% to 72%. Beyond this point of difference, the sharper divide relates to whether companies can rely on the semantics of their data—in particular, whether the data is up to date, consistently defined, trustworthy, and managed with clear ownership.

The new advantage is agent-ready data. Future-built companies are 1.4 times as likely as laggards to give agents live access to company data so that answers stay grounded in current information, and they are 1.7 times as likely to have explicit, enforceable controls over what each agent can see and do.

That advantage extends to how companies organize their data and platform foundations. Among future-built companies, 95% run a data transformation (versus 74% for laggards), and 70% run it enterprise-wide (versus 43% for laggards). Future-built companies are also twice as likely to make agent enablement an explicit objective of the transformation. Although enterprise-wide programs are a stronger vehicle for making data agent-ready, readiness requires more than adding a new technology layer. Companies need to define trusted data in a way that users can understand and implement consistently across the organization. They also need to establish a shared understanding of semantics and operating context, and redesign information flows and controls for agents.

The same pattern shows up in AI platforms. More than two-thirds of future-built companies concentrate their AI effort on a company-wide platform target, versus about one-fifth of laggards. The goal is a common architecture and control plane, not a single-vendor stack. The goal is to build shared capabilities for security, access, orchestration, observability, evaluation, and agent registration once and then reuse it across the enterprise, while allowing teams to retain flexibility across models and providers.

These findings highlight the ways in which future-built companies are ahead on tech and data today. As agents become embedded more deeply in enterprise workflows, the architecture agenda will broaden in three ways:

- **Own and protect the enterprise cortex.** As AI systems become part of how companies make decisions and execute work, traditional technology lock-in can evolve into cognitive lock-in: dependence not only on a vendor's platform, but also on the reasoning processes and operating context embedded within it. Companies should therefore keep their enterprise cortex—proprietary knowledge and IP, shared ontology, business rules, decision logic, and operating context—in a governed enterprise intelligence layer that is under their exclusive control, and they should keep it technically separable and portable, not merely contractually owned. Around that core, they should emphasize flexibility across models and providers, supported by common routing, evaluation, and fallback logic, so that they can use the best technology for different tasks without surrendering the company-specific knowledge that makes the AI useful.
- **Make AI economics visible and manageable.** Agentic systems introduce a new variable cost into the technology stack. This cost depends on the model used, the amount of context supplied, the reasoning effort required, and the number of steps an agent takes to complete a task. These criteria for gauging make traditional infrastructure cost management insufficient. Increasingly, the relevant unit is not cost per token, but cost per successful outcome at the required levels of quality, latency, and risk—including the human effort needed to review and operate the workflow. Architectures therefore must provide workflow-level visibility into consumption and value, as well as levers such as using deterministic software where a model is unnecessary, routing tasks to the smallest model that meets the requirement, and setting context reuse, caching, and stopping rules. This complements the token-management discipline discussed earlier: leaders can manage AI spending deliberately only if the architecture makes it transparent. In many instances, when assessed in this way, the cheapest model is not the right answer. A more capable model that completes a task on the first attempt can cost less in the long run than a less effective one that retries, escalates, or produces work that a human has to redo. Economics thus becomes a question of architecture and accountability rather than of procurement: every material workflow needs a named owner, an intended outcome, and explicit budget guardrails.

- **Continuously renew the technology landscape.**

Agentic AI can create complexity as quickly as it creates capability, adding new agents, interfaces, dependencies, machine-generated code, and controls on top of legacy systems. To resist accumulating complexity, companies need to treat modernization as a continuous discipline rather than as a periodic program. AI can help accelerate legacy migration, ERP modernization, documentation, and code, but only within disciplined architecture, engineering, and life cycle controls. The goal is to ensure that the rapid AI experimentation happening now doesn't rigidify into inflexible AI systems in the future.

For example, a leading European bank aimed to improve its clients' customer experience, increase productivity, and accelerate decision making with AI. But as initiatives multiplied across business units and functions, AI adoption became increasingly decentralized, and technology and data complexity created challenges to scaling. The company responded by prioritizing approximately ten key AI reshape initiatives, working with shared, agent-ready building blocks and common standards and guardrails that teams could reuse over time, rather than rebuilding for each new initiative. Teams accomplished this through a shared agentic AI platform and reference architecture aligned across stakeholders from business, technology, data, and security functions. Consolidating shared capabilities in this way improved the bank's ability to execute the initiatives and shortened time-to-market. Externally, the bank projects that it will generate annual AI- and technology-driven productivity benefits of up to 10% of its cost base within three years.

In another example, a leading global resources company had struggled to update its master data, a critical element in operating industrial assets safely and efficiently. The company fielded more than 10,000 change requests each year, but the process was manual and slow. A six-month backlog meant that some employees gave up, degrading data quality over time. To improve, the company built an AI agent to streamline the process. Briefly, an employee describes the change in plain language; then the agent pulls the details from the ERP, checks them, and generates a ready-to-submit request. The process is faster, easier, and more accurate, with just 3 steps per request (down from 68). Employees now fix data errors when they spot them, which improves data quality at the source. The company projects potential savings of millions of dollars a year in future automation phases.

Transformative Impact: Scale Functional Workflows

Sum up these efforts to combine strategic clarity with applied AI, and it becomes clear how future-built companies generate transformative impact. They win with depth and speed. Future-built companies have fewer workflows stuck in planning (just 4, versus 10 for laggards), and they run five times as many at scale (20 versus 4), from a portfolio that is only 1.5 times as large.


Future-built companies emphasize process redesign. Commercial AI tools are available for companies to deploy across existing tasks and workflows, but a growing body of BCG research indicates that companies unlock far more AI value by reshaping and inventing business processes and workstreams end-to-end. Predictably, this is where future-built companies focus their energy. Of their total pool of AI initiatives, 62% involve reshaping and inventing rather than deploying AI solutions (versus 57% for scaling companies and 44% for laggards). (See [Exhibit 10](#).) That emphasis will become even more pronounced in the future, as future-built companies project 5-percentage-point growth in reshape and invent initiatives by 2030.

This doesn't mean that AI leaders invest everywhere. To the contrary, future-built companies concentrate firepower where it counts: they scale a few high-value functions (such as R&D, maintenance, and customer service), with the widest gap in R&D at 61 percentage points. In contrast, laggards spread their resources thin across every function.

The payoff for future-built companies shows up in every industry. Wherever companies take a core workflow to scale (such as in credit decisioning in banking, claims-fraud detection in insurance, and demand forecasting in consumer goods), it delivers double-digit gains in productivity, revenue, and cost reduction.

EXHIBIT 10

Future-Built Companies Have Fewer Initiatives Stuck in Planning, Reinvent More, and Scale More



Source: Applied AI Index 2026 (n = 1,330).



A Functional View of AI Maturity

As we noted at the beginning of this report, creating value from AI requires both strategic clarity and applied AI. Companies that get both of these elements right see about five times the impact on AI value, and roughly three-quarters (74%) of that group are future-built. That same principle applies at a functional level. We looked at the core dimensions of AI value across 14 functions, to determine what separates future-built organizations from less mature competitors.

To be clear, the real value of work with AI does not arise within individual functions but emerges from end-to-end reshaping of workflows and processes that span multiple units and functions. Nevertheless, we see individual functions as a useful starting point for low-maturity-level companies to start building AI capabilities. Our analysis yielded several insights:

- Strategic clarity is the strongest of the three dimensions, and it sharply separates future-built companies from laggards. Across all 14 functions, future-built companies average an index score of 60 on strategic clarity, compared with 36 for laggards—a 24-point gap. Looking across companies overall, the highest average level of strategic clarity is reported for the IT function (64). Customer-facing, R&D, operations, and enabling functions all average between 35 and 45, signaling that many companies have the intent to upgrade their strategic clarity but not yet a clear direction for achieving that goal.
- Scale matters more than pilots. Within IT, future-built companies take a higher share of AI workflows into scale and deployment than laggards do (39% versus 7.5%) and generate more value from each (a difference of 6 to 9 percentage points). That difference is consistent across other functions, including marketing and R&D.
- Applied AI scores are higher than strategic clarity scores for operations functions (procurement, supply chain, and maintenance) and legal.

- Transformative impact shows the lowest scores, between 12 and 29, and does not exceed applied AI in a single function. Of all the functions we analyzed, legal, HR, and finance show very low scores across all three dimensions.
- Future-built companies score highest in three core functions: R&D, customer experience, and operations. This result demonstrates their focus on areas that deliver the most value. The capability gap over laggards is widest in R&D (45 points), followed by customer experience (36) and operations (33); it is narrower in enablers (24). Laggards, meanwhile, score flat across every function.

Consider a global oil and gas player that wanted to modernize its supply chain to evolve beyond slow, disconnected planning. It developed an AI-enabled digital twin—a live, virtual model of the company’s end-to-end supply chain—that included a control tower to track real-time performance, and AI-driven planning and simulation tools. More than 200 users now work with the solution daily, reducing inventory by 10% to 15% and supply-chain costs by up to 5%. Planning and purchasing processes are 80% more efficient. The project also led to greater transparency in the company’s carbon footprint.

A global automaker wanted to improve how it set prices for its cars. Like many automakers, it faced challenges in how to allocate discounts to spur sales and optimize margin across dozens of model variants and global markets. Its solution was to build an AI suite to manage the process end-to-end—forecasting demand, balancing the portfolio mix, and optimizing discounts in and across markets. The company scaled the suite to all of its top markets and achieved a boost of up to 10% in discount efficiency.




Playbook: How to Become an Agentic Future-Built Company

There is no shortcut to becoming a future-built company, and no single measure is right for all organizations. That said, one clear principle applies very broadly: becoming future-built is primarily an organizational shift, not only a technical one. In BCG's 10-20-70 model, companies focus just 10% of their AI efforts on the algorithms and 20% on the technology and data. They focus the remaining 70% on changes to people, organization, and processes.

Future-built companies track this weighting closely, putting nearly three-fourths of their top priorities in the people, organization, and processes category, with 19% coming from technology and 9% from algorithms. Their standout capabilities are human-AI augmentation, AI productization, research and ideation, AI operations, AI delivery, and the agentic AI ecosystem. (See [Exhibit 11](#).)

Leading Companies Focus on the People, Organization, and Processes Elements of Implementing AI



Source: Applied AI Index 2026 (n = 1,330).

¹Based on relative importance of capabilities (showing 5%+ relative importance only).

Beyond that, the playbook varies by a company's starting point. Organizations in the stagnating and emerging tiers that want to move up to scaling should focus on building the foundations for AI across the enterprise. Scaling companies that want to become future-built should focus on industrializing and scaling their AI initiatives. (See [Exhibit 12](#).) Here are two examples:

- A Structured Approach to Implementing AI.** A multinational consumer company wanted to reinvent core processes across functions using AI. To do so, it followed a clear sequence: map the work and tacit knowledge in detail; redesign each process, balancing AI and human tasks; deploy fast with low-code AI tools that embed enterprise knowledge; and scale the best ideas to achieve P&L impact. The journey began in marketing, where the company rolled out a suite of GenAI tools across the entire function, automating tasks that accounted for 30% to 40% of employees' time

and reinventing workflows end-to-end. As a result, human worktime spent on routine activities dropped by up to 70%, while output quality improved twofold. The company then applied the same approach across R&D and procurement, supported throughout by upskilling and value-tracking.

- Rebuilding Workflows for AI, Not Simply Automating Them.** A global insurer set out to transform its core business with agentic AI across business units and group functions. Rather than adjusting existing workflows at the margins, it rebuilt those workflows from the ground up as agentic-first, and it built a quantified business case and implementation roadmap for each one. The company then created a minimum viable product and scaled deployment, which it coordinated through a central office. The insurer expects to realize several hundred million dollars in total value from the group-wide transformation.

The Right Priorities for Agentic Success Depend on a Company's Starting Point

		Move up to scaling Build the foundations	Move up to future-built Industrialize and scale
Strategic clarity		• Bundle AI efforts into a single, multiyear program • Put clear C-suite ownership and value measurement in place	• Extend P&L value tracking beyond proxy KPIs • Build strategic clarity for AI in functions and end-to-end workflows
Applied AI	Agentic enterprise	• Keep humans-in-the-loop as the final decision maker and owner • Establish agent guardrails, audits, security, and escalation routines	• Grant agents real autonomy and ownership in bounded domains • Enhance agent decision rights across end-to-end workflows
	Agentic workforce	• Build talent, dedicate people to AI, and protect AI upskilling • Inject ambition into workforce planning and AI role reshaping	• Redesign work around agents and humans • Introduce bounded agent autonomy and decision rights
	Agent-ready architecture	• Strengthen the enterprise AI platform and data foundation • Transform data enterprise-wide, including core systems and ERP	• Build an agentic architecture, ready to reshape functions at scale • Rationalize vendor landscape, and tie spending to business outcomes
Transformative impact		• Concentrate investment on fewer, higher-value workflows • Build each workflow end-to-end across the functions it spans	• Shift toward reshape/invent; reinvest gains in AI capabilities • Optimize for agentic readiness, deployment speed, and scale
		Laggard	Scaling
			Future-built

Source: Applied AI Index 2026 (n = 1,330).

Note: ERP = enterprise resource planning.




Conclusion

Companies at lower AI maturity levels often think that the goal is to simply deploy more tools, fund more pilots, or chase the newest technical breakthrough. As our data shows, that approach won't work. Becoming an agentic, future-built organization requires a disciplined agenda anchored in strategic clarity and applied AI. The organization must embed programs into the way work gets done, govern them with confidence, and measure them by the value they create. Agentic AI raises the stakes because it transforms AI from an assistant at the edge of the enterprise into a participant in its core operating model.

The evidence is encouraging: the path to becoming future-built is open to every company. Sector, legacy, and starting point do not determine destiny; leadership choices do. Companies that concentrate their efforts, empower their people, build agent-ready foundations, and scale the workflows that matter most can turn AI from experimentation into enduring advantage. In that sense, agentic AI is more than the next wave of technology. It is an invitation for leaders to build organizations that are faster, smarter, more adaptive, and more ambitious about the future that they intend to create.

Appendix: Survey Design Methodology



We designed our 2026 Applied AI Index to focus on the AI capabilities necessary to support strategic objectives, deliver significant business value, and identify and capitalize on new market possibilities.

Compared with previous analyses, we broadened the lens in several ways. As AI progresses rapidly, a company’s position on agentic AI has become a meaningful marker of maturity, so we added a dedicated assessment of it across four dimensions: the workforce, the operating model, the technology and data foundation, and functional workflows where value is realized. We also assessed the clarity of each organization’s AI strategy—specifically, whether it defines ambition and translates it into an active roadmap, how it tracks value, and where strategic accountability sits within the leadership team.

We built our comprehensive AI maturity score on 41 foundational capabilities, each measured along four clearly defined maturity stages: stagnating, emerging, scaling, and future-built. Next, we sorted the weighted scores into four numerical ranges:

Stagnating	Emerging	Scaling	Future-built
0–25	>25–50	>50–75	>75–100

In our research, we surveyed 1,330 CxOs and senior executives who are AI decision makers in companies across 20-plus sectors, asking them to estimate their organization’s AI maturity along all 41 capabilities, including outcomes in customer experience and operations, in response to sector-specific questions. We drew insights into AI maturity and value from respondents’ self-reported data and from external metrics linked to that data. Results reflect the business area that the respondents know best and are not always applicable to the full company. Unless explicitly stated as realized (for example, as reflected in the P&L by the end of 2025), reported numbers represent expected future impact. Although we took all available precautions to achieve optimal data quality, survey responses may be subject to perception bias.

Respondents came from more than 60 countries in Asia-Pacific, Europe, and North America and from nine industries: consumer; energy; financial institutions; health care; industrial goods; insurance; public sector; technology, media, and telecommunications; and travel, cities, and infrastructure. All respondents held leadership roles in their organization: 24% were C-level executives, 58% were executives reporting directly to C-level, and 15% led a function or division. Over 50% of participating companies reported annual revenue above \$5 billion.

We sourced the report’s financial metrics (for example, TSR and EBIT) from CapitalIQ and Orbis. Of course, overall company performance has multiple drivers beyond AI. We applied multivariate statistical methods such as random forest, logistic regression, and structural equation modeling (partial least-squares regression) to calculate regression weights.

In order to compare our 2026 results with data from our 2025 Build for the Future study, we identified 33 of the 41 capabilities as measured consistently across both years. We assessed nine of these against a higher bar in 2026, reflecting how quickly expectations have risen. For year-on-year comparisons, we leverage only the 33 comparable capabilities.

About the Authors

Jessica Apotheker

CMO, Managing Director & Senior Partner
Paris
apotheker.jessica@bcg.com

Vinciane Beauchene

Managing Director & Partner
Paris
beauchene.vinciane@bcg.com

Nicolas De Bellefonds

Managing Director & Senior Partner
Paris
debellefonds.nicolas@bcg.com

Dylan Bolden

Managing Director & Senior Partner;
Global Chair, Functional Practices
Dallas
bolden.dylan@bcg.com

Sylvain Duranton

Managing Director & Senior Partner;
Global Leader, BCG X
Paris
duranton.sylvain@bcg.com

Marc Roman Franke

Partner & Associate Director,
AI and digital transformation
Berlin
franke.marcroman@bcg.com

Michael Grebe

Managing Director & Senior Partner
Munich
grebe.michael@bcg.com

Nina Kataeva

Managing Director & Partner
Zurich
kataeva.nina@bcg.com

Santeri Kirvelä

Managing Director & Senior Partner
Helsinki
kirvela.santeri@bcg.com

Djon Kleine

Managing Director & Partner
San Francisco Bay Area
kleine.djon@bcg.com

Romain de Laubier

Managing Director & Senior Partner
Singapore
delaubier.romain@bcg.com

Vladimir Lukic

Managing Director & Senior Partner;
Global Leader, Tech and Digital Advantage
Boston
lukic.vladimir@bcg.com

Amanda Luther

Managing Director & Senior Partner
Austin
luther.amanda@bcg.com

David Martin

Managing Director & Senior Partner;
Global Lead, People & Organization
Dallas
martin.david@bcg.com

Natasha Taylor

Managing Director & Senior Partner
Denver
taylor.natasha@bcg.com

Jeff Walters

Managing Director & Senior Partner
Singapore
walters.jeff@bcg.com

Rich Lesser

Global Chair
New York
lesser.rich@bcg.com

Christoph Schweizer

Chief Executive Officer
Munich
schweizer.christoph@bcg.com

Acknowledgments

The authors thank Florence Adida, Ana Chutkerashvili, Giuseppe Colombo, Soham Daruwala, Gayatri Dwivedi, Denis Guselnikov, Michael Leyh, Ekaterina Lubenskaia, Daniel McCord, Clemens Nopp, Alessia Romero, Juliana Ruseva, and Csilla Szalay for their valuable contributions to this publication.

In addition, the authors thank Matteo Coppola, Stephen Edison, Adam Job, Anne Kleppe, Matt Marchingo, and Ulrich Pidun for their guidance in creating the publication.



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.

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 [LinkedIn](#), [Instagram](#), [Facebook](#), and [X](#).

