

2026 AI Business Predictions

AI

Generative AI

[Copy link](#)

Only a few companies are realizing extraordinary value from AI today, things like surging top-line growth and significant valuation premiums. Many others are also experiencing measurable ROI, but their outcomes are often modest—some efficiency gains here, some capacity growth there, and general but unmeasurable productivity boosts. These results can pay for themselves and then some. But they don't add up to transformation.

The picture's starting to shift. It's still hard to use AI to drive transformative value, and the technology continues to evolve at speed. That's not changing. But what's new is this: Success is becoming visible. We can now see what it looks like to use AI to build a leading-edge operating or business model. From mature systems to emerging tools like AI agents, examples of impact are multiplying—across strategy, operations, workforce, trust, tech stacks, and sustainability.

Companies now have enough evidence to build benchmarks, measure performance, and identify levers to accelerate value creation in both the business and functions like finance and tax, so they can become nimbler, faster-growing organizations.

[Subscribe here](#)**Our insights. Your choices.**

Why, then, has this kind of success—the kind that drives revenue growth and opens up new markets—been concentrated in so few? Too often, organizations spread their efforts thin, placing small sporadic bets. Because AI feels easy to use, early wins can mask deeper challenges.

But real results take precision in picking a few spots where AI can deliver wholesale transformation in ways that matter for the business, then executing with steady discipline that starts with senior leadership. After success in your priority areas, the rest of the company can follow.

We've seen that discipline pay off.

Across industries and regions, our work with clients reveals how deliberate and sustained efforts can turn AI experiments into engines of growth and innovation.

Our own AI transformation at PwC has given us this perspective from the inside too. Combined with nearly a decade of research through our executive surveys and annual AI predictions, we've built a clear view of what drives success—and what holds it back.

Our forecasts are grounded in real experience and focused on practical impact—so you can take confident steps to turn AI ambition into transformative business value in 2026 and beyond.

Explore the predictions

1. **The disciplined march to value begins**
2. **Proof points and real-world benchmarks set the pace for agentic AI**
3. **Rise of the AI generalist: A new workforce emerges**
4. **Responsible AI moves from talk to traction**
5. **From vibe to value: Orchestration that accelerates impact**
6. **The demand for business returns drives AI for sustainability**



1. The disciplined march to value begins

With AI, many companies make an understandable mistake. Instead of leadership calling the shots with a top-down program, they take a ground-up approach, crowdsourcing initiatives that they then try to shape into something like a strategy. The result: projects that may not match enterprise priorities, are rarely executed with precision, and almost never lead to transformation.


Crowdsourcing AI efforts can create impressive adoption numbers, but it seldom produces meaningful business outcomes.

In 2026, we expect more companies to follow the lead of AI front-runners, adopting an enterprise-wide strategy centered on a top-down program. Senior leadership picks the spots for focused AI investments, looking for a few key workflows or business processes where payoffs from AI can be big. Leadership then applies the right “enterprise muscle”—talent, technical resources, and change management. Often, this program is executed through a centralized hub that we call an “**AI studio**.” It brings together reusable tech components, frameworks for assessing use cases, a sandbox for testing, deployment protocols, and skilled people. This structure links business goals to AI capabilities so you can surface high-ROI opportunities.

Agentic AI looks to play an increasingly important role. AI agents can go beyond analysis and automate parts of complex, high-value workflows. Especially ripe areas for agents include demand sensing and forecasting, hyper-personalization, product design, and functions like finance, HR, IT, tax, and internal audit.

What to do now

- **Have leadership pick the spots.** Top leadership picks a few areas for focused AI investments, often where business priorities, evidence of AI's value, and availability of talent and data align. Then, leadership focuses on execution.
- **Go narrow and deep.** After you identify the right high-value workflow, aim for wholesale transformation. Instead of cutting a few steps, rethink the workflow, which an AI-first approach may turn into a single step. That often starts by asking not how AI can fit into a workflow but how it can create a new one.
- **Send your A-team.** Assign top talent to those areas where you've decided to focus on AI. These business leads can both work with company leadership to define target outcomes and drive progress with process owners and AI specialists.



2. Proof points and real-world benchmarks set the pace for agentic AI

There is—rightfully—little patience for “exploratory” AI investments. Each dollar spent should fuel measurable outcomes that accelerate business value. But many agentic deployments last year didn’t deliver much value. If you looked under the hood, many weren’t using agents in ways that matter. If you asked for a demo—to see an agent at work delivering value—you often couldn’t get it because there wasn’t anything to see.

We expect that to change in 2026. We now know what good agentic AI looks like. It has proof points like benchmarks that track value that matters to the business, whether that’s financial (P&L impact), operational (market differentiation), or related to workforce and trust. Instead of siloed efforts, it has a centralized platform for deployment and oversight that draws on a shared library of agents, templates, and tools. Before each deployment, agents are tested, with flaws corrected and working demos created for future users to try—so they can offer feedback and start to trust what agents can do.

Agents are rolled out as part of all-new workflows, with clearly-articulated steps for human initiative, review, and oversight—and with people who have the training and incentives to work with agents and provide that oversight.

Built-in monitoring also includes different agents checking each other’s work, and for higher-risk scenarios, these agents come from different model providers.

Since agents can automatically document their decisions and actions, continuous monitoring can be highly effective in tracking adoption and performance, fixing errors quickly, and building stakeholder trust.

Agents today are imperfect, but new technologies generally are. Now that companies know how to proceed—with focused, centralized implementation guided by real-world benchmarks—2026 could be the year when agents shine.

What to do now

- **Create metrics that drive outcomes.** For AI that delivers the value that your business wants, set concrete outcomes for it to deliver, select suitable “hard” metrics, and stand up a capability (with a mix of tech and people) that can help make those metrics timely and reliable.
- **Follow the 80/20 rule.** Technology delivers only about 20% of an initiative’s value. The other 80% comes from redesigning work—so agents can handle routine tasks and people can focus on what truly drives impact.
- **Spell it out.** As you design a new agentic workflow, map it step-by-step, specifying where agents own the work, where people do, where people and agents collaborate, and how oversight can take place for each step.



3. Rise of the AI generalist: A new workforce emerges


AI could soon end a shift that has marked most of the industrial era—the ever-increasing specialization of work. Agents can increasingly do the specialized tasks that fill the workdays of experienced, mid-tier employees. In IT, for example, you may no longer need coders specialized in specific languages. Instead, you may want engineers who understand both tech architecture and how to manage and oversee the agents that do know these languages. In finance functions, as agents do tasks like invoice processing, purchase order matching, reconciliation, and anomaly detection, people with general finance skills can focus on growing revenue and expanding margins, engaging with vendors on payment terms, working with sales on dynamic pricing models, and conducting more scenario planning.

Across functions, demand may grow for generalists who understand a wide range of tasks well enough to oversee agents and align their work with business goals.

In knowledge work, many of these roles can be filled by entry-level employees who tend to be AI savvy. As agents—available for anyone to buy or rent—take on more “midlevel” work, differentiation comes from senior professionals who excel at strategy and innovation. With more talent concentrated at the junior and senior levels, and a smaller mid-tier, the knowledge workforce may look like an hourglass. But for front-line employee-based task work, agents could replace entry-level workers with more mid-level people needed to orchestrate and manage these agents—creating a workforce more like a diamond.

What to do now

- **Look for all-around athletes.** Evolve recruitment to look not just for people who are leading but also ones who are AI-forward and open-minded enough to be generalists and agent orchestrators.
- **Start on workforce redesign.** As agents spread, your workforce may need new skills (like agent orchestration), new incentives (aligned to business outcomes, as agents do intermediate steps), and new roles (often related to oversight and strategy). And don't underestimate the importance of having a culture that encourages change, evolution, and adoption of the future of work.
- **Measure what matters.** With agents, iterations move quickly, but you may need more to allow for the back and forth required. Still, if an outcome that once took five days and two iterations now takes fifteen iterations but only two days, you're ahead.



4. Responsible AI moves from talk to traction

Executives know what Responsible AI (RAI) is worth. In our **2025 Responsible AI survey**, 60% said that it boosts ROI and efficiency, and 55% reported improved customer experience and innovation. Yet, nearly half of respondents also said that turning RAI principles into operational processes has been a challenge.

2026 could be the year when companies overcome this challenge and roll out repeatable, rigorous RAI practices.

The acceleration of adoption leaves companies little choice, and agentic workflows are spreading faster than governance models can address their unique needs. In many cases, agents can do roughly half of the tasks that people now do—but that requires a new kind of governance, both to manage risks and improve outputs.

The good news: The proliferation of new, tech-enabled AI governance approaches brings new techniques to the challenge. Automated red teaming, deepfake detection, AI enabled inventory management, and other advancements can help make continuous assessment and monitoring a reality.

These tools are powerful and nimble, but to support effective (and cost-effective) RAI, also depends on suitable upskilling and user expectations, risk tiering (with protocols for human intervention), and clarified documentation requirements and tools. RAI can then deliver the value you want like performance, innovation, and a reduction in the costs and delays that come with governance models built for another time.

What to do now

- **Integrate early.** The faster and deeper you can align IT, risk, and AI specialists—with clear responsibilities and expectations—the easier it can be to operationalize an RAI framework that can help grow business value and stakeholder trust.
- **Explore testing and monitoring solutions.** New tech capabilities can help operationalize AI testing and monitoring. Experiment with them now to understand challenges and adapt processes to be ready for when your AI adoption takes off.
- **Add assurance.** Unless you have unlimited data science resources, independent assessments may be needed to fill gaps. For higher-risk and higher-value systems, an independent opinion can be critical for performance and risk management.



5. From vibe to value: Orchestration that accelerates impact

AI agents make possible “vibe” coding—where people write software without technical expertise—and other “vibe” work, where almost anyone can invent and test new ideas. But you usually need tech teams to “industrialize” this innovation, putting ideas into production with continuous monitoring. That’s why an **orchestration layer** is so important. Its unified “command center” view helps you catch mistakes and track and fine-tune performance. It can also help end-user innovation enhance your top-down strategy.

You can spot valuable ideas and operationalize them quickly, manage risks, and keep everything aligned with your enterprise priorities.

A good AI orchestration layer should be easy even for non-techies to use, with intuitive dashboards and commands that let you drag-and-drop agents into new workflows—even for complex, high-value tasks. It should enable you to combine AI tools from different vendors into unified processes. It should integrate real-time data and natural language. And it should be built for centralized governance and security, with integrated code reviews and tools like encrypted credential vaults and sandboxes for prototyping. Most of all, it should put you in charge, enabling you to control AI anywhere in your company.

What to do now

- **Create orchestrators.** As orchestrating agents becomes a part of people’s workdays, you’ll want your employees to know how to spot and correct agents’ mistakes, connect them into teams, and find new tasks for them to do.
- **Help IT help you.** To help run your orchestration layer and execute your AI agenda, IT likely needs new resources and skills. Agentic AI for IT can help create new capacity by automating or assisting in many common IT tasks.
- **Stay practical.** Durable, scaled, industrial-strength deployments depend on practical actions which your orchestration layer can enable, things like testing before release, constant monitoring, and protocols for patches and quick rollbacks if needed.



6. The demand for business returns drives AI for sustainability

Whether AI is a boon or a burden for sustainability in 2026 is up for grabs—but we lean toward boon. That's not to understate the challenges. Even as AI quickly gets more energy efficient, its use is growing even faster. In fact, its fast-rising efficiency—by making AI cheaper—could accelerate its use even more. That could impact emissions, water supplies, and energy prices. But companies can become more efficient in AI use by approving token usage only when it delivers significant value and using methods like carbon scheduling to further cut emissions and costs.

And, as AI starts to drive a productivity boom, more efficient operations could compensate for AI's environmental impact.

But there's another reason that AI could boost sustainability—the quest for sales growth and margin expansion. AI agents, by gathering and analyzing customer data, can identify which customers would pay what premium for what kind of sustainability in products. Or they can measure and document that sustainability to strengthen your brand and grow your markets. AI can also help manage transport and electricity use to lower travel and power bills. Its

simulations can show how to grow resilience against natural disasters. And it can cost-effectively help trace products across your value chain to reduce both environmental impacts and costly product recalls. These and other AI solutions can create financial value for you while making your operations more sustainable.

What to do now

- **Be deliberate—and use what you have.** If you design AI with sustainability value as a goal, that value can come quickly. Since sustainability data is business data, for example, you can cost-effectively add AI use cases for Scope 3 indirect carbon emissions as you modernize supply chain data for AI.
- **Follow your customers.** AI to personalize products, marketing, and pricing for customer sustainability wishes can deliver quick returns. You may even discover that you're already meeting those expectations—you just haven't been communicating or pricing for it.
- **Act now to fend off rising costs.** As the grid struggles to meet AI-driven demand, you may face higher energy bills and even scarcity. Prepare to diversify energy sources or build your own, with renewables often the most affordable long-term option. You can also realize cost savings by integrating sustainability (such as carbon scheduling or protocols for when to use AI) into architecture.



Artificial Intelligence

Lead with trust to drive outcomes and transform the future of your business.

[Learn more](#)

What can AI do for you?

AI is already transforming business. Contact us to learn more about this rapidly evolving technology — and how you can begin putting it to work in a responsible way.

Get in touch

Related content

NEW

AI · Generative AI

AI native engineering for software delivery transformation

5 min. | Jul 14 | Blog

AI

Scaling AI with agentic architecture and enterprise governance

10 min. | Jun 12 | Publication

AI · Generative AI

Modernizing controls testing with agentic AI

6 min. | Jun 4 | Blog

AI · Generative AI

Coordinating enterprise AI agents at scale

4 min. | May 28 | Blog

- Privacy
- Data Privacy Framework
- Cookie info
- Legal
- Terms and conditions
- Site provider
- Site map
- Your Privacy Choices