

Generative AI Leadership Playbook

Idea In Short

Business leaders don't need another explainer on what generative Artificial Intelligence [AI] is; they need a decision framework for where to deploy it and how to contain its risks. The clearest signal from recent enterprise surveys is that adoption has outpaced value capture: most organizations use generative AI somewhere, yet few have scaled it past pilots or tied it to measurable profit. The immediate action is narrow and specific. Pick one workflow with a defined bottleneck, choose build, buy or blend based on data uniqueness and timeline and put trust and safety guardrails in place before the tool reaches a customer. Everything else, including talent development and long-term governance, follows from getting that first deployment right and measured.

Generative Artificial Intelligence [AI] has moved from a novelty in the boardroom to a line item in the budget and the gap between the two speeds, how fast it spreads and how slowly it pays off, is now the central management problem for business leaders. The technology creates new text, images, code and designs from existing data instead of simply scoring or classifying it, which is what makes it useful across sales, operations and product development at once. That same generative range is also why it resists a single playbook: a tool that drafts a marketing email and one that screens a loan application carry entirely different risk profiles. This article sets out what business leaders need to decide, in what order, to capture value from generative AI without absorbing its worst failure modes.

What Generative AI Actually Does

Generative AI differs from the predictive and analytical AI most businesses already run because it produces new content rather than a score, a ranking or a forecast. Fed a data set, product specifications, customer feedback transcripts or design briefs, it can generate a product description, a press release or a prototype sketch in the same session. A company launching a solar-powered bicycle, for instance, could use it to draft marketing copy, social posts and a press release within minutes rather than starting each from a blank page, though a human still needs to check the output before it goes out. That review step matters

because speed without verification just moves errors faster.

Adoption has already reached the point where it shapes daily workflows rather than sitting in a pilot folder. McKinsey's most recent state-of-AI survey found that nearly nine in ten organizations now use AI regularly, though most have not scaled it past isolated pilots and only 39% report any measurable profit impact from it so far¹. That gap between usage and payoff is the strategic problem this article addresses directly:

broad adoption without a deployment framework produces activity, not advantage

Where Generative AI Changes the Work

Executives speaking at recent industry forums, including Seamus Blackmore, chief technology officer at Trudell Medical International and Leila Pishdad, machine learning lead at Borealis AI, have pointed to a consistent set of functions where generative AI is already changing daily output rather than just promising to. Their observations, drawn from sessions at the Ivey Impact Live event, cluster around five areas that recur across industries.

1. Sales personalization at scale, where the tool drafts outreach tailored to a prospect's public profile, letting sales representatives spend their time on the relationship and the close rather than the first draft
2. Customer support triage, where it answers routine questions and navigates policy lookups fast enough to cut wait times measurably across a support queue
3. General productivity, where it summarizes research, drafts code or turns raw data into a written brief, freeing skilled staff for the judgment calls machines cannot make
4. Creative and design work, where architects, engineers and automakers use it to generate draft visuals and prototypes faster than a traditional design cycle allows
5. Operational tasks such as scheduling, document handling and preliminary cybersecurity monitoring, where marginal time savings compound across a large organization

These are not speculative use cases; they are the ones already showing up in enterprise deployments, which is precisely why business leaders can no longer treat generative AI as optional exploration.

Five Decisions Before Deployment

Getting generative AI into production well requires five decisions and getting them in the right order matters as much as getting each one right individually.

Identify the Use Case First

Every organization's highest-friction task is different, whether that is a hospital system parsing patient records, a retailer setting dynamic prices or a bank screening transactions for fraud. The starting discipline is asking vendors for examples from comparable organizations rather than assuming a generic tool will fit a specific workflow. A use case chosen for its visibility rather than its fit tends to produce a flashy pilot and no lasting adoption.

Decide Whether to Build or Buy

Building a proprietary model can make sense when an organization holds genuinely unique data or in-house expertise, but it demands specialized talent and a long runway. Buying, particularly customizing a pre-trained foundation model with an organization's own data, is usually faster and cheaper. Gartner frames this less as a binary choice than a portfolio decision: distinguished analyst Hung LeHong has argued that the most effective enterprise approach blends embedded AI already inside existing software, independently adopted tools chosen by individual teams and centrally built or blended systems, all under one governance layer². Leaders should weigh timeline, integration cost with existing systems and total cost of ownership before committing to any single path.

Invest With a Clear Return in Mind

Deploying generative AI costs money: specialized hardware, skilled staff and secure data infrastructure all carry a price and pre-trained models only reduce that bill, they don't eliminate it. The return has to come from a specific, named efficiency or innovation gain, not from the technology's novelty. Thomson Reuters' 2026 professional services survey found that only 18% of organizations currently track return on generative AI investment formally, with another 40% unsure whether it is being measured at all³. That measurement gap, more than the technology itself, is what turns a promising pilot into a stalled budget line.

Build Trust and Safety Into the Design

Generative AI can hallucinate, producing false or biased output with the same fluent confidence as an accurate one and that failure mode carries real reputational risk. Practical guardrails include topical limits that keep a chatbot inside its intended domain, output filters that catch unsupported claims before a customer sees them and access controls that protect the underlying data. Harvard Business Review has catalogued the hallucination problem alongside three related risks, unclear accountability, manipulative output and unresolved liability, as the core reasons leaders report feeling disoriented by generative AI even as they face pressure to adopt it⁴. Building guardrails before launch, not after an incident, is what turns that anxiety into a manageable risk register.

Measure Success Against a Defined Goal

None of the previous four decisions matter if success is never defined. Business leaders should set a target, hours saved, resolution time cut or customer satisfaction lifted, before deployment, then track it against that baseline rather than against a vague sense of improvement. A tool that nobody is measuring is a tool nobody can defend at the next budget review.

The Risks That Come With the Opportunity

Generative AI's opportunities carry a matching set of risks and leaders who treat the two as separate conversations tend to underinvest in the second.

High upfront cost remains real, even as it declines with broader access; staffing and data security account for a larger share of the bill than most first estimates assume, particularly for organizations running systems on premises rather than in the cloud. Hallucinated output can mislead customers and damage a brand's credibility, especially where training data carries embedded bias the model reproduces at scale. Skilled AI talent, people who can integrate, govern and evaluate these systems, remains scarce relative to demand, which is one reason many organizations are training existing staff rather than competing for a thin external pool.

Overhype is its own risk. Traditional software or a simpler analytics tool solves most business problems more reliably than a generative model bolted on for its own sake, so the discipline is assessing fit before defaulting to the newest option. Compliance adds a further layer: the European Union's Artificial Intelligence Act now requires organizations to classify

AI systems by risk level, disclose AI-generated content and, for higher-risk applications, document testing and ongoing monitoring, with compliance timelines running twenty-four to thirty-six months depending on classification⁵. Leaders operating across borders need to treat that regulatory timeline as a planning input, not an afterthought.

A Five-Step Blueprint for Deployment

A workable blueprint for prioritizing and deploying generative AI follows a sequence rather than a checklist and skipping a step tends to show up later as rework.

1. Focus on impact by identifying the small, recurring frustrations, slow content drafts, repetitive document review, that erode employee and customer satisfaction, then target those first
2. Customize with data by starting from a pre-trained model and layering an organization's own information on top, whether that produces a virtual assistant for event attendees or sentiment analysis for a bank's service desk
3. Boost productivity by embedding generative AI into tools employees already use, such as code generators or analytics platforms, rather than introducing a standalone system that adds friction
4. Develop skills deliberately, pairing technical training in data handling with softer skills like bias evaluation, since a tool is only as good as the people reviewing its output
5. Innovate with oversight, testing thoroughly, setting explicit confidence thresholds for automated decisions and keeping a human in the loop wherever the cost of an error is high

Each step depends on the one before it:

a blueprint built around impact but skipped on data customization produces a generic tool nobody adopts, while skills development without oversight produces confident but unchecked automation

Setting the Conditions for Long-Term Success

Beyond individual deployments, generative AI performs only as well as the organizational conditions surrounding it. Data has to be integrated, current and governed, since a model trained on fragmented or duplicated records will reproduce those flaws in its output. IBM's research into data exposure under generative AI describes this as a shift from infrastructure-focused security toward data-centric governance, given that employees can now input sensitive information directly into a model or leak it through an unmonitored output⁶. Cloud infrastructure, rather than legacy on-premises systems, tends to handle the computing load more sustainably and at lower marginal cost as usage scales.

Security, privacy and hallucination risk all require ongoing governance rather than a one-time review; a risk assessment done at launch and never repeated leaves an organization blind to how a model's behavior drifts as it processes new data. Starting with small, goal-oriented pilots and expanding only after they prove out builds both organizational confidence and a data-driven culture that survives the initial novelty phase. None of this requires abandoning ambition, but it does require treating generative AI as infrastructure to be governed rather than a feature to be switched on.

Where This Leaves Business Leaders

Talent and skills sit at the center of how this plays out over the next several years. The World Economic Forum's most recent talent and technology outlook found that more than half of business executives globally expect AI to displace existing jobs, while only about a quarter expect it to create new ones on net, a gap that depends heavily on how quickly workforce skills development keeps pace with AI capability⁷. That is not a reason to slow adoption; it is a reason to fund the retraining alongside the technology rather than after it.

Fredrik Odegaard, who moderated the Ivey Impact Live discussion referenced earlier in this article, put the responsibility squarely on leadership.

It's up to us to shape the AI revolution

Generative AI for business is neither a shortcut nor a threat by itself; it is a capability that rewards the leaders who sequence their decisions correctly, use case first, then build-or-

buy, then investment, then guardrails, then measurement and who fund the people development that has to run alongside it. Leaders who skip that sequence tend to get the pilot; leaders who follow it tend to get the return.

- 1The State Of AI: Agents, Innovation And Transformation
- 2Deploying AI: Should Your Organization Build, Buy Or Blend
- 32026 AI In Professional Services Report
- 4Generative AI-Nxiety
- 5EU AI Act: First Regulation On Artificial Intelligence
- 6Governing Data Exposure In The Age Of Generative AI
- 7Four Ways AI And Talent Trends Could Reshape Jobs By 2030

Summary

Generative AI rewards leaders who treat it as an operating decision rather than a technology purchase. The organizations pulling ahead pick a specific use case, choose a build, buy or blend path deliberately and measure results against defined goals instead of chasing novelty. The risks, hallucinated output, data exposure, talent gaps and tightening regulation, are manageable when governance is built in from the first pilot rather than bolted on afterward. Generative AI will not replace strategic judgment; it will keep demanding more of it, as leaders decide where automation helps and where human oversight has to stay in the loop.