HARVARD BUSINESS REVIEW

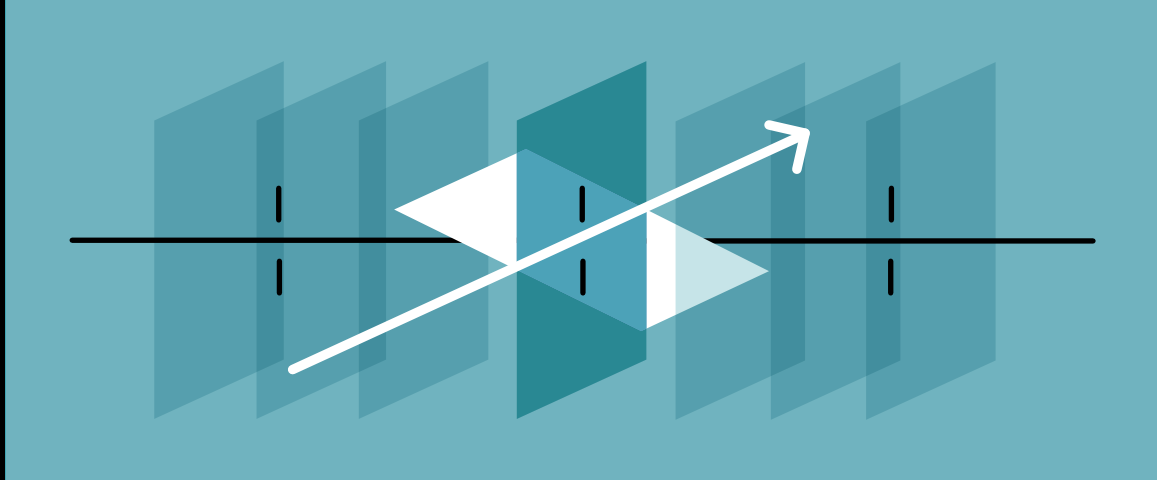

Moving Generative AI from Experimentation to Operation

HBR Analytic Services **BRIEFING PAPER**

Sponsored by

Hitachi Digital Services

Beyond the Hype: Making Generative AI Work for Your Business

The rapid advancement of generative AI (gen AI) is shifting from a phase of experimentation to full-scale business integration, redefining operational efficiencies, cost structures, and innovation strategies. While consumer applications like ChatGPT have captured public attention, enterprise adoption requires a more structured and strategic approach. Organizations must navigate complex challenges, including identifying the right use cases, cost, ROI, data quality, data privacy scalability, security, and workforce integration, to ensure that artificial intelligence (AI) delivers tangible business value.

Hitachi Digital Services recognizes that the operationalization of gen AI in an enterprise is not just a technological shift but a fundamental business transformation. Companies that succeed in moving beyond proof-of-concept deployments stand to unlock significant competitive advantages. From demand forecasting, context-based maintenance procedures, and safer workplaces to contract management and intelligent automation, gen AI is poised to reshape industries by enhancing decision making, optimizing processes, and uncovering new revenue streams.

However, as with any transformative technology, AI adoption is not without its hurdles. Many enterprises struggle with identifying the right use cases, sourcing data, ensuring quality, complying with regulatory requirements, and developing a compelling business case, including alignment with business objectives. The key to success lies in a well-defined AI strategy, robust governance models, leveraging ecosystem partners, and a collaborative approach that involves cross-functional teams. IT leaders must work alongside business executives to ensure gen AI implementations align with long-term corporate goals.

This report explores the evolving landscape of gen AI operationalization, highlighting best practices, emerging use cases, and key considerations for successful deployment. Drawing insights from industry leaders, analysts, and practitioners, we examine the steps organizations must take to transition from pilot projects to enterprise-wide AI adoption.

For Hitachi Digital Services, the focus is on leveraging AI not as a standalone technology but as an enabler of IT-OT convergence, digital transformation, and industry-specific innovation. By embedding AI within the broader digital ecosystem, businesses can drive efficiency, enhance customer experiences, and position themselves at the forefront of the AI revolution.

As enterprises take their first steps toward operationalizing gen AI, they must adopt a mindset of continuous learning and adaptation. The future of AI-driven business transformation will depend not only on technological advancements but also on how effectively organizations integrate AI into their core operations by complementing internal expertise with technology partners. This report serves as a guide to navigating this journey, offering practical insights to help businesses harness the full potential of AI while mitigating risks and maximizing ROI.

Premkumar Balasumbramanian
Chief Technology Officer,
Head of Artificial Intelligence,
Hitachi Digital Services

Rajesh Rajappan
Chief Strategy Officer,
Hitachi Digital Services

Patrick Corcoran
Global Lead,
Analyst and Advisor Relations,
Hitachi Digital Services

Moving Generative AI from Experimentation to Operation

Successfully moving generative AI (gen AI) projects into business operations—a process Stamford, Conn.-based technology research and consulting firm Gartner refers to as operationalization—can generate significant business benefits in terms of improving efficiency, effecting cost savings, and creating new revenue streams and business opportunities. Gen AI systems can even support new business processes by combining data sets and language models that simply weren't possible before.

IN A MCKINSEY & CO. global survey of 1,363 IT executives and business unit leaders titled "The State of AI in Early 2024: Gen AI Adoption Spikes and Starts to Generate Value," published in May 2024, 65% of the respondents indicated their companies are regularly using gen AI. That figure is nearly double the percentage from another McKinsey survey conducted just 10 months prior.[1]

However, Tom Gruber, cofounder of and head of design for Siri, Apple Inc.'s renowned digital assistant, sees chatbots like ChatGPT as just one example of gen AI. "What the consumer sees with ChatGPT is a lot different than what most enterprises want to do with it," says Gruber, who is also an artificial intelligence (AI) consultant to several companies, including Pleasanton, Calif.-based Cognizer.ai and Bilbao, Spain-based Sherpa.ai. "You can operationalize it at an individual level, but if you have an important business problem, that's not solved out of the box by the ChatGPTs of the world." He points out, however, that the complex gen AI models that support ChatGPT are capable of addressing more-complicated business problems.

In fact, according to a January 2025 survey of 460 business unit leaders by Cambridge, Mass.-based HFS Research, companies have made varying levels of progress applying gen AI platforms to a range of business functions and use cases. Thus far, though, just customer service chatbots are a widespread use case, with 28% of the respondents currently using gen AI to

HIGHLIGHTS

Establishing a **comprehensive framework of metrics** can help organizations ensure their generative AI (gen AI) projects are on track to deliver the anticipated value.

In enterprise environments, businesses are finding **gen AI delivering business value in a variety of use cases** that go far beyond chatbots.

Future adoption and **efficient operations** will ensure gen AI platforms are constantly evolving and everyone using those systems will provide that continuity.

"You can operationalize [ChatGPT] at an individual level, but if you have an important business problem, that's not solved out of the box by the ChatGPTs of the world," says Tom Gruber, cofounder of and head of design for Apple Inc.'s Siri and an independent AI consultant.

support chatbots and 50% expected to do so within the next two years. Demand forecasting is another strong function, with 17% currently using gen AI to support forecasting efforts and 68% planning to over the next two years.[2] FIGURE 1

As companies prepare to make the move from experimentation and proof of concept (POC) to operation with gen AI, it's important for executives to accurately identify the use case or business problem, determine the appropriate scale, and ensure the entire organization is both informed and involved, says Edward Hanapole, managing director and chief AI officer at New York City-based Alvarez & Marsal Public Sector Services. "During the initial efforts of prototyping, it is essential to understand the art of what is possible," he says. "Some of these projects are going to be valid and useful. Some of them, you might need to take a step back and take a look again."

Hanapole emphasizes the importance of involving the entire team in the process. "This isn't just an engineering task. It requires collaboration across functions—finance directors, business unit managers, and other key stakeholders all working with a shared understanding of the objective and a clear approach to solving the problem," he says.

Meanwhile, moving a gen AI platform into operationalization is ultimately still an IT project. "You're going to have the same issues as any IT program, like security and scalability from an infrastructure standpoint," Hanapole explains. "Then you're going to have other required skill sets around

FIGURE 1

Adoption Levels Vary for Different Tasks

Companies are making varied progress in applying generative AI to business functions

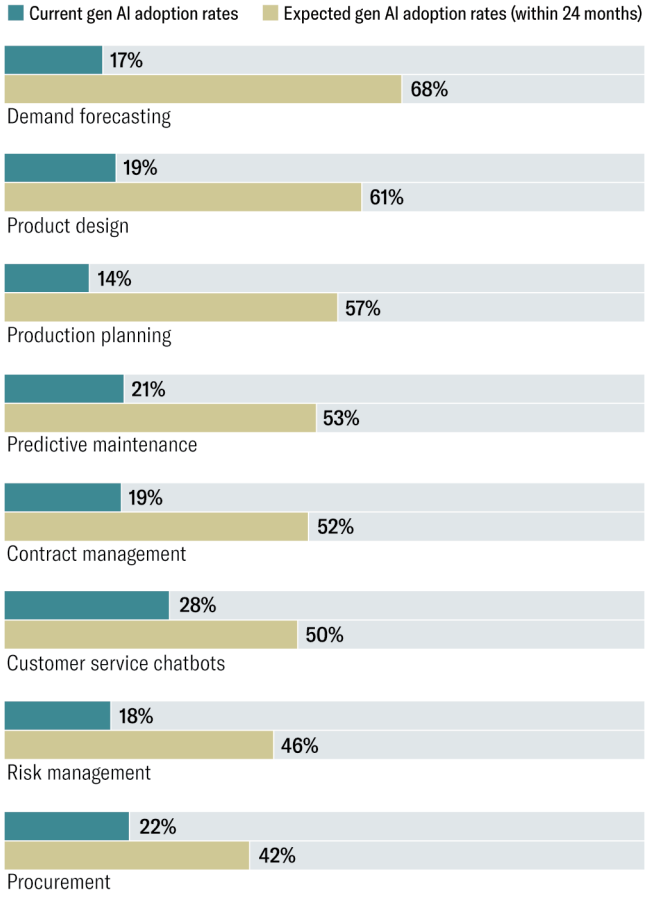

Source: HFS Research survey, January 2025

advanced data analytics to get your data where it needs to be to power your AI instance."

This report will help guide companies toward a successful transition and help executives make their plans and preparations. It will also ensure executives are tapping into the right skill sets, whether that means relying on internal teams or engaging with an external partner.

Getting Ready for the First Steps

ChatGPT and similar chatbots may have thrust gen AI into the collective consciousness, but the technology has been around for years. Many people have already been using gen AI-powered functions, possibly without even realizing it, Hanapole says. "There are things people use on a day-to-day basis that leverage gen AI," he says. "Recommendation engines on Netflix are using machine learning-based AI. Navigating on Google maps involves using AI principles of predictive analytics and reinforcement learning."

Hanapole compares the rapid emergence of gen AI to that of cloud computing. "2023 was transformative for gen AI as a technology not because it was new, but because of the explosion of consumer-facing applications like ChatGPT. It's similar to what we saw play out in cloud computing, when cloud computing ultimately became 'consumerized,'" he says. "Gen AI has been around for a few years in deep learning-based architectures, but foundational elements like natural language processing and neural networks date back decades. Once you put it in front of users, the fascination occurs and it becomes clear what the power is and what they can do with it."

After gen AI captured such immediate widespread attention, many enterprises began experimenting with the technology. "It became top of mind for all our clients. [They all] set up internal task forces," says Akiba Stern, a partner in the New York City office of international law firm Loeb & Loeb, where he focuses on outsourcing contracts, technology-enabled business transactions, e-commerce, technology transfers, and licensing. "What is this new technology? How do you safely and properly implement it? How do you utilize it? What are the fundamental use cases to consider?" he asks.

Among other things, gen AI has the potential to help organizations get more value out of the data they already have. "Today you can have AI that understands every sentence in every document. What would you do with that in the enterprise?" asks Gruber, the AI consultant. "You can now have a deep understanding of the content of the documents and can extract truly structured data from unstructured data."

Taking those first steps when putting gen AI into business operations requires preparation, however. Alvarez & Marsal Public Sector Services' Hanapole advises executives to ask several questions as they craft their gen AI plans. "How are you going to adopt it? How will you continuously measure it? How are you going to scale it into production?" he asks. "As with traditional IT projects, you start with these things under consideration."

Establishing a comprehensive framework of metrics can help organizations ensure their gen AI projects are on track to deliver the anticipated value. "How do we know if it's delivering repeatable results?" asks Hanapole. "How are we going to measure whether it goes off the rails or not? Can we in fact confirm that it is performing accurately and predictably and delivering the types of responses we would expect?"

That initial preparation is critical, he says. "It becomes difficult to go from experimentation to full implementation if you haven't already thought of those things. Understand what the possible pitfalls and drawbacks are when going from POC into actual implementation," Hanapole explains. "During the earliest stages, it can be more technology-driven—a solution looking for a problem rather than solving a defined need."

Loeb & Loeb's Stern agrees that identifying the right business process is crucial to ensure success. "To go beyond the POC, [identify] the use cases that will provide significant benefits," he says. "Conduct a benefit analysis around those use cases and determine what guardrails need to be put in place."

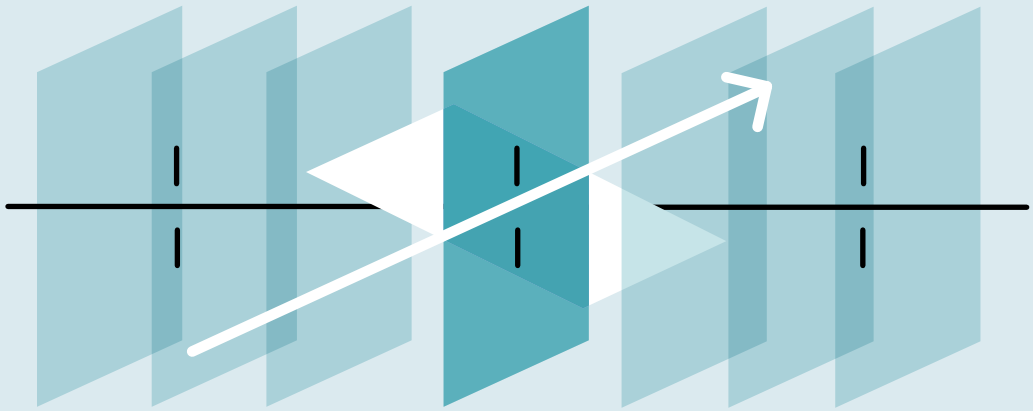

"You can now have a deep understanding of the content of the documents and can extract truly structured data from unstructured data."

Tom Gruber, independent AI consultant

Use Cases, Business Benefits

In enterprise environments, businesses are finding gen AI delivering business value in a variety of use cases that go far beyond chatbots. "Where is this untapped value of gen AI in the realm of document understanding? There are all kinds of interesting use cases," says Gruber.

He cites processing legal documents and contracts as one significant use case. "Legal documents are the most expensive document word for word and have the most risk of getting it wrong. And one sentence in that document can change a lot about a business," Gruber says.

Applying gen AI to legal documents has broad implications across a range of industries. "[Legal documents] are involved in all aspects of business—sales contracts, service contracts, partnerships, and M&A deals," says Gruber. "Legal contracts are extremely important, yet they are still written in this medieval language, not written as structured data. We can tackle that with language models, and [those models] have gotten better over time."

Executives would do well to step back and take a look at which business processes could be expedited with automation, he continues. Contracts are a good starting point. "Looking at underserved business processes, either [companies automate] inefficiently or don't do it at all because it costs too much," Gruber says. "Things like legal analysis—right now, company A buys company B. They pay a ton of money for some consultancy to pay lawyers to analyze and integrate all of the legal obligations. That's expensive and tedious. It's not practical to do exhaustively without some serious gen AI, but that's an example of where you can do new things."

As a corporate attorney, Stern clearly appreciates the data privacy and security issues that need to be accommodated when implementing gen AI systems within a law firm. "That's issue number one for law firms," he says. "In order to be able to run a gen AI system against client data, a law firm needs to run it in a manner that keeps the provider from having access to the data ingested by the system, whether on premises or within instances in the cloud where the provider and others cannot view that information."

Stern sees gen AI as particularly useful for performing initial review and markup on legal contracts. "How do you build a tool to go through a particular type of contract and do an initial markup?" he asks. "We would use it to spend less time, as billable hours are how we charge. More-efficient service delivery reduces the amount of lawyer hours and lets us focus more on high-end advice. We could reduce work like preparing initial drafts of massive sets of documents. If we can give that to the client at a lower price point, that's very helpful because we are in a competitive environment."

Besides employing gen AI platforms to generate rough contracts or perform initial contract reviews, Stern also uses gen AI to perform deeper, more sophisticated searches of previous legal documents. "Another use is looking for a particular provision or some way we've handled something for a previous transaction," he explains. "I know we did it, and I could use the search tools we have today, but it's a tedious process. [With gen AI], I could search the database and also clean up the data and ensure we split the documents into various components." Stern again points to the time and cost savings that can then be passed on to his clients.

Beyond the legal use case, Gruber sees how gen AI can also help bring together and extract value from large data sets in ways that were previously impossible or at the very least impractical. "I'm working with a company [whose clients include] a giant financial services firm and a giant telco. They obviously can't share customer data with each other, but it turns out there's a valuable cross section," he says. "There's a lot of value if you can combine [the data] to help predict if someone is going to default on a loan, for example. [They can] make a privacy-protected version of an aggregate model so gen AI knows in a distributed way about both models, without any data leaving its protected silo. It used to be impossible to build aggregate models like this given the rules of the game, but privacy-protected machine learning is changing those rules."

Another potentially life-changing use case Gruber sees is using AI to combine data sets to address infant, orphan, and rare diseases without compromising patient privacy. "There's not enough data in any one hospital to build an AI model to diagnose those diseases, but if you pooled all the data from all the hospitals, you would have enough. But it's too expensive and no one is going to do it for privacy reasons," he says. "Now, with [privacy-protected machine learning], you'd be able to start addressing this class of diseases. You can think outside your normal level of constraints. Gen AI can overcome these constraints."

Bringing gen AI to bear on more-complicated corporate use cases will require more thoughtful consideration, however. "Other use cases can be a bit more complex, like using [gen AI] to do financial analysis or planning transportation," says Hanapole, citing the extent and quality of the data required and the accuracy and sophistication of the models needed to perform such analysis. "It's still gen AI-based, but you're going to want to have a bit more rigor around the way you establish the models."

When considering the primary business benefits of bringing gen AI systems into operation, time savings and cost savings leap to the top, which is probably why support of call center operations is such a major use case. "The [use cases] that are really ripe for gen AI are customer care, customer support, and even agent support," Hanapole asserts. "[Gen AI] can quickly help an agent find information to present a more satisfactory answer. You can reduce costs and resolve questions faster." He can envision a point at which human intervention is reserved for exceptions. "You can get the agent involved only when gen AI lacks sufficient confidence in its response," he notes.

But there are deeper, more complex benefits that can lead to strategic differentiation for an enterprise. "It's ultimately helping you scale your ability to find patterns, to find trends within your own data. [Gen AI] extends beyond the limitations of your human analysts," says Hanapole. "It's like having a thousand analysts at your disposal to be able to perform tasks and identify patterns to provide actionable insight."

FIGURE 2

Generative AI Benefits Customer Experience Most

Operationalizing it benefits revenue growth, too, among other things

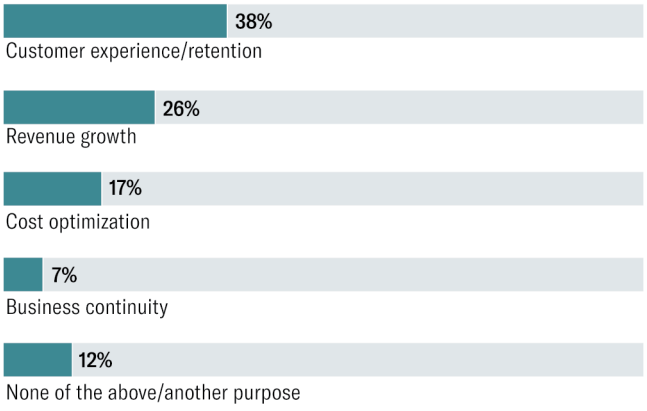

Source: Gartner survey, 2023

A 2023 Gartner survey of more than 2,500 executives found that 38% of the respondents listed customer experience and customer retention as the primary business benefits they expect to see from their gen AI platforms. Several other notable business benefits that executives anticipate from bringing gen AI platforms into the organization include revenue growth (26%), cost optimization (17%), and business continuity (7%).[3] FIGURE 2

Regardless of the intended use case, Hanapole advises applying gen AI platforms across as broad a spectrum as possible. Limiting a single gen AI platform to a single use case can restrict the cost savings and performance potential because if an organization has already developed and trained a gen AI platform to perform a specific function, adjusting the language models that power the gen AI platform can help an organization apply that platform to other functions

> **“[Gen AI] can quickly help an agent find information to present a more satisfactory answer. You can reduce costs and resolve questions faster.”**
>
> Edward Hanapole, managing director and chief AI officer, Alvarez & Marsal Public Sector Services

without having to start from scratch. “How can you shift [the gen AI platform] to other domains?” he asks.

He cites a large education services organization that implemented a gen AI-powered virtual guidance counselor to help high school students prepare for college. The chatbot interviews the students to help them understand what they need to do to get into the colleges of their choice. It asks what courses they’ve taken, which tests they’ve taken, and what other interests they have. “The gen AI is conversational, so the student feels natural. It feels like they’re texting with someone [who’s] asking them questions,” says Hanapole.

After several rounds of “conversation,” the system returns its guidance. “Based on what you told me, here’s what we think your safety schools are. Here are your target schools,” says Hanapole. “Based on what you told me, here are the chances of you getting into these schools.”

The system has been so successful that this organization has extended the model to other areas, like financial aid and eligibility. “They’re now looking to leverage a single platform architecture for multiple use cases [beyond career guidance],” he says.

Potential Implementation Pitfalls to Avoid

To ensure a successful gen AI implementation, there are some relatively common mistakes that are essential to avoid, says Hanapole, with the biggest being a lack of preparedness. It’s important to have a plan for how to support the system once it’s up and running.

Hanapole also advises thoroughly defining success metrics well in advance of putting any gen AI system into production. Among significant oversights, he says, are “not knowing how you’re going to gauge success [and] not having a human in the loop in terms of understanding transparency of what models are doing.”

Other potential pitfalls abound with the gen AI platform architecture optimization. “You might be overusing certain AI components that are actually quite expensive,” Hanapole reports. “Based on how you’re going to pay for AI usage and API [application programming interface] processing, the cost might be considerably higher based on poor architecture. [You could] offload some of what gen AI handles and use some other technique. [Carefully examining system architecture] before you get to gen AI could lower the costs.”

Stern sees the particular need for guardrails around how gen AI systems interact with data. “What’s the provenance of your data? When you use your data, what purposes are you using it for? You have to have a solid set of rules around [the data],” he says. “Specify your requirements and obligations. That is what companies want. They do not want to be in the press.”

Beyond privacy and security, ensuring data quality is equally critical. “Gartner and the like will tell you everybody is doing a lot of pilots and POCs and that is as far as it goes,” says Stern. “One of the reasons is that the data isn’t clean. You need clean data to take full advantage of [gen AI].”

Hanapole also expresses the need to prepare the data that will drive gen AI. Consider what data will be fed into the language models. How will that data be sourced, curated, and refined? And what metrics will be applied to measure success? “You really have to start with a focus on the data and data science,” he says. “At the front end, it’s all about sourcing the right data.”

Gen AI models require rigorous examination, he advises. “Fine-tuned and foundational models can help control costs while maintaining performance,” Hanapole says. “You don’t want to be prevented from scaling because you didn’t architect [the platform] in a satisfactory manner.”

At a fundamental level, adopting gen AI is still an enterprise IT initiative and should be subject to the same best practices. “It’s the discipline and maturity that should be part of any core IT program,” he says. “How to handle privacy, scalability, and security—you don’t throw that out the window just because you’re going to use gen AI.”

On the human side, there are overarching concerns about job loss. “The short headline is automation replaces jobs.

> **“Responsible AI is not just an afterthought.”**
>
> Alvarez & Marsal Public Sector Services’ Hanapole

And it will [do so] for jobs that are easily replaced, [such as an] entry-level call center operator,” says Gruber. “The [call center] industry will be one of the better examples of outright automation replacing jobs, because humans—at least the entry-level workers—are doing information-processing jobs that can be performed as well or better by gen AI.”

That replacement of lower-level jobs is just one aspect of the potential impact gen AI may have on the human workforce, Gruber says. “In the vertical [use cases] with high-value documents like legal contracts or in health sciences, those are places where doing knowledge work is the bottleneck,” he says. “Things like drug discovery will be well augmented by AI tools. It’s going to be the same scientist using those tools. AI is not going to replace the scientist; it’s just that there’s so much low-hanging fruit to [be able to] automate the grunge parts of their work.”

Approaching gen AI in a responsible fashion is another aspect of operational success. “[That responsible approach] will help ensure you’re being fair, not just in terms of how you’re using AI, but that your practices are fair. You’re helping humans understand how AI operates, identifying potential biases, ensuring ethical deployment, and understanding inconsistencies in your treatment of information,” says Hanapole. “Responsible AI is not just an afterthought.”

Essential Skill Sets Needed

There will ultimately be the expectation that AI is a core competency within the enterprise at all levels, according to Hanapole. Future adoption and efficient operations will ensure gen AI platforms are constantly evolving and everyone using those systems will provide that continuity. “The most important thing for executives is to be curious and start using [gen AI]. Just use the tools and understand what it is and what it isn’t,” he says.

Once there is a thorough understanding of the challenges the business is trying to address, that comfort level will help executives determine the most valuable place to leverage gen AI, he says. “In some companies, it might be financial efficiency or providing more-responsive budgets,” Hanapole explains. “Based on those [use cases], you’ll be able to map out the subject-matter experts you’ll need to ensure AI projects are successful.”

Gruber stresses that involvement at all levels of the organization is essential. “Both [knowledge workers and executives] need a good sense of what [gen AI] does and how it does it,” he says. “They can calibrate the business value of the technology by knowing what it can and can’t do.”

Continued use, refinement, and training of the platform will ultimately determine its success. “These [gen AI platforms] do not know anything, despite the language people use. They have been trained—which is the proper, accurate word—and they’ve been shown information,” Gruber says. “They don’t know in the way that we know. They don’t reason from first principles. They don’t think from theory. They just imitate thinking they’ve seen in the form of human writing.” Once workers and executives realize the parameters gen AI platforms adhere to, only then, Gruber believes, can they understand both how to use and how not to use gen AI.

“For example, you should know what [a gen AI] hallucination feels like. When those things hallucinate, there’s no intent behind it. They can’t help themselves,” says Gruber. An AI hallucination is a false or misleading AI-generated response that is presented as truth, often the result of an inadequately trained language model. “They lie and tell the truth with equal facility. Ask them to show only evidence that is true, and they’ll answer you. They’ll make up stuff that looks accurate, but you can’t tell. They’re not doing it as [if they were] a used car salesman.”

It’s important, too, for executives to approach gen AI platforms and models from a similar perspective to that with which they approach their staff. “From a manager’s point of view, you already know how to delegate,” says Gruber. “You know how to assess whether or not an employee is doing a good job. You know how to

assign tasks calibrated to suit the skills of the employee. Those are managerial skills, and those are the skills you need to use AI."

When it comes to outsourcing the move to getting gen AI into operations, Gruber recommends a combination of outsourcing and keeping work in-house, although he believes it is more important for larger enterprises to engage a partner. "Enterprise apps [require much more work] than just making a demo work," he says. "Solve the important problems first. For example, modernize your enterprise to be more agile at partnering with other companies to solve strategic problems, empowering R&D to think more creatively and flexibly. You can use gen AI to solve these problems that are traditionally harder to solve."

Hanapole similarly recommends a balance of outsourcing and in-house expertise for getting gen AI projects into operation. "[An outsourcing partner can] help accelerate the adoption and potential use cases and help clients avoid some of the pitfalls that have become common in getting from the prototyping stage to production, [such as] costs, performance, security issues, or other ways of architecting an AI solution," he explains.

Stern is a strong advocate of outsourcing operationalization. "As long as you have the right parameters, you can and should contract out. Everybody does it, even banks and large pharma companies, which are among the most highly regulated enterprises," he says. "The trend continues to be to move these types of functions to providers."

"As long as you have the right parameters, you can and should contract out. Everybody does it, even banks and large pharma companies, which are among the most highly regulated enterprises," says Akiba Stern, a partner in the New York City office of international law firm Loeb & Loeb.

On the Horizon for Gen AI

Gruber sees more opportunities for gen AI to solve increasingly complex problems in areas like the life sciences and materials science. He cites the work that Demis Hassibis has already accomplished in protein structure prediction, which can have a significant impact on drug discovery for the pharmaceutical industry. Hassibis, who won a Nobel Prize for his work in 2024 (along with John M. Jumper), is the CEO and cofounder of Google DeepMind and Isomorphic Labs, both based in London.

Researchers are also making breakthroughs in the areas of manufacturing and materials science. When companies are able to discover and develop new materials, it can lead to manufacturing and product innovations. "In materials science, when you discover new materials, that's how you get batteries and storage and transmission and energy," says Gruber.

He also sees call center use cases expanding with gen AI. "We're going to see the call centers of the world rapidly changing," he says. "Most call centers have had really stupid chatbots, so people just stopped using them." As chatbots improve, however, Gruber expects a cultural shift from customers being reluctant to approach call centers because of their assumption that they'll have a frustrating and inefficient experience with call center chatbots. As customers begin to realize gen AI-powered chatbots are more responsive, more accurate, and more efficient than prior iterations, customers will become more comfortable using chatbots to get their customer service problems resolved.

Hanapole anticipates the emergence of gen AI platforms preconfigured for specific tasks. He describes these as agents—gen AI platforms configured to accomplish a specific task—similar to a chatbot. "We're going to start to hear more about the rise of agents," he says. "Intelligent agents are basically just pretrained gen AI models that will be able to handle

Gartner anticipates that gen AI will become a "general-purpose technology with an impact similar to that of the steam engine, electricity, and the internet."

things specific to a particular domain, whether it's housing benefits, procurement, or whatever. By year-end, you will see a lot of [organizations] leveraging gen AI to solve specific needs within their environment."

Indeed, more sophisticated and efficient gen AI platforms can help on numerous fronts. Hanapole cites sustainability as one of those. "I hope to see more focus on sustainability of AI solutions that comes with creating a smaller footprint in terms of energy consumption," he says. "Once we start getting more sensitive to architecture, design, and using [gen AI] models, that will ensure [gen AI] is as effective and efficient as possible." As their use becomes more widespread, Hanapole believes gen AI platforms will also help companies develop a higher degree of strategic differentiation.

Gartner anticipates that gen AI will become a "general-purpose technology with an impact similar to that of the steam engine, electricity, and the internet."[4] For gen AI platforms to truly generate business value, though, organizations need to move them from the experimentation and proof-of-concept stage and get them into business operations. Only when those gen AI platforms are directly supporting business operations can they generate business value.

"The technology is being adapted to actual problems to deliver maximum value," says Hanapole. "You have to roll up your sleeves and determine how it's going to help you amplify your strategic differentiation."

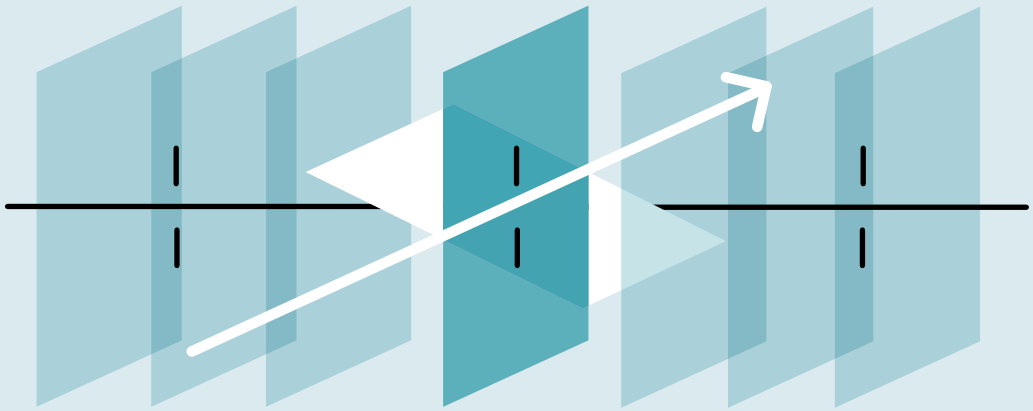

For gen AI platforms to truly generate business value, though, organizations need to move them from the experimentation and proof-of-concept stage and get them into business operations.

Endnotes

1 Michael Chui, Bryce Hall, Alex Singla, Alexander Sukharevsky, and Lareina Yee, "The State of AI in Early 2024: Gen AI Adoption Spikes and Starts to Generate Value," McKinsey & Co., May 30, 2024. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-2024.

2 KS, Krupa, "Agentic AI Era Begins: GenAI Will Surpass Traditional AI in Most Supply Chains Areas," HFS Research, January 16, 2025. https://www.hfsresearch.com/research/agentic-ai-surpass-supply.

3 Gartner, "Gartner Experts Answer the Top Generative AI Questions for Your Enterprise," 2023. https://www.gartner.com/en/topics/generative-ai.

4 Ibid.

VISIT US ONLINE

hbr.org/hbr-analytic-services

Harvard Business Review Analytic Services is an independent commercial research unit within Harvard Business Review Group, conducting research and comparative analysis on important management challenges and emerging business opportunities. Seeking to provide business intelligence and peer-group insight, each report is published based on the findings of original quantitative and/or qualitative research and analysis. Quantitative surveys are conducted with the HBR Advisory Council, HBR's global research panel, and qualitative research is conducted with senior business executives and subject-matter experts from within and beyond the *Harvard Business Review* author community. Email us at hbranalyticservices@hbr.org.

CRE61110425