snowflake®

GLOBAL RESEARCH

# THE RADICAL ROI OF GEN AI

Enterprises are winning big with generative AI.
Here's how they're doing it.

# TABLE OF CONTENTS

# GEN AI DOES NOT DISAPPOINT

## Early adopters are enjoying early — and significant — ROI

There's no need to sugarcoat this or be clever about it. But we can put it in bold: **Generative AI works.** Fifty-seven percent of business and IT professionals we surveyed worldwide — 1,900 of them — have deployed gen AI solutions. We asked how that's working out for them, and 92% told us that not only are they happy with their efforts, but their investments have already paid for themselves.

Just 2½ years since gen AI became the subject that dominated every tech conversation, these early adopters are not only taking advantage of opportunities today, they're accelerating toward a transformative tomorrow. As quickly as gen AI has flourished in the enterprise, the technology keeps leaping forward. In January, just as this research was concluding, China's DeepSeek AI assistant debuted, disrupting the stock market and beating out ChatGPT in Apple's app store a few days later. And later this year, autonomous agents, capable of executing complex, multistep tasks with little or no human intervention, may make as big a splash. At Davos in January, Salesforce CEO Mark Benioff said that today's CEOs would be the last to manage only a human workforce.

Major technology providers, from Salesforce to Microsoft, Google, OpenAI and others, have either begun launching agentic services or are expected to release them this year. And the promise of AI that doesn't merely help a human employee work smarter or faster, but independently takes on whole tasks, workflows or roles, stands to dramatically improve many companies' competitiveness. The next few years may really shake up a lot of industries.

In the pages that follow, we dig into what those 1,900 professionals told us and explore use cases by domain, from IT to procurement to HR. We'll look at their strategies around LLM adoption and tuning, their data platforms and more. Finally, we'll provide highlights of the results for six major industries and the nine countries our research covered.

But first things first. Just how happy are the companies that made the earliest moves into the fresh green field that is generative AI?

### A Global Research Project

The Enterprise Strategy Group, an Informa TechTarget company, surveyed 1,900 business and IT leaders — at companies in nine countries in North America, Europe and Asia — who are actively using generative AI for one or more use case.

For full details, see Methodology, in the appendix.

More than nine in 10 early adopters say that their gen AI investment is in the black. For the 1,268 respondents who said that they have specifically quantified the ROI of their gen AI initiatives, the average return is 41%. Meaning that for every $1 million sunk into a gen AI project, the company is seeing $1,410,000 in business value, through increased revenue, reduced costs or both.

A major question for gen AI deployment is whether to use it internally to help employees, or externally to improve customer interactions. The latter is often seen as riskier, because of potential negative customer experiences and higher risk of exposing sensitive data. The early adopter community seems to be split on the two approaches:

- 55% have prioritized employee-facing solutions to improve productivity and efficiency.
- 44% started with customer-facing solutions to elevate customer experience, satisfaction, etc.

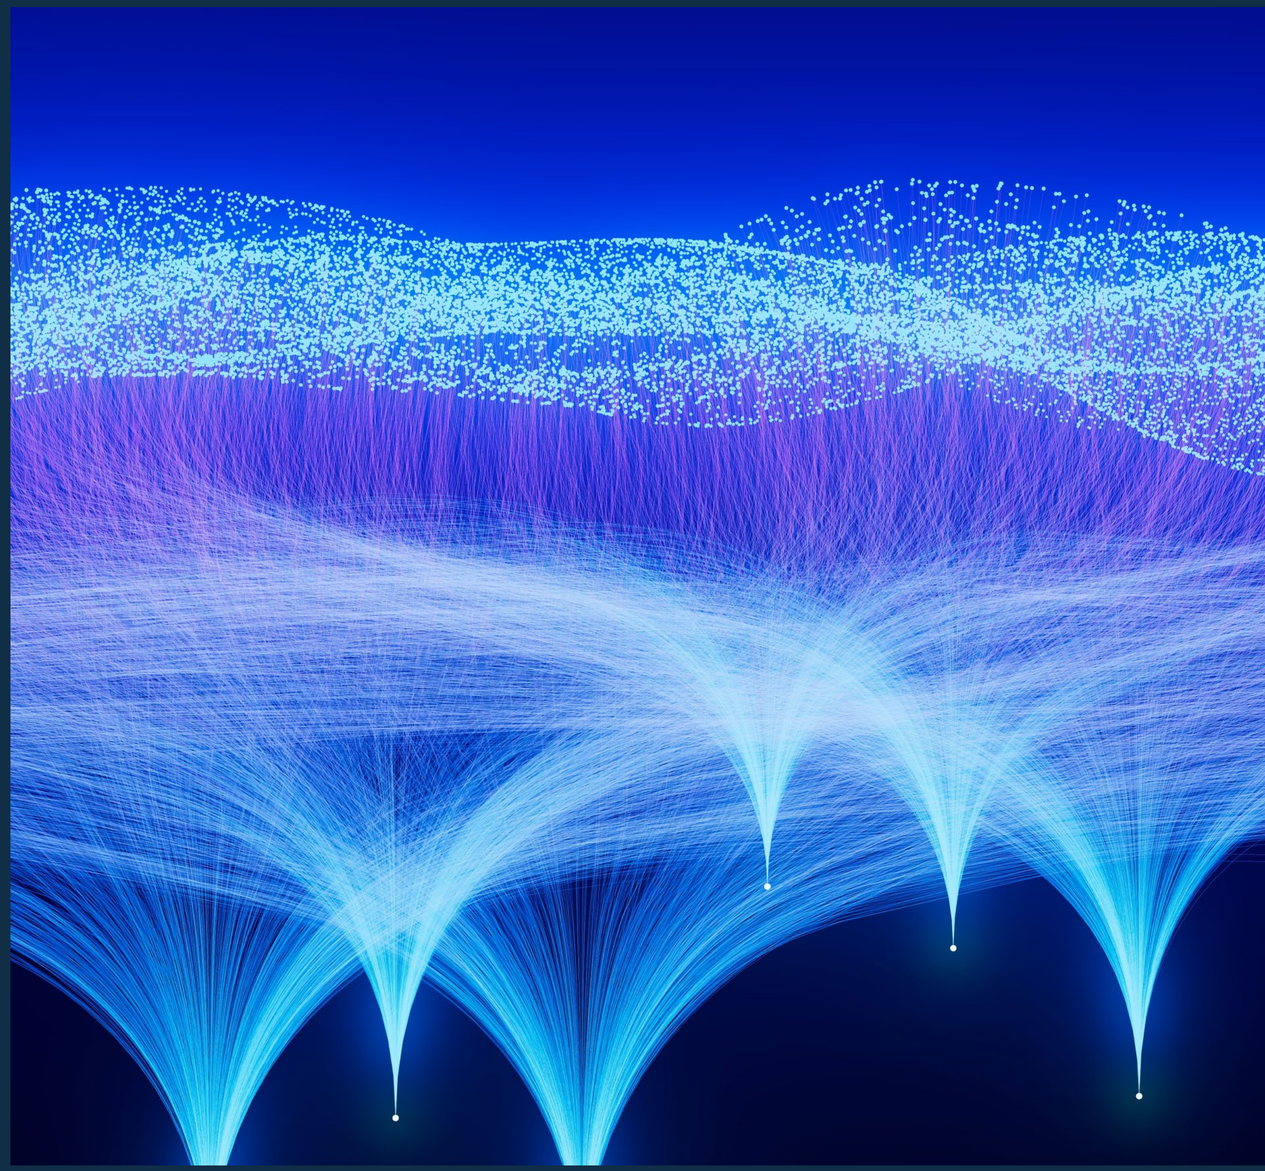

Operational efficiency is among the leading motivations for gen AI adoption, with 51% citing it as a top-three driver. Early adopters are definitely reporting great results:

- 88% report that they've seen a material improvement in efficiency as a result of their gen AI efforts.

Another top driver, called out by 43% of the early adopters, is the desire to improve customer experience. Again, the results are very promising:

- 84% report that their gen AI initiative has led to measurable CX improvement.

The third-biggest motivator — top three for 40% of respondents — is improving innovation outcomes. The numbers look great:

- 84% again report that gen AI is already accelerating innovation outcomes.

All said, 93% of early adopters say that their gen AI initiatives have been very (40%) or mostly (53%) successful.

**92%**

**of early adopters say their gen AI investments are already paying for themselves**

## THINKING BIG

There is a correlation between company size and who organizations are targeting with their gen AI initiatives. Large enterprises (5,000+ employees) were the most likely to report targeting end customers:

- **49% of large enterprises prioritize customer-facing AI initiatives.**
- **42% of large midmarket organizations (1,000–4,999 employees) focus on customers.**
- **38% of small/midsize enterprises (500–999 employees) do so.**

# UNSTRUCTURED DATA IS AI ENEMY NO. 1

Let's not get carried away here. Implementing gen AI is not a trivial project, and the challenges are real. Later in the report, we'll dive deeper into where even the most successful early adopters have struggled. But one challenge stands out and merits discussion up front.

The vast majority of data is unstructured — 80%–90% by many estimates. And what good is a tool that uses your data to uncover hidden trends, to provide smart answers to queries, if it's working with less than a quarter of your data?

Only 11% of the early adopters say that more than half their unstructured data is ready to be used in LLM training and tuning. And that's from a group that was already out in front on AI adoption. Diving into the problems at the data platform level, we found:

- 55% are hampered by time-consuming data management tasks such as labeling.
- 52% struggle with data quality — including issues of error, bias, irrelevance and timeliness.
- 51% say data preparation is too hard.
- 50% cite issues with data sensitivity.
- 42% say they lack the needed range or diversity of data.

And only 8% say that none of these factors are a problem. Must be nice.

Organizing unstructured data in a way that allows automated systems to extract meaning has been a challenge for years. As generative AI and other LLM-related technologies become not only the cutting edge but the enterprise norm, organizations will have to grapple with how* to draw more signal from that noise.

*We can help.

**How much of your unstructured data is AI-ready?**

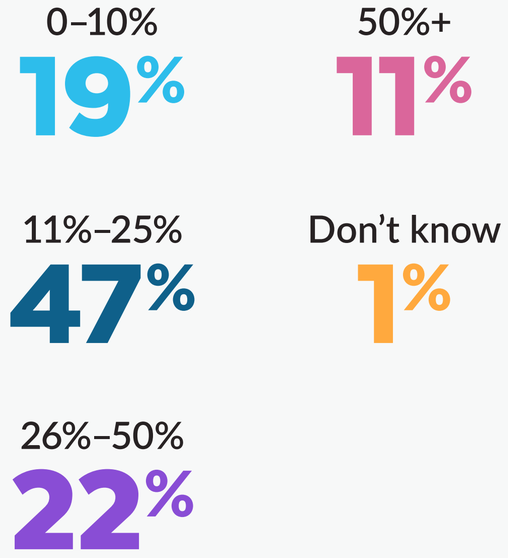

# GEN AI: FACING IN AND LOOKING OUT

Most organizations — 69% — apply generative AI where they think it will do the most good. For some, that's internal, employee-facing tools such as coding copilots for developers; for others, it's customer-facing solutions such as chatbots. That said:

- 18% figure that customer-facing projects would have the strongest business impact but prioritize employee-facing projects, for three main reasons:
  - Insufficient infrastructure and data management capabilities.
  - Security concerns.
  - Worries about AI accuracy.
- 13% say that employee-facing projects would probably have the biggest kick, but are starting with customer-facing options. The top two reasons:
  - They've identified existing solutions that fit their customer use cases.
  - The potential ROI for customer-facing use cases may be lower, but also seems more certain.

All told, the leading gen AI use cases include IT operations (70%), cybersecurity (65%), customer service and support (56%) and software development (54%). However, we note that IT pros consistently identify less use of generative AI than their business counterparts. This may mean that shadow IT has been joined by shadow AI — possibly in the form of online consumer gen AI tools.

For instance, among early adopters:

- 69% of marketers use gen AI for web copy development, ad copy, SEO, etc., yet only 42% of the IT professionals seem to know about it.
- 73% of HR pros use gen AI to screen resumes, train employees, et al., but just 44% of IT are aware.
- 70% of respondents in manufacturing and supply chain operations use gen AI; or 47%, according to IT.
- And 72% of sales orgs report using gen AI, but only 37% of their IT colleagues seem to know about it.

So what's to be done? No magic wand will wave shadow IT — including shadow AI — away. Still, 81% tell us that their organization will establish a gen AI center of excellence at some point in the next 12 months. Such organizations can not only standardize AI rollouts, but move faster and better understand where gen AI is most needed.

**Which of the following lines of business are using generative AI technologies?**

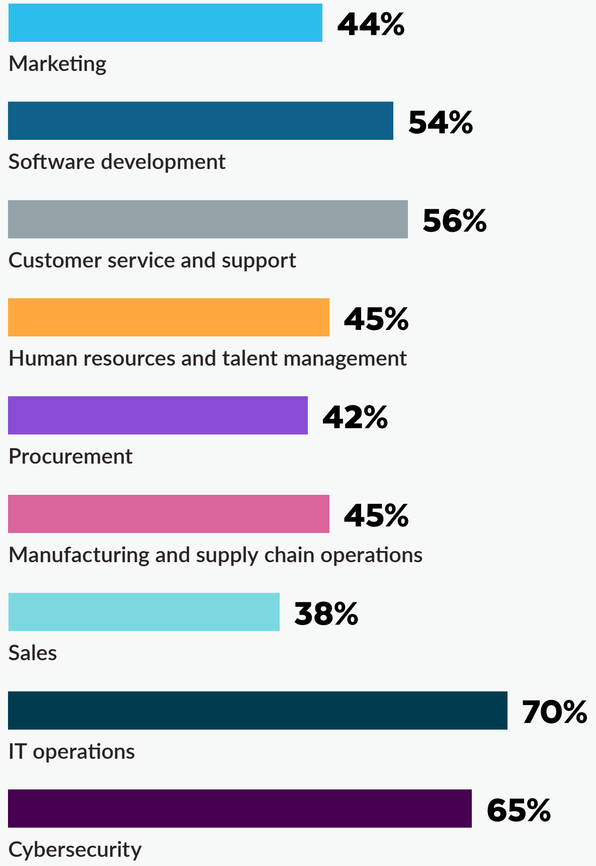

# AN EMBARRASSMENT OF OPPORTUNITIES

Selecting the right use cases for gen AI is a high-stakes game for both companies and their leaders. The ability of generative AI to write credible copy and surface insights from large pools of data has many potential applications in the enterprise. Most organizations have found that they need to be selective about which possible projects they pursue:

- 71% of early adopters agree they have more potential use cases that they want to pursue than they can possibly fund.
- 54% agree that selecting the right use cases based on objective measures like cost, business impact, and the organization's ability to execute is hard.

Limited resources — even where funding is generous — and a need to focus efforts require difficult choices. The promises and proven capabilities of LLMs and generative AI are changing week by week. Enterprises feel like the competitive pressure to keep up with technology advances has been cranked up to 11, and yet, there's so much low-hanging fruit. It's what a military strategist might call a target-rich environment.

Yet the easy answer of "So just pick something and get on with it" isn't easy at all:

- 71% acknowledge that selecting the wrong use cases will hurt their company's market position.
- 59% of respondents say that advocating for the wrong use cases could cost them their job (26% "strongly agree").

That's an application that can manifest in different ways across the many departments of a sprawling enterprise. In the following pages, we'll examine more specifically where and how generative AI is succeeding in various domains within the enterprise.

## GEN AI TRANSFORMED ANALYTICS

Nearly all early adopters say that their gen AI initiatives have:


**Accelerated the generation of analytics**

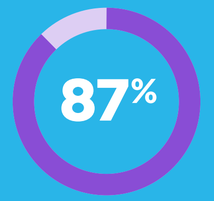


**Helped users better understand analytics**

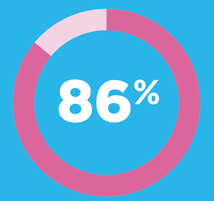


**Made analytics more self-service**

# GEN AI ACCELERATES IT, DEV AND CYBERSECURITY

One of the earliest public use cases for gen AI was as a coding assistant to developers. And one of the earliest gloomy reactions to the technology was its implication, in the wrong hands, for cybersecurity teams. Across the board, our early adopters in software development, IT operations and cybersecurity are enthusiastic about the boost their AI-powered tools are delivering.

**In software engineering, 54% of development teams are using gen AI. Of them, 83% say gen AI is either game-changing (38%) or significant (45%).**

- **Leading use cases:** Analytics and reporting (71%), code generation (63%), and code reviews/QA (60%).
- **KPIs improved:** Code quality (62%), detection and resolution of bugs (56%), and vulnerabilities (54%).

**In IT ops, where 70% of orgs use gen AI, 85% of those early adopters tell us that the overall impact of gen AI is either game-changing (37%) or significant (48%).**

- **Leading use cases:** Analytics and reporting (62%), infrastructure optimization (57%), and cloud cost analysis (56%).
- **KPIs improved:** Costs of operations (47%), uptime (47%), faster response time to tickets (46%), and mean time to detect (45%) and repair (44%) incidents.

**For cybersecurity, 65% of orgs are using gen AI, and 84% of these early adopters say that the overall impact gen AI is having on security operations is either game-changing (39%) or significant (45%).**

- **Leading use cases:** Security hygiene and posture management analysis and prioritization (65%), analytics and reporting (58%), internal knowledge base growth (57%), and data/asset classification (51%).
- **KPIs improved:** Uptime (57%), mean time to detect (54%), mean time to resolve (54%), and manual toil reduction (53%).

Our CIO often notes that code review is a developer's least-loved activity, and any CISO will tell you that basic security hygiene is more than half the battle. Seeing generative AI make such a difference in these areas, as well as the overall productivity enhancements, is particularly encouraging.

**Has gen AI made a measurable improvement in these software development KPIs?**

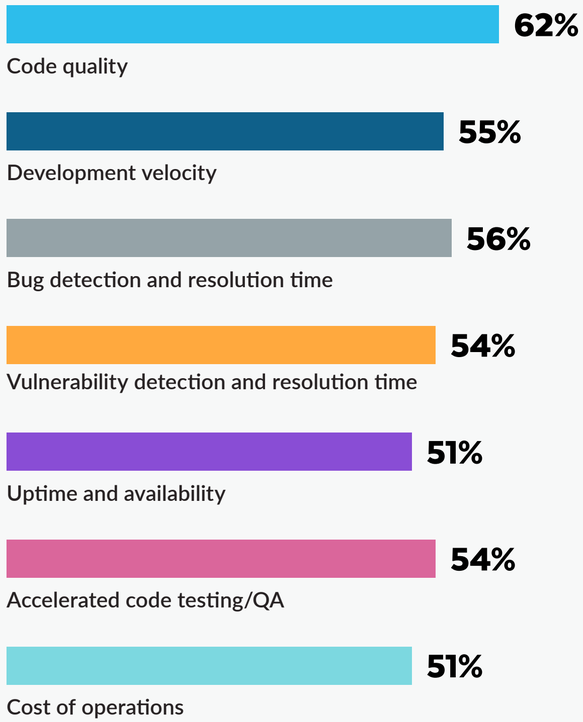

# GEN AI SUPERCHARGES SALES, MARKETING AND CUSTOMER SERVICE

In all the ways that businesses want to reach their potential customers, generative AI is improving outcomes.

**About 38% of early adopters say that their sales teams are using gen AI. Among that group, 77% say that the overall impact of gen AI on sales is either game-changing (30%) or significant (47%).**

- **Leading use cases:** Analytics and reporting (65%), recommendations based on forecasting (56%), and CRM data enrichment (52%).
- **KPIs improved:** Revenue growth (56%), forecast accuracy (55%), and cost of operations (50%).

**In marketing, 44% of early adopters say that they're using gen AI. Of that group, 80% say it's either game-changing (30%) or significant (50%).**

- **Leading use cases:** Generation of personalized marketing content such as offers and recommendations (62%), customer-facing chatbots (60%), and social monitoring and engagement (52%).
- **KPIs improved:** Customer engagement (63%), click-through rates (55%), and cost of marketing operations (54%).

**For customer service and support, 56% of early adopters report using gen AI, and 82% of that group says that the overall impact is either game-changing (32%) or significant (50%).**

- **Leading use cases:** Chatbots for customer service and support (66%), analysis of customer feedback and sentiment (62%), and help managing customer knowledge bases, including FAQs, documentation, troubleshooting guides (57%).
- **KPIs improved:** Customer satisfaction scores (66%) and average first-response time to customer queries (53%).

When we see relatively low adoption rates paired with very high satisfaction rates, we know two things: the early adopters were right, and they're pulling ahead fast.

## Has gen AI improved these marketing KPIs?

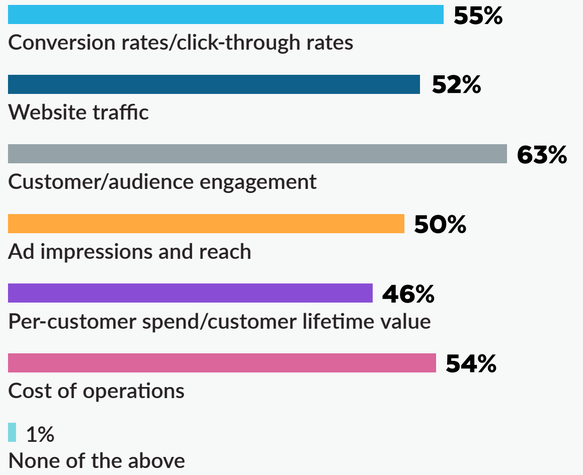


## What sales use cases are you pursuing with gen AI?

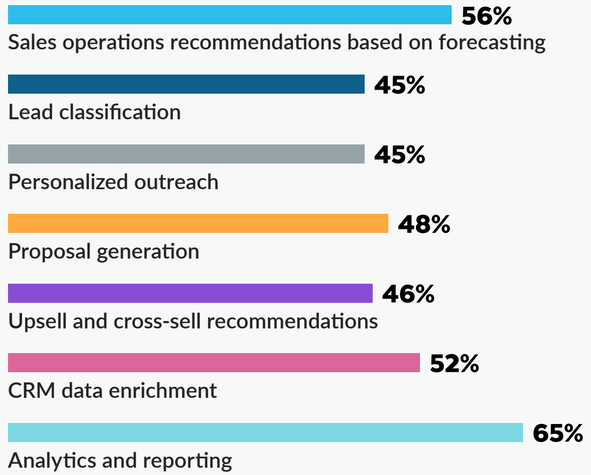

# GEN AI TACKLES HR, PROCUREMENT AND MANUFACTURING

Generative AI is making further inroads — and significant contributions — across nontechnical departments, including in human resources, procurement and manufacturing.

**For HR, 45% of early adopters say that they're using gen AI, of whom, 74% say the overall impact of gen AI is either game-changing (30%) or significant (44%).**

- **Leading use cases:** Gen AI is being applied to both the recruitment operations and ongoing talent management: 63% say their HR teams apply gen AI to performance management processes like reviews, feedback and goal setting; 60% use gen AI to assist with analytics and reporting; 56% use the technology to enable more automated onboarding; and 54% say it's helping to generate job descriptions and interview questions.
- **KPIs improved:** Roughly three-fifths of respondents report improved employee engagement (58%), higher-quality hires (60%), and reduced costs of operations (58% for each).

**Forty-two percent of organizations use gen AI for procurement processes, and 76% of them describe the overall impact as either game-changing (27%) or significant (49%).**

- **Leading use cases:** Analytics and reporting (74%), documentation creation (64%), and contract management (56%).
- **KPIs improved:** Decision making (80%), velocity of operational processes (72%), and supplier relationships (60%).

**About 45% of respondents use gen AI for manufacturing/supply chain. Of that group, 79% say that the overall impact is either game-changing (31%) or significant (48%).** (Note that manufacturing, here, includes all businesses whose processes include some amount of manufacture. For manufacturing as an industry, see page 22.)

- **Leading use cases:** Inventory management recommendations based on demand forecasting (58%), generating maintenance schedules based on equipment data (57%), and analytics and reporting (55%).
- **KPIs improved:** Reduced costs (59%), more production uptime (54%), and demand forecast accuracy (53%).

These strong results point toward continued gen AI adoption across many lines of business. There's still a lot of growth to come.

### Which HR use cases are you pursuing with gen AI?

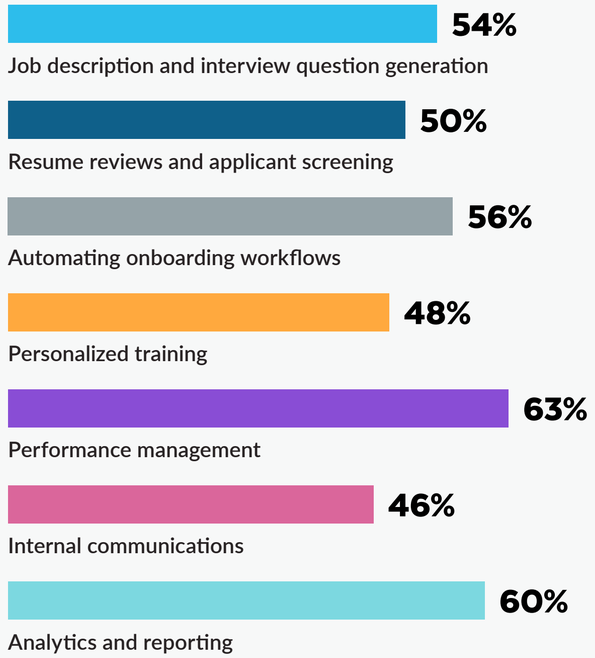

# UNDER THE HOOD

A look at the strategies and challenges organizations report in deploying generative AI and optimizing the data that powers it.

# THE INVESTMENTS DRIVING GEN AI SUCCESS

Just as there are a healthy range of use cases, based largely on the needs of the individual enterprise, organizations also adopt an array of approaches to implementing their gen AI projects. As we'll see, certain fundamental values — and some stubborn challenges — persist.

Let's start with the bad news: 96% of early adopters report that one or more components of their gen AI solutions have cost more to date than was initially anticipated, and 78% say that half or more of their gen AI use cases have cost more than expected to get into production.

- 64% cite compute cost overruns.
- 61% report that supporting software, such as for logging, monitoring and observability, guardrails, model optimization, and developer tools have often cost more than anticipated.
- 58% say data collection, labeling and processing have been more costly than expected.

In the larger picture, those numbers might not entirely count as bad news. Note that these struggles contrast with the 92% of early adopters who report satisfaction — and return on investment. And rather than pull back because they feel stung, these organizations are accelerating adoption to reap further success, increasing their investments in:

- Data (according to 81%)
- Large language models (78%)
- Supporting software (83%)
- Infrastructure (82%)
- Talent (76%)

Overall, 98% of organizations expect their budget for generative AI projects to increase in one or more areas spanning data, models, supporting software, infrastructure or staffing.

**Which best describes your org's gen AI budget in these areas for the next 12 months?**

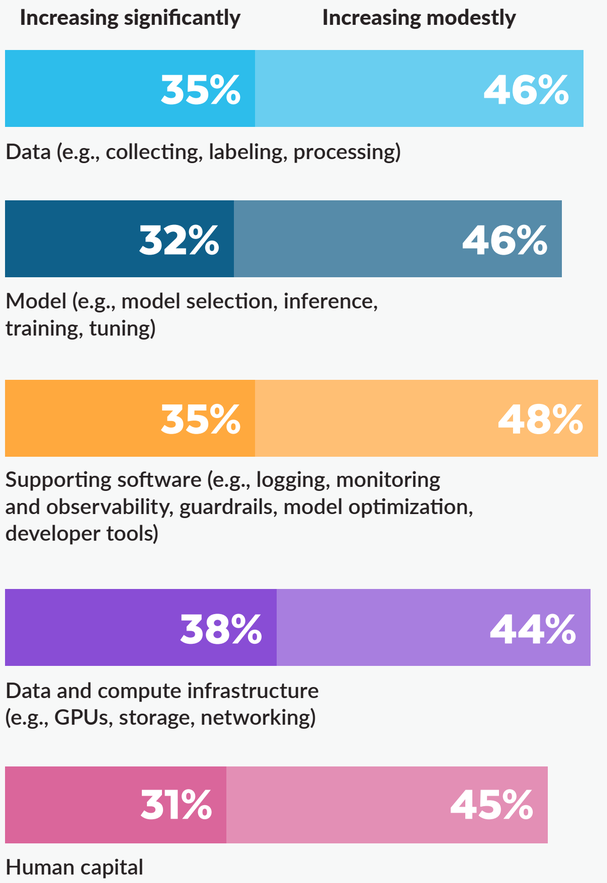

# LLMS (PLURAL): THE MULTI-MODEL MODE

There's no one size to fit all use cases for the LLMs that power generative AI applications. Primarily, the available models break into three categories.

- **Consumer-grade LLMs** (e.g., free versions of ChatGPT, Gemini): Inexpensive, but often come with limited features or customization options. Still, they can work for experimentation and simple uses.
- **Commercial LLMs** (e.g., Azure OpenAI, Amazon Bedrock): Enterprise-class support, frequent updates and more advanced capabilities all contribute to better performance and security. Though well-suited to enterprise-scale applications, commercial LLMs can be expensive and may limit control over the model.
- **Open-source LLMs** (e.g., Meta LlamA, BLOOM): Maximum transparency, full customization, and community support and innovation. Cost-effective, suitable for long-term use and when control over training data is important. But on the other hand, they require more technical expertise to implement and manage.

Though the advantages and downsides of each type of model are clear, an organization's needs don't tend to fall neatly into a single bucket. We found that most early adopters are taking a diversified approach, with commercial models most popular — incorporated by 91% of early adopters.

A diversified, multi-model approach is expected to continue, as well; 59% of early adopters expect that their organization will have deployed three or more LLMs in 12 months, and 93% expect at least two LLMs will be in use.

To manage the resulting complexity, organizations tend to use an LLM gateway. A gateway is an interface that manages LLM interactions, optimizes performance and enhances security. Sixty-eight percent expect to buy their gateway from a vendor, and 28% say they'll attempt to build it themselves. Just 3% say they don't need an LLM gateway.

The diversification we see in models carries through to model infrastructure. Fully 80% of early adopters report using a mix of cloud-hosted and on-premises infrastructure to run their models. Notably, organizations were more than twice as likely to report being mostly reliant on cloud infrastructure versus being primarily reliant on on-premises infrastructure.

**Which of these LLM approaches is your org using?**

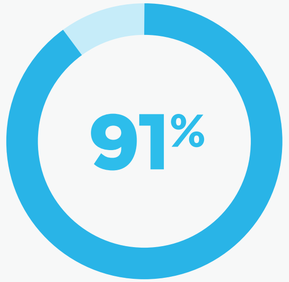


Commercial
(e.g., Azure OpenAI, Gemini, Amazon Bedrock)

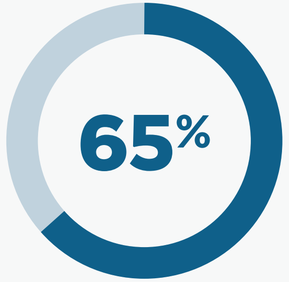


Open source
(e.g., Meta LlamA, Hugging Face BLOOM)

However one sources their LLM(s), they don't come ready to use straight out of the box. Organizations generally face the potentially complex challenge of training, tuning and augmenting an LLM with additional data sets, such as internal or third-party data. Organizations will train or augment LLMs to customize them for specific uses, to enhance relevancy and accuracy of outputs with domain-specific knowledge, and to improve decision-making capabilities.

Training the LLM is all but mandatory; 96% of early adopters say they're training, tuning or augmenting their commercial and open source LLMs.

- 80% are fine-tuning models with their own proprietary data. While this enhances the relevance and accuracy of outputs, it brings potential challenges in terms of data quality, bias amplification, and privacy concerns around both proprietary business information and customer personal data leaking into outputs.
- 71% use RAG. Retrieval-augmented generation enhances the LLM's performance by incorporating relevant external information into the prompt, which enhances the outputs by informing them with greater context and providing answers that are more easily verified. This is achieved without retraining the LLM.

Organizations reported that their augmentation efforts are done to improve user productivity by improving model accuracy and contextual awareness (according to 65% of early adopters) and to improve task-specific knowledge (56%); 46% say they're training their model to better withstand attack or resist misuse. Across the spectrum, the work of training or tuning models is no small task.

- 71% of organizations have found that effective model augmentation requires multi-terabytes of data (or several million documents).

**Which best describes the volume of data your org will use to train, tune or augment its pre-trained commercial and open source LLMs for current use cases over the year?**

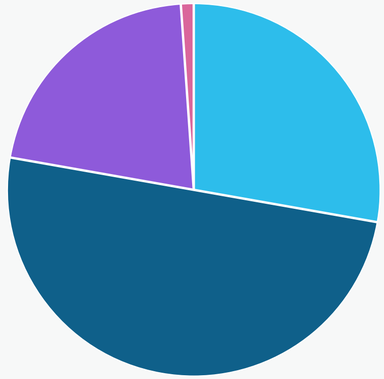

**28%**
Up to hundreds of gigabytes (typically up to 1M documents)

**21%**
Many terabytes and above

**50%**
Up to a few terabytes (typically up to 10M docs)

**1%**
Don't know

# THE DATA MANAGEMENT STRATEGIES DRIVING GEN AI SUCCESS

With many different use cases for generative AI, and so many ways to source and augment large language models, we might worry about a resulting variety of case-specific, siloed data strategies within any given organization. Fortunately, the opposite is true among the early adopters of gen AI. While organizations may employ different models in different scenarios, this diversity of needs and approaches has not fragmented the enterprisewide data strategy:

- 88% of early adopters affirm that they need data strategies and tools that span all generative AI use cases.

A good follow-up question would be: What does that kind of comprehensive, AI-supporting data strategy look like? The research data suggests a preference for a comprehensive platform approach, with respondents saying the tools and solutions they employ must:

1. **Break down data silos:** 64% of early adopters say integrating data across sources is challenging today.
2. **Integrate governance guardrails:** 59% say enforcing data governance is difficult.
3. **Measure and monitor data quality:** 59% say measuring and monitoring data quality is difficult.
4. **Integrate data prep:** 58% say making data AI-ready is a challenge.
5. **Efficiently scale storage and compute:** 54% say it's difficult to meet storage capacity and computing power requirements.

Digging deeper into the challenges, the research points toward the data fueling gen AI success, with early adopters mentioning data quality (45%) and quantity (38%) most often among an array of issues. That's in addition to the data discussed earlier in this report showing that a slim 11% of organizations say the majority of their unstructured data is ready to train or augment their LLMs. Estimating the percentage of unstructured data that's AI-ready, the median was 20.5%.

And in the aggregate, the challenges of implementing gen AI have hurt organizations' deployment timelines:

- 77% of respondents say half or more of the gen AI use cases they've pursued to date have taken longer than expected to get to production.

**How difficult are the following challenges in preparing data for gen AI?**

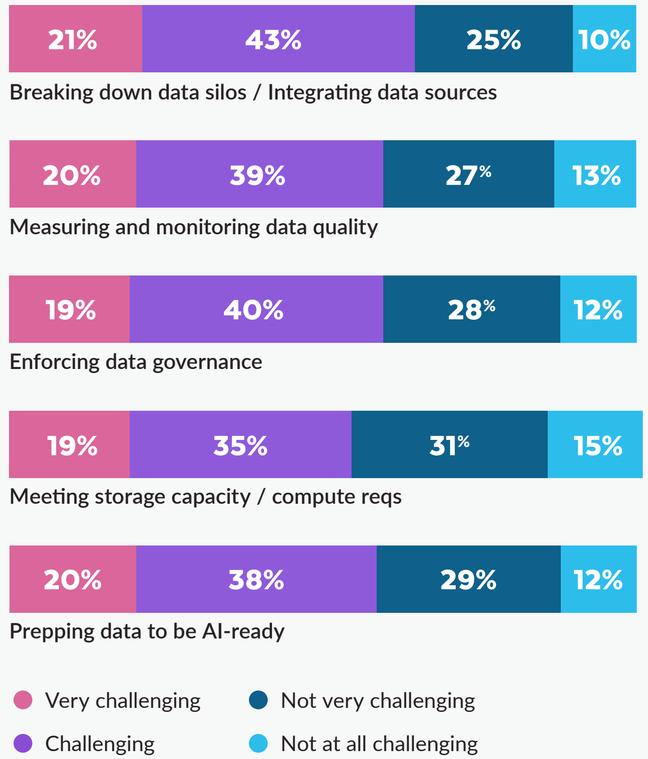

We noted previously that early adopters are using a mix of cloud and on-premises solutions, with a decided tilt toward the cloud. No surprise there — cloud technologies have revolutionized enterprise IT over the past decade-plus. And today, they offer greater scalability and cost control, as well as access to the high-performance computing resources and advanced services and tools needed for gen AI initiatives.

While a plurality of early adopters say they're using an even mix of cloud and on-prem platforms for data warehousing, cloud ultimately comes out ahead:

- 49% use cloud-hosted technologies, with 35% mostly and 14% exclusively using cloud-based platforms.
- Just 5% rely mostly/entirely on on-premises data warehousing.

That cloud preference, mirroring overall global technology trends, is expected to continue, with 81% of early adopters saying that they'll increase their investments in cloud-based data warehousing solutions over the next 12 months. On average, organizations expect a 24% increase in that spending.

When evaluating an investment in cloud-based data warehouses, key features emerge:

- **Security:** 46% rate this critical with a further 38% saying it is important.
- **Advanced AI functionality:** 39% say it's critical, and a further 45% say it is important, to have LLM and ML capabilities under the hood.
- **Integrated analytics capabilities:** 39% rate this critical with an additional 45% saying it's important.

An implication of the majority's desire for advanced AI functionality from their data warehouse provider is that businesses are looking for vendors to provide not just AI-friendly solutions, but platforms that provide the AI capabilities they might otherwise have to build themselves.

### How important are these capabilities when selecting data platforms for gen AI uses?

| | Critical | Important | Somewhat Important | Not Very Important | Not At All Important |
|---|---|---|---|---|---|
| Scalability of storage and compute resources | 34% | 47% | 16% | 2% | 0% |
| Ability to integrate data in either real time or batch | 35% | 46% | 16% | 2% | 0% |
| Analytics capabilities | 39% | 46% | 13% | 2% | 0% |
| Ease of data sharing | 34% | 46% | 18% | 2% | 0% |
| Security and governance controls | 46% | 38% | 13% | 2% | 0% |
| Rapid experimentation and iteration | 25% | 50% | 21% | 4% | 0% |
| Cost monitoring/management tools | 34% | 43% | 20% | 3% | 0% |
| Integrated advanced AI functionality (e.g., LLMs, ML models) | 39% | 45% | 14% | 2% | 0% |

# AGENTS OF RAPID CHANGE

Generative AI is not a single wave of change. Having found that more than half of organizations surveyed have deployed gen AI, and more than 90% of those score it as a success, you could imagine that what follows is steady expansion — leaders adding use cases, laggards catching up. But already in 2025, the next wave of change is upon us.

The first offerings of autonomous agents are in the marketplace. Capabilities and use cases will only increase, and they'll probably increase at the same rapid rate as gen AI has over the past 2½ years, if not faster. Today's early adopters are walking with gen AI today, and they'll be better positioned to run with autonomous agents tomorrow. (Seriously, *tomorrow.*)

In a nod toward the very near-term advent of autonomous AI agents:

- 72% say it is likely that tasks currently executed by my direct reports will be taken over by autonomous agents.
- 70% say they'll be reviewed, at least in part, by AI systems.

And yet, 67% say they've only operationalized a quarter or less of their organization's existing gen AI use cases. Leaders of generative AI initiatives are themselves change agents, and they're very clear-eyed about the amount of change bearing down on us all.

Let's get ready.

**72%**

**of early adopters expect autonomous agents to take over some tasks by the end of 2025.**

# INDUSTRY OUTTAKES

Looking closer at what the data reveals about gen AI uptake in key sectors: Finance, healthcare, manufacturing, marketing/advertising, retail and technology.

## FINANCIAL SERVICES

# FINANCIAL SERVICES FIRMS FOCUS ON SUPPORT AND SECURITY

Though financial services firms are in a data-sensitive, heavily regulated industry, they're ambitious in their pursuit of generative AI. The most commonly cited driver among early adopters is to improve the bottom line.

- 43% of financial services early adopters cite improving financial performance as the key driver, versus 30% of early adopters in aggregate.

Leading use cases reported by financial services firms include customer support (63% versus 56% of all early adopters) and cybersecurity (70% versus 65%).

- **Customer support:** 53% say that gen AI has helped increase the number of tickets handled (versus 48% generally).
- **Cybersecurity:** 62% report that gen AI has helped them grow their internal knowledge bases (versus 57%).

With an eye toward the future, there are a few requirements financial services firms report at higher levels than their peers:

- 87% report that they'll increase their investments in cloud-hosted data warehousing over the next 12 months, versus 81% in the aggregate.
- They also more frequently note (88% versus 84%) that the integrated security and governance controls of these cloud data platforms are critical or important to their ultimate purchase decision.

The financial services industry is highly competitive and a leader in digitalization of services. Given both competitive pressures and the capacity for generative AI to improve digital customer experience and target compelling offers tailored to consumer need, the industry is likely to remain a leader in gen AI adoption.

### What are your org's primary business drivers for gen AI initiatives?

| | Average | Ad/Media | Retail/CPG | Finance | Health/LS | Manufacturing | Tech |
|---|---|---|---|---|---|---|---|
| Operational efficiency | 51% | 43% | 48% | 43% | 53% | 63% | 50% |
| Decision making | 32% | 45% | 32% | 31% | 35% | 34% | 31% |
| Customer experience | 43% | 47% | 52% | 43% | 35% | 41% | 38% |
| New products / services | 40% | 32% | 35% | 36% | 42% | 44% | 44% |
| Cost reduction | 32% | 21% | 37% | 26% | 28% | 37% | 33% |
| Talent management | 26% | 40% | 23% | 27% | 27% | 20% | 28% |
| Risk management | 26% | 36% | 33% | 45% | 36% | 35% | 41% |
| Financial performance | 30% | 31% | 28% | 43% | 32% | 22% | 27% |

## HEALTHCARE & LIFE SCIENCES

# HEALTHCARE FOCUSES ON IT, HR — AND STRONG ROI

Like financial services, healthcare and life sciences is a highly regulated industry with a lot of reasons to be cautious where data security is concerned. But also like finance, healthcare and life sciences companies have embraced generative AI. Early adopters in this industry report higher than average ROI on gen AI spend — 44% versus 41% in the aggregate. The industry's adoption of HR and IT use cases particularly stands out:

- 53% are using gen AI for HR, versus 45% of early adopters generally. In particular, these organizations are bringing gen AI to bear on application screening processes (62% versus 50%), and slightly more often, they credit gen AI with accelerating their time to hire (50% versus 47%). Overall, healthcare and life sciences organizations are more likely to say that gen AI technology is having a game-changing impact on HR (36% versus 30%).
- IT teams in this industry are adopting gen AI a bit more aggressively. 76% say that their IT operations team uses gen AI, compared to 70% in the aggregate. In particular, IT teams in healthcare and life sciences are more often using gen AI to help them detect incidents faster and with greater accuracy (58% versus 49%). These IT teams also say gen AI technology is explicitly helping to accelerate first-contact resolution rates (48% versus 39%), improve resource utilization rates (46% versus 38%), and lower cost of operations (50% versus 47%).

Separately, while they still tend to use multiple LLMs, they expect their approach to LLMs to be a bit less diversified than their counterparts in other verticals, with 51% reporting that they will be using three or more LLMs in the next 12 months, compared to 59% of respondents in the aggregate.

**How do you characterize your org's ROI on gen AI investments to date?**

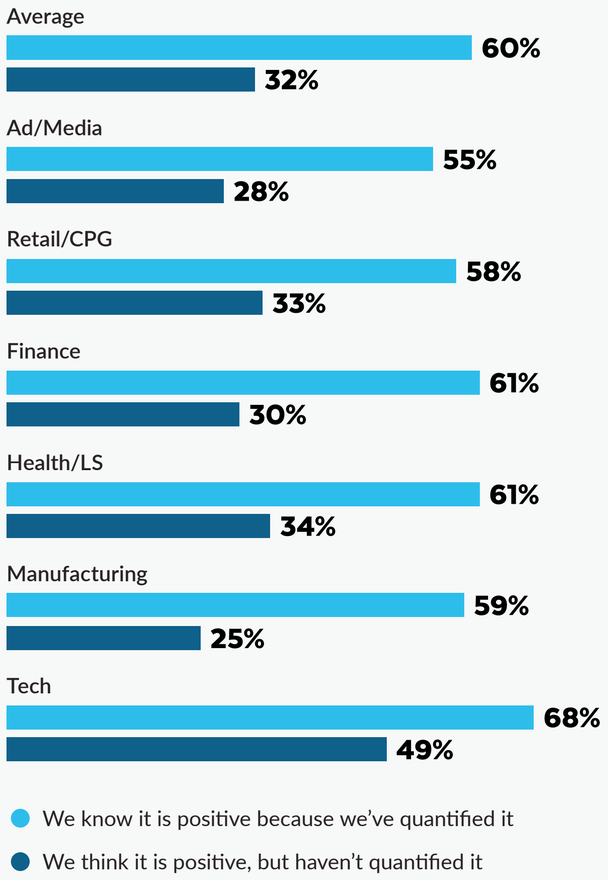

## MANUFACTURING

# MANUFACTURERS TARGET EFFICIENCY WITH VENDOR-LED STRATEGY

The research delivered two sets of data for the manufacturing segment. We asked our entire early adopter cohort about their gen AI use cases, and a number of businesses that are not considered traditional manufacturers still listed “manufacturing and supply chain” as an area to which they’re applying gen AI. As noted above, 79% of those respondents say gen AI has been either game-changing (31%) or significant (48%).

Among manufacturing companies proper, there’s an apparent lag on gen AI adoption compared to other industries.

- Manufacturers most often report using gen AI for their initial use cases (45% versus 36% among all qualified respondents) as opposed to using the technology in several areas.
- Driving their initiatives, manufacturers were the most likely industry to say they’re targeting operational efficiency (63% versus 51% generally) and cost reductions (37% versus 32%).
- And it’s the least likely industry to focus on transformative business results like increased revenue and market share growth (22% versus 30%).

Manufacturers also seem to be relying more heavily than other enterprises on vendors to provide out-of-the-box AI services. More manufacturers say they need both a vendor-provided LLM gateway (81% versus 68%), and that their data warehousing solutions must include strong analytics capabilities (46% versus 39%) and integrated LLMs and ML models (45% versus 39%).

Where manufacturers lead: on deploying gen AI technology to their production/supply chain management teams (71% versus 45%). Further, they are more often applying the technology to assist with inventory management (64% versus 58%) and creating quality inspection protocols (60% versus 54%).

**45%**

**of manufacturers are working only with initial gen AI use cases, versus 36% of all early adopters.**

## MARKETING/ADVERTISING

# MARKETING/AD/MEDIA FIRMS ARE TRAILING — AND THEY KNOW IT

Here's a shocker: The marketing, advertising and media sector is struggling to make gen AI work. Yes, 83% report positive ROI — which is great, but noticeably trails the 92% cross-industry average. These respondents were more likely to report that half or more of their gen AI projects have taken longer than expected to get to production (84% versus 78%). They also more often cite the difficulty of identifying compelling gen AI use cases (77% agree that it's hard, versus 54% generally).

While respondents seem motivated to improve, they do so with a considerable measure of fear: 45% strongly agree that failing to pick the right use cases could cost them their jobs (versus only 26% generally).

The industry is slightly more likely than average to pursue gen AI deployments for internal use (60% versus a 55% aggregate) than for external customers (40% versus 44% generally).

- Notably, about a fifth of respondents in this sector said they thought gen AI had the most potential for customer interactions but were pursuing employee-facing use cases instead.
- Asked why, the marketing/ad/media sector was the most likely (64% versus 50% globally) to cite concerns about accuracy of AI, or just not trusting current gen AI tech for customer interactions.

Looking at marketing use cases, this sector was in line with the cross-industry average in terms of using gen AI for writing copy, analyzing social engagement and running analytics. Where they led the average was in visual creative design (53% versus 48% across industries) and employee-facing chatbots (68% versus 50%). Reflecting the previous paragraph, deployment of customer-facing chatbots and assistants lagged dramatically (45% to the aggregate 60%).

On the plus side, advertising, marketing and media are reporting smoother gen AI deployments in key areas, being more likely than the cross-industry average to say that it has been less challenging than expected to manage usage, such as tracking prompts and responses (27% versus 19%) and ensure uptime (28% versus 22%).

**Which of these challenges have you encountered in curating unstructured data to train, tune or augment LLMs?**

| | Average | Ad/Media | Retail/CPG | Finance | Health/LS | Manufacturing | Tech |
|---|---|---|---|---|---|---|---|
| Data quality | 52% | 52% | 43% | 46% | 50% | 53% | 58% |
| Lacking data range/diversity | 42% | 47% | 33% | 42% | 42% | 38% | 47% |
| Data sensitivity | 50% | 54% | 42% | 50% | 49% | 47% | 53% |
| Data management tasks | 55% | 68% | 45% | 57% | 53% | 49% | 58% |
| Data prep tasks | 51% | 62% | 46% | 55% | 50% | 47% | 50% |
| None of these | 8% | 1% | 13% | 6% | 13% | 13% | 5% |

## RETAIL/WHOLESALE/CPG

# CUSTOMER SUPPORT DRIVES SUCCESS, WITH ROOM TO GROW ON SUPPLY CHAIN, ROI

Looking at the broad retail industry, including wholesale and consumer packaged goods, the data shows some stark differences from the cross-industry average. Retail uses fewer LLMs, with 50% of early adopters using one or two, compared to 39% of the aggregate. Retailers are about half as likely as the average to be using three or more LLMs (12% versus 22%).

It's no surprise that retail would more often include improved customer experience as a top gen AI driver (52% versus 43% on average) and with noteworthy success: 87% say they have positively impacted customer service/support via gen AI projects — tied with financial services and technology for highest figure across industries.

And the hits keep coming. Gen AI adopters in retail:

- Are least likely to report trouble with unstructured data quality (43% of retail/wholesale/CPG versus 52% across industries).
- Have fewer complaints about time-consuming unstructured data management tasks (45% versus 55%).

However, relative to other sectors, it appears CPG companies and retailers may have an untapped opportunity to affect change in procurement and supply chain. To date, just 20% say gen AI has had a game-changing impact on procurement (versus 27% in the aggregate data). For supply chain management, the figure stands at 25% versus 31%. This may be part of the reason that the retail sector reports the lowest quantified ROI: 30% versus 41% across all industries.

### Has gen AI led to material improvements in these areas? (Percentage of "Yes" responses)

| | Average | Ad/Media | Retail/CPG | Finance | Health/LS | Manufacturing | Tech |
|---|---|---|---|---|---|---|---|
| Operational efficiency | 88% | 80% | 85% | 89% | 83% | 91% | 90% |
| Decision making | 84% | 87% | 83% | 82% | 87% | 81% | 86% |
| Customer experience | 84% | 79% | 87% | 87% | 80% | 81% | 87% |
| Innovation | 84% | 79% | 78% | 83% | 82% | 84% | 87% |
| Cost reduction | 81% | 75% | 81% | 81% | 78% | 84% | 80% |
| Talent management | 75% | 73% | 73% | 75% | 73% | 72% | 78% |
| Risk management | 81% | 77% | 77% | 86% | 82% | 82% | 82% |
| Financial performance | 79% | 79% | 74% | 81% | 81% | 78% | 81% |

## TECHNOLOGY

# DAZZLED BY GEN AI, TECH FIRMS LEAD IN EARLY ADOPTION

You'd expect technology companies to be leading on gen AI adoption, and you'd be right. But the downside seems to be that the sheer range of possibilities is sometimes daunting.

Tech companies have gone in head first:

- 70% report using gen AI for multiple use cases versus 64% among all early adopters.
- 33% say they use three or more LLMs versus 22% across all industries.
- And 43% say they built their own LLM gateway, compared to 28% in the aggregate.

As to drivers, tech companies were the most likely to report targeting innovation (44% versus 40% broadly) and applying the technology to their software engineering operations (66% versus 54%). And when gen AI is used by development teams, those at tech companies are more often using solutions to generate (68% versus 63%) and debug (66% versus 60%) code. Tech companies also most often see developer velocity accelerate thanks to gen AI technology (64% versus 55%).

However, leadership in gen AI is not without challenges:

- 76% say they have more use cases to pursue than their budget can fund (versus 71%).
- 61% say it's hard to identify the right use cases to pursue based on factors like cost, business impact, and ability to execute (compared to 54%).

The irony being, of course, that the industry with the greatest ability to spot the potential of gen AI ends up being slightly penalized for that deeper insight. Yet all told, the sector that brought us generative AI continues to be a trailblazer in its application.

**It is hard to identify the right use cases to pursue based on factors like cost, business impact, and our ability to execute.**

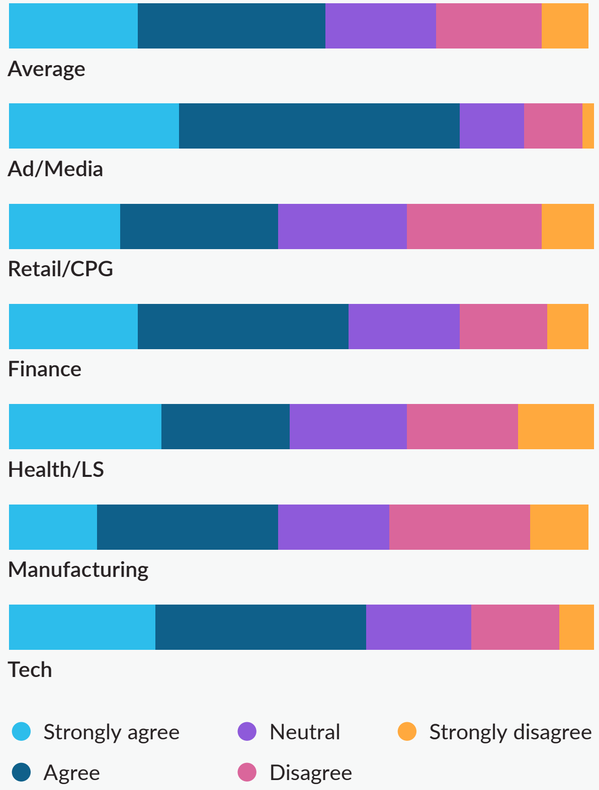

# GLOBAL PERSPECTIVE

Data highlights from eight regions: Australia and New Zealand, Canada, France, Germany, Japan, South Korea, the United Kingdom and the United States

## AUSTRALIA AND NEW ZEALAND

# DESPITE DATA AND TALENT CHALLENGES, CUSTOMERS ARE HAPPIER DOWN UNDER

The data from respondents in Australia and New Zealand (ANZ) paints a nuanced picture, with early adopters here differing from the global aggregate in terms of intention, results and challenges.

**Focus:** Compared to the global average, ANZ organizations more often cite enhancing customer satisfaction, delivering personalized experiences and improving customer engagement as key goals (53% versus a 43% global average). Given that focus, it's no surprise that ANZ respondents were more likely to emphasize gen AI projects for end customers (53% versus 44%) and less weighted toward employee-facing efforts (47% versus 55%). And given that more public approach, ANZ more often stresses the goal of improving security and compliance (55% versus 39%).

**Results:** ANZ's early adopters were more likely than the global average to report their efforts with gen AI were enabling the organization to make better, faster decisions (91% versus 84% globally) and reported a slightly higher than average ROI attributed to gen AI spend (a 44% return versus 41% on average globally). And that customer focus pays off with 85% reporting improvement to customer experience.

**Challenges:** Respondents in ANZ more often reported certain challenges in their gen AI initiatives:

- **Many competing priorities:** They often struggle to identify the right use cases to pursue (71% versus 54%).
- **Unstructured data:** ANZ more often cites a lack of data diversity/range (56% versus 42%), time-consuming data management tasks (62% versus 55%) and data preparation (59% versus 51%) as difficult areas. And they more often say that it's hard to break down data silos (76% versus 64%).
- **Staffing up:** The cost of staffing up for gen AI was more often higher than expected (63% in ANZ versus 48% in the aggregate).

**Outlook:** ANZ organizations tend to allocate more funding to gen AI, with 32% putting more than a quarter of their tech budget for the next 12 months toward gen AI (versus 25% globally).

**Where is your org putting more emphasis with current and near-term gen AI projects?**

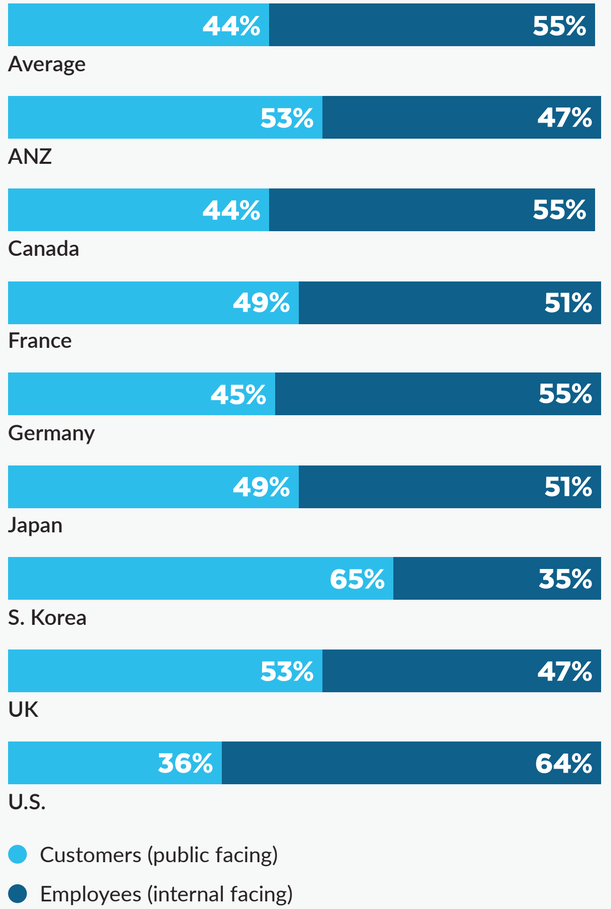

## CANADA

# CANADIANS HESITATE ON GEN AI, BUT LOG GOOD OUTCOMES

The overall story in Canada is a struggle to catch up and keep up.

**Focus:** Suggesting that Canadian companies are earlier in their gen AI journeys, respondents were notably more likely to say that they're pursuing only an initial use case (45% versus the 36% global average). Use of gen AI tended to be highly focused within technical teams, the only lines of business where Canadian respondents outpaced the global average were IT (74% versus 70% globally), cybersecurity (69% versus 65%), and software development (56% versus 54%). Conversely, Canadian respondents less often reported gen AI penetration in marketing, HR, procurement or sales.

**Results:** Those technical teams did report some outsized results to date:

- **More agile IT**: In Canada, IT teams have more often reduced both their average first response time to tickets (53% versus 46% globally) and total handle time (46% versus 40%).
- **Faster cybersecurity:** Canadians were more likely to say that security had accelerated its mean time to detect (61% versus 54%).
- **Accelerated development:** Respondents in Canada beat the global average, 65% to 54%, on using gen AI to accelerate development teams' code testing and QA.

**Challenges:** Canadian respondents are more apt to cite issues with data quality (49% versus 45% globally), legacy integrations (36% versus 31%) and costs (38% versus 30%), which may complicate their efforts to drive adoption into additional lines of business.

**Outlook:** Reflecting their slower start, Canadian companies more often say they'll be using only one or two LLMs 12 months out (51% versus a 39% global average). In fact, Canada trails the global average for expectations to significantly increase budgets for data (31% versus 35%), models (30% versus 32%), infrastructure (33% versus 38%) and people (25% versus 31%).

Overall, Canadians are less likely to say that gen AI will represent 25% or more of tech budgets (20% versus 25% globally). Gen AI proponents should hope that strong initial results noted above drive further investment.

### How are you using gen AI and LLMs right now?

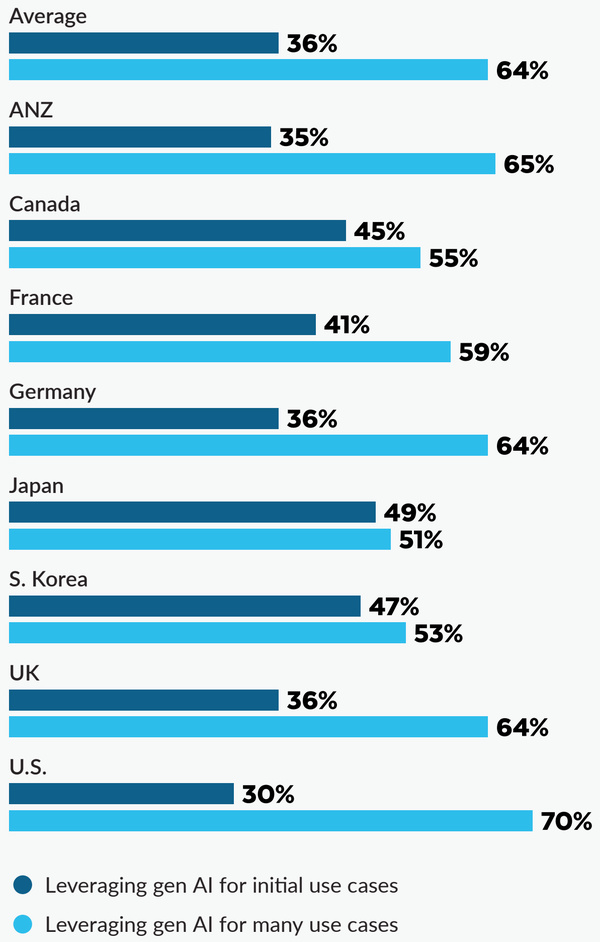

## FRANCE

# FRENCH COMPANIES START MODESTLY, RAISE THE BAR ON INNOVATION

The French market appears to lag slightly in gen AI maturity, with lower overall ambitions and more off-the-shelf gen AI deployments. The data shows a relatively low degree of challenges, but also more muted business impacts.

**Focus:** Respondents in France more often say they're still working on an initial gen AI use case (41% versus a 36% global average). And their strategy for gen AI seems less mature: French respondents less often use methods to train or augment the LLMs with proprietary data, such as retrieval-augmented generation (59% versus 71%). They're also less likely to enable end users to easily query databases with text-to-SQL services (52% versus 66%).

**Results:** All this said, French respondents were global leaders on achieving material improvement to organizational innovation (91% versus an 84% global average). If such results motivate decision makers to double down on gen AI, more data challenges may come to the fore. Certainly French organizations will need to up their investments: currently just 9% of them will allocate more than 25% of their technology budgets to gen AI, compared to 25% of organizations in the aggregate.

**Challenges:** France's less ambitious use of proprietary data sources in the context of gen AI LLMs is probably why French respondents report fewer and less acute data challenges associated with their gen AI projects. The French complain less about:

- Lack of unstructured data diversity/range (33% versus 42%), time-consuming unstructured data management tasks (35% versus 55%) and difficult data prep (41% versus 51%).
- The challenge of breaking down data silos (42% versus 64%).
- A struggle to measure and monitor data quality (40% versus 59%).

**Outlook:** Competitive pressures, gen AI advances and their own positive results should drive French organizations to do more with generative AI. And with such efforts will come greater awareness of the challenges with unstructured data, data silos, data quality and governance.

### How much of your unstructured data is AI-ready?

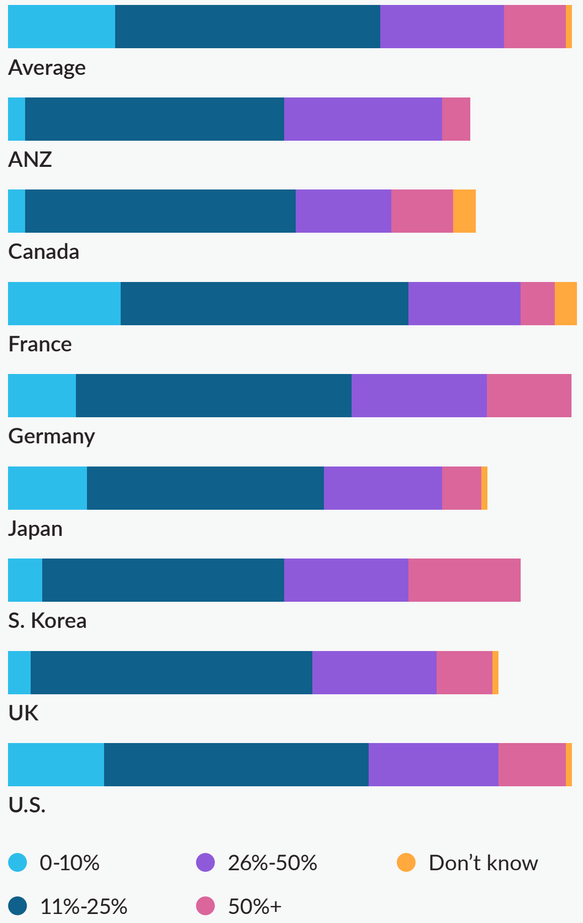

## GERMANY

# GERMANS SHOW DATA DEXTERITY BUT NEED INFRASTRUCTURE AND BUDGET

Germany has an interesting mix of strengths and obstacles: a more sophisticated approach to data and gen AI projects that's slowed by the availability and cost of compute and data infrastructure.

**Focus:** Germans more often say that they're fine tuning their LLMs with proprietary data (88% versus an 80% global average). And when it came to unstructured data, Germans suffer less from time-consuming data management tasks (46% versus 55%) and data preparation (45% versus 51%), and they say that more unstructured data is AI-ready (39% say that more than 25% of their data is versus 33%). All of this reflects a strong commitment to getting value from their data.

**Results:** The data's ready, but are their lines of business? Germans trail the global average on bringing gen AI to most of their internal organizations (see chart). While their ROI to date is only just below par (with 89% saying their initiatives are in the black, near the 92% global average), German organizations are also:

- Less likely to report improved innovation thanks to gen AI (78% versus 84%).
- Less likely to report better talent management outcomes (70% versus 75%).
- Lower quantified ROI (34% versus 41%).
- More likely to say that it's hard to identify the right use cases (62% versus 54%).

**Challenges:** Infrastructure blockers abound: 69% struggle to meet gen AI's storage and compute requirements (versus 54% in the aggregate). Not surprisingly, then, German respondents most often strongly agree that they have more use cases to pursue than their budgets can support (35%, against a global average of 29%).

**Outlook:** They are the most likely to say that their organization will leverage more than three large language models within the next 12 months (34% versus 22%). Looking ahead, the data implies German organizations should be focused on exploring efficient ways to bring their high-quality approach to generative AI to end users across internal teams.

### Which of these lines of business are using gen AI today?

| | Average | ANZ | Canada | France | Germany | Japan | S. Korea | UK | U.S. |
|---|---|---|---|---|---|---|---|---|---|
| Marketing | 44% | 44% | 40% | 49% | 36% | 53% | 45% | 42% | 44% |
| Software development | 54% | 52% | 56% | 50% | 39% | 46% | 55% | 62% | 58% |
| Customer service & support | 56% | 58% | 56% | 52% | 49% | 45% | 53% | 61% | 60% |
| HR | 45% | 52% | 42% | 37% | 35% | 45% | 38% | 44% | 48% |
| Procurement | 42% | 53% | 41% | 39% | 41% | 31% | 30% | 41% | 44% |
| Manufacturing & supply chain | 45% | 44% | 45% | 39% | 41% | 47% | 49% | 41% | 46% |
| Sales | 38% | 38% | 34% | 39% | 35% | 37% | 31% | 40% | 40% |
| IT operations | 70% | 63% | 74% | 61% | 63% | 57% | 65% | 74% | 75% |
| Cybersecurity | 65% | 59% | 69% | 67% | 57% | 56% | 59% | 69% | 69% |

## JAPAN

# JAPAN WALKS THE GEN AI WALK — AND PREPARES TO RUN

Japan seems to be lagging, with respondents there listing fewer challenges, but also more muted business impacts.

**Focus:** Japanese organizations seem to be taking a slower, more cautious approach to gen AI. They're the most likely to report pursuing only an initial use case (49% versus 36% in the aggregate). They trail on model training and augmentation, being less likely to report using RAG (65% versus 71% globally), tuning with proprietary data (75% versus 80%), and providing text-to-SQL functionality (53% versus 66%).

Strategically, respondents in Japan are the least likely to report a focus on customer service and support (30% versus a 43% global average) and financial performance (18% versus 30%), but the most likely to be employing gen AI for cost cutting (43% versus 32%).

**Results:** Less maturity and a more tactical mindset may explain gen AI's relatively low ROI in Japan, where organizations that have measured the return put it at 30%, compared to 41% globally. A brighter point: Respondents in Japan were the most likely (93% versus 88% globally) to report improved operational efficiency tied to gen AI projects, meaning results are well-aligned to priorities.

**Challenges:** Japanese organizations more often cite a lack of use cases (29% versus 18% overall), employee skills (49% versus 35%) and a belief that existing solutions are immature (31% versus 21%). In terms of data-centric challenges, respondents in Japan were the least likely to say it's "very challenging" to:

- Break down data silos (5% versus 21%).
- Measure and monitor data quality (5% versus 20%).
- Enforce data governance (3% versus 19%).

This may change as their gen AI efforts grow more ambitious.

**Outlook:** Acceleration and ambition seem likely: Respondents in Japan are the most likely to say that practical use cases are emerging and understanding of gen AI is rising (45% versus a 26% global average). As organizations in the country advance their projects, both challenges and outcomes should increasingly align with those of other countries.

### What quantified ROI has your org seen to date on gen AI investments?

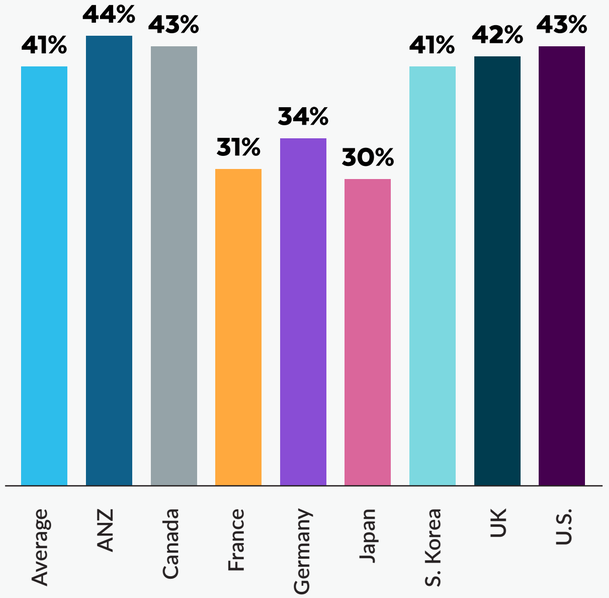


Two-thirds of respondents (66.7%) said their org had quantified gen AI ROI. The number represents the ROI to date. An ROI of 41 means that for every dollar spent, $1.41 has been earned through cost savings, increased revenue or both.

## SOUTH KOREA

# AMBITIOUS SOUTH KOREANS ROLL UP THEIR SLEEVES

South Korean respondents display remarkable ambition in terms of their work with LLMs and data, yet they have identified fewer use cases, and deployed gen AI across fewer lines of business.

**Focus:** South Korean businesses really want to get hands-on. They're the most likely to report use of open source models (79% versus a 65% global average) and much more apt to report augmenting models via RAG (82% versus 71%), fine-tuning embeddings (81% versus 72%), and by leveraging text-to-SQL services (74% versus 66%).

**Results:** South Korean organizations' results seem to reflect a high data fluency:

- South Korean organizations were the most likely to report both a high degree of data management expertise for unstructured data (35% versus 20%) and that the majority of their unstructured data is "AI-ready" (20% versus 11%).
- They're also less likely than the global average to find it very challenging to measure and monitor data quality (10% versus 20%), enforce data governance (11% versus 19%) or prep data for gen AI (10% versus 20%).

**Challenges:** South Korean organizations more often say that they struggle with gen AI's technical complexity (39% versus 34% generally). They also are more often challenged by a lack of use cases (26% versus 18%) and (more often than any other country) with coordination across their organizations (31% versus 23%). This may be why they trail the global average in deploying gen AI across many lines of business. (See chart, p. 30)

**Outlook:** South Korea's early adopters seem determined to press on. There's a lot of optimism in this group, which is notably more likely to predict using more than three LLMs within the next 12 months (32% versus 22%). And respondents in South Korea were the most likely to report they will use "many" terabytes of data (30% versus 21%) to train, tune and augment models. Expect all this hard work to pay off with the widespread use case adoption that has so far eluded them.

### In which of these ways is your org enhancing its LLMs with proprietary data?

| | Average | ANZ | Canada | France | Germany | Japan | S. Korea | UK | U.S. |
|---|---|---|---|---|---|---|---|---|---|
| Retrieval-augmented generation (RAG) | 71% | 71% | 66% | 59% | 74% | 65% | 82% | 73% | 73% |
| Fine-tuning the model with proprietary data | 80% | 79% | 78% | 73% | 88% | 75% | 76% | 76% | 82% |
| Fine-tuning embedding models | 72% | 72% | 74% | 78% | 72% | 76% | 81% | 68% | 70% |
| Text-to-SQL services | 66% | 76% | 63% | 52% | 68% | 53% | 74% | 70% | 67% |

## UNITED KINGDOM

# THE UK TARGETS SUCCESS AND PREPARES TO BUILD BIG

The UK data shows targeted success. Early adopters there are applying gen AI to specific use cases at elevated rates, achieving specific outcomes with great effect.

**Focus:** UK respondents beat the global average in citing both operational efficiency (57% versus 51%) and innovation (46% versus 40%) as primary business drivers for gen AI initiatives. (And not to spoil the results section, but it's going well: 93% see efficiency gains, versus 88% in the aggregate, and 89% report greater innovation, versus 84% globally. See chart, p. 35)

Keys to success include UK organizations being the most likely to report that they're training, tuning or otherwise augmenting their LLMs (98% versus an impressive 96% global average). And they tie (with Germany) as the most likely to use mostly or entirely cloud-hosted infrastructure to run their LLMs (47% versus 40%) which may ease adoption and utilization across lines of business.

**Results:** Getting into specific use cases, we see more concrete examples of how UK organizations are driving greater operational efficiency. UK early adopters are the most likely to report:

- Using gen AI to aid software development (62% versus a 54% global average) and to free up dev resources by applying gen AI to code reviews and debugging (69% versus 60%).
- Applying gen AI technology to customer service and support (61% versus 56%) — and more often shrinking their typical first response time as a result (59% versus 53% globally).
- Using gen AI technology to help cybersecurity teams (69% versus 65%) — and to have logged less spent on manual tasks (64% versus 53%) and lower overall cost of operations (56% versus 47%).

**Challenges:** It's not all sunshine and roses. UK respondents were least likely to report a high degree of capability at using unstructured data for gen AI (10% versus 20% on average). They more often note time-consuming data management as a top struggle with unstructured data (63% versus 55%).

**Outlook:** The question is: Will UK respondents' strengths overcome their weaknesses? Data shows that they're calling for backup: 93% expect to budget even more for cloud-based data warehousing solutions for gen AI (versus an 81% global average), and 47% (compared to 39% globally) say it's critical that any data warehousing solution for gen AI must have integrated LLMs and ML models.

**Is your org training, tuning or augmenting its commercial and open source LLMs with proprietary data?**

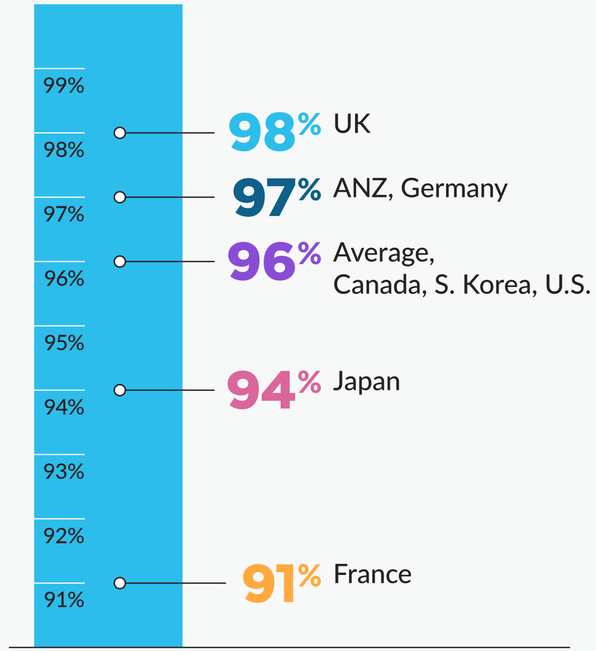

## UNITED STATES

# U.S. PULLS AHEAD ON GEN AI AND INVESTS TO KEEP IT THAT WAY

The United States seems to have the most gen AI momentum, with 70% of U.S. early adopters reporting that their organization leverages generative AI not only for an initial proof of concept, but in multiple use cases (versus a global average of 64%). Perhaps they're driven by a heightened sense of urgency: These U.S. early adopters were also the most likely to strongly agree that failing to pursue the right use cases could cost their company its market position (35% versus 30% globally).

**Focus:** U.S. respondents beat the global average on deploying gen AI technologies for IT operations (75% versus 70%) and cybersecurity (69% versus 65%).

**Results:** U.S. companies cited improved financial performance (i.e., increases in market share, revenue and profitability) as a driver for gen AI more often than the global average (35% versus 30%), and were the group to most often report that gen AI is delivering on that promise (84% versus 79%).

U.S. early adopters, more often than in any other country, say that they've been "very successful" at operationalizing gen AI to achieve their business goals (52% versus a 40% global average) and were also the most likely to say that their ROI was both explicitly quantified and positive (71% versus 60% globally).

Drilling down, they were also more likely than average to claim a game-changing impact from their IT ops (44% versus 37%) and cybersecurity (43% versus 39%) efforts.

**Challenges:** When asked about five particular data difficulties (breaking down silos, measuring/monitoring data quality, enforcing data governance, meeting storage and compute requirements, and prepping data for gen AI), U.S. respondents were 3-5 percentage points more likely than the global average to deem them "very challenging."

Fifty percent of U.S. respondents strongly agree that their organization needs to adopt data strategies and tools that can span all gen AI use cases, compared to 44% globally.

**Rate your org's success in operationalizing gen AI to achieve its business goals.**

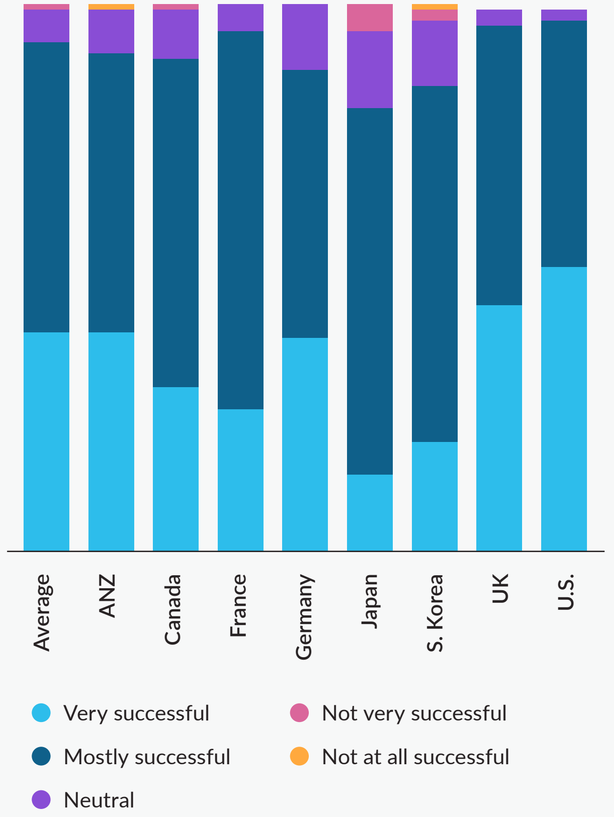

**Outlook:** The U.S. seems poised to keep leading. U.S. respondents were the most likely to say that it's "very likely" that their organization will establish a generative AI center of excellence in the next 12 months (47% versus 39%). They also led the globe in planning significant gen AI investments in:

- Data management (41% versus 35%)
- Models (38% versus 32%)
- Supporting software (42% versus 35%)
- Infrastructure (43% versus 38%)
- Human capital (36% versus 31%)

U.S. companies significantly outstripped the average in saying that they expect to invest more in cloud data platforms, 88% to 81%. Other countries were below that average, except for UK respondents, which beat even the U.S. cohort, at 93%.

Finally, U.S. respondents seem to have the most faith in autonomous agents, with 36% saying they expect some tasks currently done by direct reports to be handled by agents by next year, compared to 29% globally. (Only Germany, at 31%, also exceeded the average.)

## Has gen AI led to material improvements in these areas? (Percentage of "Yes" responses)

| | Average | ANZ | Canada | France | Germany | Japan | S. Korea | UK | U.S. |
|---|---|---|---|---|---|---|---|---|---|
| Operational efficiency | 88% | 85% | 87% | 89% | 87% | 93% | 85% | 93% | 88% |
| Decision making | 84% | 91% | 81% | 85% | 84% | 75% | 71% | 81% | 87% |
| Customer experience | 84% | 85% | 81% | 81% | 83% | 73% | 73% | 89% | 88% |
| Innovation | 84% | 83% | 85% | 91% | 78% | 74% | 76% | 89% | 85% |
| Cost reduction | 81% | 76% | 83% | 76% | 84% | 83% | 77% | 82% | 82% |
| Talent management | 75% | 81% | 67% | 71% | 70% | 65% | 62% | 76% | 80% |
| Risk management | 81% | 81% | 80% | 83% | 77% | 79% | 73% | 87% | 82% |
| Financial performance | 79% | 75% | 79% | 73% | 77% | 70% | 73% | 82% | 84% |

# APPENDIX

How we did it, and what to do next.

## APPENDIX

# METHODOLOGY

In preparing to create this report, Enterprise Strategy Group conducted a comprehensive online survey, which was fielded between Nov. 21, 2024, and Jan. 10, 2025. All respondents represent organizations with 500 or more employees. About 39% of respondents represent organizations with more than 5,000 employees; 47% represent orgs with 1,000–4,999 employees; and 14% come from companies of 500–999 workers.

The intent of this report, and the research that underpins it, is to better understand the experiences enterprise organizations are having with generative AI technologies. As such, the survey sought to find organizations meeting two key characteristics: First, organizations already using gen AI to augment and execute business processes in production. Second, the gen AI tech being used must go beyond consumer-grade "off-the-shelf" solutions — such as subscriptions to ChatGPT — to commercial and open-source models that can be trained and tuned with proprietary data to improve accuracy and relevance.

The process of finding this cohort of early adopters allowed us to observe how broadly gen AI has been adopted by enterprises to date. Of 3,324 respondents, 1,900 (57%) said they are using commercial or open source generative AI solutions.

The 1,900 early adopters who completed the survey were drawn from IT and cybersecurity (52%), software development (13%), data operations (8%), and other lines of business (e.g., marketing, customer support, manufacturing, 27%). To qualify, respondents must have reported both knowledge of and influence over their organization's AI strategies. A range of seniorities are represented, from C-level executives to senior individual contributors. Additionally, the research includes responses from across the globe, including the United States (45%), Canada (8%), the UK (8%), France (8%), Germany (8%), Australia and New Zealand (8%), Japan (8%), and South Korea (8%).

The margin of error for this sample size is +/- 2 percentage points at the 95% confidence level. The totals presented in figures and tables throughout this report may not add up to 100% due to rounding.

# JOIN THE AI DATA CLOUD

Generative AI is moving the enterprise into a new era, and early adopters are finding tremendous success — and positive ROI — already.

The research shows that businesses want a data platform that is truly easy to use, trusted by thousands of successful customers and able to connect your entire data estate. Explore how Snowflake can automate platform management, bring AI to your data, and unify crucial security and governance features.

Control hidden costs, unsilo your data, and tap into unstructured data like never before. See how the AI Data Cloud securely connects businesses globally to deliver AI, power applications and more.

Learn more

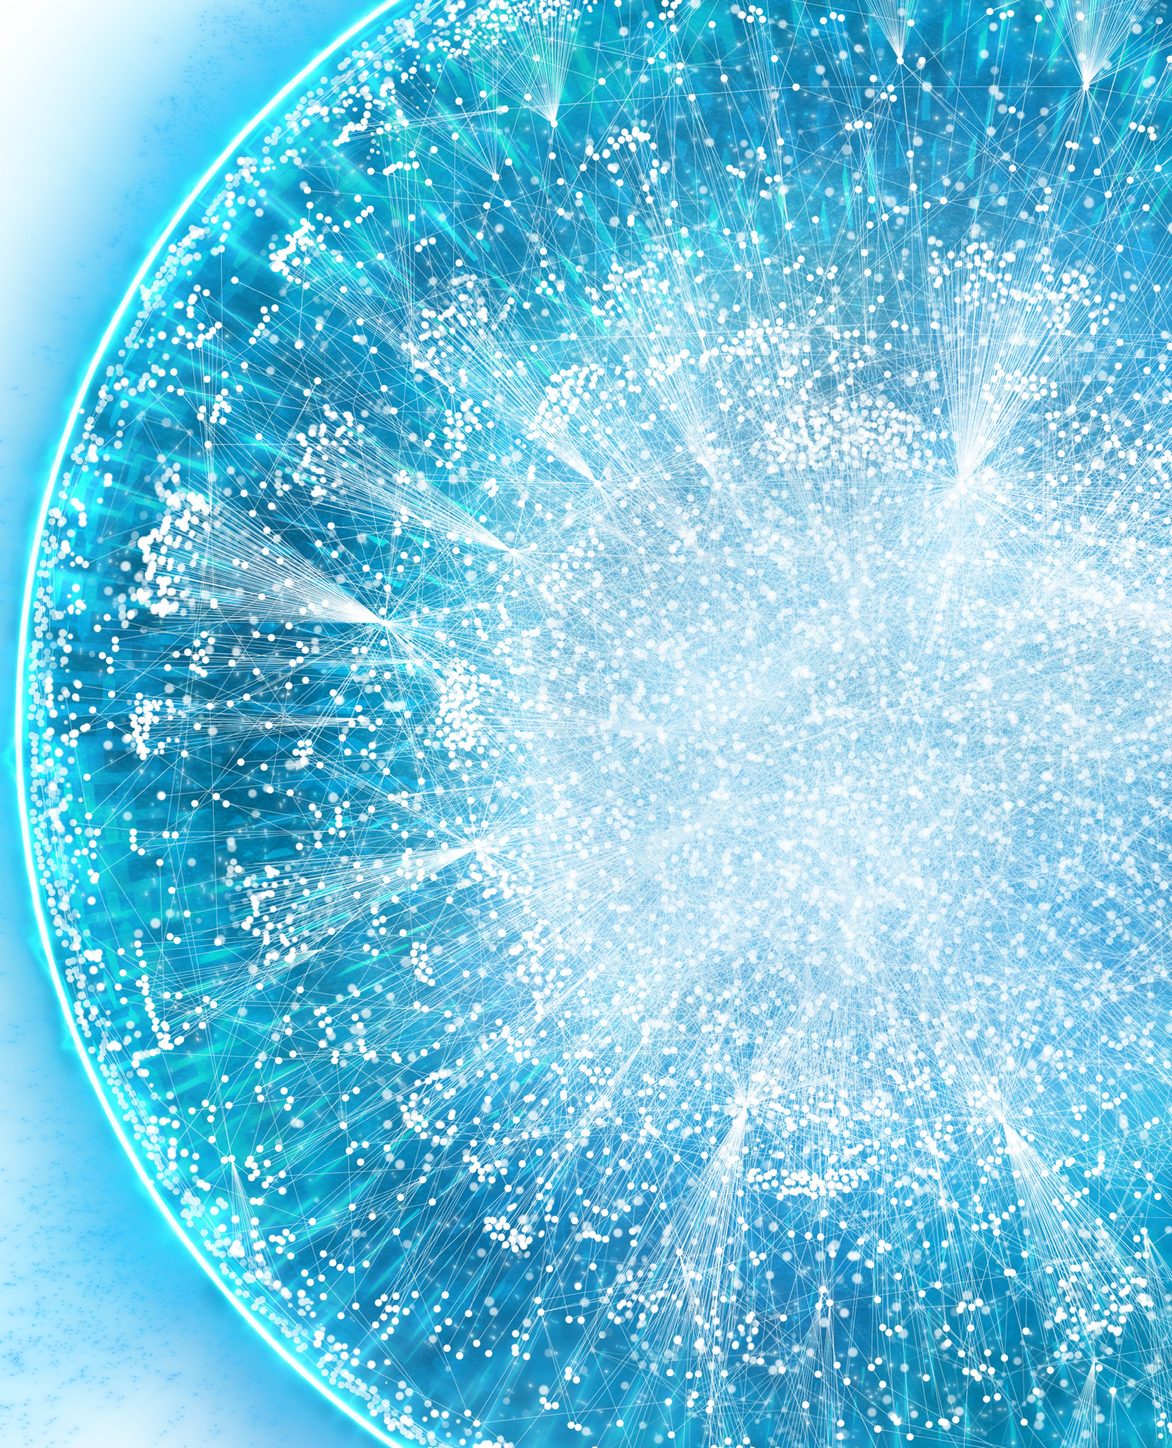

Snowflake makes enterprise AI easy, efficient and trusted. More than 11,000 companies around the globe, including hundreds of the world's largest, use Snowflake's AI Data Cloud to share data, build applications, and power their business with AI. The era of enterprise AI is here.

Learn more at **snowflake.com**

(NYSE: SNOW)