snowflake®

AI + DATA PREDICTIONS 2026

The rise of agentic AI will create a new kind of enterprise and a new kind of worker

The Year of Agents and Ecosystems

It seems like every new year is labeled a “breakthrough year for enterprise AI.” And pretty much every year is. Just looking at the recent past, 2024 was the year businesses got creative, experimenting with explosively hot generative AI technologies to varying levels of initial success. By 2025, the focus was on ROI and proving that all the hype and FOMO wasn’t getting in the way of solid business value. (Countless surveys, including one published by Snowflake, found that many AI-forward organizations were measuring solid return on their gen AI investments.)

So here comes 2026, and another evolution in enterprise-class artificial intelligence. Two things, broadly, are on the horizon: extending ROI from the level of individual projects to comprehensive, strategic AI ecosystems, and agentic AI. The former reflects a maturing of the enterprise itself, as the CDO and other leaders integrate individual successes into a data and AI strategy that helps every team and contributor perform better. The latter reflects a maturing of the technology, as large language models evolve into large reasoning models, and as these LRMs become increasingly capable of reliably taking action with less human supervision.

This year’s report considers the ever-shifting state of AI technology, the outlook and opportunities for enterprise adoption, and the challenges, particularly around security and data governance. It draws from in-depth conversations with more than a dozen leaders and experts at Snowflake. Click through the summaries for an expanded view of these topics:

The AI landscape

Choke points and paths of progress

- Feedback loops will improve agents
  - EU regulations may advance, not hinder, innovative AI initiatives
- A dominant AI protocol will facilitate agentic development and prevent vendor lock-in
- Open-source foundation models will break the hold of the handful of giants
- Agents will start small, stitching together ‘micro-agents’ into effective tools
- Context windows and memory will be the keys to better AI agents
- Postgres will be a foundational technology for agentic AI

Enterprise uptake

- Data strategy will determine AI readiness — and AI outcomes
- What's *not* in your data will limit agentic decision-making

The AI-augmented workforce

- Workers will (have to) master human-AI collaboration and communication
  - AI coding assistants should make developers 33% more productive by 2027
- Everyone will (have to) become a strategic thinker
  - *What* junior engineers?
- Workers will (need to) develop cross-functional and orchestration skills
- For workers and the enterprise to thrive, leaders must provide continuous, contextual learning capabilities
  - AI-created applications will be the new general-purpose spreadsheet
- Humans will remain in the loop as interpreters and quality controllers

AI drives cybersecurity (for better and worse)

- Agents as cyberweapons will help human attackers refine and scale their attacks within a year
- Cybercriminals will use increasingly sophisticated 'dark AIs'
- In three years, AI agents and tools will finally close the security talent gap

Human creativity will persist

- Generative AI will accelerate and improve human creativity — but it will take smart, skillful approaches to avoid it being a crutch

Key industry predictions

- Retail: Gen AI is flipping customer 360 on its head
- Financial services: Shifting back to a data-first mindset
- Manufacturing: AI adoption takes center stage

The AI Landscape

In 2025, trying to understand the state of AI was an exercise in industrial-scale cognitive dissonance. Last summer, Anthropic CEO Dario Amodei said that generative AI technologies were developing so quickly that massive layoffs across the global workforce were nigh. About two weeks later, Apple reported that its research with large reasoning models showed that these ostensibly brighter cousins of LLMs couldn't reason their way through a puzzle that might bamboozle a high school sophomore — which might cast doubt on the near-term potential of agentic AI.

So which is it? Is the AI industry flying down the freeway at 90 miles an hour, or is it (and are we) stuck in stop-and-go traffic while obnoxious radio ads implore us to *ACT NOW! DON'T MISS OUT! on this once-in-a-lifetime opportunity*? The general consensus at Snowflake is, "a little bit of both."

"It's like my commute home from the office," says Anupam Datta, AI Research Lead at Snowflake. "It goes pretty fast, but there are choke points. That's AI in the coming years: In general, it works well, but there are choke points we have to invest in to get AI agents to scale."

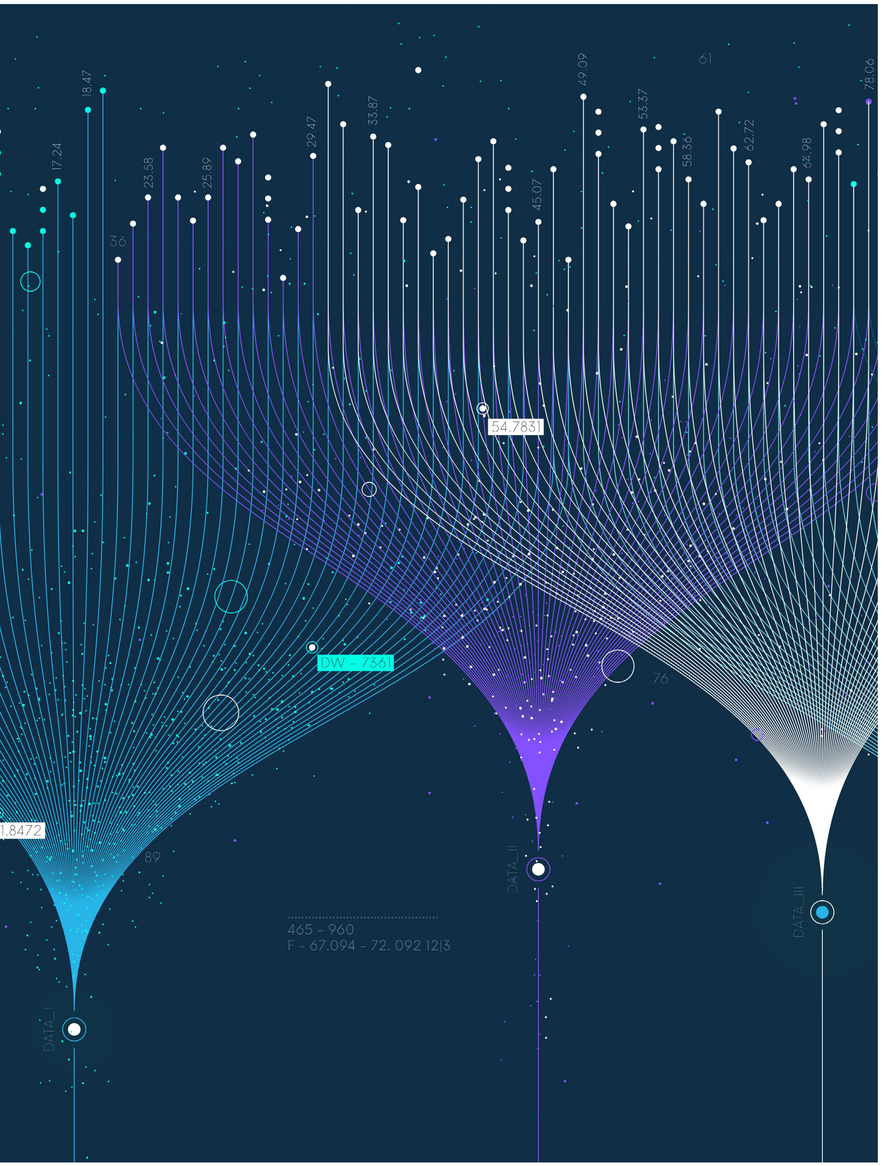

Choke points and paths of progress

Datta, who co-leads Snowflake's AI research team, breaks the choke points slowing the adoption of LLM and LRM systems into a three-part challenge. "To make agents more reliable and trustworthy, three factors are going to be at play: evaluation, prompt-based and post-training optimization, and AI safety and security."

Evaluation, or validation of a model's success, is perhaps the key 2026 challenge on the mind of Snowflake Director of AI Infra Mona Attariyan.

"It will be very hard to rely on agents if we don't have a way of systematically measuring their accuracy," she says. "There are ideas for how to do this, but it's not well established yet. And the main way to do it is to test it against queries you know the answers to. But for every agent, or every domain, there's a different corpus of knowledge involved, so it's complex."

"Enterprises are already demanding that reliability be quantified, because that's what it takes to succeed in the enterprise," notes CEO Sridhar Ramaswamy. "I mean, there is exactly one answer to 'How much money did Snowflake make yesterday?' That's not a matter of doubt or opinion."

> **Enterprises are already demanding that reliability be quantified, because that's what it takes to succeed in the enterprise."**
>
> **—SRIDHAR RAMASWAMY**
> CEO, Snowflake

Feedback loops will improve agents

Feedback loops, Ramaswamy says, will be essential to quickly improving agentic outcomes. “At first, Google’s search algorithm ranked pages based on who pointed to whom,” he says, “but the secret sauce was when they factored in which pages users actually clicked on, and for which query. Similarly, the more feedback loops that are set up in the natural way that AI products operate, the more powerful they’re going to become.”

Coding copilots are already getting tons of feedback as users accept, modify or reject outputs, he adds. “Once there are feedback loops in every product that you use, more complicated use cases become possible.”

Reinforcement learning, benchmark testing and other validation methods will continue to evolve, as will the actual technologies they validate. Dwarak Rajagopal, VP of AI Engineering and Research at Snowflake, says agents will be developed to rate the performance of other agents.

“With longform workflow agents, we’ll evolve self-verification for each step within the agent,” he says. “Another agent will grade the agent at each step, like grading a student, and you’ll have to set, basically, the minimum acceptable grade-point average for the entire workflow.”

Datta adds that carefully crafted evaluations are already becoming part of production AI pipelines to catch mistakes and identify blind spots.

“The evaluation-to-optimization loop will be much, much more sophisticated and become part of production pipelines in a more streamlined fashion over the next year or two,” he says.

EU regulations may advance, not hinder, innovative AI initiatives

Even fans of regulation — those who accept that a certain requirement or guardrail ensures a necessary public good — may agree that such laws tend to limit innovation. Certainly that has been true in 2025 in the United States, where there have been proposals to ban regulation of AI for up to a decade specifically to avoid stifling an innovative industry. Principal Data Strategist Jennifer Belissent, who is based in France, says the perspective is different in the European Union.

“We have GDPR and the EU AI Act, but the business perspective here is not negative,” she says. “While there has been a lot of hand-wringing about the regulation, I see that much less from Europeans, and more from Americans looking in.”

She says that many European customers and colleagues find that existing regulations open up opportunity, providing reassurance that data is well-governed, which can allow an organization to adopt new AI use cases securely and efficiently.

“The mandated transparency has sparked collaboration across the organization because different teams are seeing what others are doing,” she says. “They either want to do it themselves or they want to borrow the model, and the resulting reuse is a further incremental benefit. Several of the people that I’ve spoken to here say, ‘It’s not stifling innovation, it’s accelerating it.’”

A dominant AI protocol will facilitate agentic development and prevent vendor lock-in

Agents will work across multiple systems and databases, so a big challenge will be making sure that all these AI systems can communicate with each other. One part of the solution is open data formats, but Snowflake Co-Founder Benoit Dageville says that even more important than open formats and APIs will be the emergence of a dominant AI protocol.

"There are multiple attempts right now to create the defining protocol," he notes. "The acceptance of that winning protocol will be super important."

Much as TCP/IP provides the framework to connect billions of internet-enabled devices, AI agents will need a protocol to standardize communication among agents and with other systems. At present, three notable entrants in the AI protocol sweepstakes are MCP, the Model Context Protocol from Anthropic; A2A, Google's Agent-to-Agent protocol; and ACP, the Agent Communication Protocol championed by IBM and the Linux Foundation.

While it's too soon to declare which protocol or protocols will become widely adopted standards, expect consensus in the next year or so. The need to make complex AI systems work in an effective, unsiloed way, is not the only driving force. There's also enterprise customers' exhaustion with proprietary systems.

"Enterprises have learned the hard way that vendor lock-in is difficult and expensive to break," Anupam Datta says. "Cloud-vendor lock-in led to cloud-agnostic services, including Snowflake, and data format lock-in led to Apache Iceberg and other open standards getting a lot of play. With agents, enterprises have realized early that lock-in to a specific agent-building framework is not a good idea, and they are insisting on interoperability. Protocols, such as MCP, are a response to that explicitly articulated customer need."

> **"There are multiple attempts right now to create the defining protocol. The acceptance of that winning protocol will be super important."**
>
> **—BENOIT DAGEVILLE**
> Co-Founder, President of Product, Snowflake

Open-source foundation models will break the hold of the handful of giants

Because it takes very scarce talent and very, very large budgets to create ever-improving foundational models, the assumption has been that a tiny handful of giants will corner that market. The cost issues seem unlikely to change, and 2025 saw a considerable escalation in competition for AI talent. So we'll probably always be able to count the makers of a Claude or a ChatGPT on one hand.

> **A lot of innovation is happening in the post-training phase — after the early training phase has produced a pretty good base model — as new, specialized data is fed into it."**
>
> **—SRIDHAR RAMASWAMY**
> CEO, Snowflake

But early in 2025, researchers in China debuted Deepseek, which used innovative training methodologies to achieve remarkable reasoning capabilities, competitive performance — and potentially lower cost and greater transparency. Suddenly the cash-consuming "build bigger to be better" approach wasn't the only game in town. This helped convince Ramaswamy that the overall model market would be bigger than initially thought.

"A lot of innovation is happening in the post-training phase — after the early training phase has produced a pretty good base model — as new, specialized data is fed into it," Ramaswamy says. "Many will start with the standard open source model and then go ahead and build more custom ones. So my thinking has shifted from 'models are going to be the realm of only four or five companies.' I do think that there will be a broader democratization over the next year."

Datta agrees. "Lots more players can play in that space," he says. "The market opportunity is *really* big."

The agentic evolution

The promise of agentic AI is systems capable of problem-solving, rather than following rigid, predetermined workflows; more like coworkers than tools. Vendors in 2025 have hyped various tools as "agents," combining deterministic automation flows with statistically driven probabilistic reasoning, but the fantastic leap beyond generative AI assistants hasn't broadly materialized yet.

Snowflake Head of AI Baris Gultekin notes that while definitions of agentic AI are varied and evolving, solid examples have emerged.

"Customer support is a very complex process for many companies, and AI can significantly reduce cost and improve customer experience," he says. "Agentic AI is working quite well there. End-to-end solutions exist already, and even large companies with very bespoke processes can easily tweak AI to learn from documentation. They'll still have humans in loop, but with substantially increased overall efficiency."

He notes that agents are also proving effective in the insurance industry with the complexities of claims processing. In terms of agentic development, he points to a study released in March by research nonprofit METR that found that agentic AI's ability to automate hours of work is doubling every seven months. The report concludes that while agents today can't carry out substantive projects unattended, "in under a decade, we will see AI agents that can independently complete a large fraction of software tasks that currently take humans days or weeks."

Agents will start small, stitching together 'micro-agents' into effective tools

While generative AI's big public debut was amazing because it could generate long discourses on virtually any topic (with varying but steadily improving accuracy), expect AI agents in the enterprise to start granularly, with simple tasks and human oversight. Over time, they'll handle greater complexity with less human intervention.

CIO Mike Blandina likens the steady buildup of agentic capabilities to microservices, the idea of individual components that can be used to build more flexible and adaptable software systems.

"In the next couple of years, we'll see 'micro-agents' that do a task or a few small tasks really, really well," he says. "Then we'll combine those agents like Lego blocks to do bigger tasks."

Through that process, agentic systems will develop with abilities comparable to a capable coworker, rather than an intern who can handle only tightly defined assignments. But don't expect a single agent with a wide range of capabilities.

Instead, hundreds of specialized AI models will be optimized for specific tasks like coding or performing web research. When a user demands a complex task or asks for a wide variety of competencies, an orchestration layer will select the right sub-agents to execute the broader, high-level assignment one step at a time.

"Verticalized, bounded agents will be easier to productionize, especially in a high-trust environment," says Dwarak Rajagopal. "Multiple bounded agents will excel at specific tasks, and you'll have an orchestrator on top of them to route queries to the correct agents. That makes it much easier from a verification perspective."

Context windows and memory are the key to better AI agents

Agents will continue to rapidly improve, says Vivek Raghunathan, Snowflake's SVP of Engineering and Support. In the near term, he expects there to be perhaps a few dozen tools that can handle somewhat complex instructions successfully. Their development will rely on improvements to two things: memory and context windows.

"Context windows — which define the number of tokens or amount of data a model can work with at any time — will grow exponentially, from thousands to millions of tokens," he says. This will lead to more dynamic, comprehensive problem-solving abilities and will fundamentally change how AI systems can approach complicated tasks.

Secondly, improvements around memory mean that systems will be better able to determine what information to store, when to retrieve it and how to apply it to a situation. "It's a more human-like capability, to be able to remember the larger context of a situation to solve the problem at hand," he says. This will dramatically improve an AI agent's ability to work autonomously on complex tasks.

All of this will be amazing in the laboratory, but the more tokens, the more context involved in an agentic undertaking or a gen AI conversation, the higher the cost. A significant challenge will be making such technological breakthroughs affordable for the enterprise or end user.

Postgres will be a foundational technology for agentic AI

As AI models evolve from conversational tools into agents that reason and take multistep actions with limited human oversight, they'll need tools suited to the job(s). In particular, they'll need transactional databases suited to fast-paced, real-time data and memory, and that's where Postgres comes in.

Postgres will be valuable for agentic AI platforms because it provides a fast and reliable database optimized for the transactional and real-time demands of AI agents. Postgres is an online transactional processing (OLTP) database geared to high-frequency, low-latency read-and-write operations. This is vital because agents generate a continuous stream of small, real-time data points that must be stored and accessed instantly.

"Postgres is the lingua franca for structured data and a great experimentation layer," says Dwarak Rajagopal, Snowflake's VP of AI Engineering and Research. The developer community's embrace of Postgres and its rich ecosystem of extensions makes it ideal for building agentic applications. "With agents writing more code and doing more things, Postgres is just an easier tool for agents to use."

Enterprise Uptake

AI's capabilities are increasing rapidly. But exactly how, and how rapidly, businesses will adopt gen/agentic AI depends not just on the tech, but on the vision, data strategy and budget commitment of the enterprise.

"It's very easy in the tech industry to get caught up in all the new developments," says Snowflake Principal Data Strategist Jennifer Belissent. "But when we feel like we're racing into the future, a lot of other businesses and consumers haven't even gotten onto the road yet."

More than a decade into the explosion of cloud and SaaS solutions, many businesses are still dealing with legacy systems and the "cloud migration journey." On the AI front, adoption of traditional machine learning — for predictive analytics, process automation and more — is still underway. So we're probably a long way from gen AI and agents being as ubiquitous across the business world as email and spreadsheets, but early adopters may be the first to benefit from the greater data insights and efficiencies these technologies promise.

The steady growth of agentic capabilities that CIO Mike Blandina noted will require new levels of management and preparation. Belissent says that successful organizations will prioritize the management of agents, not just their rollout.

"You can't just define the job — you have to define a job well done, and measure it against expectations, outcomes and ethical guardrails," she says. "Constructive feedback loops will help agents refine their responses and perform better, or call in human intervention when necessary."

Data strategy will determine AI readiness — and AI outcomes

As generative and agentic AI systems continue to evolve, organizations with well-governed, high-quality data will be best positioned to harness these technologies for meaningful business outcomes, says Anahita Tafvizi, Chief Data and Analytics Officer at Snowflake. These leaders will achieve earlier — and ultimately greater — competitive differentiation.

"The pace of AI innovation is extraordinary, with new capabilities emerging every week. But most enterprises are still struggling to translate that innovation into impact," Tafvizi says. "Gaps in data readiness, governance frameworks, and organizational skills are holding many companies back. Based on what I'm hearing from customers across industries, the next wave of AI adoption won't be evenly distributed — it will create clear divides between companies, and even entire industries, that are prepared to scale versus those still stuck laying the groundwork."

That's not just about "breaking down silos" to make enterprisewide data available to AI systems. Mike Blandina notes that while increasingly capable and accurate AI systems will drive adoption efforts, governance and security challenges will necessarily restrict most companies to a measured pace.

"When AI delivers an accurate answer, you also have to be sure that private or proprietary data isn't being exposed," he says. "Should the user have the permissions to see this answer? Is your marketing chatbot giving out employees' Social Security numbers and customers' credit card numbers? That's not about the AI, that's about how you govern and secure your data."

And while CDOs and CIOs will be on the front line of developing and enforcing data strategy, it's not a discrete, small-team job.

"Responsibility for AI-ready data cannot lie only with a single individual or department," says Jennifer Belissent. "It has to be a shared, cross-functional effort with multiple stakeholders, from business leaders to technical teams, data owners and end users."

What's *not* in your data will limit agentic decision-making

Generative AI delivers increasingly reliable information as models and post-training techniques improve. But agentic AI is meant to not only find answers, but to turn reasoned outputs into independent action. Reaching that point, says Streamlit Co-Founder Amanda Kelly, requires not only getting your data in order, but clarifying your thinking, as well.

"A lot of companies don't really know how they make decisions, and you can't just apply AI on top of that to fix the underlying problem," she says.

Tafvizi agrees. "The logic behind higher-order business decisions is nuanced, often informal, undocumented or context-dependent," she says, "and until that changes, agents will work best in tactical, structured scenarios, like marketing budget optimization for a specific campaign, or prioritizing your lead list for outreach to customers."

Kelly says that mastering the challenges around decision-making will give considerable advantage in the next three to five years — the horizon on which more capable AI agents are expected to arrive. "The companies who have these really good processes, who understand their decision-making, who know what they need to do, and are able to apply agents and AI in this tight framework are going to be big winners," she says.

> **"The pace of AI innovation is extraordinary, with new capabilities emerging every week. But most enterprises are still struggling to translate that innovation into impact."**
>
> **—ANAHITA TAFVIZI**
> Data and Analytics Officer, Snowflake

The AI-augmented workforce

The macro effect of AI on employment is unclear. In the micro, certainly some jobs will be lost, other jobs will be created, and still more will begin to require new, AI-aligned skills. The rapid adoption of AI technology has led some to voice alarm at the potential for job loss faster than society can absorb. But as enterprise adoption actually lags behind the dizzying speed of technological development, the dominos may not fall as quickly as tech prognosticators fear.

Generative AI is already significantly affecting a number of jobs, providing specific examples from which to assemble a larger picture. Everyone, from front-line workers to senior leaders, will have to upskill, reskill and make some cognitive shifts as AI becomes a more common tool, and even a partner, throughout the enterprise.

Being successful in the AI-infused enterprise will require new skill sets of every worker, from key technical roles, such as software developers and data engineers, to nontechnical workers and leaders all the way up to the C-suite. What follows are our predictions for the successful workers and those fortunate enough to have leadership that understands the needs and opportunities of the AI era.

Everyone, from front-line workers to senior leaders, will have to upskill, reskill and make some cognitive shifts.

Workers will (have to) master human-AI collaboration and communication

When ChatGPT first scared and surprised everyone all at once, there was talk of a new job category: prompt engineer. Now, a few years later, it is understood that everyone needs to know how to interact with generative and agentic AI. It's not a job, it's a basic skill.

Software engineers exemplify this shift most clearly. AI coding assistants will write code described in natural language. That means that engineers will need to be able to talk a good game.

"Rather than being good at writing code, the engineer has to be good at describing very clearly what they want built," says Snowflake Head of AI Baris Gultekin. "This is a new paradigm, and even here at the very beginning, it's massively transformative."

"The ones who turn out to be good at describing code to AI may not be the ones who are great today at coding in, say, a syntax-laden language like Java," adds Mike Blandina. "They really are different skills."

Snowflake Co-Founder Benoit Dageville, who still does plenty of hands-on coding, has found that his style makes collaborating with AI more effective. "When I program, I write a lot of comments — I like to explain what I'm going to write," he says. "My AI assistant is very effective at suggesting code because I put in those comments. I see other developers having less success because they don't comment as much. We all have to learn how to use AI, how to describe what you want to solve and how you want to solve it."

The same skills will be essential to the marketer or sales rep telling a gen AI app to write an email, or tasking an agent with achieving a campaign outcome. It will be vital to the HR specialist or finance leader who "asks the data" for insights — and needs to ask the right question to get the right answer.

> **"The ones who turn out to be good at describing code to AI may not be the ones who are great today at coding in, say, Java."**
>
> **—MIKE BLANDINA**
> CIO, Snowflake

AI coding assistants should make developers 33% more productive by 2027

In the nearest term, ever-improving coding assistants are going to make a tremendous difference in IT, says CIO Mike Blandina.

"In the next year or so, you're going to see a big productivity gain — and in my world, big is 33%," he says. He's quick to add that predictions of even greater gains are probably overblown. "I'll tell you why it's not higher. The AI can write generic code quickly, but it doesn't know all the standards and compliance requirements of your company and industry. AI writing generic code probably only buys you 25% efficiency, because there's already a lot of generic open source code out there that a developer could search for."

While the makers of agentic coders promise ever-improving tools that will eventually take more of that context and complexity into account, Blandina thinks it'll be a while before such capabilities exist and are customized across industries and individual companies.

Everyone will (have to) become a strategic thinker

When we all get good at telling AI what to do, how will we know what we actually want it to do? Rather than focusing only on technical specialties or executing tightly assigned tasks, workers will have to understand the big picture to successfully hand off mundane tasks to AI tools.

Data engineers are a great example. As AI permeates the enterprise, there will be more pipelines to build and maintain, with greater complexity. The only way to scale, says Snowflake VP of Product, Data Engineering Chris Child, is with AI. Data engineers will spend less time creating pipelines and more time orchestrating them to drive business success.

"They'll need to be thinking at higher levels of abstraction, more about the business outcomes they're driving," Child says. "Traditionally, data engineers run the pipes. We move the data, and how that ultimately gets used isn't our problem. That's already changing."

Same goes for software developers, says Amanda Kelly. "Engineers will think more about outcomes and use AI to work backward."

"There are two parts to software engineering," says Director of AI Infra Mona Attariyan. "Writing the code one line at a time, and the more global understanding of what you are building and why. AI can really help with the first part, but not the second. At least not yet."

CIO Mike Blandina says this will be the case for tech leaders as well. He notes that a CIO's metrics, and often their thinking, is dominated by the tickets that measure work done. Those metrics will still matter, he says, but forward-thinking IT leaders will adopt a "solution mindset."

"If Legal comes to the CIO with a problem, in the past, the CIO might say, 'Well, there are three vendors with a SaaS solution for that,'" Blandina says. "With AI, I might wonder whether I can find an AI that does what they need. Maybe I bring together three different LLMs and create a solution that solves not only that problem today, but a bunch of future problems, too. So that's where broader solution thinking comes in as well."

What junior engineers?

It's often said that agentic AI is going to take the low-level tasks that junior engineers used to do, to let senior engineers do higher-level work. But without junior engineers, where will future senior engineers come from?

"It's a big question," says Snowflake Co-Founder Benoit Dageville. "You used to train up an intern or entry-level worker. And now you're going to use AI assistants instead. AI has arrived so quickly that we don't yet know how the world is going to reorganize itself."

Mona Attariyan says that education will change. Tomorrow's grads will need more of a business education that instills an ability to think strategically.

"Junior engineers will come in with less coding skill and more focus on higher-level aspects around design and architecture," Attariyan says. "Senior engineers will have that perspective, as well as a better understanding of the company's specific architecture and code base."

Workers will (need to) develop cross-functional and orchestration skills

Part of the reason that workers at all levels will need a better grasp of the overall business strategy is that people won't be snug in their task-limited siloes anymore.

"Jobs and personas are blurring," says Baris Gultekin. "Now an analyst can do some of what a data scientist does, and a marketer can do some of what an analyst can do. Product managers can now code. Now engineers need to do more PMing to understand what to build."

AI's blurring of traditional role boundaries will require workers to develop broader skill sets, both in understanding other departmental objectives and the soft skills of effective cross-team collaboration.

The result, says Chief Data Analytics Officer Anahita Tafvizi, is a flatter organization. "AI will compress org layers and amplify generalists who can think across domains," she says. "Specialized skills won't vanish, but the value shifts to those who can orchestrate across tools, data and AI."

This orchestration extends to managing AI agents themselves. Jennifer Belissent suggests that every worker will become a manager of sorts, overseeing multiple AI assistants and agents optimized for that worker and for specific tasks.

"Workers will have to be trained to manage these digital subordinates," she says. "We'll see the CDO's team working closely with HR to make sure we govern our agents appropriately and train people to work effectively with them."

For workers and the enterprise to thrive, leaders must provide continuous, contextual learning capabilities

The pace of AI advancement requires ongoing adaptation rather than one-time training. "Companies can't afford to outsource this learning," Tafvizi says. "Training must be embedded, contextual and ongoing. The winners will treat AI upskilling as strategic infrastructure, not a side program."

This learning must be hands-on and practical. Senior leaders, Tafvizi says, must be "using AI tools themselves, sponsoring real use cases and tying AI to business priorities — not just approving budgets."

AI-created applications will be the new general-purpose spreadsheet

Just as the spreadsheet application moved from its primary role as a finance tool to become something every sales rep, marketing wonk and HR specialist uses to keep track of things, in the next year or so, more and more of us will add making on-the-fly applications to our basic toolkit.

"I see agents becoming a core part of how everyone does their work," says Baris Gultekin. "Everyone building custom applications for their specific needs will be a very common thing."

Benoit Dageville says that if you can clearly verbalize the application you need, you'll be able to get an AI assistant to spin it up. "Every one of us knows how to use an application to interact with something," he says. "So you should be able to describe to an AI what you want the system to do, and how you want to interact with it."

Humans will remain in the loop as interpreters and quality controllers

Even as AI democratizes access to data and analysis, human expertise remains crucial — and not everyone who chats to a chat app is an expert. Amanda Kelly notes that if you gave your sales team access to an AI agent that will create dashboards on demand, it would be a mistake to assume that means you don't need your analyst team anymore.

"Anyone can ask for a dashboard or insight, but does everyone actually understand what to ask for and how to interpret it?" she asks. "Do they actually know correlation from causation?"

She suggests that when you democratize access, you may not also be democratizing understanding. The result could actually mean you need more analysts, not fewer.

And just as we've mentioned that many organizational decisions today are made by imperfect human instinct, Chris Child observes that AI is similarly black boxed and will never be completely free of error. So rather than replacing "gut instinct" with data-driven AI insights, workers from the CEO on down will need to loop in their instinct when reviewing AI outputs.

"AI models will have a deep understanding of your data — and not just metrics approximating your data, but the actual data," Child says. "But you'll still have to know when to doubt, when to ask deep follow-up questions before taking action."

"Don't discount how much data is not in the database, and therefore cannot be in the AI," Kelly adds. "The data can be wrong, incomplete or out of date. A manufacturing exec told me, 'An AI won't know that we just stopped production because there was a spill on the floor, until someone logs that event — which might never get logged at all.'"

Therefore, she says, there will always be knowledge and perspective stored only in our analog brains.

AI drives cybersecurity (for better and worse)

Earlier in this report, Anupam Datta outlined three challenges around agentic AI adoption, which apply as well to the ongoing integration of generative AI. We've covered the evaluation and optimization challenges, leaving us the third category: AI safety and security.

Some types of agents or gen AI tools will be built and entirely used in-house, which reduces some of the risks, but many systems will interact with the public, such as a company's customers, or with external websites, databases or tools. That, Datta notes, is a riskier proposition.

"When you connect agents and their data, you improve performance, but you also significantly expand the attack surface," he says. "As you move into a distributed setting where you're calling AI tools built by others, you're working with data created and stored outside your perimeter. All of this heightens the need for product research into capabilities that guard against adversarial attacks."

Snowflake CISO Brad Jones says that as companies introduce agentic AI, security teams will have to work to balance guardrails around agents' capabilities and behavior with sufficient room to experiment and innovate. And they'll have to pay even greater attention to data governance.

"There are likely to be many documents or data sets in a company that don't have permissions correctly locked down," he says. "If you feed that into generative or agentic AI, the tool may expose data that it shouldn't."

Agents as cyberweapons will help human attackers refine and scale their attacks within a year

Cybersecurity's classic arms race of offensive and defensive tools and techniques continues, with even higher stakes, in the AI era. Add in agents capable of researching, devising and executing attacks and you're looking at a seriously upleveled adversary. Jones says he expects agentic cybercrime to be a problem in 2026.

"The ability of agents to scale up the capabilities of human attackers is just around the corner," Jones says. "You'll see agents that will look at code, find a vulnerability and custom-build an exploit kit to attack and exfiltrate data, launch a ransomware attack, etc."

Key threats with AI agents include prompt injection — ways of tricking the agent into doing things it shouldn't, despite guardrails — and hallucinations. Getting an AI system to generate a sales document that promises capabilities or security classifications your product doesn't actually have could lead to major legal problems, for instance. Jones says that security teams will have to do a lot of red-teaming, in collaboration with data science teams, to prevent such possibilities.

And that's just while agents are accelerating conventional attacks. It can definitely get worse.

"Right now, AI can execute human attacks faster, better and with increasing autonomy, but they don't do novel things," Jones notes. "When they reach the point of being able to execute completely novel attack strategies, that's a whole new level of trouble."

> "The ability of agents to scale up the capabilities of human attackers is just around the corner."
>
> **—BRAD JONES**
> Chief Information Security Officer, Snowflake

Cybercriminals will use increasingly sophisticated ‘dark AIs’

The foundational models that drive most generative AI are designed with guardrails to prevent misuse. Such protections are not always difficult to evade, but at least they’re there. But security vendors have been describing specific chat tools for creating phishing scams, malware and more as “dark AI” since 2023. Such tools as HexStrike AI, FraudGPT and WormGPT — uncensored versions of large language models specifically geared for malicious activities — are already generating phishing emails, malicious code and social engineering attacks.

Consider it the dark side to the open source movement. Many malicious AI tools are created with open source foundational models (like GPT-J-6B) that are simply deployed without the ethical guardrails of commercial models.

Such AI-enhanced tools have entered the supply chain of the industry known as cybercrime-as-a-service. The underground economy features global businesses built not on attacking targets, but on providing the data and tools that attackers need — with subscription models and customer support. The tools are out there, and even expensive and cutting-edge AI will be increasingly bent toward malicious goals.

“Financially motivated threat groups now have the resources of nation-states,” Jones says. “They are increasingly leveraging purpose-built LLMs unhindered by such constraints as ethical guardrails or regulations.”

In three years, AI agents and tools will finally close the security talent gap

It’s not all bad news when you talk to the security folks. Generative and agentic AI, alongside traditional machine learning, will also improve the security operations center’s ability to fight back. The most persistent problem in the SOC has been the shortage of human analysts. Talent is hard to find and hard to retain. The promise of advanced AI is to supplement human analysts to take not their jobs, but those that CISOs have never been able to fill.

“It will still be hard to find good analysts,” Jones says, “but they’ll be supplemented by AI agents and tools robust enough to finally make the SOC feel fully staffed and resourced.”

And just in time, judging from the odds against them.

Human Creativity Will Persist

The subject of AI and creativity is most often addressed from the perspective of the arts. If every chucklehead who can type “Jet-pack ferret in samurai armor with a freakin’ big sword” can trigger an AI-generated image, will anyone draw anymore? And will anyone get paid for it? Will people write? Will even our podcasts be AI-generated?

These questions are beyond the scope of this report. But more firmly within the enterprise space, there also have been questions about whether a human *je ne sais quoi* will have a place in a business world dominated by data-crunching AI systems. The conclusion of our expert panel is: Yes.

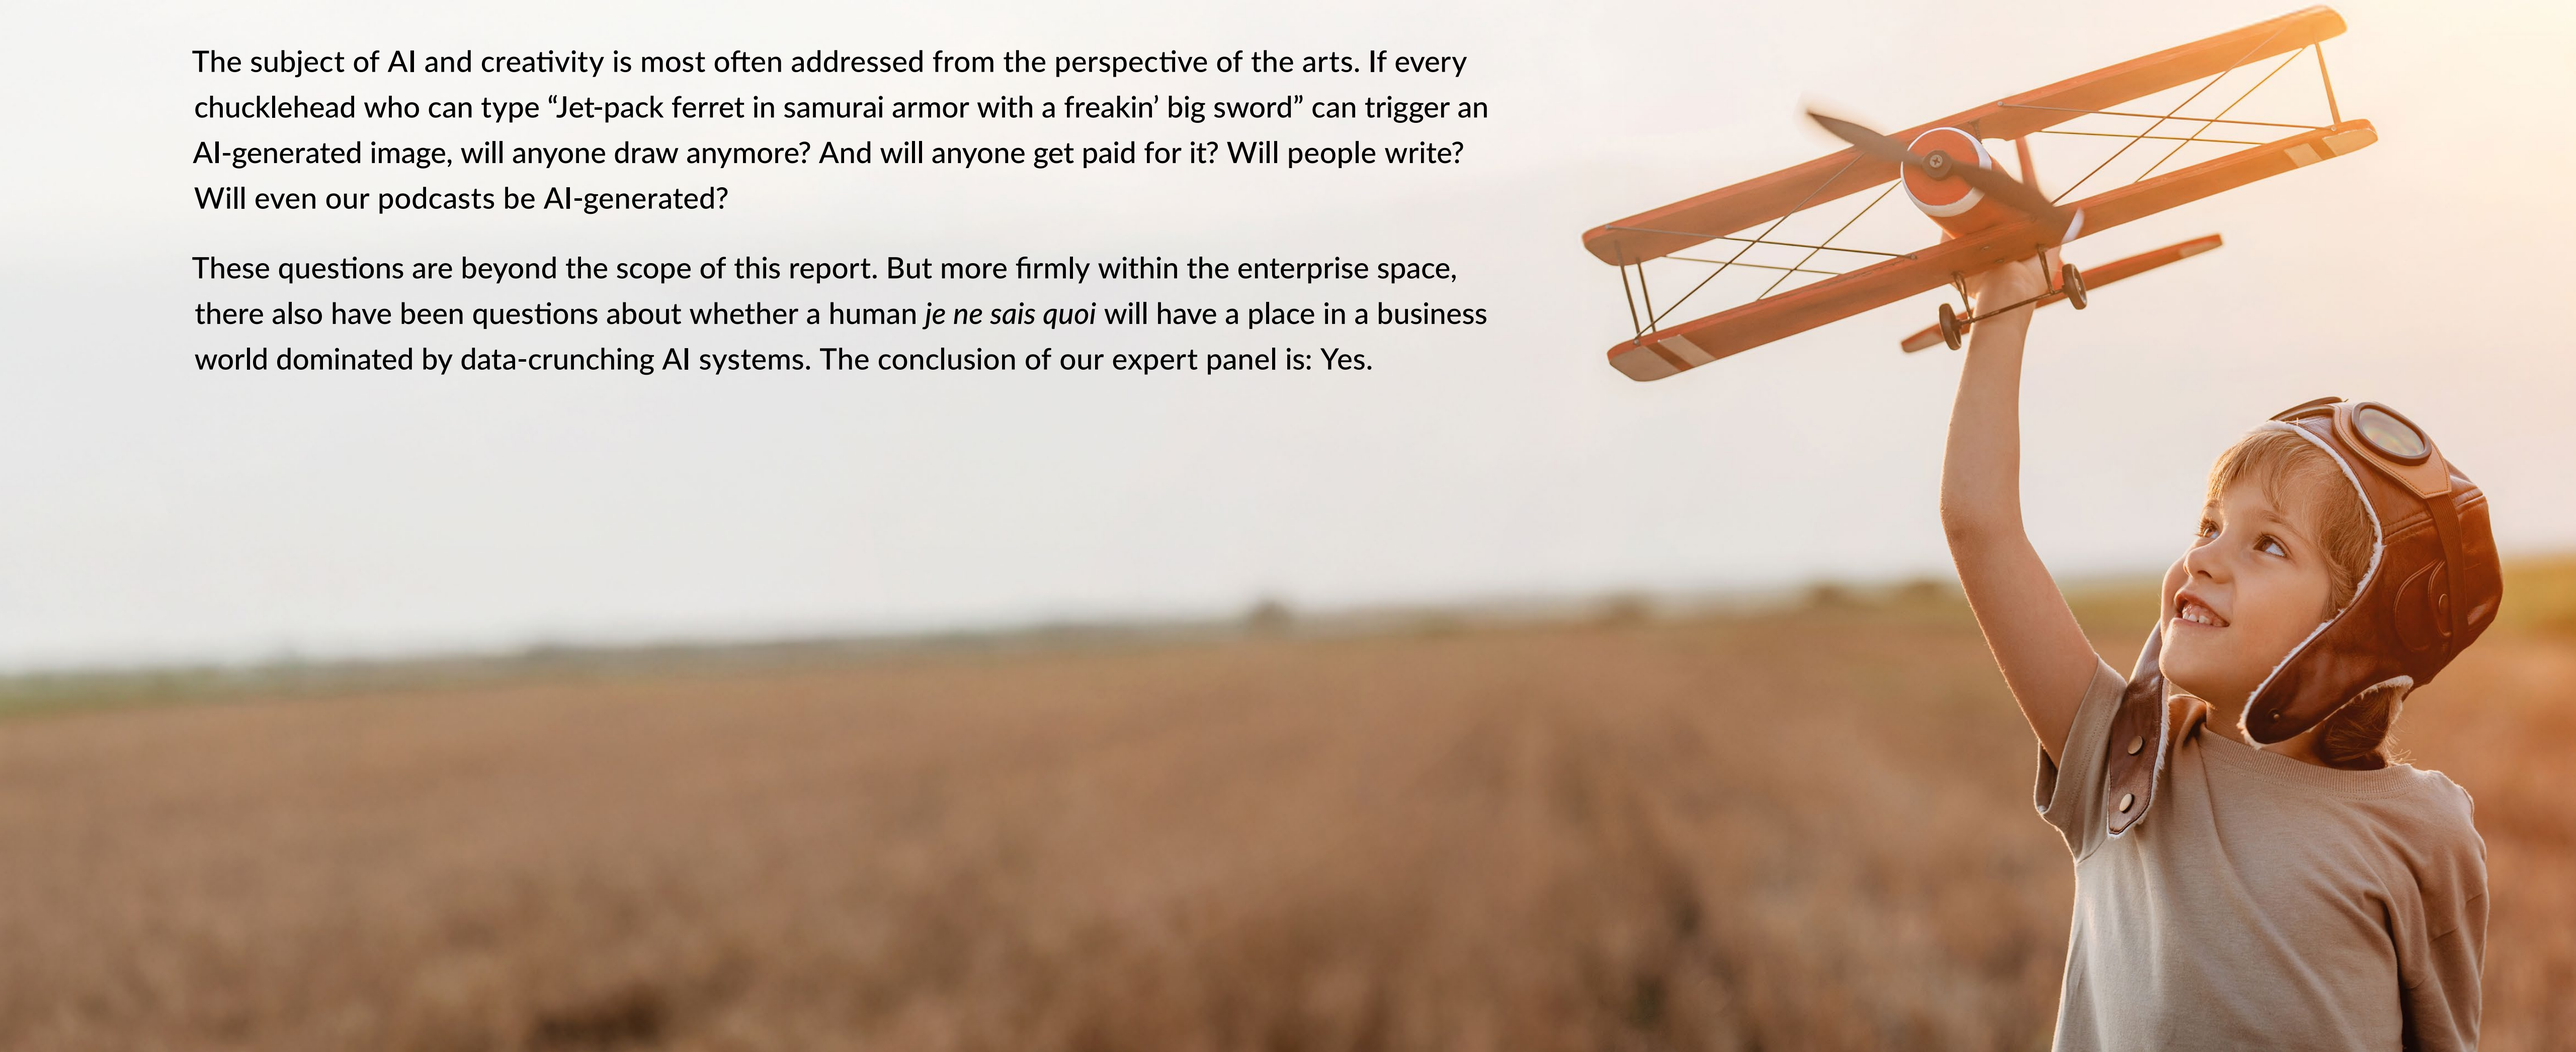

Generative AI will accelerate and improve human creativity — but it will take smart, skillful approaches to avoid it being a crutch

In addition to the need for human instinct to gut-check AI outputs, AI tools present creative opportunities far beyond the AI slop that may be choking your social feed. Of course, there are those who say (and studies that underscore) that using gen AI can actually be a crutch that makes us lazier or less intelligent. Sridhar Ramaswamy says that creativity is up to the user, not the tool.

"True, it's tempting to take the easy route and use AI to generate a bunch of content that is low quality, but you can also use AI as a foil to push your own creativity," he says. "I believe in humanity. I think that, when given a powerful tool, there will be enough among us who will use it well and push it, and push ourselves. People will continue to be very, very creative. If anything, AI models vastly increase the number of choices that are available to creative people."

Baris Gultekin says that the difference is in part a learning curve, an adaptation to new technologies.

"Initial reaction to ChatGPT in academia was to ban the thing, and now we're seeing the opposite, people embracing it, figuring out how to make it a core part of the learning experience, so there's a balance there," he says. "We can easily outsource critical thinking to our detriment, but if you're able to use it as an effective tool, it is a very effective tool."

Ramaswamy agrees. "I was talking to a colleague recently, and we were saying, 'We are going to be limited by our ideas, not by our ability to get things done.' Which is both fun and scary."

INDUSTRIES

A closer look at three global sectors

RETAIL

Gen AI is flipping customer 360 on its head

After years of observation, retailers are moving full steam into implementing gen AI. Initial wariness has given way to widespread fear of missing out on the powerful promise of generative and agentic AI.

"Finding niche products or optimal flights is commonplace on ChatGPT," says Rosemary DeAragon, Global Retail & CPG Lead at Snowflake. "To ensure discoverability of long-tail products, retailers and consumer brands must get their product data ready for gen AI."

With new features like ChatGPT's instant checkout, retailers now recognize that generative AI is already a part of the shopper's journey. In 2026, she says, customers will be providing deep, personal behavioral information to large language models. And thus, the industry practice of collecting customer 360 data will be radically changed.

> "To ensure discoverability of long-tail products, retailers and consumer brands must get their product data ready for gen AI."
>
> **—ROSEMARY DeARAGON**
> Global Retail & CPG Lead, Snowflake

Data-rich large language models will fuel a new level of customer 360

Since customers are willingly sharing personal details with large language models, retailers and brands will race to embed gen AI features into their customer experience. Retailers already used traditional machine learning for fraud, personalization, dynamic pricing and supply chain management. The big change with generative AI is the bidirectional customer engagement, with customers volunteering personal information in exchange for discovering products they didn't know existed. Looking forward, generative and agentic AI will enhance traditional AI to provide deeper personalized recommendations than ever before.

"The strategic moat for retailers and brands in the age of AI, where applications are built in seconds, is how much rich or unique first-party consumer data they have, and how that data fuels their value to customers," DeAragon says.

Customers will be able to shop for items through a gen AI application such as ChatGPT, which will connect with retail and payment processing partners to source and purchase the item. Retailers will also have their own agents, helping customers within the retailer's platform. It's the personalized interaction that will help determine success.

"Applications and even LLM models will become more and more commoditized," she says. "The strategic differentiation lies on the user input side, namely, in human-generated data. The more that retailers and brands are able to become a trustworthy digital entity, the more the customers will provide deep insights, accelerating a flywheel of rich personalization."

Distrust of AI slop will elevate influencer marketing

DeAragon says that consumers increasingly trust influencers and brands more than mass advertising or text reviews. And that's partly because of a mass amount of AI-generated content, commonly known as AI slop. "Digital natives are starting to not trust what they see on the internet, because so much content is now AI-generated," she says. "They're more likely to trust a human influencer."

AI-generated summaries of reviews have become commonplace, but faith in individual product reviews has dimmed, DeAragon says. "The trust in product reviews is declining because consumers don't know whether it's written by a human or AI."

With a lot of knockoffs in the market, consumers turn to social media, relying on influencers to tell them which brands they should buy and which they should avoid. "Digital natives are especially skeptical of what they see, leading consumers to seek trusted human sources."

Reliance on social influence is especially true for Gen Z. "Forty-one percent of Gen Z rely on influencers to show them a product for them to buy that product," DeAragon says. And Gen Z's habits are significant because spending power is shifting from baby boomers to millennials and Gen Z.

"It seems novel that we're changing from big brand advertisements to pouring that money into influencer marketing," she adds. "But for the foreseeable future, this will be the new norm for big marketing budgets and how even the largest companies reach customers."

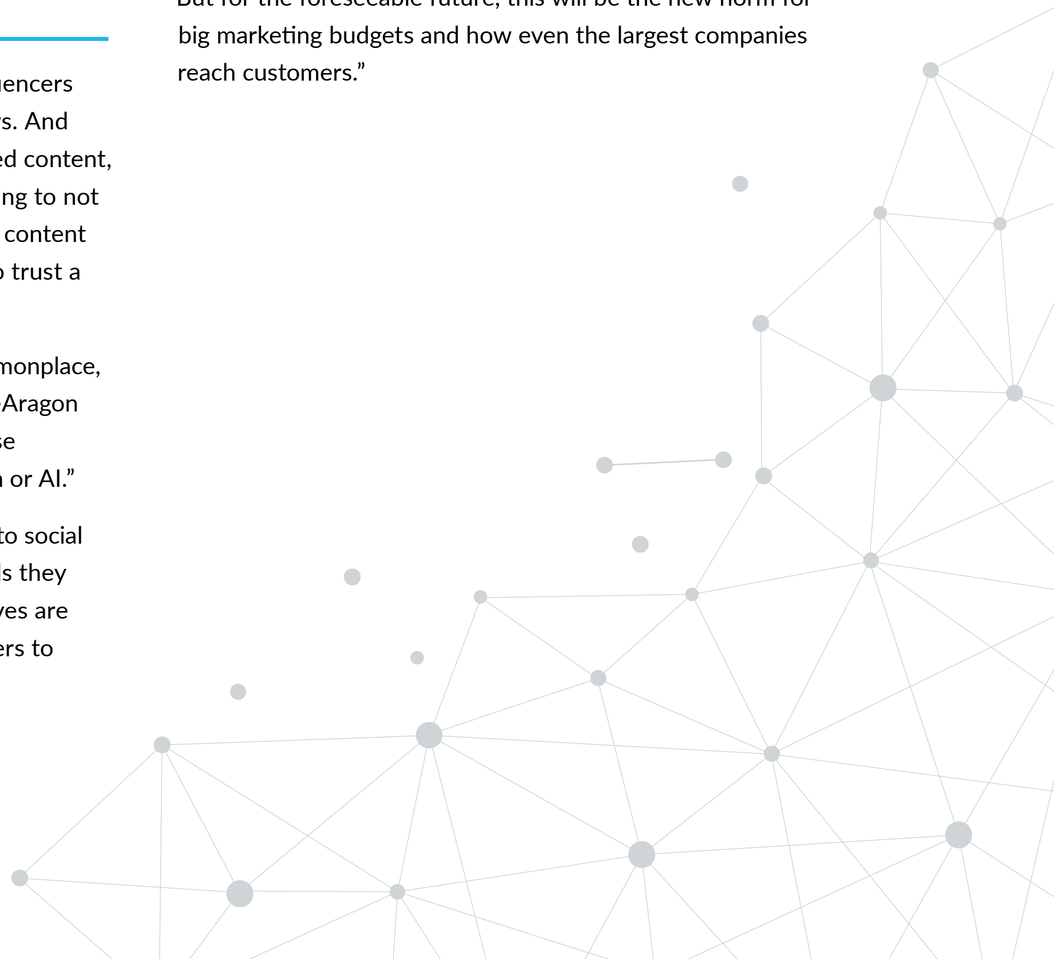

Privacy approaches will drive the success of AI shopping agents

People already turn to generative AI for shopping ideas, and the next frontier is agents that will do the shopping for you. As AI technology develops, consumers will still have to trust the tools and agents they're offered. DeAragon says that how brands deal with privacy concerns will play a major role in building that trust.

She notes that Apple, which is sometimes seen as late to the AI race, has spent the last few years building a reputation for taking stronger positions in favor of consumer privacy. That, combined with its devices and operating system in so many hands, make it a strong contender as a core provider of AI tools for retail consumers.

"Apple knows who I talk to, what's on my calendar, where I go, and more, so it can create smart agents that really know me," DeAragon says. And she expects the other major phone platform, Google's Android, to pursue the same idea. AI that shops for you will be everywhere. "Within the next two or three years, you're going to see truly automated agentic AI workflows at the B2C level."

As personal chat and agentic tools become more common, retailers will be challenged to handle personal data in a way that's both optimized for AI shopping and consistent with consumers' and regulators' expectations around privacy. Brands that adopt an ethics-first approach to private data — and make that approach clear to their customers — will gain trust and loyalty. "That is going to be table stakes, but it's way easier said than done," she adds.

These challenges were already evident as generative AI first exploded into public consciousness, of course. The difference now, DeAragon says, is that retailers are facing the challenges head-on and taking the lead in infusing AI into the consumer experience. After a couple of years of hesitation, the industry is finally on board, and the train has definitely left the station.

> **Apple knows who I talk to, what's on my calendar, so it can create smart agents that really know me."**
>
> **—ROSEMARY DeARAGON**
> Global Retail & CPG Lead, Snowflake

FINANCIAL SERVICES

Shifting back to a data-first mindset

The craze for AI experimentation in financial services is fading as firms focus on ROI and remember that value starts with data.

"Twelve months ago, firms were running around trying to build chatbots and copilots, but they weren't thinking about the architecture and the underlying data," says Rinesh Patel, Snowflake's Global Head of Financial Services. "Now firms are looking at access to structured and unstructured data, multimodal data."

Data has always been a primary source of value for financial services firms. Understanding market trends, reducing risk and targeting the right offer to the right customer all rely on extracting insight from data. Rather than seeing AI as an exciting new imperative of its own, Patel says, firms understand that AI offers new ways to get at those insights, but only if first the data is in order.

"That has been the evolution of data and AI strategy across financial services in the past year," Patel says. Organizations still want to keep up with the pace of innovation, including the development of agentic AI, but they face both risk-averse boards and the eternal imperative to maximize profit.

Showing the ROI of AI initiatives will be a top priority for firms

Patel says that experimentation with AI will not be about exploring the possible, but the pursuit of likely profit. "There's less or no interest in AI for AI's sake," he says. "Firms are trying to solve the same problems and pursue the same metrics that have always been there."

He says that, as with many industries, financial services firms will transform as they infuse data throughout their businesses. They'll just do it incrementally, with clear returns for every step.

"Financial service organizations want to augment AI into the entire data lifecycle, infusing the way they use data from ingestion all the way to distribution," Patel says. "They're getting there by focusing on use case specificity, prioritization and measuring value."

Regulation concerns will increase

While the United States has shown a strong taste for deregulation, Patel says that Europe will continue to put more restrictions on how data and AI are used. The EU's AI Act provisions have been enacted steadily since 2024, with the complete regulatory regime, as currently envisioned, to be live from 2027. The AI Act is a risk-based approach to AI that will continue to develop as AI and data become increasingly interconnected. Global firms will have to contend with evolving international regulations as they figure out how to use AI.

"How the EU continues to regulate data will have implications on how AI is used," Patel says, adding that European requirements are likely, in many cases, to become global defaults.

Agents will be a hard nut for financial firms to crack

As agentic AI develops, firms will wrestle with AI security and regulatory compliance. AI adoption is moving from automation and augmentation toward autonomy with AI agents. So in the near term, firms will have to grapple with how and when to integrate agents, and manage the resulting risks.

Patel says that tech-forward financial firms especially will be wary of agentic AI, for fear that more autonomy means more room for error.

"When you think about what AI agents are meant to do, there's going to be a greater focus on risk," he says, "and not just limited to hallucination and ethics. There will also be questions around data residency, operational resiliency and more."

Even with the increased risk, firms know they'll need to evolve and develop ways to incorporate agents before too long.

And incorporate them they will, Patel adds. AI, whether traditional ML, generative or agentic, offers a route to efficiencies and insights too profitable to pass up.

"Financial services firms are in the business of making money," he says. "It's their north star for everything they decide when it comes to data technology and AI."

> **"When you think about what AI agents are meant to do, there's going to be a greater focus on risk."**
>
> **—RINESH PATEL**
> Global Head of Financial Services, Snowflake

MANUFACTURING

AI adoption takes center stage

The manufacturing industry is ready to debut the next phase of AI. Manufacturers have been automating and investing in AI/ML for years, for everything from predictive maintenance to computer vision and the foundational data strategies that allowed them to incorporate traditional machine learning. So, now they're positioned to adopt more advanced AI solutions, including generative and agentic AI.

The AI prep work may help U.S. companies under pressure to bring more manufacturing back to the United States.

"We don't have the U.S. workforce to support the manufacturing needs of today," says Tim Long, Snowflake's Global Head of Manufacturing. We're short 500,000 pipe fitters in the current market. For every other type of tradesperson, we're 100,000s short of the talent needed in the U.S., he says.

To make up for the lack of qualified workforce, manufacturers will have to perform differently and more efficiently by making significant process improvements. "They'll need to use AI to streamline the production and supply chains as much as possible," Long says.

In 2026, manufacturers will use AI to streamline performance, monitor quality control and improve production outputs.

> **"[Manufacturers] will need to use AI to streamline the production and supply chains as much as possible."**
>
> **—TIM LONG**
> Global Head of Manufacturing, Snowflake

Chatbots will be the next big thing in generative AI advancement

"Manufacturers are realizing that they can use generative AI to get to the next level of performance," Long says. While still in the experimentation phase, chatbots are gaining traction to support manufacturers' business objectives.

"The most important metric in manufacturing is overall equipment effectiveness, because it captures efficiency of the production process and quality of the products that are produced," Long says. Soon, chatbots will summarize this effectiveness using natural language to diagnose equipment issues, better parse customer sentiment and understand service logs. "I've seen great examples of chatbot interfaces being pursued, and I expect them to be a common use case in the next year."

AI will accelerate quality control and production

As generative and agentic AI are hogging the spotlight, manufacturers will continue to use computer vision to observe workers and identify efficiency gains on bespoke production efforts. Observing sequences of how things are put together can suggest ways to modify the steps assembly techs take for improved productivity.

And computer vision applications will only improve in the year ahead. It will predict maintenance needs and product defects with increased precision. "This enables manufacturers to better see when equipment isn't running at full capacity and determine whether it's due to the operator, training or product steps," Long says. "They can also determine where a defect is first detected."

Creatives will not be replaced with generative AI (yet)

AI may lead to reductions in creative staff, Long says, but it will not be able to replace human creativity entirely. "We'll need a workforce with the competencies to not just use an AI agent, but with the innovative and creative experience to invent process improvements and products themselves."

Generative AI for ideation, digital prototyping and product design hasn't proven itself yet, so it will remain a future capability, he adds. "AI that can be creative and simulate, then evaluate, concepts is not something I've seen yet."

Agents will start to make a real business impact

AI agents have high potential to improve product production and logistics. "We'll likely start with agents making decisions on how to categorize products, pulling products out of a process for inspection or, potentially, expedite decisions in the production process," Long says.

Additionally, "When we talk about large language models, we'll start to see agents that are specialized for specific industries," he says. For manufacturing, agents will be equipped to decide to expedite one product over another to accommodate shipping schedules and delivery data.

Letting agents act will of course be a headache for legal and IT teams. But Long is undeterred, predicting that industry-specific AI agents will emerge and gain popularity nonetheless. And AI will be crucial for rebalancing global manufacturing, offsetting labor costs and improving efficiency.

New technologies always present challenges. But the industry's advantage is the strategic data work it has done already. "Manufacturing is well positioned to overcome new challenges," Long says.

Contributors

Mona Attariyan

Mona is Snowflake's Director of AI Infra, and holds a Ph.D. in computer science from the University of Michigan. Previously, she served as Head of Analytics and Machine Learning at Redfin and held various technical roles at Google.

Jennifer Belissent

Jennifer is Principal Data Strategist at Snowflake. She spent over a decade at Forrester Research as an internationally recognized expert in data sharing, the data economy and data leadership, including best practices for building world-class data organizations. She has an M.A. and Ph.D. in political science and organizational theory from Stanford.

Mike Blandina

Mike joined Snowflake as CIO in 2025. Before that, he had been CIO of Payments at JP Morgan Chase since 2020. He was also CEO at Bakkt, a SaaS provider in the crypto space, as well as VP of Engineering at PayPal and Engineering Director, Payments, at Google.

Chris Child

Chris is VP of Product, Data Engineering, at Snowflake, where he supervises data engineering, open lakehouse, open source and developer products. He was Snowflake's VP of Worldwide Sales Engineering, and previously led teams at Segment, Salesforce, RelateIQ and Foursquare. He studied computer science and electrical engineering at MIT, and has an MBA from Stanford.

Benoit Dageville

Benoit is the Co-Founder of Snowflake, and today serves as the company's President of Product. Before co-founding Snowflake in 2012, Benoit spent more than 16 years as a data architect and developer at Oracle.

Anupam Datta

Anupam is a Principal Research Scientist and Snowflake AI Research Team Lead. He joined Snowflake via the acquisition of TruEra, where he was Co-Founder, President and Chief Scientist. He has published over 100 research papers and was on the faculty at Carnegie Mellon University from 2007 to 2022. He received his Ph.D. from Stanford and currently teaches a course on trustworthy AI there.

Rosemary DeAragon

Rosemary is Snowflake's Global Head of Retail & Consumer, and previously worked on data strategy at Walmart, on machine learning for Amazon, and at market intelligence firm Numerator. She was also the cofounder and executive director of EmpathyFX, a nonprofit that builds schools and provides healthcare in Ghana.

Baris Gultekin

As Head of AI, Baris leads AI and ML initiatives and drives Snowflake's AI product roadmap and strategy. He joined Snowflake following its acquisition of nxyz in August 2023, where he was co-founder and CEO. Before that, he spent over a decade at Google across various product leadership roles.

Brad Jones

Brad is the Chief Information Security Officer at Snowflake and VP of Information Security. Prior to Snowflake, he spent more than six years as CISO and VP of Information Security at Seagate. He has actively participated in a number of customer advisory boards and is currently part of the CISO Advisor Council at NightDragon.

Amanda Kelly

Amanda is Co-Founder of Streamlit, which was acquired by Snowflake in 2022. Amanda has also led product and operations in autonomous vehicles at Zoox and for several stealthy Google X projects around Google's natural language processing capabilities. She holds an MBA from Stanford.

Tim Long

Tim is Global Head of Manufacturing at Snowflake and a technical leader with education and experience practicing advanced analytical methods including ML, data mining, optimization, simulation, natural language processing, small and big data warehousing, and the creation of compelling data visualizations.

Rinesh Patel

Rinesh is Global Head of Financial Services at Snowflake. With over 20 years of experience in financial services, he has held several leadership roles across enterprise data and analytics, consultancy and investment banking.

Vivek Raghunathan

Vivek is SVP of Engineering at Snowflake. He co-founded Neeva (acquired by Snowflake in 2023) where he pioneered the development of a next-generation AI search engine for consumers. He also spent over a decade at Google as a VP of Engineering in various technical leadership roles.

Dwarak Rajagopal

Dwarak is a distinguished technology leader with over two decades of experience at some of the world's most innovative companies, including Google, Meta, Uber, Apple and AMD. As Vice President of AI Engineering at Snowflake, Dwarak oversees the AI and ML engineering organizations and helps Snowflake's AI Research Team drive innovation.

Sridhar Ramaswamy

Sridhar was Co-Founder of Neeva, acquired in 2023 by Snowflake. Not quite a year later, he was named Snowflake CEO. He spent more than 15 years at Google, where he started as a software engineer and rose to SVP of Ads and Commerce. Sridhar earned a Ph.D. in computer science from Brown University.

Anahita Tafvizi

Anahita joined Snowflake in 2024 as Chief Data & Analytics Officer. She previously served as VP of Data Science & BizOps at Instacart and held leadership roles at Google, YouTube and eBay. She serves on the boards of Candid, Make-A-Wish and Raya AI. Anahita holds a Ph.D. in physics from Harvard University.

Snowflake is the platform for the AI era, making it easy for enterprises to innovate faster and get more value from data. More than 12,000 customers around the globe, including hundreds of the world's largest companies, use Snowflake's AI Data Cloud to build, use and share data, applications and AI. With Snowflake, data and AI are transformative for everyone.

Learn more at **snowflake.com**

(NYSE: SNOW)

report-ai-data-predictions-2026-r04