prove ai

From Pilot to Production:

The State of AI Adoption, 2025

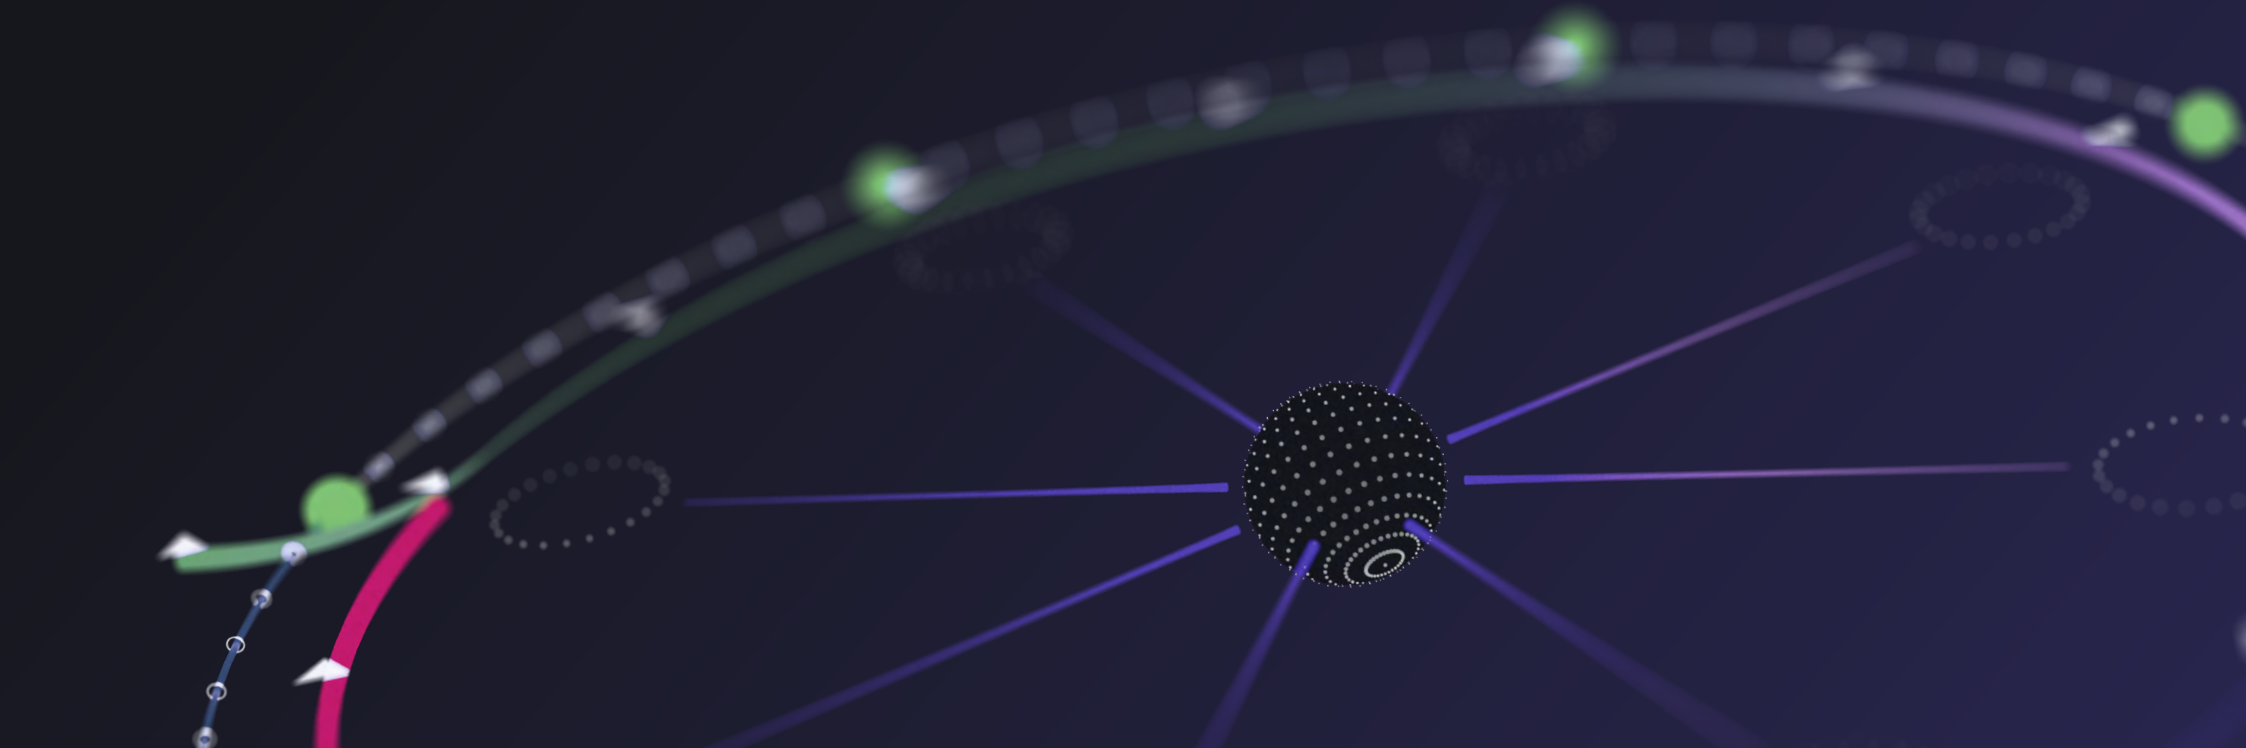

Table of Contents

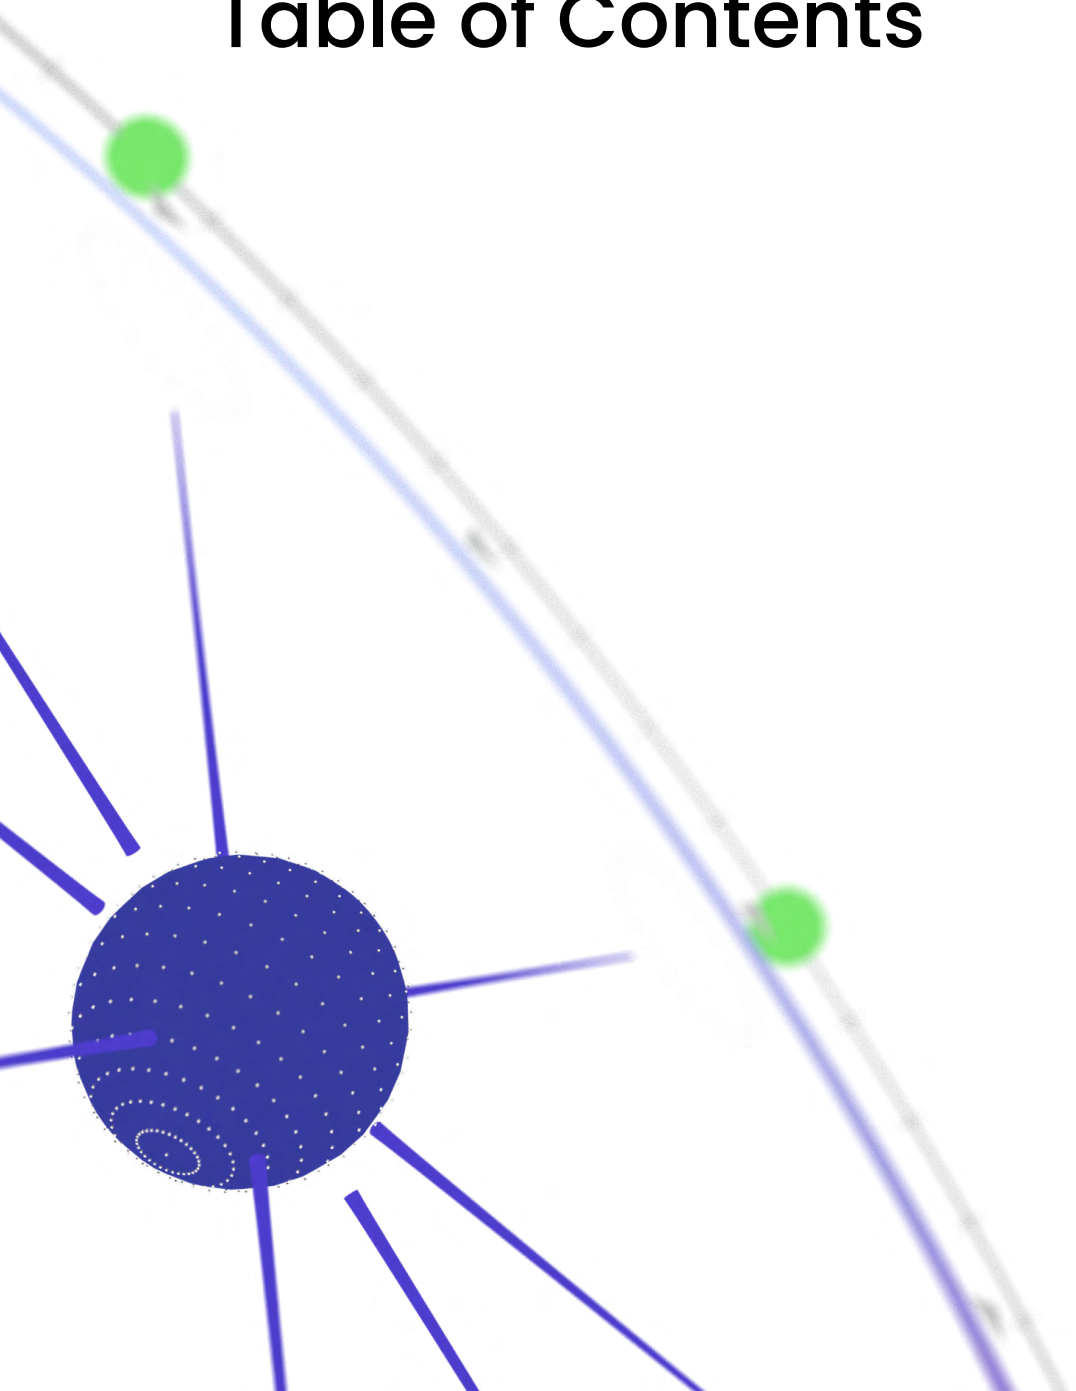

Letter from Prove AI

After years of hype, AI has firmly established itself as a foundational technology for businesses.

AI is no longer confined to R&D labs or innovation teams; it's now embedded into core business operations, including customer service, coding, fraud detection and more – and at increasing scale.

The findings from our latest report, which polled 1,000 business leaders across the U.S. and Canada, bring two paradoxical truths to the forefront: AI adoption continues to surge, but tangible, AI-driven progress remains elusive for many.

While companies are overwhelmingly adopting AI for its vast potential, they are, by and large, still grappling with its unbounded risk and struggling to integrate it within existing technology stacks.

Given the unprecedented nature of the transformation that AI promises, it's not surprising to see such a deliberate shift. Organizations remain bullish on AI – most respondents in this survey reported investing at least $5 million annually on AI initiatives – but struggle to show a clear return on those investments.

While reasons for this disconnect are myriad, several key themes emerged from the data. Among other factors, a lack of transparency around the data that trains AIs, challenges integrating a still-evolving "AI stack" with legacy technologies and difficulties attracting qualified AI experts from a limited talent pool are colluding to keep AI risk prohibitively high for many organizations.

And yet, nearly all signs point towards organizations making steady progress on each of those fronts. Respondents remain overwhelmingly optimistic that their respective organizations are on the right track to deploying AI-driven applications and workflows, and that the right tooling and talent exists.

Ultimately, AI's promise rests on trust. Increasing trust in AI means building greater transparency and accountability into AI systems – a process that is now well underway. The decisions we make in the coming months and years will have far-reaching implications. It's critical, then, that we not overlook this critical moment in AI's evolution.

Mrinal Manohar
CEO, Prove AI

Data Highlights

In 2025, AI has moved from proof of concept to production-ready… but hurdles remain

Business leaders tasked with overseeing AI deployments report considerable progress; more than two-thirds of respondents report moving beyond the pilot phase and into active production. Paralleling this trend, an overwhelming majority **(86%)** of organizations have moved to appoint a Chief AI Officer.

At the same time, **more than half** of respondents cited the continuous training, or "fine-tuning," of AI systems to be more difficult than initially anticipated. This has caused additional challenges around data quality, availability and model validation, hindering AI's current output.

beyond the pilot phase and are now in active production of a custom AI solution.

of those already in active production will add to their active solutions in the coming year.

Top business use cases for AI include: [Q3]

55% **Chatbots and personal assistants**

54% **Software development and maintenance**

52% **Predictive analytics like demand forecasting and fraud detection**

Most organizations have an annual AI budget of $5–10 million. Within that budget:

86%

of organizations have a Chief AI Officer.

83%

cited data sovereignty as their top priority when deploying AI software, followed by establishing data lineage, implementing clear guardrails, and audit readiness.

Around 9-in-10 respondents stated:

- Their organization is **effectively managing AI policy**
- Their organization can **establish AI guardrails**
- Their organization can **trace AI data lineage**
- Their **AI data is secure** and **honors all user terms and conditions**

Two-thirds of organizations want to deploy AI models on premises, citing efficiency and security as the main drivers over cloud deployments.

Demographics and Methodology

Zogby Analytics conducted an online survey of 1,000 business decision makers in the U.S. and Canada, with many representing organizations with 2,001 to 5,000 employees.

Using internal and trusted interactive partner resources, thousands of adults were randomly invited to participate in this interactive survey. Based on a confidence interval of **95%**, the margin of error for **1,000 is +/- 3.1** percentage points.

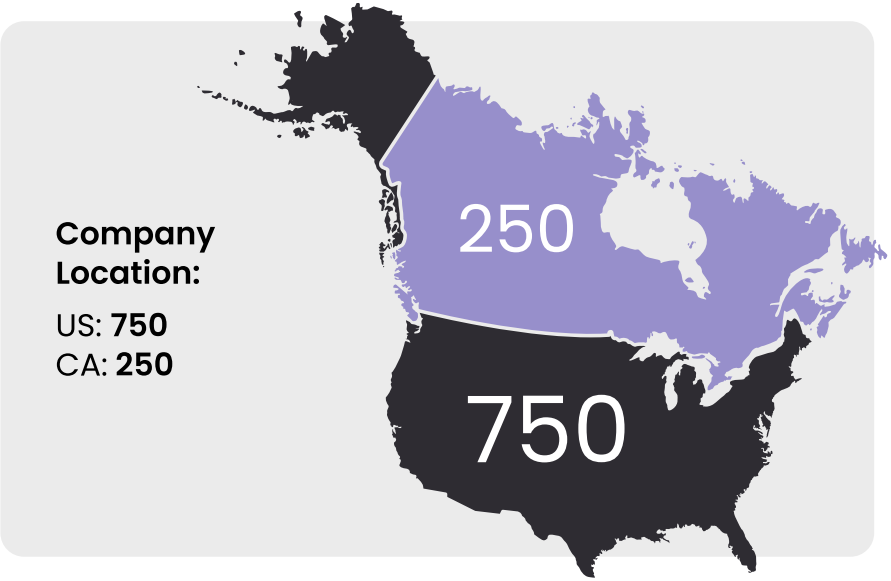

CHAPTER 1

From Pilot to Production, AI is Table Stakes for Businesses Today

AI is a core component of long-term business strategies across the board. Organizations are dedicating significant financial resources to supporting AI initiatives on a sustained, long-term basis.

Foundational applications like chatbots and analytics remain popular, although as more companies move into active deployment, additional use cases – most notably, software development – are also gaining momentum. Businesses that are newer to the AI space are more interested in prioritizing generative AI over traditional machine learning.

AI maturity is growing, and quickly:

Nearly **two-thirds (63%)** of respondents have had AI deployments live for at least 6 months.

Of those, **16%** are at **2 years** or more of maturity with live AI deployments.

Businesses are moving beyond pilot AI stages:

68% of organizations are out of a pilot phase and in active production of a custom AI solution, with **43%** of those looking to add to their active solutions in the coming year.

Businesses are taking a multifaceted approach

More than half (55%) of organizations have at **least three** live AI use cases, with **18%** having **more than 5.**

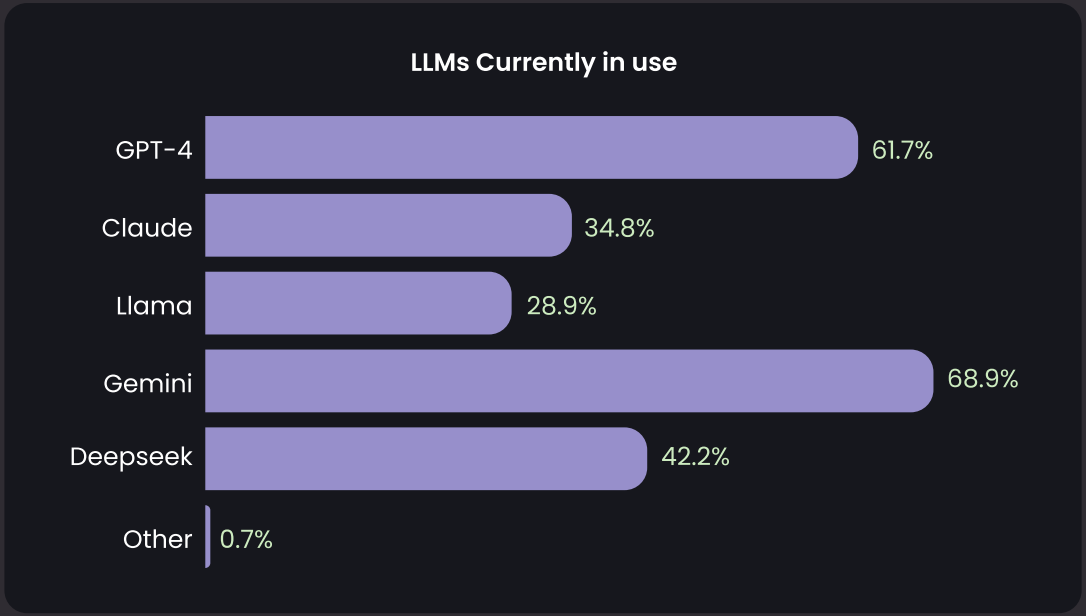

SECTION SIDEBAR:

What types of LLMs are powering companies' AI usage?

Most (59%) report using 2-3 LLMs:

- Gemini, ChatGPT-4, and Deepseek are most commonly used
- Claude, Llama, Deepseek are most likely to be added within the next year

While foundational applications like chatbots **(55%)** and predictive analytics **(52%)** remain common, a growing number of organizations are now applying AI to software development and maintenance **(54%)**, signaling a shift toward more technical, behind-the-scenes use cases.

Marketing applications, once a primary entry point for AI, are seeing relatively less focus than in years past.

Strategically, the majority of organizations **(57%)** report a stronger emphasis on generative AI, although a significant segment **(36%)** are balancing generative AI efforts with traditional machine learning, particularly among those earlier in their adoption journey.

Financial commitment to AI is equally telling. A vast majority of organizations **(81%)** have annual AI budgets of at least **$1 million**, with over a quarter **(26%)** investing **$10 million** or more.

So far, most organizations **(43%)** have already spent between **$1 million** and **$3 million** on AI deployments, underscoring the transition from exploratory spending to scaled investment.

CHAPTER 2

As Specialized AI Strategy Grows, Foundational Hurdles Remain

Organizations are establishing dedicated leadership roles, often in the C-Suite, that focus specifically on AI strategy, implementation and ongoing maintenance.

Many organizations are also taking a direct role in training and fine-tuning their custom AI models. This allows for more tailored solutions, but it also introduces complexity.

Despite increasing maturity and centralized management tools, organizations still face challenges securing high-quality, traceable data and ensuring that their AI systems can safely and effectively integrate with their previously existing technologies.

- The majority of organizations **(86%)** have someone with the title of Chief AI Officer.
- Nearly the same percentage of Chief AI Officers / Head of AI **(42%)** as CEOs **(43.3%)** are responsible for setting AI strategy, showing the critical importance of AI to overall business goals

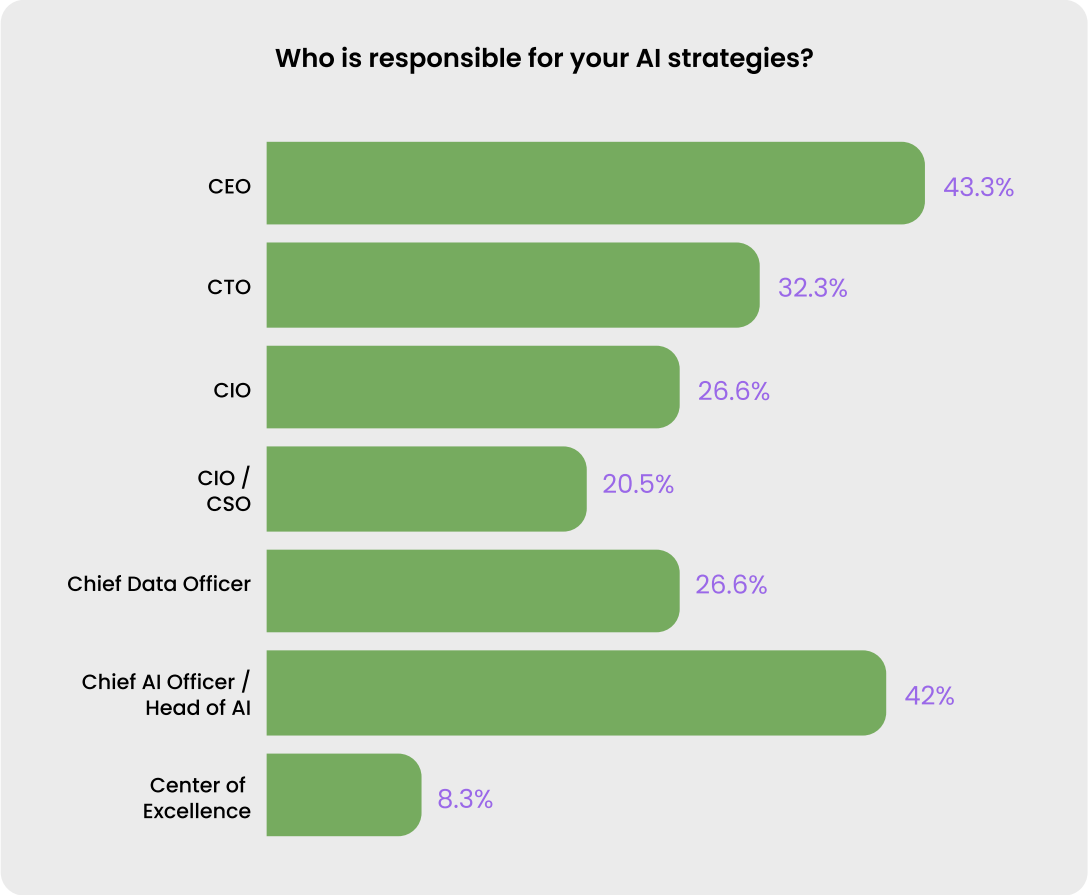

As AI initiatives mature, organizations are building internal infrastructure to support more sophisticated development and oversight. Nearly **three-quarters (73%)** report having a dedicated machine learning operations (ML Ops) team, and an even larger share **(84%)** have implemented a centralized dashboard to monitor all AI activity across the organization—an essential step in reducing silos and improving oversight.

The majority of companies **(86%)** have taken the step of training or fine-tuning their own AI models, signaling a strong commitment to customization and control. However, this process has proven more complex than many anticipated. **More than half (55%)** found that training or tuning models was more difficult than expected, citing common challenges such as safety concerns, limited data availability, and the high cost of infrastructure and software needed to support these efforts.

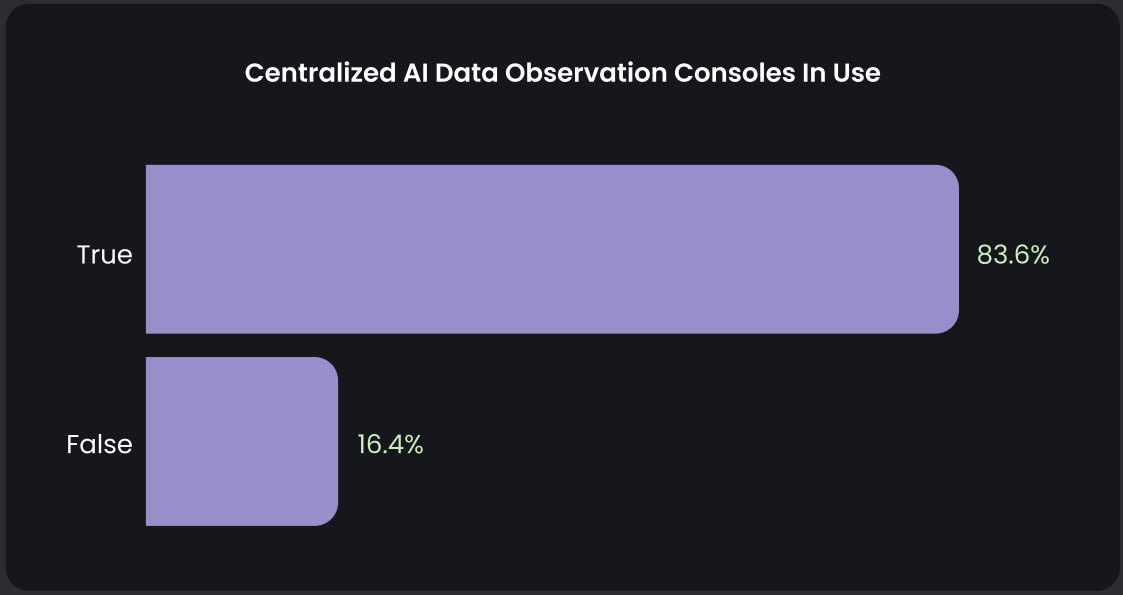

SECTION SIDEBAR:

A closer look at the benefits of having a centralized view into AI data—a single dashboard that can reduce silos and oversight.

CHAPTER 3

More Data Meets Established Complexity in AI Deployments

Business leaders are broadly expressing confidence in their companies' ability to manage AI policies and scale data. But even with specialized strategies and leadership in place, many organizations still face delays in their AI rollouts.

The ability to verify data, audit changes, and establish governance and guardrails are key components to effective deployments. Challenges with data quality, availability, and validation persist, complicating model training.

91%: The vast majority of organizations both believe they are effectively managing AI policy **AND** are confident they can scale data with their current tooling.

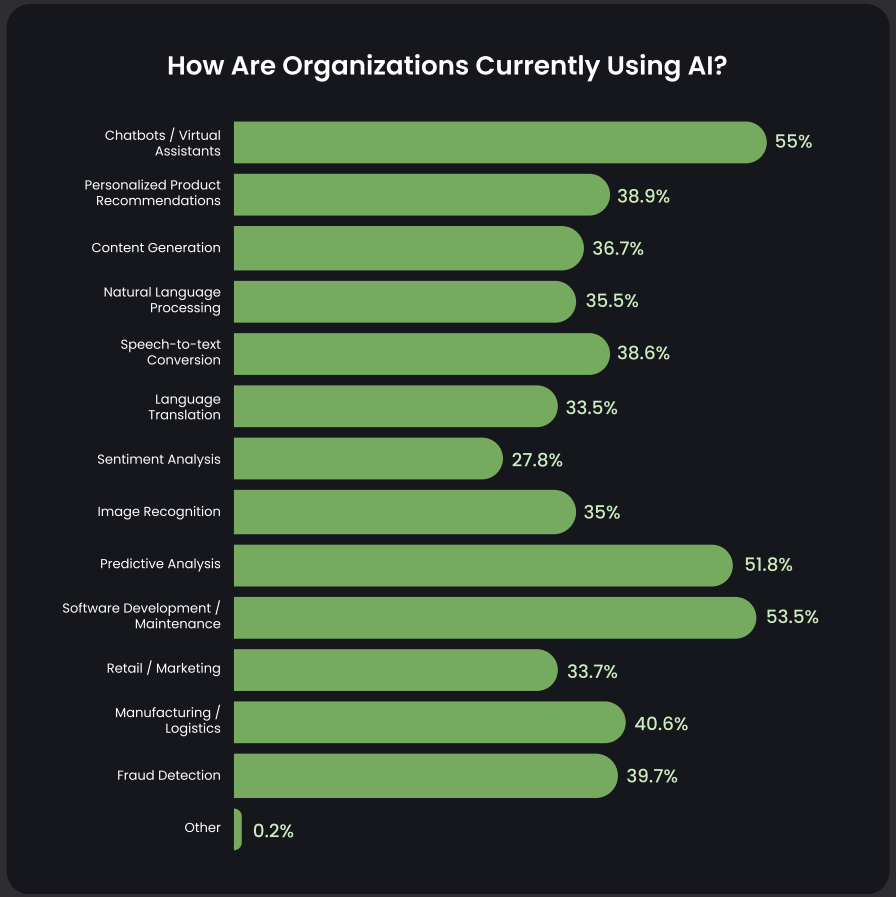

SECTION SIDEBAR:

How is your organization currently using AI?

>50% of organizations are using AI-power chatbots, software development, and predictive analytics

SECTION SIDEBAR:

How Prove AI collects and verifies data for clean deployments

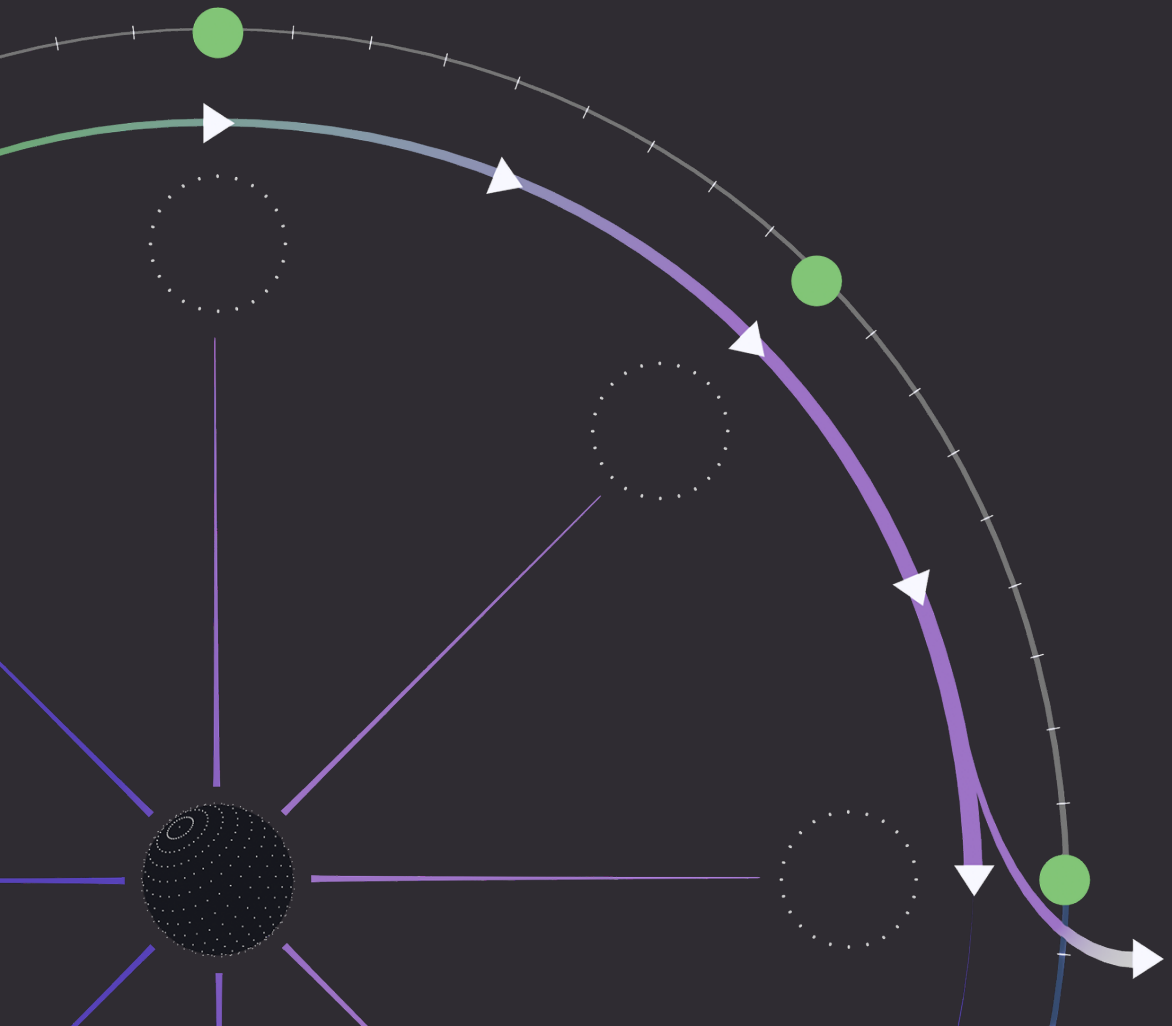

Despite growing investment and strategic focus, AI deployments are not without friction. Nearly **70%** of respondents reported that at least one of their organization's planned AI initiatives is currently behind schedule. The most commonly cited reasons for these delays include issues with data quality and availability, a shortage of skilled talent to manage AI deployments, and difficulties integrating AI systems with legacy infrastructure.

When it comes to deployment priorities, data sovereignty emerged as the leading concern for **83%** of respondents, underscoring the critical importance of knowing where data resides and who controls it. Other top priorities included establishing data lineage, implementing clear guardrails, and ensuring audit readiness.

Although most organizations rated their AI lifecycle management as effective, the processes of data labeling, model training, and data validation stood out as areas with comparatively lower favorability. These foundational steps remain persistent hurdles in the push toward scalable and reliable AI systems.

CHAPTER 4

Balancing Confidence With Control in AI Infrastructure

Despite the existing obstacles with data, business leaders trust their current AI security frameworks. Companies express high levels of confidence in their own ability to manage AI guardrails, ensure data compliance, and maintain traceability.

Still, strategic decisions about AI infrastructure demonstrate that security remains top of mind. Many leaders view non-cloud environments as more secure and efficient than cloud environments, prompting a shift to hybrid or on-premises solutions to better control data and align with evolving privacy expectations. The next 12 months will likely see more transitions away from cloud-native AI deployments.

High confidence in the traceability, security, and compliance of AI deployments:

- **97%** are confident in their ability to establish AI guardrails; of these, **60%** noted they are extremely confident
- **84%** are confident that their AI data remains within the jurisdiction of their SaaS provider
- **91%** are confident in their organization's ability to trace AI data lineage
- **90%** believe that their AI data is secure and honors all user terms and conditions

Despite this confidence, organizations still recognize the central importance of security and privacy in AI deployments, a driving factor behind future plans to shift AI operations to hybrid or on-premise infrastructures. Nearly **9-in-10** respondents stated that their organizations manage at least part of their AI infrastructure in the cloud.

However, citing data privacy concerns, security concerns, cost predictability, and better integration with SaaS AI, two-thirds of respondents believe:

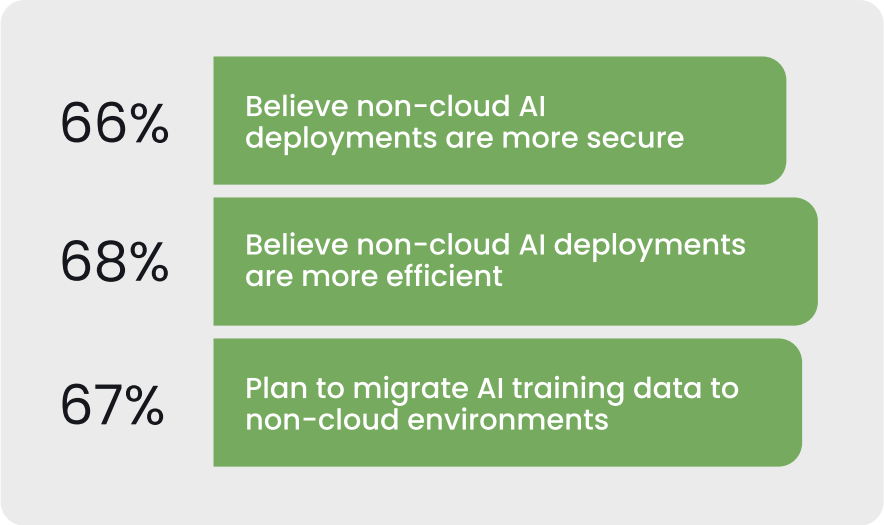

Conclusion

AI is no longer a distant horizon—it's embedded in the present tense of enterprise strategy. Organizations across North America are moving past experimentation, with **nearly 70% of companies actively deploying custom AI solutions** and **planning to scale further in the year ahead.** AI is no longer isolated to chatbots and forecasting; it's powering software development, redefining workflows, and reshaping internal infrastructure.

And yet, ambition alone isn't enough. The majority of organizations report delays in AI rollouts, most often due to data quality issues, skill shortages, and constraints from legacy systems. These hurdles signal that even as AI strategy matures, execution remains complex.

The confidence is real. So is the caution.

Executives overwhelmingly believe their AI data is secure, their policies are effective, and their infrastructure is ready to scale. However, they're also **pivoting away from overreliance on cloud-based tools,** favoring hybrid and on-premises solutions in pursuit of **greater security, sovereignty, and control over performance.** These shifts reflect a new maturity: **a desire not just to deploy AI, but to govern it meaningfully.**

The path forward is clear. As AI becomes central to how organizations operate, so too must AI governance become a foundational layer of enterprise infrastructure. **Transparency, traceability, and trust are no longer aspirational—they are structural requirements.**

Prove AI exists to meet that moment, ensuring your AI systems don't just work, but prove their worth every step of the way.

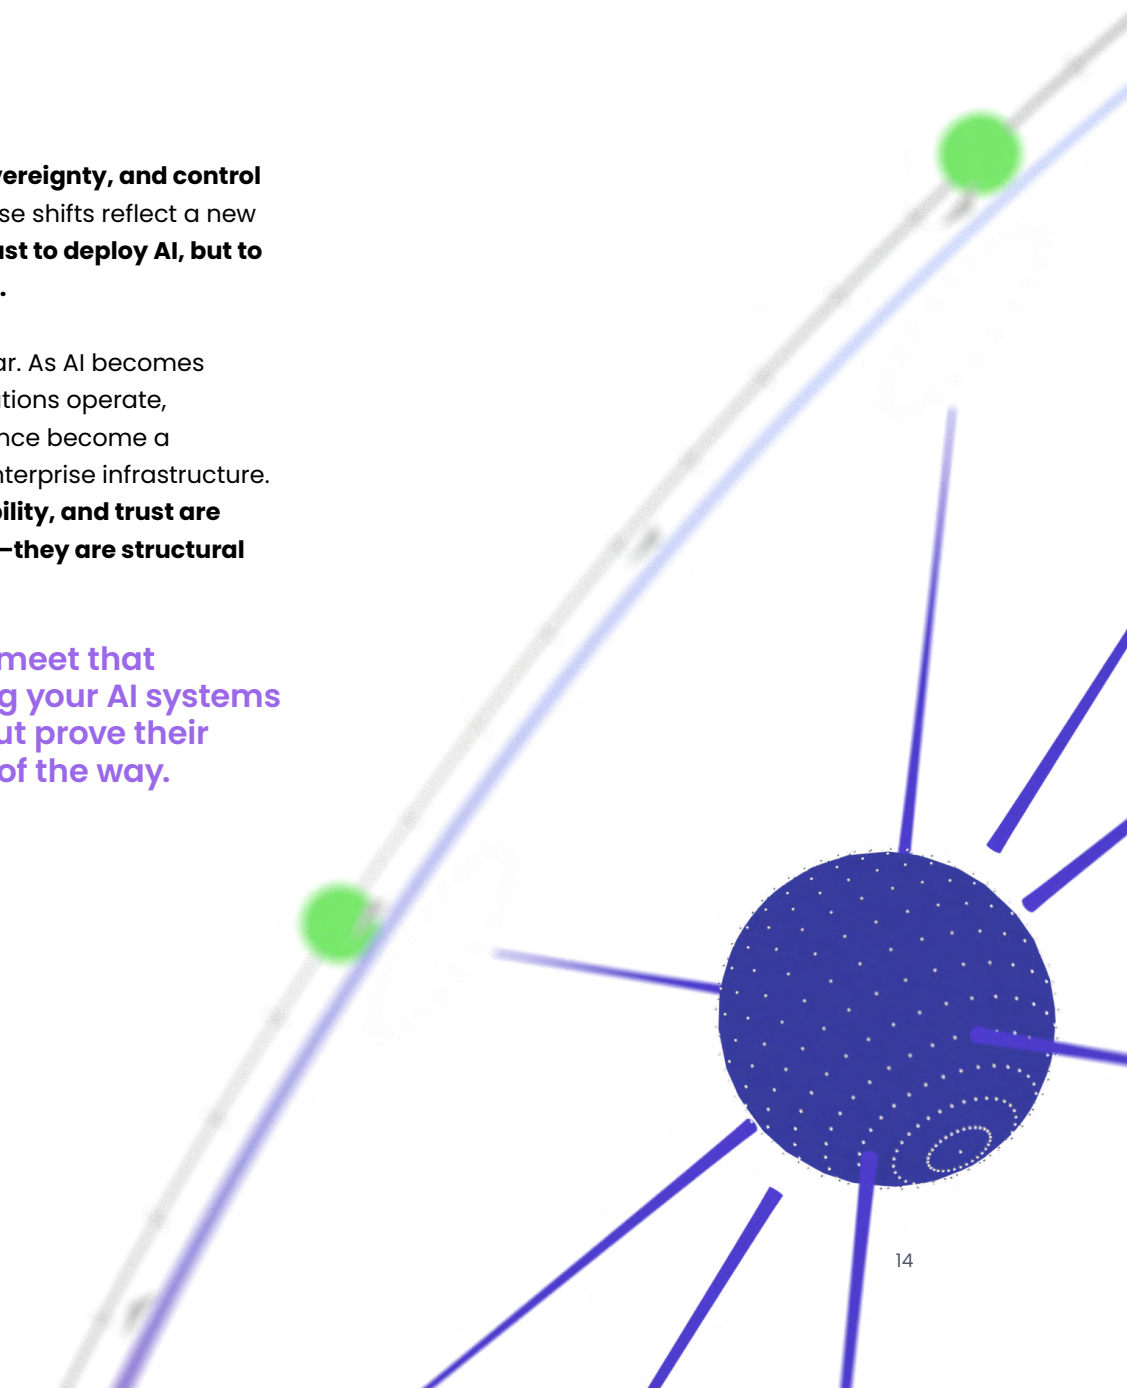

prove ai

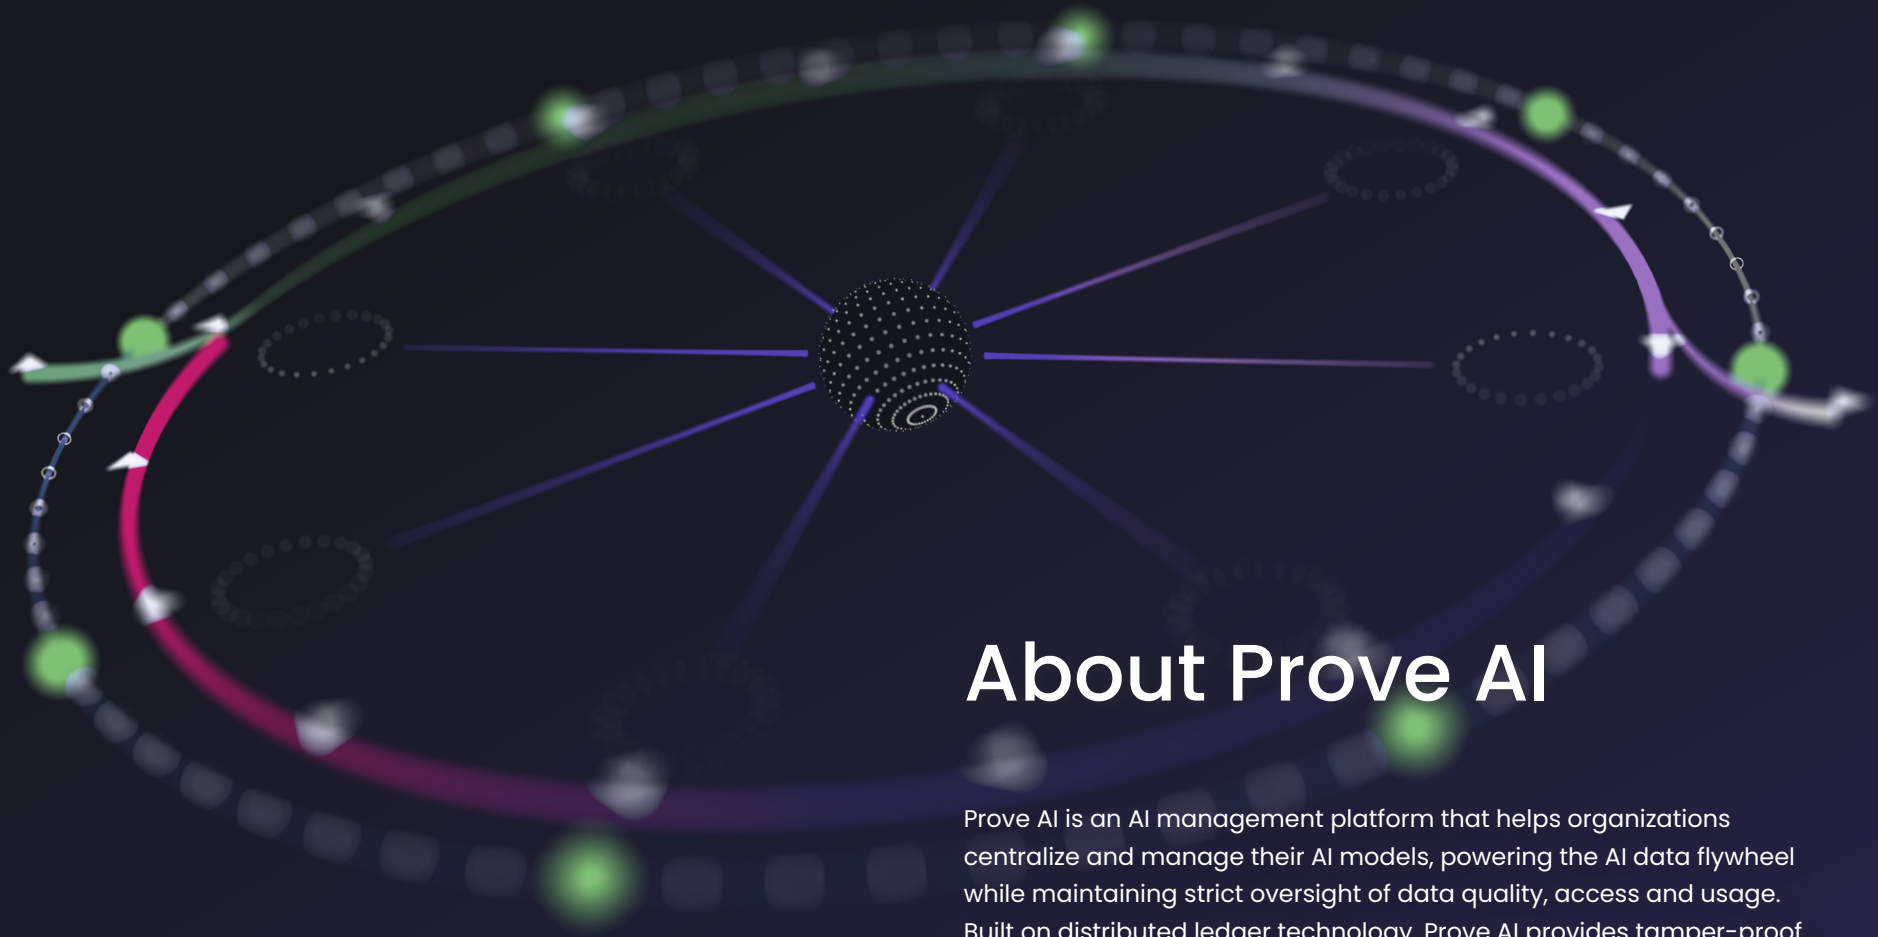

About Prove AI

Prove AI is an AI management platform that helps organizations centralize and manage their AI models, powering the AI data flywheel while maintaining strict oversight of data quality, access and usage. Built on distributed ledger technology, Prove AI provides tamper-proof oversight, ensuring organizations can set and enforce guardrails to mitigate risk and optimize AI performance.