snowflake®

THE DATA EXECUTIVE'S GUIDE TO EFFECTIVE AI

Best Practices from Data Executives for an AI Transformation Journey

TABLE OF CONTENTS

PREFACE

"The Data Executive's Guide to Effective AI" provides a roadmap to help data, analytics and AI leaders effectively implement and scale the AI initiatives across their organizations. To produce this report, I drew on in-depth interviews with data executives at Snowflake customer companies conducted in January and February 2024, as well as discussions with other data leaders working to transform their organizations with new technologies.

Snowflake would like to especially thank the following executives for their time and insights:

MICHELINE CASEY
Chief Data & Analytics Officer, **Siemens Energy**

BETH QUINTON
Vice President, Data, **Air Canada**

PERRY PHILLIP
Chief Data Officer, **Entain**

VIJAY KOTU
Chief Analytics Officer, **ServiceNow**

AMAN THIND
Chief Technology Officer, **State Street Alpha**

PETE WILLIAMS
Director of Data, **Penguin Random House**

BENOIT REID
Director of Data and Analytics, **Altitude Sports**

KIM MACAULAY
Chief Information and Data Officer, **the International Air Transport Association (IATA)**

ANDREW CURRY
Central Data Office Manager, **ExxonMobil**

— JENNIFER BELISSENT
Principal Data Strategist at Snowflake

SETTING THE STAGE

AI is now everywhere. And this time it's really here to stay. Yes, public excitement has waxed and waned as people celebrated innovations but quickly lost interest. Yet this time the current resurgence of AI seems less transitory. What's brought it to the forefront is that it's now within everyone's reach. AI now speaks our language(s). While classical AI, also known as machine learning, could predict sales revenues or a customer's propensity to spend based on a limited set of data, and deep learning could predict foot traffic in a store based on telecom network traffic patterns or video footage, extracting these insights from the data required technical expertise. Generative AI now provides tools that speak in plain language, not code. The renewed excitement has accelerated all AI projects across the enterprise.

In many organizations, gen AI entered through the back door. Less than a year after the highly publicized debut of ChatGPT, most millennials report that they've used a generative AI tool, including for work purposes, and over half report using them regularly. The adoption of gen AI was, as many now refer to it, the "iPhone moment" of AI as a whole.

Yet, according to the **CDO Agenda 2024**, many data leaders didn't know this adoption was happening on their watch. In the summer of 2023, fewer than 30% of respondents reported that employees were experimenting at an individual level, just over 10% reported pilots, and well under 10% reported production deployments. Most respondents were doing little to address growing adoption. Less than one quarter reported preparing data for AI use cases, and just over 10% had acquired additional training data to augment their internal sources.

Data leaders might have been initially slow off the mark to expand AI adoption. However, many are now in a dead sprint to bring not only their technology but their people and processes up to the task. To understand how CDOs are actually prioritizing opportunities and tackling challenges, we asked them directly. Snowflake interviewed about a dozen CDOs who are driving AI innovation within their organizations to capture their strategies and best practices for "The Data Executive's Guide to Effective AI."

As these data-focused leaders explained the challenges they faced in launching AI initiatives, a journey consisting of five distinct stages emerged. This report organizes their experiences around these five milestones, from evangelizing and experimenting with these new tools, to operationalizing their use for greater scale across organizations, expanding use cases and beneficiaries, and eventually embedding them into the organization's DNA to transform the business.

THE STATE OF AI TODAY

While generative AI has been the catalyst for the recent excitement, AI comes in multiple forms. Early AI tools matched job applicants to appropriate jobs or predicted outcomes like fraudulent activity, often based on what would today be considered smaller data sets.

Most early AI tools can be divided into four main categories: matching, predicting, recommending and assisting.

MATCHING

Determine best fit for job or academic placement, clinical trials including overlap analysis for ad placement on compatible profiles on data apps.

PREDICTING

Predict outcomes or the likelihood of specific outcomes or henomena, such as win or loss in sports, disease outbreaks, fraudulent activity, or maintanence needs.

RECOMMENDING

Recomment specific products or services, medical treatments, learning paths, or other actions based on information specific to a segment or individual.

ASSISTING

Support decisions across use cases by identifying general patterns and correlations to streamline choices and facilitate decision making.

With access to larger data sets and the ability to analyze even unstructured data, such as documents, images or video, models deliver more accurate predictions and offer prescriptions to guide better decision-making. Which is the "next best offer" for a particular customer? Or which substitute product could be delivered to replace an out-of-stock online order? Patterns identified in the data offer insights to assist decision-makers — either human or machine.

With the release of generative AI models, we're now not only predicting and prescribing, we're able to access vast amounts of information and create new content. This isn't just Alexa telling us the weather or Siri explaining what causes a rainbow. These are tools that can dig into enterprise records to inform all sorts of critical business decisions. At Morgan Stanley, a new tool designed for its financial advisers and support staff, **Morgan Stanley Assistant**, provides access to more than 100,000 research reports and documents, putting information at their fingertips and increasing employee productivity. A global logistics company fielding shipping requests from existing customers has used an AI tool to extract pertinent information from contracts to accelerate responses.

In addition to information retrieval, these new generative assistants can draft the email response to a customer or the content of a marketing campaign — or even the code to derive insights or build a new application. These copilot tools are some of the most widespread enterprise uses of generative AI, and **early reports** indicate positive productivity gains.

CONCERNS ABOUT AI

The explosive adoption of generative AI has raised awareness of the promise of all AI tools but also of the concerns associated with their widespread use. Among the most prominent concerns are:

Data leakage

The growing use of public generative AI tools raises security concerns among data executives. While it is possible to **configure the settings** to prevent ChatGPT from using conversations to train and improve models, ChatGPT itself will tell you to **avoid sharing sensitive or confidential information**.

That's why the data team at a global specialty food manufacturer went on a sleuthing mission to uncover internal use of ChatGPT by monitoring its network traffic. Lo and behold, they found that roughly 10% of employees had visited and tested ChatGPT in the first quarter of 2023. Yet one team was responsible for 60% of the ChatGPT traffic, and in fact it was only one employee — an intern learning to code side-by-side with her AI copilot. While this is a common use case among developers and younger employees, it's one to monitor to ensure usage is in compliance with company policy and employees are instructed on how best to use these tools.

For a VP of Data, Analytics and IT Innovations at a manufacturer of recreational products, the risk that proprietary or sensitive data, shared into an open or public platform, could be misused in violation of regulatory requirements is too high. That fear of data exfiltration is the driver for a very controlled access policy regarding the use of open source gen AI. "We need to ensure that employees know how to use these new tools properly," he says. "Establishing a companywide policy is the top priority."

> **"We need to ensure that employees know how to use these new tools properly. Establishing a companywide policy is the top priority."**
>
> — VP of Data, Analytics and IT innovations at a manufacturer of recreational products

Capture of historical bias

The potential for data bias is well known among data leaders. Stories of institutional bias in housing, hiring, loans and policing are not new, and will likely continue to surface, particularly as they are captured by AI models trained on historical data. If an HR department wants to identify a profile for a specific role in its organization, for example, using only internal data would capture the characteristics of past employees in that role.

These potential internal biases exist across all decisions. If a retailer wanted to expand into a new market, forecasting demand with only internal data would produce results based on sales in its current market. A move from the Nordics into Southern Europe, for example, wouldn't be well informed without external data about the new market.

Most data leaders understand how to mitigate these risks, but should keep an eye on the evolving regulations. It's concerning enough that incorrectly forecasting demand in a new market might result in bad business decisions, but discriminatory hiring practices could be a regulatory violation with more significant consequences.

Inaccuracies and fakes

A Forbes Advisor survey found that 76% of consumers are concerned about misinformation from AI tools such as Google Gemini, ChatGPT and Bing Chat. At the same time, 54% of those surveyed said they think that they can tell the difference between content written by a human and that generated by automated chatbots. But have you tried? The New York Times published a quiz to identify AI-generated images, and it's not easy. Now new tools, such as AI or Not, can help determine whether an image has been generated by artificial intelligence or a human. But when we interact with chatbots, we assume the information provided is correct. And businesses may be held accountable for the information their chatbots deliver.

> **"As people leave an organization, those who start on the journey might not be the ones who finish it. Use cases and models will get orphaned and passed around."**
>
> **— PERRY PHILLIP**
> CDO, Entain

Job disruption

According to the Forbes survey, one of the biggest concerns among the broader population is job loss: AI will take our jobs. Some tasks will be replaced by AI, like the time spent by a paralegal finding precedents in previous cases. For many others, the new arsenal of AI tools will transform jobs as we know them. They will replace certain tasks and accelerate others. Imagine life before your mobile phone or even a computer. But AI has limitations: It lacks judgment, doesn't understand context, and is only as good as the data it's trained on. And that's good news for humans. We might not be completely replaced — at least for creative, innovative thinking. That's what humans are good at. But to get the most out of AI tools, employees will need to be retrained.

At a CDO-level networking event in London earlier this year, data leaders emphasized the need to address that foundational skills gap. The proliferation of gen AI accelerates the need to improve data literacy of current employees, but it also raises the question of who is responsible for providing that retraining, and even incorporating data and AI training into future onboarding? Is it HR or the data team, or a collaboration among stakeholders? At ExxonMobil, CDO Andrew Curry reports that the data literacy program is owned by their Central Data Office, which develops curricula for a mix of business and technical personas.

Organizational turnover

It's not only about job loss. Data leaders also worry about people loss. According to Perry Philipp, CDO of Entain, one of the biggest risks during the adoption of AI technologies is turnover. "As people leave an organization, those who start on the journey might not be the ones who finish it," Philipp says. "Use cases and models will get orphaned and passed around."

The lack of continuity can derail new projects and jeopardize business benefits. Project management and governance will need to fill in the gaps.

The increased AI adoption and awareness of concerns has galvanized data executives to accelerate policies and coordination across their organizations. For many, strategy and processes remain a work in progress. However, best practices are emerging.

THE JOURNEY TO EFFECTIVE AI

The challenge with these new AI tools and technologies isn't to identify useful applications of them. Recent history has demonstrated that users are eager to put them to use. For enterprise leaders, the real challenge is to figure out how to enable their use effectively, responsibly and at scale.

At Siemens Energy, the data teams have spent a lot of time with business areas to refine their strategies, and have introduced an AI maturity model to track progress in their implementation. A plethora of maturity models illustrate the potential evolution of AI within the enterprise. No one size fits all; CDOs report pulling from many to create a model that fits best. While some models suggest that the journey is unidirectional, as companies achieve each successive level of maturity, that's not usually the case. Data leaders report that stages overlap, and activities that characterize earlier stages persist into subsequent stages. Especially in large organizations, not everyone will be at the same place at the same time.

Consider the different stages illustrated in the diagram below. Data leaders wouldn't abandon evangelism when they begin to experiment, or forego experimentation when they begin to operationalize AI. A transformation is never really complete. Think of this journey as ongoing, a continuous evolution with new technologies to evangelize and use cases to test. While the stages might seem basic, they provide a useful organizing construct by breaking down a complex transformational process. Data leaders use these for planning purposes, and to enable scale as adoption grows. It is important to recognize that transformation is a journey across the whole company and even across an ecosystem, not just in pockets with ad hoc projects.

For some organizations, the transformation requires pulling together existing yet disparate initiatives, with the formalization of the CDO role or its expansion into a Chief Data and AI Officer (CDAIO). For U.S. federal government agencies, designating a Chief AI Officer **is now required**. Proposed legislation in the UK would require that all organizations using AI **designate an AI officer**.

Conversations with CDOs across industries provide insights into the journey, and how they have navigated challenges to deliver on the promise of AI. Their experiences offer lessons learned and some best practices for others embarking on the journey.

Stages in the Journey to Effective AI

Use the Stages to Plan, Execute, and Scale

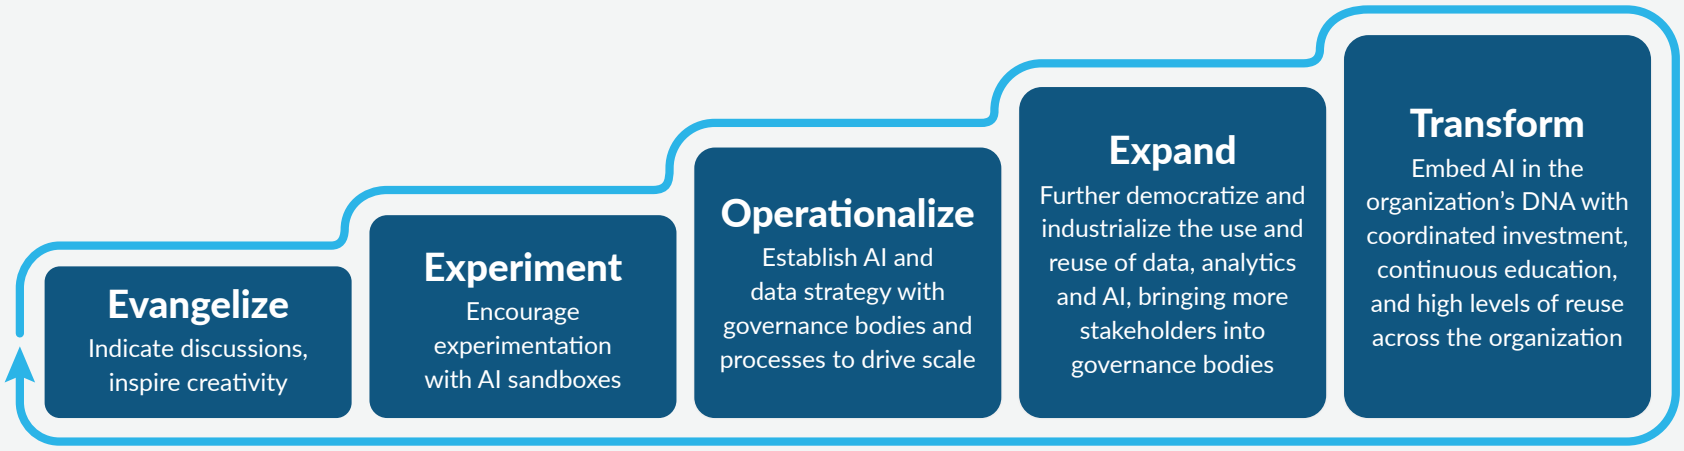

STAGE ONE: EVANGELIZE

While many projects enter the enterprise through the back door, there might not be back doors into all areas of the organization. The goal of evangelism is to drive awareness systematically across the entire organization. Cashiers need to understand that the data they capture when ringing up a sale influences inventory and replenishment orders at the end of the month. Or when field technicians close out a work order, they must know that the information they've captured will be analyzed. The time spent on the repair will determine equipment replacement or predict future maintenance needs. Evangelism is required from the top floor to the shop floor. Growing awareness is also a source of new ideas.

Communicate clear and simple messages relevant to the business. While water cooler conversations might spark serendipitous collaboration, AI evangelism must be more systematic and cross-company. It must also be audience specific. At Air Canada, CDO Beth Quinton has presented to a lot of different audiences, starting with high-level content to build awareness at the department level. At a senior leadership summit, the content was less technical and more about change management and the attendees' roles as leaders. While some executives have an understanding of how they need to adapt and are willing to make the necessary changes, others push back, saying they're too busy, or don't want to deal with automation. Speak their language, she says, and make sure the message is clear.

"We want to be on the winning side," Quinton says, "and we think those who don't experiment and invest will fall behind."

Repeat messages to set expectations and drive real change. A continuous dialogue reinforces the message and eventually wins converts. That ongoing evangelism is key to driving change — and setting expectations. As Benoit Reid, Director of Data and Analytics, Altitude Sports, observes, "People think that AI is a magic potion that will solve everything. But you can't just put data into the machine and expect it to do everything." He goes on to note that, "It's not about doing it once. It requires being consistent with reminders and training on the power and risks." AI training is now embedded in ongoing security training in many organizations.

Extend evangelism to customers and partners. When AI initiatives affect customers, external evangelism and engagement can alleviate concerns and earn their clients' confidence. At State Street Alpha, educating customers was also part of the company's Responsible AI Framework, says CTO Aman Thind. State Street Alpha offers asset servicing, investment management, market research and trading services to institutional investors worldwide. As AI is incorporated into its platform, clients need to know about the changes and what to expect. "We need to win our clients' confidence," Thind says.

A specific client advisory board on AI allowed State Street Alpha to get client input up front, as well as solicit ongoing feedback. In addition, client-facing webinars delivered updates on new technology and use cases, and provided a forum for questions from an even broader audience. Users were encouraged to forward webinar invitations to other client stakeholders within their firms.

> **"We want to be on the winning side, and we think those who don't experiment and invest will fall behind."**
>
> **— BETH QUINTON**
> CDO, Air Canada

EVANGELISM OF AI AND DATA MUST EDUCATE ALL EMPLOYEES

Everyone in an organization plays a role in AI and data transformation — capturing data, protecting it, analyzing it or making decisions based on it. But not everyone initially realizes that. One data executive shared a story about the early days of digital transformation at a food services management company. The data team saw a spike in the purchase of breakfast sausages in one of their corporate cafeterias. Then suddenly, no one was buying anything else: no croissants, for example, and this was in France. Certainly, a curious insight into the culinary habits of the employees at this site. But was it?

It turns out that the change in the data corresponded with the adoption of point-of-sale terminals with individual buttons for each item. The design was intended to capture the item sold and facilitate inventory management. But breakfast sausage was the button of choice. There was no malice involved; sausages and croissants were the same price, and the cashier merely wanted to move customers along quickly. If the site managers had used this data to determine future food orders, however, customers would soon have found their choices limited to breakfast sausages, with no croissants on the menu. AI transformation requires education across the entire organization. Everyone plays a role.

STAGE TWO: EXPERIMENT

As word spreads, the curious will want to experiment with AI. Experimentation explores the art of the possible and generates ideas for ongoing projects. However, that's not all. The value of experimentation isn't only in the models built, but in the experience itself. In the long run, will most companies build their own AI models? Air Canada's Beth Quinton doesn't think so.

"We could build everything," she says, "but we shouldn't. We're an airline. However, the experience of building models has helped us develop firsthand knowledge to better evaluate potential purchases."

The bottom line: Experimentation is a good thing — when done correctly.

Encourage broad yet managed experimentation. While leaders admit the use of new AI tools started ad hoc, some advise to not just let things happen organically. Yet it's important to encourage experimentation and get started, stresses Micheline Casey, CDO at Siemens Energy. To that end, her company decided to open a generative AI sandbox in 2023 to anyone who wanted to play in it. To encourage broad experimentation, traditional business case approaches were done away with. As Casey argues, the time and requirements for scaling, along with potential value creation, are generally unknown at the ideation phase with gen AI. Instead, users have to fill out a use case form, and a stage gate process formalizes the testing with structured steps. "We get much clearer about the potential value, use case targets and scalability parameters with this approach," says Casey.

Host a hackathon to expand ideation. A hackathon — sometimes called a datathon when building AI or analytics applications — helps expand and accelerate ideation. Where experimentation hasn't happened organically or evenly across the organization, leaders might want to provide incentives or find a solution for a specific use case. A hackathon is an intensive, goal-focused event, often held over a full day or weekend. These events promote collaboration and creativity, often more with the goal of finding novel approaches to problems than of developing fully usable applications. A recent **hackathon hosted by the International Organization for Migration**, "Bridging Climate Change and Human Mobility," explored the feasibility of modeling the links between different drivers of human mobility. Participants used 70 different data sets and new machine learning tools to analyze migration flows and identify new indicators.

At another recent **hackathon for the UK utility sector**, 30 participants from different organizations came together to address persistent challenges. Several data providers offered their data free of charge for participants to come up with different innovative solutions. Two industry experts served as judges for the competition. The resulting collaboration offered solutions for better identifying vulnerable households facing risks associated with energy costs, developing intelligent demand-side management solutions, and ensuring the resilience of the utility grid as it transitions to multiple energy sources.

> **"We could build everything, but we shouldn't. We're an airline. However, the experience of building models has helped us develop firsthand knowledge to better evaluate potential purchases."**
>
> **— BETH QUINTON**
> CDO, Air Canada

A Prioritization Matrix Guides Use Case Selection

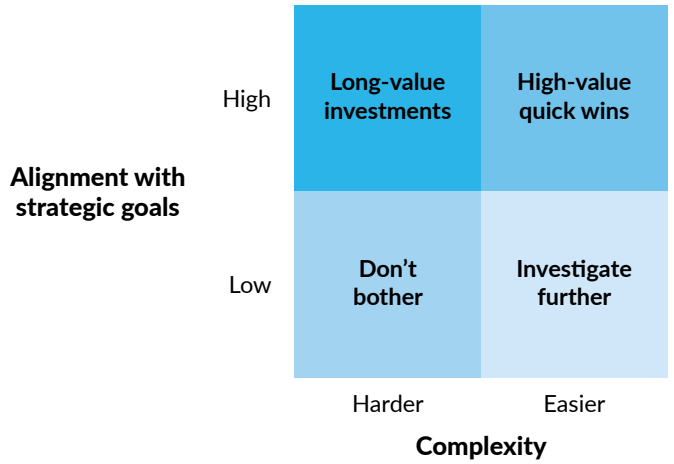

Prioritize selected use cases. Be deliberate. Not all ideas will survive. Yes, you want to ideate and experiment. You want to cast a wide net and capture all ideas, but you're not going to fund them all. You need to be very intentional about what you invest in for the long term. At Siemens Energy, the 300 projects in its AI sandbox were whittled down to 30, based on strategic business objectives.

Across organizations, the prioritization happens either centrally, but coordinated, or in a distributed manner, with each domain prioritizing its own initiatives. Some focus on improving performance or reducing costs of existing processes. Data quality improvements is a common priority.

At State Street Alpha, AI anomaly detection models are used to manage data quality, emphasizing a circular dependency between AI and data. That focus on data makes sense, Thind says. "Data is more valuable than the AI model itself. Models are converging. It is the nuances in the data that differentiate the outcomes."

Other more down-to-earth use cases often feature augmenting humans with automation but don't feature the lifelike AI characters of sci-fi movies. In its report, "The Executive's AI Primer," IT research firm *Forrester* advises against"marquee AI use cases" that oversell and end up relying on a wizard behind the curtain. Keep it simple to start.

> **Data is more valuable than the AI model itself. Models are converging. It is the nuances in the data that differentiate the outcomes."**
>
> **— AMAN THIND**
> CDO, State Street Alpha

ENTAIN PRIORITIZES PLAYER PROTECTION USING AI TO IDENTIFY RISKY BEHAVIOR

Sometimes prioritization reflects core values that might seem counterintuitive to short-term business interests. Entain, one of the world's largest sports betting and gaming groups, created its **Advanced Responsibility and Care (ARC) program** to address the best interests of its users, which meant signaling unhealthy behavior and putting guardrails in place to limit play. Developed in collaboration with leading experts, such as Harvard Medical School faculty, the model identifies 26 signals of harmful behavior — such as play frequency, intensive betting, increases or variability in wager amounts, the adoption of riskier gambling positions, and increased losses — with an accuracy rating of over 90%, according to Entain. Once identified, Entain will alert the player and offer preventive measures, such as loss limits or temporary lockout.

At the end of 2022, ARC was rolled out to 22 jurisdictions globally. That same year, Entain reports, there were 3.7 million interactions and interventions globally, representing 670,000 unique players, and resulting in a 36% reduction in harmful behaviors following interventions. The program didn't automatically translate to positive Net Promoter Scores, or add to the bottom line. As Entain CDO Perry Philipp observes, "We can be happy that ARC reflects our corporate and personal values. Longer term, it's a key sustainability driver for the business."

Any short-term penalty to the bottom line is overshadowed by the well-being of its customers and business sustainability by reducing the risk of noncompliance with industry regulation. Moreover, Philipp says, safer gaming can be a competitive advantage in the long run.

"More than 95% are gaming for entertainment, but there are customers who are at higher risk," he says. "We're pleased to be able to help our customers enjoy the entertainment of our products, but intervene at crucial points on their journeys where necessary."

And they do, with 97% of higher-risk customers and more than 73% of medium-risk customers using gambling controls following interventions, the company reports.

STAGE THREE: OPERATIONALIZE

The way to expand AI use effectively — and increase the value delivered from it — is to operationalize and industrialize. That requires alignment between the technology and the business, coordination of resources, and standardization of policies and processes.

Align your AI and data strategies to your business. In the immortal words of Arthur Conan Doyle's Sherlock Holmes, "Data! Data! Data! I can't make bricks without clay!" Or as Frank Slootman, Snowflake's former CEO, says, "There is no AI strategy without a data strategy." Moreover, successful data leaders know to align their data and AI strategies to their business objectives. If AI is the fuel for growth, as many are saying, data is the fuel for AI.

In true data and AI alignment, data leaders use AI to ensure the quality and freshness of data. As one data leader observed, "While everyone in the room is excited about the benefits that we can all imagine we're going to get from it, there are certain foundational blocks that you have to have in place before you even attempt that journey."

> **"We're rolling out AI/ML Ops with a registry to facilitate re-use and quality monitoring of models. That's how we'll be able to move fast. Reusability is one of the KPIs that we report to the board."**
>
> **— MICHELINE CASEY**
> CDO, Siemens Energy

Get that first data step right, and the rest will be easier — and potentially less expensive. That's not to say you have to wait until all the data is clean and in place. The foundation is about operationalizing: getting the policies and processes in place to accelerate out of the turn.

Define standard operating procedures collaboratively. To operationalize is to put the strategy into practice, and that requires policies and standard processes. But standardization requires coordination and collaboration. A first step for many is to bring together the technology teams and risk or privacy management. At Air Canada, the first order of business was to bring the privacy team into the process. Similarly, at State Street Alpha, CRO and CIO coordination ensured that development has been done responsibly. Specifically, a model risk management team under the CRO established a responsible AI framework. Then the bionics team under the CIO takes that framework and builds technology guardrails to ensure compliance — no data leakage, no data privacy issues. Responsibilities are well established, creating checks and balances.

Establish a continuous operating model. When implementing AI at scale, standard processes must be part of an operating model. In addition to the prioritization process mentioned earlier, companies are establishing MLOps frameworks for streamlining the development and deployment of predictive and generative models. A robust MLOps framework is designed to accelerate processes, facilitate cross-functional collaboration, increase scalability and ensure transparency.

For many, it's still a work in progress. For some, an AI factory is in the works; for others, it's a more distributed approach with appropriate governance. A first step for many companies has been the launch of a model registry to track all AI models, the data they're using, their performance, and the value they are delivering.

"We're rolling out AI/ML Ops with a registry to facilitate re-use and quality monitoring of models," says Siemens Energy CDO Micheline Casey. "That's how we'll be able to move fast. Reusability is one of the KPIs that we report to the board."

Similarly, ServiceNow has built an application internally to govern its AI models. Others use their existing PMO to prioritize, structure and manage AI initiatives. There is no one-size-fits-all method, but the basic approach answers simple questions:

- **Which data is the model using?**
- **Is appropriate governance in place?**
- **Is the model running?**
- **Did the model deliver value?**

Source and use external training data responsibly. More data sources ensure data diversity, mitigating the risk of AI hallucination and bias. But sourcing of external training data requires coordination and must take an integral role in AI model ops. At the International Air Transport Association, industrywide initiatives, such as traffic forecasting and a CO2 calculator, require sharing data with airlines; manufacturers such as Boeing, Airbus and Bombardier; and industry data providers such as OAG. Four new airlines came onboard earlier this year and more are expected over the course of the year. Where sensitive data is involved, IATA needed a trusted research environment. For example, data clean rooms can facilitate privacy-preserving data collaboration. Or as another CDO referred to it, "answer sharing," rather than data sharing. Parties to the collaboration can ask questions of the data and receive answers that don't reveal the underlying data.

Data from multiple sources, however, brings some complexity. "When you get data from partners, the quality and thoroughness is not the same," says Altitude Sports' Benoit Reid. "You must set standards and monitor the incoming data against them. Don't be afraid to push back on poor quality. It's OK to say no — you won't accept the data if it doesn't meet the quality standard."

> **"Don't be afraid to push back on poor quality. It's OK to say no — you won't accept the data if it doesn't meet the quality standard."**
>
> **— BENOIT REID**
> Director of Data and Analytics, Altitude Sports

Define and track performance metrics, starting with model output. Operationalizing AI means knowing how the models are doing. A first step is to monitor the output of the model. Does the model perform as expected from a technical perspective? The team at ServiceNow has instituted quarterly performance reviews with specific metrics, chosen before the model is deployed. If those metric requirements are not met, then the model needs more work. These output metrics include:

- **Hallucination rate, or accuracy. Does the model deliver expected answers?**
- **Precision and recall, or a standard confusion matrix. Does testing reveal true positives and true negatives?**
- **Average error of forecasts. How accurate is forecasting?**
- **What value is the model delivering to the company?**

Air Canada establishes test criteria and the ongoing performance metrics for each system put into production. Yet that's not always easy. Performance metrics for traditional AI models are known. They are quantitative. For example, models are used to predict on-time performance. Actual flight times are compared to predictions; a model is expected to have 98% accuracy. But what about gen AI?

Air Canada has applied that same process to generative AI, but it's not always possible. Different systems might prioritize different metrics, preferring accuracy or consistency or both. Take a knowledge management use case where a guided flow helps agents communicate the policies to customers. There is a right answer to the customer's question, "How much does it cost to send a large dog from Toronto to Colombia?" Agents need to have the same answer (consistency) and the right answer (accuracy). But depending on the model and tools you use, you don't always get the right answer. If the answer is in a table, for example, the tool might not read it correctly. That requires further testing. Air Canada is now part of an academic consortium exploring these kinds of issues.

Also track business outcomes and attribute value to the data and AI applied. The holy grail of all data initiatives is to deliver business value. Demonstrating that value, however, is often a tall order. The first step is to benchmark existing programs to ensure the ability to identify the impact of applying new AI models. That's part of the process at ExxonMobil. Its Central Data Office works with business units to identify goals and requirements for data products.

"The Central Data Office guides the business with a formal process," says ExxonMobil's Central Data Office Manager, Andrew Curry. Data products are built according to the needs of specific business capabilities, such as supply chain or operations. "The business units determine their objectives in their domain, and identify which data they need, the quality of the data, etc.," added Curry.

This helps ensure the resulting data product provides the value that the business needs. ExxonMobil's Central Data Office also looks across business units to coordinate requirements and identify possible re-use, which increases the value generated by a single data product. As ExxonMobil's Curry explains, the supply chain team wants operations data, but the central team recognizes that others might want that data product in the future. The objective is to achieve data gravity, in which the more the data product is used, the more others will want to use it, delivering a greater ROI.

Establish performance management, combining automation and human monitoring. Most data leaders recognize the importance of human-in-the-loop. A combination of automated performance monitoring and human feedback on outcomes provides the best performance management. Rather than perceiving AI as a mere tool, the advice of experts is to envision AI as a remote team member — an additional employee capable of improving business operations or customer experience. Doing that changes the dynamics of the relationship. We know how to manage an employee, and we have tools and processes in place to do so.

- **Define the role with an explicit "job description":** Does it chat with customers, interpret longer documents, analyze data, flag anomalies? Set expectations.
- **Identify the right AI model "candidate," either to be bought or built:** Different AI and gen AI tools have different capabilities.
- **Invest in "onboarding" by training models with appropriate data:** Prepare your models to perform their specific function, with continuous training to keep them up to date with job requirements.
- **Establish output and outcome performance metrics:** Having defined the "job," define a job well done and how you'll measure it.
- **Monitor performance against benchmarks to attribute value generation:** Benchmarking performance against a baseline established before the arrival of this new "employee" will help justify investment in data and AI.

These methods ensure models are kept up to date and continue to deliver value.

STAGE FOUR: EXPAND

This next stage in the journey is to further democratize the use and re-use of data, analytics and AI. Expansion means more users and requires the ability to scale through "industrialization," or building of AI factories. As Siemens Energy CDO Micheline Casey says, "We want to democratize AI — but do it responsibly. We are using this opportunity to rethink the approach to the operating model, steering and governance necessary to support expansion of AI use."

Drive and manage further change inclusively. Effective AI is not a solitary pursuit. At Air Canada, while the first order of business was to bring the privacy team into the process, now the Operating Committee (OpCo), which meets biweekly, is even more inclusive. Initially the meeting was open to anyone interested. As interest grew, the meeting was split into two: the OpCo for those executing AI initiatives and driving outcomes, and an InfoForum for those who are interested in learning more about AI initiatives at Air Canada. Similarly, at IATA, the generative AI working group includes not only data scientists but also data enthusiasts from departments without data teams. These broad forums allow continuous evangelism and education.

A Center That's Not Too Centralized

A data and AI center of excellence doesn't have to be central command. Rather, it can be a hub that coordinates across business units to establish strategy and collaborate on shared responsibilities. The spokes may define use cases and business needs that the CoE helps address.

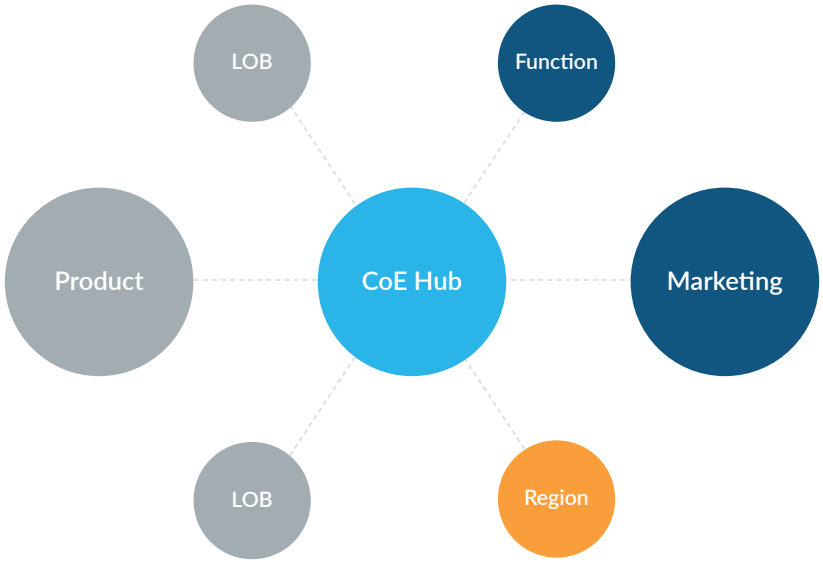

Inclusivity at the executive level also increases the likelihood of broad support. Air Canada started with a monthly executive steering committee that included representation from the CIO, CFO, the Chief Talent Officer (HR) and the Head of Digital, Data and Marketing. Since then they have moved to regular updates and quarterly program updates incorporated into the sustainability governance framework.

Coordinate use and policies with an AI Center of Excellence. Somewhat ironically, establishing a center of excellence (CoE) is one of the best ways to build a broader base of knowledge and understanding of data and analytics. It's not just an ivory tower. In fact, a CoE is not necessarily a central organization but rather a constellation that includes a hub to coordinate across business units, domains and other entities within and beyond the organization. The spokes are different business units, functions or more granular data domains.

A CoE dedicated to data and AI can be organized in a number of ways. Some organizations might have more capabilities centralized, while in others the spokes are more data-mature. There are gray areas. The role of the CoE is to capture requirements, coordinate initiatives and galvanize resources. By creating a CoE, leaders highlight the strategic nature of data and AI within their organizations.

> **"Smaller models that do specific things can be well-tested and understood. Combining those capabilities allows a better understanding of what's going on behind the scenes."**
>
> **— AMAN THIND**
> CTO, State Street Alpha

Several Snowflake customers have set up AI CoEs, with varying distributions of responsibility. For ServiceNow's Chief Analytics Officer, Vijay Kotu, coordination was critical to connect the dots at the company level, eliminating duplication and inconsistency. A centralized data team establishes the enterprise schema and the data models, creates foundational data products, and manages the tools that enable federated teams to perform additional insights tasks. At IATA, the company's AI and data center of excellence reaches across the aviation ecosystem with use cases around safety and sustainability.

A CoE is just good practice, says State Street Alpha's Aman Thind: "A hub-and-spoke model creates good discipline of which models to use and how to govern them."

Reinforce data product thinking for scale and explainability. Data leaders emphasize the importance of establishing data products as a cornerstone of their data strategy. Coordinating the development of foundational data products and AI models improves performance, and ensures the re-use and efficiency required to scale.

At State Street Alpha, the centralized bionics team is mandated to create reusable models to be used across the organization. They have streamlined the process of picking models, validating models, making sure they fit into the overall risk framework, and making them available to the rest of the community. One advantage of composability is explainability, Thind says.

"Smaller models that do specific things can be well-tested and understood," he says. "Combining those capabilities allows a better understanding of what's going on behind the scenes."

The bottom line is that coordination makes models better. "Everyone benefits from each other," Thind says. "If 100 teams built them, they'd have 100 different models that don't learn from each other. But with a centralized model, they all learn from each other."

The way to expand efficiently and effectively is to collaborate across the organization.

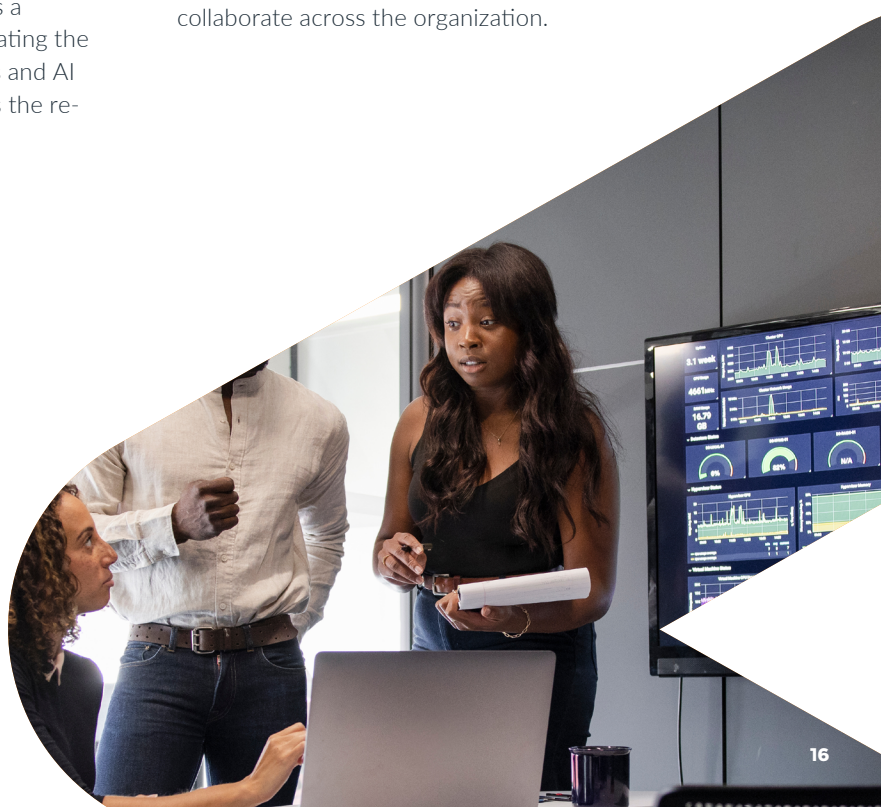

STAGE FIVE: TRANSFORM

True transformation means that AI is embedded in the organization's DNA with coordinated investment, continuous education, and high levels of collaboration and re-use across the organization. You "want AI to be a program, as opposed to a technology," says State Street Alpha's Aman Thind. "Only thinking about it as a technology is limiting. Programs have multiple pillars and stakeholders, all standing on the same platform and making decisions together."

This kind of collaboration, he says, "allows you to be more fearless about AI adoption. Usage is blessed not only internally but by clients. And you can adopt it such that it becomes part of your DNA, not separate, but part of how you think and what you do. That's the only way AI within a firm can be successful."

Coordinate and justify ongoing investment. The excitement around AI has driven some unsustainable practices. At Air Canada, Beth Quinton worries about resourcing, among other things. "We can't just have people do it off the side of their desks," she says. "We do have curious people on the cutting edge, but that alone doesn't scale. This is something new that wasn't in the plan 12 months ago. We're adjusting the structure of the organization and the budget to enable people to work on this."

Reflecting on the mission and the vision of the Central Data Office — to maximize the value of ExxonMobil's data — Andrew Curry, Central Data Office Manager, has focused on how to calculate that value. "I dislike the phrase 'data is an asset,'" he says. "I prefer to think of data as a resource, which in turn raises the question of how you invest in that data." For ExxonMobil, investment is justified by the value estimated through robust AI performance management.

Build a program for communication and community building. To build the necessary support — and a strong data culture — data leaders have launched communications initiatives. At IATA, Kim Macaulay launched "Falling in Love with Data" for the month of February, putting a spotlight on demystifying generative AI. What can you use the tool for? What do you need to be aware of? To extend communication to a broader ecosystem, IATA has planned a series of webinars with the airlines and two in-person industry events. The aviation-exclusive Data and Tech Summit in 2025 will offer an "Ask Our CDOs in Airlines" booth to encourage questions.

At ExxonMobil, Curry heads up a data culture team that sponsors a series of events for communities of practice in different cities. These events include "A-chats" — internal show-and-tell sessions that bring internal speakers to present various data trends and AI topics — along with showcasing projects. In addition, the "A-Games," quarterly hackathons, had more than 500 people participating across 21 countries last year, with six finalists chosen by a panel of experts. This program not only provides excellent training opportunities but offers professional recognition, with gamification making it fun.

Build an Organization of Data A-C-E-S

A framework for building comprehensive data literacy

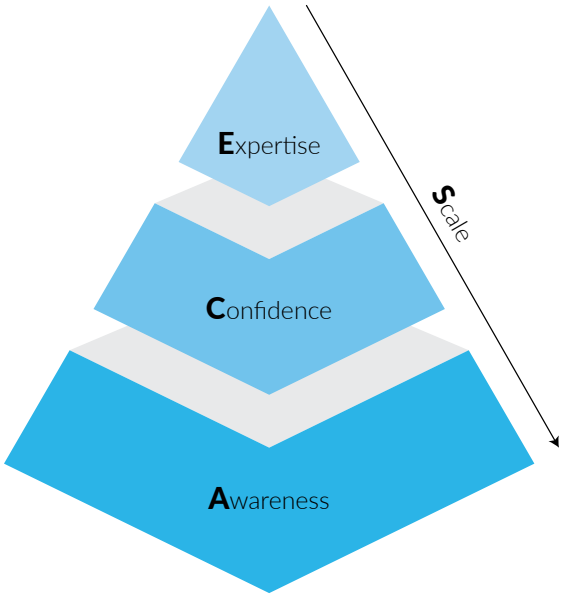

> **The purpose of AI is to improve decision-making in every role, from CEO to account executives, engineers and marketers. What capabilities would we need to put into place for them to benefit?"**
>
> **— VIJAY KOTU**
> CAO, ServicesNow

Offer comprehensive education to uplevel the organization. Not all employees can participate in a hackathon. In fact, most employees are not data experts. Keeping in mind the wide range of knowledge and even basic understanding, AI and data literacy programs must educate all levels. A comprehensive program would addresses four goals:

1. **Raising awareness across the broader organization**
2. **Increasing confidence of decision-makers who will take action based on insights**
3. **Enriching expertise of data and AI experts with new tools and data**
4. **Enabling scale by encouraging data teams to share their expertise and disseminate project information and results**

At ExxonMobil, the data team offers an extensive data literacy program. To understand their baseline, the team has developed a self-assessment tool available to anyone in the organization.

Tailor education to specific audiences — including company leaders. For ServiceNow's Vijay Kotu, "The purpose of AI is to improve decision-making in every role, from CEO to account executives, engineers and marketers. What capabilities would we need to put into place for them to benefit?"

That then begs the question of how to train each of these roles in how to use the AI that has been developed for them. Education programs must be tailored to personas such as the data novice, the data leader, the data expert, or to other specific roles in the organization. At State Street Alpha, Aman Thind says, the data team has held multiple town halls including show-and-tells for the developer, executive and even board level to ensure that there is a good understanding.

Siemens Energy's Micheline Casey agrees that educating the board is critical to setting realistic expectations. "Board members love AI, for many good reasons. But it's not a Magic 8 Ball," she says. "You can't shake it and get an answer. We continue to educate our board on the capabilities and skills required to realize the benefits."

> **Board members love AI. But it's not a Magic 8 Ball. You can't shake it and get an answer. We continue to educate our board on the capabilities and skills required to realize the benefits."**
>
> **— MICHELINE CASEY**
> CDO, Siemens Energy

The data team at ExxonMobil has focused on the best way to educate company leaders within the organization. Senior leaders aren't typically available to attend a three-day AI workshop, yet they're responsible for big investment decisions. The goal is to drive a mindset in which company leaders ask the right questions about data as a resource: Who else could use that data? Is there another piece of data we should collect? What is the quality required?

Sometimes you might only have 15 minutes to deliver the message. Is that in the classroom or through videos? To reach a busy audience, a multimedia strategy with diverse content, delivered on repeat, is often required. This is not a one-and-done proposition. Some teams have brought in external consultants, others deliver content with internal resources, and others use a mix of both.

The bottom line, says ExxonMobil's Andrew Curry: "If you can't get your leaders to talk about the value of data and understand, you're in trouble."

Follow training with opportunity to ensure retention. A recent study revealed that on average a data scientist will **switch employers after just 1.7 years**. That's not a long time, particularly if the company has invested in training. Retaining key talent must be a priority for AI transformation. At a multinational networking and telecommunications company, thousands of employees have completed their data literacy training, but the initiative doesn't stop there. Yes, one of their key performance metrics tracks how many employees go through literacy readiness training. But as their CDO points out, "It's not just about how many complete the training; it's about giving them the opportunities to apply what they've learned in their work. In our experience, if you train people but don't give them the opportunity to apply what they've learned, they leave the organization."

It seems that others have had a similar experience. Data leaders must continuously review attrition rates and evaluate their ability to retain top performers. The advice: Don't stop at training; ensure that you are retaining those critical to driving your transformation by giving them innovative, mission-critical work.

Ensure insights turn into actions and deliver business value. Ultimately, data is worth nothing unless it is used, and that means only if actions are taken that improve the business. "I'm a big believer that data by itself doesn't produce value," says ExxonMobil's Andrew Curry. "The business must act and do something different with the data."

That's why executive education is so important, and measuring the business impact is critical to the success of an effective AI practice. It's not about theoretical formulas, but a concrete methodology of measuring the impact of a particular data product in a specific use case or business initiative. Knowing the value delivered helps determine how much to invest and the eventual ROI. Perhaps with these best practices, we will finally be able to answer the perennial question, "How much is my data worth?" — but only if the data is used, the insight acted on and the impact measured.

In short, follow Curry's advice: "Don't measure based on insights delivered, but on business actions taken."

> **"Don't measure based on insights delivered but on business actions taken."**
>
> **— ANDREW CURRY**
> Central Data Office Manager, ExxonMobil

TOP TRAVEL TIPS

The emergence of gen AI has led data leaders down a bumpy road. Pete Williams, Director of Data at Penguin House UK, worries that "gen AI has corrupted the conversation." Penguin Random House, one of the world's largest book publishers, considers unauthorized use of content to train AI models to infringe copyright. But that's only one aspect of the story. At Penguin, Williams tries to navigate the obstacles by helping the organization differentiate between the risks and the opportunities. And the key to focusing on the opportunities is to get the foundation in place for well-governed use of the data.

Like all the data executives interviewed, Williams stressed the need to get the foundation right. "Our well-structured data platform puts us in a good place," he says. "Doing the governance up front allows us to better take advantage of the opportunities."

Though adopting generative AI presents many new challenges, the process of digital transformation is not exactly new to most data leaders. Rather than view the road to advanced AI as filled with obstacles, look at it as an exciting journey and prepare accordingly. Here are a few tips for getting into a good starting place and enjoying the ride:

1. **Pack what you'll need (hint: it's data):** It's important to ensure access to more of your data. Make sure you build the foundation you need for a steady supply of quality data. Break down internal silos to improve access. Transform previously inaccessible data to enable use. Imagine being an electric car owner and starting the journey with your battery only 10% charged. IDC estimates that 90% of data in the enterprise today is unstructured.
2. **Understand the rules of the road:** Start by ensuring that your data is well governed. Check under the hood to make sure it's all in good working order. But also keep in mind the underlying goals. Keep asking, "Where are we going?" — or in this case, "What is the problem we're trying to solve? How do we define success? How do we measure success? Did we make progress?" Those common guardrails need to be applied here, just as with other technology journeys. This is likely not the f irst you've undertaken.
3. **Plan the route but embrace the unexpected:** The prioritization process sets an itinerary for the journey, but interesting detours might be worth the investment. Ideation and experimentation keeps you open to pleasant surprises along the way.
4. **Bring a companion along for the ride:** A road trip is certainly more fun if shared with others. Similarly, the AI journey should not be a solo pursuit. Enlisting partners, customers and data partners into your AI ecosystem will help deliver richer insights and mitigate risks of hallucination and bias.
5. **Learn what you'll see along the way:** Knowing what you, your colleagues and your peers are getting yourselves into helps smooth the ride. Start evangelizing early. Then maintain continuous education across all roles. Develop a comprehensive communications plan — a travelog — to share where you've been and what you learned.
6. **Keep a scrapbook to remember the highlights (and learn from the obstacles):** Document the journey by capturing the details from which data was used, to the models employed, the challenges faced along the way and the outcomes achieved. Just like a travel log, success stories encourage others to re-use the data and models and increase the ROI.

When done correctly, you'll be planning your next trip before you finish the first. It is an ongoing journey. Bon voyage!

JENNIFER BELISSENT

As Principal Data Strategist, Jennifer advises Snowflake customers on data and AI strategy and best practices in building world-class data organizations. Previously, she spent over a decade as a Forrester Analyst. She has held management positions in tech sales and marketing, designed urban policy programs, taught math as a Peace Corps volunteer, and earned a Ph.D. in political science from Stanford.

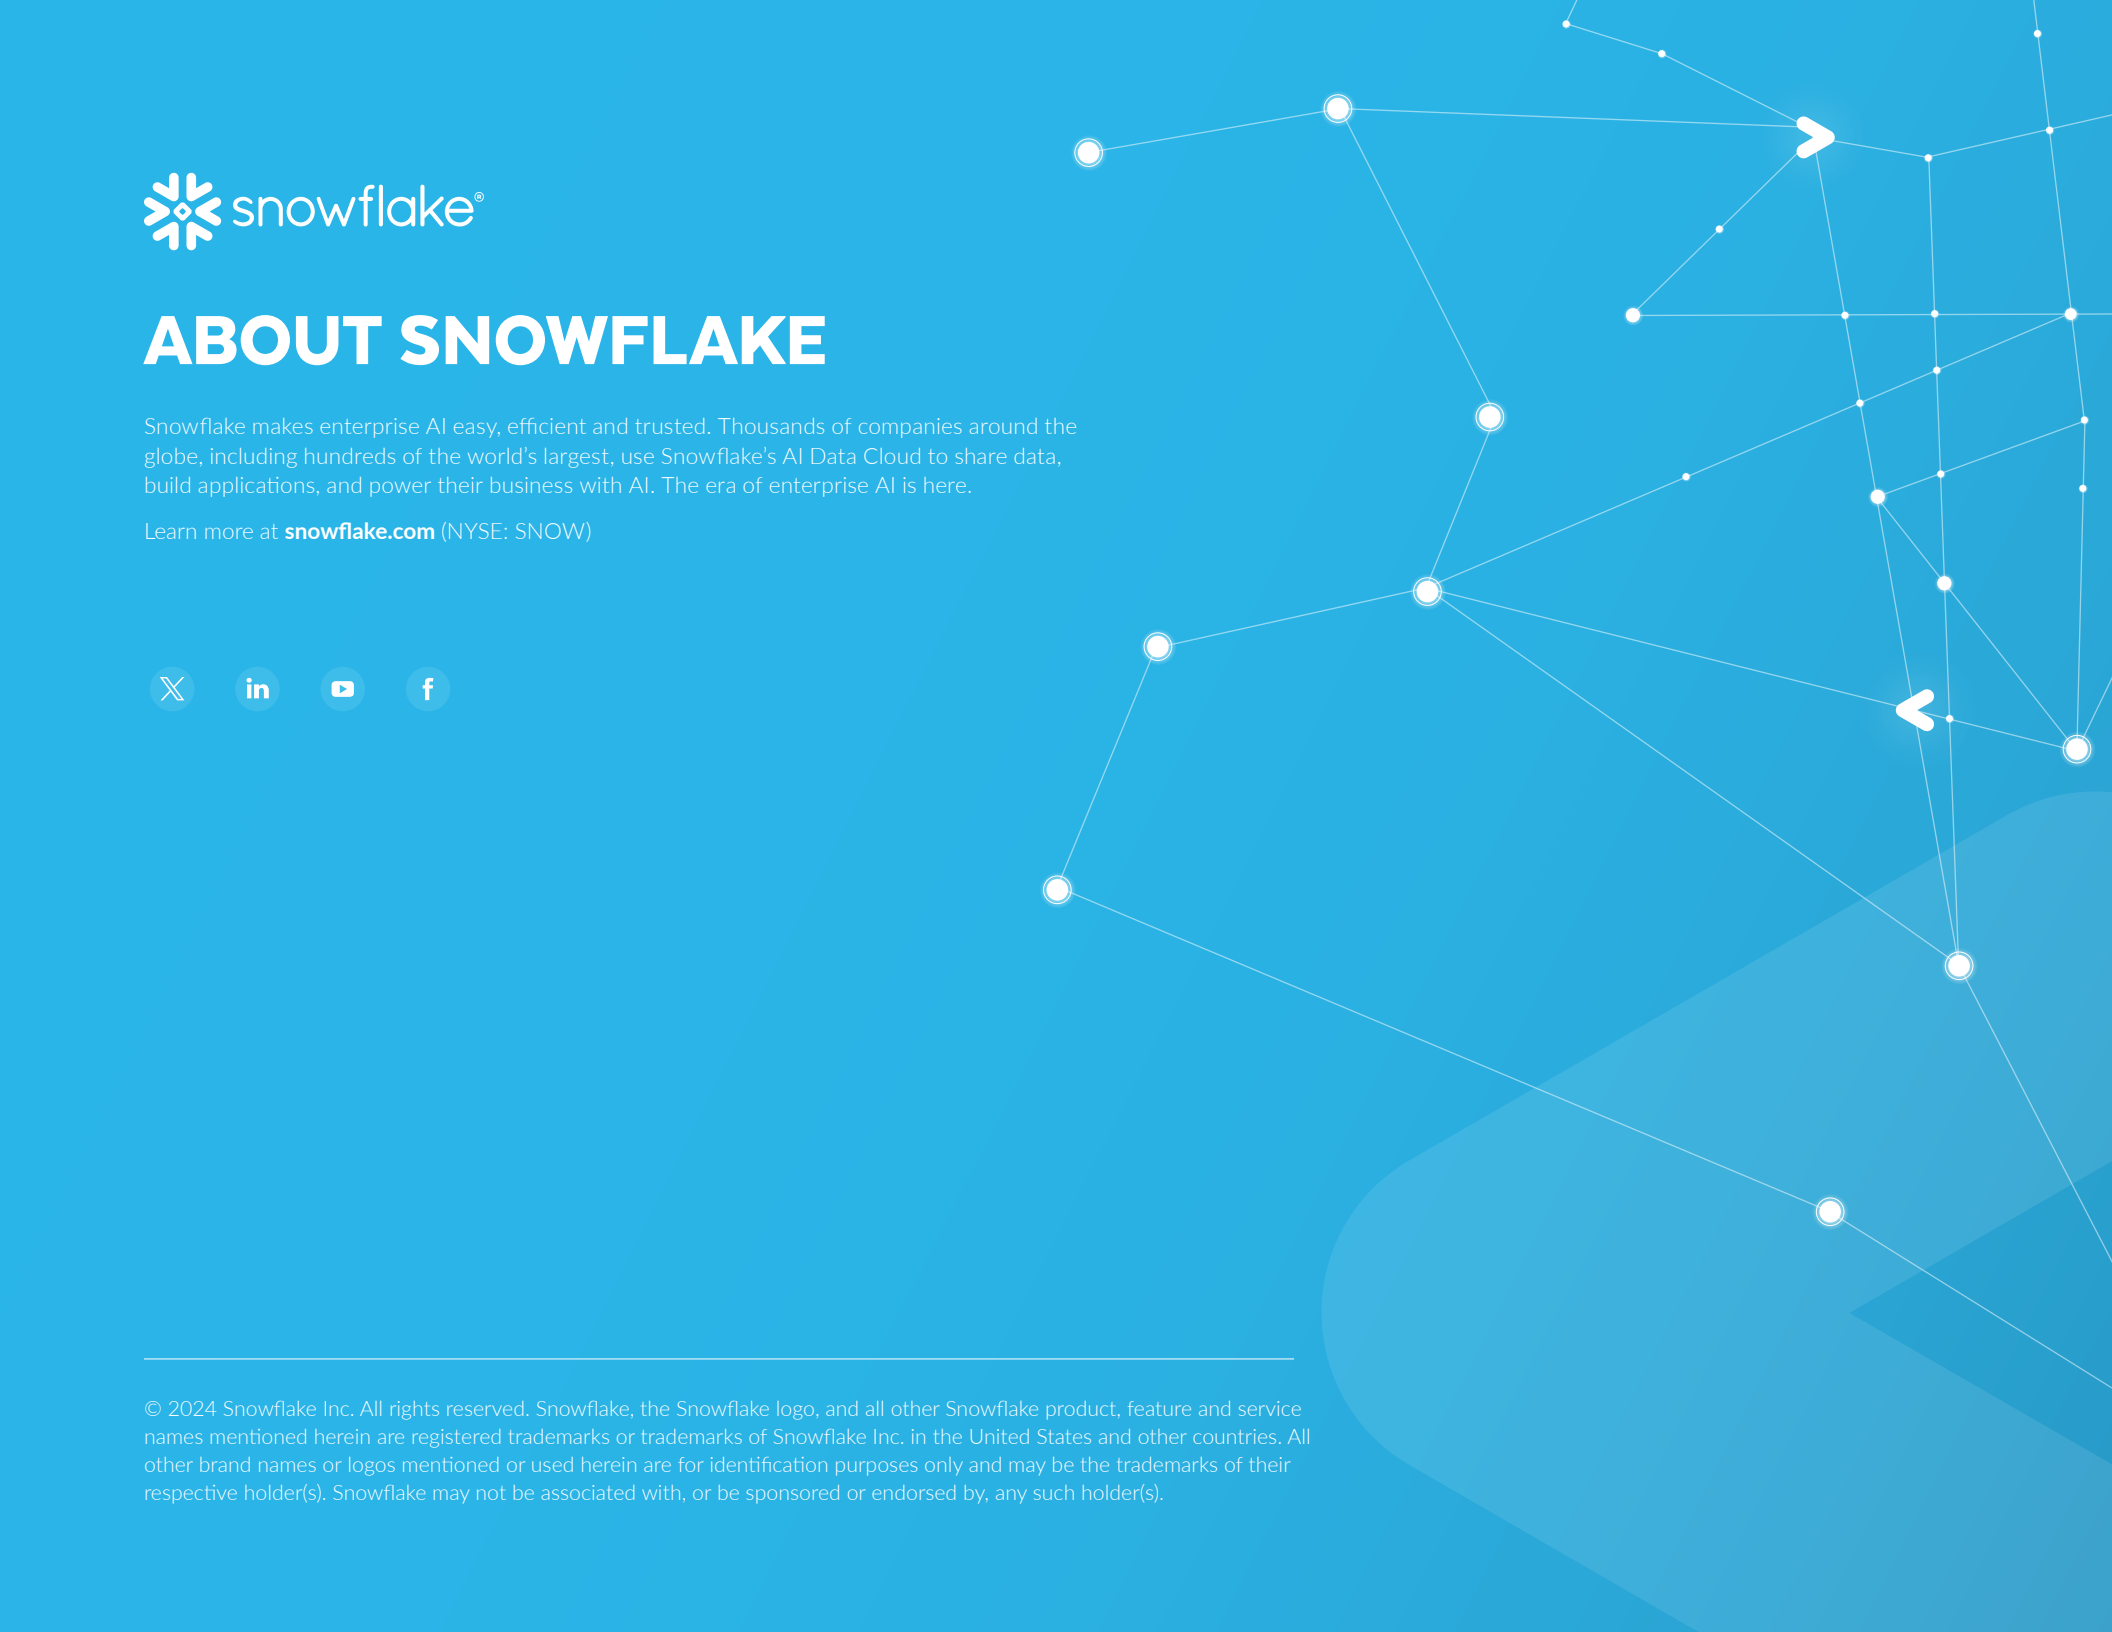

ABOUT SNOWFLAKE

Snowflake makes enterprise AI easy, efficient and trusted. Thousands of 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)