Co-sponsored by: zoominfo

BEST PRACTICES REPORT

Q1 2025

Creating an AI-Ready Organization



By Fern Halper, Ph.D.

tdwi TRANSFORMING DATA WITH INTELLIGENCE™

Creating an AI-Ready Organization

By Fern Halper, Ph.D.

Table of Contents

Research Methodology and Demographics 3
Executive Summary 5
Introduction: The Journey to Artificial Intelligence 6
The Current State of AI 8
- Traditional AI Use Cases. 8
- Generative and Agentic AI Use Cases. 9
- AI Is Providing Value. 12

Data, Tools, and Platforms 13
- AI Applications 16

Challenges with AI 17
- In Their Own Words: Overcoming the Challenges 18
- Intelligent Tools and Platforms. 21

Organizing to Execute 21
- New Leadership Structures for AI 21
- Important New Roles for AI 22
- Organizational Models 24
- AI Governance 25
- Cultural Considerations 26

Considerations for AI Success 27
- Data Foundations to Support AI Are Critical 27
- CDO and C-Suite Presence Are Key 27
- No One Organizational Model Wins 28
- Plan to Upskill and Hire New Talent 28
- Literacy Enablement Teams Are Also Important 29
- Measure the Impact 30
- AI Governance Will Become Critical 30
- Make Use of AI-Infused Tools and Platforms. 30

Recommendations 31
Research Co-Sponsor: ZoomInfo 33

© 2025 by TDWI, a division of 1105 Media, Inc. All rights reserved. Reproductions in whole or in part are prohibited except by written permission. Email requests or feedback to info@tdwi.org.

Product and company names mentioned herein may be trademarks and/or registered trademarks of their respective companies. Inclusion of a vendor, product, or service in TDWI research does not constitute an endorsement by TDWI or its management. Sponsorship of a publication should not be construed as an endorsement of the sponsor organization or validation of its claims.

This report is based on independent research and represents TDWI's findings; reader experience may differ. The information contained in this report was obtained from sources believed to be reliable at the time of publication. Features and specifications can and do change frequently; readers are encouraged to visit vendor websites for updated information. TDWI shall not be liable for any omissions or errors in the information in this report.

About the Author



FERN HALPER, Ph.D., is vice president and senior director of TDWI Research for advanced analytics. She is well known in the analytics community, having been published hundreds of times on data mining and information technology over the past 20 years. Halper is also coauthor of several Dummies books on cloud computing and big data. She focuses on advanced analytics, including predictive analytics, machine learning, AI, cognitive computing, and big data analytics approaches. She has been a partner at industry analyst firm Hurwitz & Associates and a lead data analyst for Bell Labs. She has taught at both Colgate University and Bentley University. Her Ph.D. is from Texas A&M University.

You can reach her by email (fhalper@tdwi.org) and on LinkedIn (linkedin.com/in/fbhalper).

About TDWI Research

TDWI Research provides industry-leading research and advice for data and analytics professionals worldwide. TDWI Research focuses on modern data management, analytics, and data science approaches and teams up with industry thought leaders and practitioners to deliver both broad and deep understanding of business and technical challenges surrounding the deployment and use of data and analytics. TDWI Research offers in-depth research reports, commentary, assessments, inquiry services, and topical conferences as well as strategic planning services to user and vendor organizations.

About the TDWI Best Practices Reports Series

This series is designed to educate technical and business professionals about new business intelligence technologies, concepts, or approaches that address a significant problem or issue. Research for the reports is conducted via interviews with industry experts and leading-edge user companies, and it is supplemented by surveys of business intelligence professionals.

To support the program, TDWI seeks vendors that collectively wish to evangelize a new approach to solving business intelligence problems or an emerging technology discipline. By banding together, sponsors can validate a new market niche and educate organizations about alternative solutions to critical business intelligence issues. To suggest a topic that meets these requirements, please contact Fern Halper (fhalper@tdwi.org).

Acknowledgments

TDWI would like to thank many people who contributed to this report. First, we appreciate the many users who responded to our survey, especially those who agreed to our requests for phone interviews. Second, our report sponsors, who diligently reviewed outlines, survey questions, and report drafts. Finally, we would like to recognize TDWI's production team: Lindsay Stares and Rod Gosser.

Sponsors

Cube, Snowflake, SAP, and ZoomInfo sponsored this report.

Research Methodology and Demographics

Report purpose. As AI continues to transform industries, businesses must develop the right organizational models, build effective teams, and implement the necessary data and analytics technologies to support AI initiatives. While previous TDWI research indicates that many organizations are taking steps toward AI readiness, significant gaps in strategy, skills development, and governance remain. This report will examine these challenges and outline the key factors that drive AI success, offering a road map for organizations to address the critical elements necessary for a successful AI deployment.

Roles

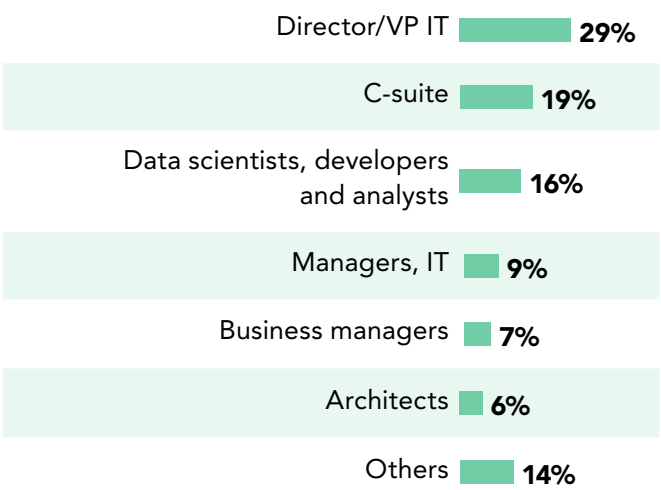

Survey methodology. In November 2024, TDWI sent an invitation via email to the analytics and data professionals in our database, asking them to complete an online survey. Over 300 respondents answered the survey; 197 met quality standards for completeness and validity. This group is primarily used for analysis.

Research methods. In addition to the survey, TDWI conducted interviews with technical users, business sponsors, and data experts. TDWI also received briefings from vendors that offer products and services related to these technologies. Data was analyzed using a variety of methods including statistical techniques such as chi-squared analysis and other tests for significance.

Survey demographics. Respondents act in a variety of roles. These include corporate executives (19%), VPs/directors of IT (29%), data scientists and analysts (16%), architects (6%), and others.

Respondents came from a range of industries including information technology (34%), financial services (15%), healthcare/life sciences (11%), manufacturing (10%), and others. Most survey respondents reside in the U.S. (77%). Respondents come from enterprises of all sizes.

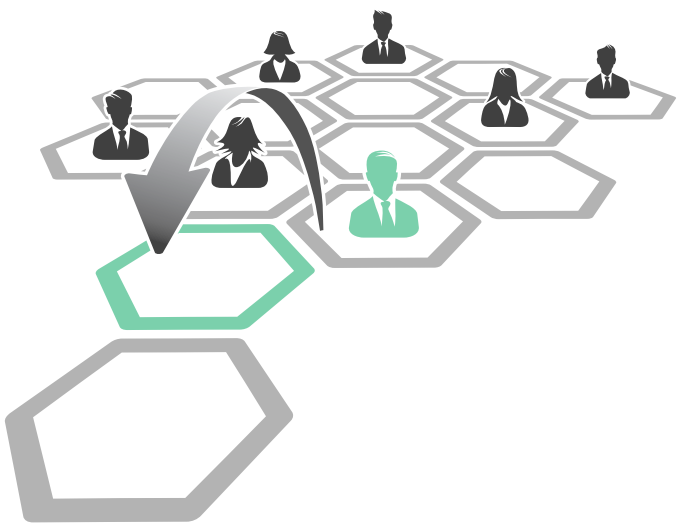

Industry

Information technology	34%
Financial services	15%
Healthcare/life sciences	11%
Manufacturing (non-computer)	10%
Public sector	7%
Retail/CPG	7%
Construction/engineering	5%
Insurance	4%
Other	7%

("Other" consists of multiple industries, each represented by less than 4% of respondents.)

Company Size by Revenue

Less than $100 million	22%
$100 million to $499 million	23%
$500 million to $999 million	12%
$1 billion to $4.9 billion	14%
$5 billion to $9.9 billion	9%
$10 billion or higher	11%
Don't know	9%

Geography

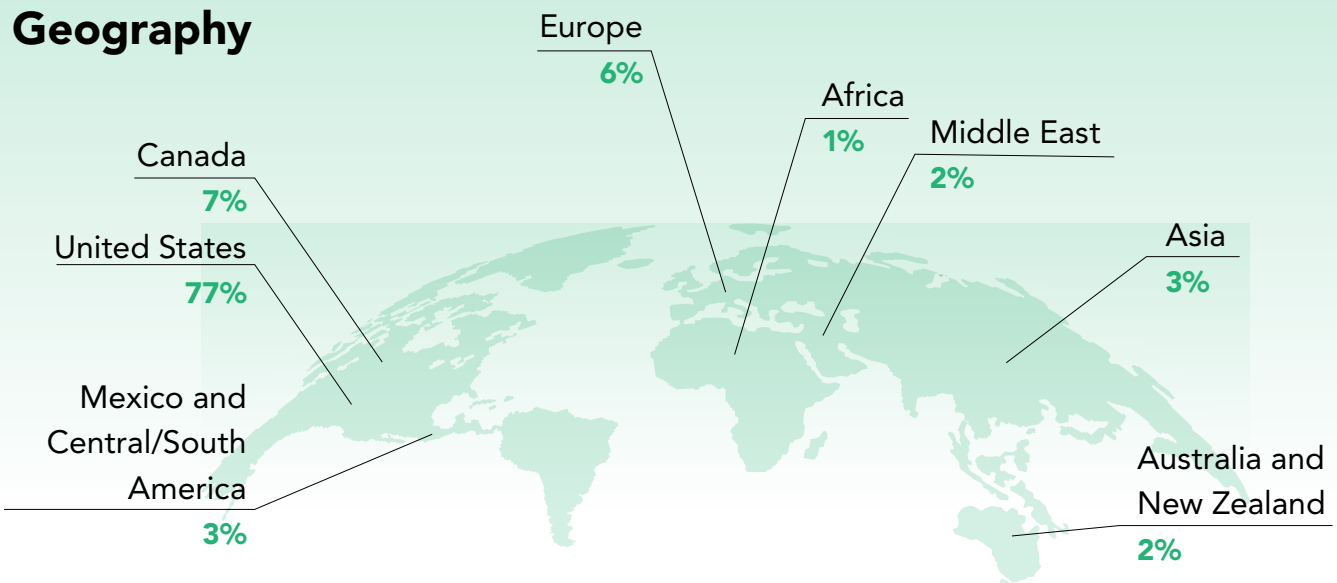

Based on 197 respondents.

Executive Summary

Artificial intelligence (AI) has become a driving force behind innovation, capturing the attention of organizations eager to unlock its transformative potential. With advancements in technologies such as generative AI, businesses are prioritizing AI to help with decision-making, improve efficiency, and gain competitive advantages. While AI represents transformative potential, realizing its benefits requires a holistic approach that combines technological capabilities, skilled teams, a supportive organizational culture, and the right leadership.

This TDWI Best Practices Report examines the current state of AI including the challenges organizations are facing with AI, how they are organizing to execute, and considerations for success. The top challenges include poor data quality (28%), lack of AI skills (25%), and insufficient AI literacy (24%). Addressing these challenges requires investments in tools, comprehensive training programs, and cross-functional collaboration.

A unique feature of this report is its examination of the characteristics of companies that have achieved impact with AI—as defined by examining the extent to which AI has contributed to key business goals such as revenue growth, cost reduction, or customer satisfaction. The report looks at three categories of respondents: those that have achieved a high impact with AI, those with a moderate impact, and those with a low impact. Impact also corresponds with maturity in analytics and utilizing AI, reinforcing what TDWI has seen before—success begets success in a virtuous circle.

Key findings include:

- **Technology foundations:** High-impact organizations are more likely to manage diverse data types (including text, images, and IoT data) and use modern platforms such as cloud data lakes and lakehouses. Data fabrics are emerging as a solution for hybrid environments as well, although no single approach (physical vs. logical) dominates.

- **Skills and roles:** Organizations achieving higher AI maturity and impact have dedicated roles such as MLOps engineers and AI application developers to operationalize and scale AI efforts.

- **AI literacy programs:** Data literacy and AI literacy programs are also critical, with 54% of high-impact organizations offering these programs.

- **Leadership:** Having a chief data officer (CDO) or chief analytics officer (CAO) in the C-suite significantly improves AI outcomes.

- **AI governance teams:** Organizations with dedicated AI governance teams are better equipped to manage risks such as bias, security, and explainability. AI governance is separate from data governance.

- **Measuring impact:** Mature organizations are significantly more likely to measure AI's impact using ROI and metrics tied to business outcomes. Respondents reported efficiency gains of approximately 25% and cost savings of 20% for specific AI use cases.

Introduction: The Journey to Artificial Intelligence

There is no doubt that the market is excited about artificial intelligence (AI); it is the buzzword du jour. Yet AI is certainly not new. The idea that machines could act "intelligently" has been around since the ancient Greeks. Back in the 1950s, when John McCarthy of Dartmouth College coined the term, he described it as, "Making a machine behave in ways that would be called intelligent if a human were so behaving."[1]

At TDWI, we think of AI as an umbrella term encompassing myriad methodologies and techniques. It leverages advances in mathematics, computer science, computational linguistics, cognitive sciences, and robotics, among others. Popular AI technologies include machine learning, natural language processing, and neural networks, which collectively drive modern AI systems (see next page).

> *In a 2024 TDWI survey, generative AI ranked higher than machine learning as a priority for analytics.*

Today of course, the buzz is all around generative AI and agentic AI, which also fall under the AI umbrella. Generative AI has catapulted AI once again into the spotlight, and organizations appear to be on board. In a 2024 TDWI survey, for instance, generative AI ranked higher than machine learning as a priority for analytics.[2] In fact, it ranked in the top four! In a current 2025 survey, generative AI is a top investment priority for analytics.[3] There is no question that generative AI represents a leap forward in AI capabilities.

In TDWI surveys, we typically see that over 90% of respondents indicate they have plans for generative AI. However, organizations need to be thoughtful as they take steps toward AI and generative AI. We've been tracking AI adoption for more than a decade.

Back in 2017, about 40% of respondents to our surveys were utilizing AI technologies such as machine learning and natural language processing.[4] Another 25% claimed that they would be starting to use these technologies in the next year. If users were able to stick to their plans, close to three-quarters of organizations we survey would be making use of these technologies already.

The plans haven't matched the reality. Typically, in surveys, we see about 25% of respondents making use of AI technologies and another 25% starting to use them. The rest are still using only self-service BI tools (albeit perhaps with AI infused into the product). There are numerous factors that come into play in terms of succeeding with AI.

This TDWI Best Practices Report examines what it takes to create an AI-ready organization: a combination of technology, skills, processes, and leadership/culture. In this report, we examine different groups of respondents including those who are mature versus not mature with AI technology as well as those who have measured high, moderate, and low impact with AI. These groups are used to help identify best practices for AI.

[1] See: https://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html.
[2] Unpublished 2024 TDWI Data and Analytics Survey.
[3] Unpublished 2025 TDWI Data and Analytics Survey.
[4] See 2017 *TDWI Best Practices Report: Advanced Analytics: Moving toward AI, Machine Learning, and Natural Language Processing*, online at tdwi.org/bpreports.
[5] Unpublished 2024 TDWI Data and Analytics Survey.

AI Technologies

The actual definition of AI has been debated since the 1950s. Some definitions include the ability of a computational device to learn, reason, and interact. Other narrower, task-oriented definitions include cognitive tasks that humans do easily, such as image recognition or natural language understanding. Some researchers talk about general versus applied AI. Others use the term weak or narrow AI for a system that performs specific tasks versus strong or true AI, where machine intelligence is equivalent to human intelligence. That is why we prefer to use AI as an umbrella term that includes a number of technologies, including:

Machine Learning. This method of data analysis originated in the field of computational science. In machine learning, systems learn from data to identify patterns with minimal human intervention. The computer learns from examples, typically using either supervised or unsupervised approaches (although different approaches also exist). In supervised learning, the system is given a target (also known as an output or label) of interest. The system is trained on these outcomes using various attributes (also called features). In unsupervised learning, there are no outcomes specified.

Machine learning includes its subset, **deep learning**, which uses algorithms to learn functions that can classify complex patterns such as images. Some deep learning algorithms, such as artificial **neural networks**, have input nodes and a number of hidden layers that act as a kind of "black box" to model one or more output nodes.

Natural Language Processing (NLP). NLP involves analyzing, understanding, and generating responses ultimately to enable interfacing with systems using human rather than computer languages.

Generative AI. These AI systems are designed to create new outputs such as images, music, text, or other forms of output based on the input data that they were trained on. Generative AI builds on NLP by using advanced models, such as transformer models (another kind of deep learning), to create coherent and contextually relevant text, such as articles, dialogues, or code. This connection allows generative AI to produce human-like language outputs based on patterns learned from vast amounts of text data.

Agentic AI. *Agentic* is a nod towards *agency*, meaning that these are AI systems designed to act autonomously or on their own and proactively toward specific goals or objectives, often without needing continuous direct human oversight.

The Current State of AI

TDWI has been tracking AI readiness for years. For instance, in 2024, we introduced the TDWI AI Readiness Assessment, which measured five key dimensions of enterprise fitness and readiness to use AI operationally: organizational readiness, data readiness, skills readiness, operational readiness, and governance readiness.[6] The overall median score for all enterprises responding to the assessment in 2024 was 62 out of 100, which put respondents on average at the Standardizing stage of readiness.[7]

During the Standardizing stage, organizations are developing an AI strategy, often supported by preliminary use cases and proofs of concept, with leadership recognizing AI's potential impact but hesitating to fully commit resources like technology, talent, and training. They begin collecting diverse data types beyond structured data—such as social media, machine-generated, geospatial, and real-time data—while working to make this data accessible using data warehouses or lakes on-premises or in the cloud.

Predictive modeling and machine learning efforts typically start at this stage, supported by data scientists and a growing recognition of the need for data engineers and operations roles to operationalize AI. Governance and architecture for supporting data growth and AI initiatives are in early consideration. In the 2024 AI readiness assessment, 24% of respondents were making use of AI technologies; another 25% were starting to make use of predictive analytics and machine learning.

[6] See: https://tdwi.org/pages/assessments/adv-all-tdwi-ai-readiness-assessment.aspx.
[7] See the *State of AI Readiness report*, online at https://tdwi.org/Research/2024/07/ADV-ALL-2024-State-of-AI-Readiness.aspx.

In this survey, we asked the maturity question in a slightly different way; we asked how mature survey respondents *thought they were* specifically in terms of their use of AI technology. Forty-five percent said that they are mature or very mature (Figure 1). Very mature is defined as an organization that makes use of tools and technologies such as machine learning, AI, and generative AI. These models are embedded into decision processes. The organization also builds applications and monetizes them, with the right team to make this happen. Mature in this report refers to an organization that uses tools and technologies such as machine learning and other AI technologies to make decisions, and acts on their insights. In the mature organization, organizational culture supports data-driven decisions. The rest of the respondents to the survey were either somewhat mature (using self-service dashboards and analytics tools) or not mature.

Traditional AI Use Cases

Organizations are deploying or planning to deploy many kinds of AI use cases. These include marketing use cases such as recommendation engines, next best offer, and predicting customer churn. They include operational use cases such as fraud detection, predictive maintenance, scheduling, and supply chain optimization. They also include vertical use cases such as disease detection or crop management.

In the survey for this report, between 20-30% of respondents were deploying AI marketing and operational use cases now; another 40-50% were planning to implement them (not shown). In an open-ended question, we asked specifically about AI use cases. Examples of verbatim responses are shown in Table 1, grouped by type of use case.

Figure 1

Which statement best describes your organization's maturity with AI?

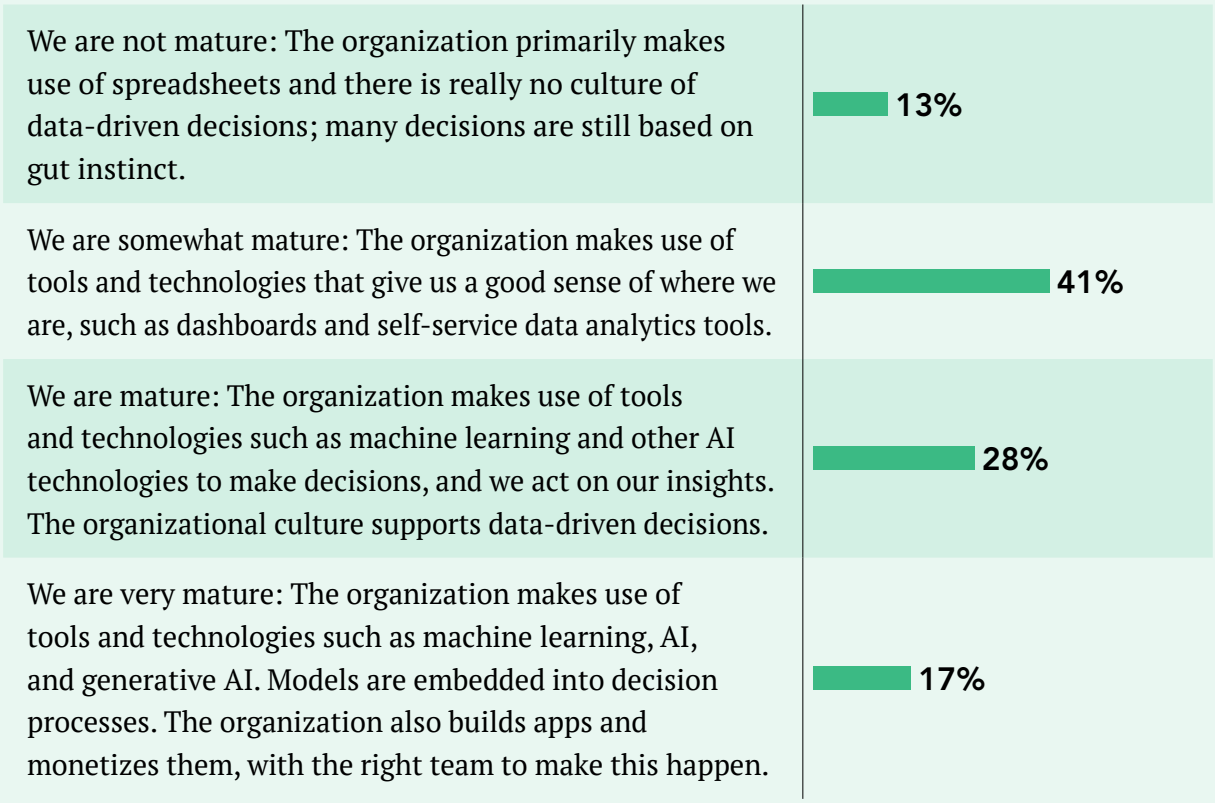

Based on 197 respondents.

These quotes illustrate diverse AI applications across industries, highlighting innovations in personalized customer engagement, operations optimization, fraud detection, predictive analytics, quality control, and healthcare. They provide some indication of how AI can help to transform business processes, improve outcomes, and drive strategic advancements across key functional areas.

Of course, many of these are the kinds of use cases that have been deployed for some time and from which companies are already generating value (see the section about AI generating value).

Generative and Agentic AI Use Cases

We also asked respondents whether they were making use of generative AI and even agentic AI. Close to 90% of respondents to this survey said that they were either deploying generative AI now, were experimenting with it, or were planning to use it (Figure 2). In other words, even those who hadn't yet implemented traditional AI techniques are using or planning to use generative AI. Twenty-six percent said that they had generative AI in production; we've typically

Area	Responses
Marketing, Customer Experience, and Personalization	• *"We are implementing a 'recommendation engine' that looks at life stage and customer engagement to recommend products or provide advice."* • *"Churn prevention, customer segmentation for targeting."* • *"Optimizing marketing campaigns for best interactive results."* • *"Ability to target leads effectively for conversion through fine-tuned campaigns."*
Operations and Process Improvement	• *"Supply chain optimization that would take into account locations, delivery, and storage capabilities."* • *"Identify process bottlenecks and recommend improvements."* • *"We use AI to detect fraud by monitoring outlying and duplicate requests compared to approved entries."* • *"We use it to improve our server security in detecting malware and fraudulent activities."* • *"We are exploring and trying AI for security and data protection."* • *"Predictive Fault Detection & Preventive Maintenance."*
Specific Vertical Use Cases	• *"Triaging and scheduling specialists and surgeons."* • *"...Using specific sensors drills approximately 4000 holes an hour. AI determines if a hole was drilled correctly or flagged as problematic."* • *"Crop yields forecasting demand of commodities from product demand."*

Table 1. Verbatim responses to an open-ended question about AI use cases.

been seeing a number around 10-15%. About a third were experimenting with it now and another third were planning to implement it in 2025.

In terms of agentic AI, adoption in this respondent group is limited but evolving. The majority of respondents indicated that they are not currently using agentic AI, with many citing early adoption stages, lack of maturity, or ongoing experimentation as reasons for this. However, a few respondents did mention examples of agentic AI in use, such as autonomous customer service bots, virtual assistants that handle complex tasks independently, and systems with limited autonomy for customer engagement and operational efficiency. Several respondents cited plans or interest in exploring agentic AI in the future, while others highlighted barriers including security concerns and the need for clear use cases.

Respondents *did* cite numerous use cases for generative AI. Many of these are what TDWI refers to as "phase 1" implementations. Some make use of company data (some do not), but not always the structured data that data and analytics professionals typically associate with analytics.

Chatbots are popular. In 2024 TDWI research, we saw that chatbots were the most popular use case for generative AI.[8] In the survey for this report, chatbots were also a top use case in open-ended

[8] 2024 Unpublished TDWI Data and Analytics Survey.

Figure 2

Is your organization currently implementing generative AI use cases?

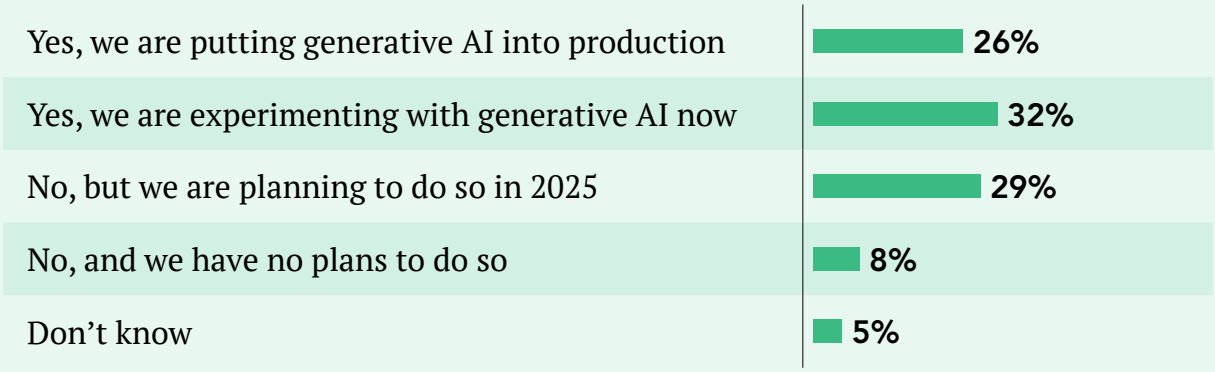

Based on 197 respondents.

responses. Respondents mentioned use cases such as "Customer self-service through natural language prompts; ability to automate customer service requests resolution," help desk and HR responses, and personalized interactions with customers via chatbots. Chatbots might be used to handle routine questions while customer service agents provide help for more complex questions. For instance, one respondent noted, "Introduce a hybrid support model where chatbots handle initial queries, but seamlessly escalate more complex issues to human agents. The chatbot can be trained to detect when customers are frustrated or when inquiries require human intervention."

Some of these examples might use product information such as product specifications, pricing, and FAQs. For instance, one respondent mentioned an "Internal LLM for generic questions and policy answers based on company standards." Use cases might employ articles and solutions for common customer issues or even policy documents about warranties, returns, refunds, and terms of service. However, other use cases will require using traditional customer data as part of the bot experience. This might include previous customer conversations, complaints, and resolutions. It will include information about the customer such as their account information, products and services used, and spending information.

Code generation. Another popular use case cited by respondents to this survey was code generation. For instance, respondents described using generative AI to suggest code, to write code, and to generate functional specifications from legacy code. Generative AI for coding and programming assistance can potentially help accelerate development by generating code snippets, automating repetitive tasks, and helping debug errors. It can also potentially help teams quickly prototype solutions, improve code quality with suggestions based on best practices, and reduce barriers for less experienced programmers through natural language interfaces.

Content creation. Respondents also talked about creating content; this was a top use case mentioned in TDWI's 2024 research. This might include commentary creation. A person can provide prompts or input data that outline the topic, tone, and style of the commentary they need. They might input a summary of an event, data sets, or articles, and the AI can generate insights, opinions, or analyses aligned with the desired context. More sophisticated use cases (that use traditional data) might generate personalized sales customer interactions and marketing campaigns based on customer segments, past history, and other data.

Business process improvement. Respondents also provided examples of how generative AI might help improve business processes. These include document preparation, automating complex report creation based on extensive text analysis, and supply chain optimization. For instance, TDWI has seen use cases where a company might use generative AI to standardize résumés that they receive from applicants. Other use cases include using a natural language front end to help supply chain managers better understand inventory levels and other supply chain components. Some of these examples use customer data, others do not.

> *Those who are already putting generative AI models into production are either mature or very mature in analytics.*

Data analysis. Using generative AI as a front end for data analysis has been cited as a popular use case in TDWI surveys. While commercial vendors provide AI-infused products with natural language interfaces to help democratize analytics, we're hearing more about use cases that analyze unstructured data that may not be stored in a company's data management systems. For instance, a popular example of this is to use generative AI to analyze a company's written call center notes and trouble tickets in order to classify the kinds of issues customers are experiencing. This data might be stored in a content management system or an S3 Bucket. Companies are also using generative AI to analyze spreadsheet data, such as survey data or other operational data. In fact, there are generative AI plug-ins for this now. Respondents mentioned generating insight summaries on operational situations that may make use of internal or external data sources. We expect this area to continue to grow.

The good news here is that those who claim to have generative AI models in production are "very mature" meaning they've been using other AI technologies (74%). Likewise, organizations that are "mature" also show a moderate level of generative AI in production (29%). Organizations that are "somewhat mature" and "not mature" have lower levels of generative AI in production (14% and 0%, respectively, see Figure 3). This relationship is statistically significant.

AI Is Providing Value

We asked an open-ended question about how much AI has improved efficiency and cost savings. The results indicate that AI has significantly contributed to efficiency, although the extent of impact varies based on use cases and implementation maturity. On average, respondents using AI reported efficiency gains of about 25% and cost savings of 20% for certain

Figure 3

Relationship between overall AI maturity and whether generative AI is in production.

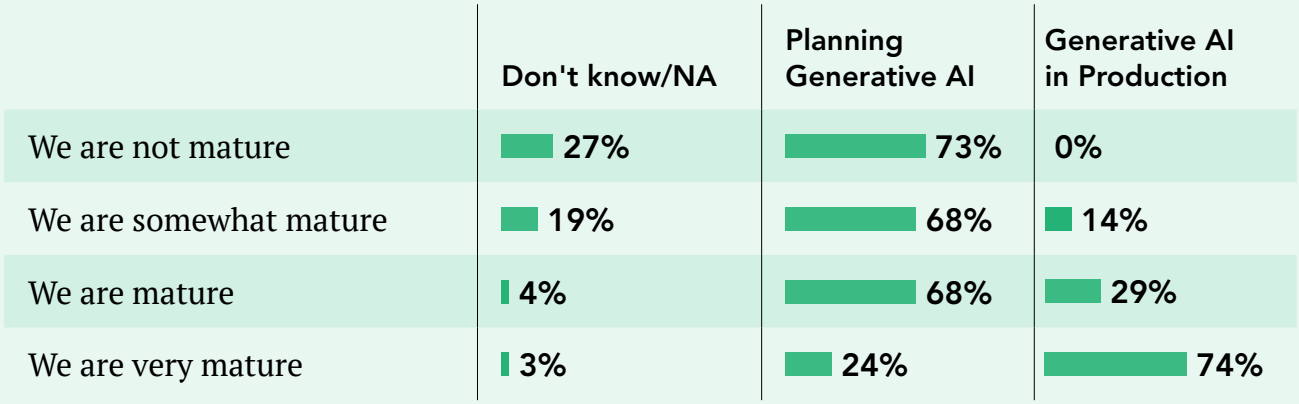

Based on 197 respondents. Rounding may cause some rows not to add to 100%.

use cases (not shown). For example, automation, such as AI-powered chatbots and workflow optimization tools, has reduced manual workloads and enhanced operational efficiency. Despite widespread acknowledgment of benefits, some organizations are in the early stages of adoption, with tangible results yet to be realized.

We also asked respondents to what extent AI has contributed to achieving key business goals (e.g., revenue growth, cost reduction, customer satisfaction). The responses were grouped into the three buckets mentioned previously: high impact, moderate impact, and low impact (32%, 28%, and 40% of respondents respectively). Not surprisingly, those respondents that are mature or very mature are more likely to say that AI has contributed to achieving key business goals (high impact). We will come back to these impact groups throughout this report.

Data, Tools, and Platforms

At TDWI, we see that a top reason organizations are evolving their data infrastructure is to support modern analytics such as AI and machine learning. Many organizations are moving to cloud platforms to manage the data diversity as well as the performance and scalability needed for AI workloads. Scalability will be needed for AI, where model training can be computationally intensive and often large diverse data volumes are needed to support model training and building. Scalability ensures the platform can support growing AI models and complex real-time applications without compromising speed. In this survey, we asked respondents what kind of data they were currently managing, analyzing, and utilizing (Figure 4). Slightly more than half of the respondents are primarily using structured data from databases as well as SaaS application

data. This includes CRM and ERP data. It also can include newer data types such as firmographic and intent data. Over 40% are already making use of structured data along with newer data types such as text data. About 20% of those are also using images, video, and audio.

Those organizations that are more mature in AI tend to analyze more diverse data, including newer data types such as text, image, video, or IoT data. Less mature organizations are more likely to focus predominantly on structured data, such as that found in traditional databases. This suggests that AI maturity correlates with an organization's ability to analyze a broader and more complex spectrum of data types, reflecting advanced technical capabilities and data strategies. Likewise, those respondents that collected and analyzed more diverse data are also more likely to achieve their key business goals with AI (both not shown). This is a virtuous circle that we have seen in other TDWI research. As companies achieve success with their analytics programs, the success tends to build on itself.

In terms of platform types to support AI, respondents are using a wide variety of platforms (Figure 5). These include on-premises (16%) and cloud only (46%) deployments. Cloud deployments include the growing popularity of the cloud data lakehouse. In a recent TDWI survey, the cloud data lakehouse was a top investment priority for 2025.[9] Respondents are also using hybrid deployments (31%) that utilize both on-premises and cloud platforms, unified using an approach such as a data fabric. The term *data fabric* has been used to describe a way to bring together disparate data in an intelligent fashion. The data fabric maps and connects relevant application data stores with metadata to describe data assets and their relationships. This approach can help unify data that might be on different platforms.

In other words, respondents are using both physical and logical approaches to managing

[9] Unpublished 2025 TDWI Data and Analytics Survey.

Figure 4

What kind of data does your organization manage, analyze, and utilize?

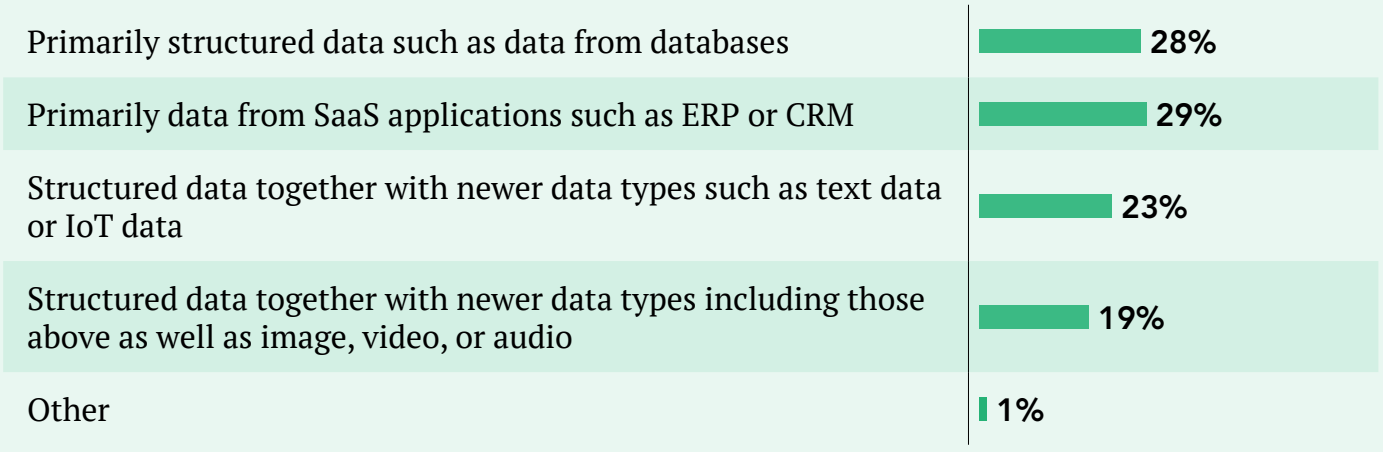

Based on 197 respondents.

Figure 5

Which statement best describes your organization's data management environment?

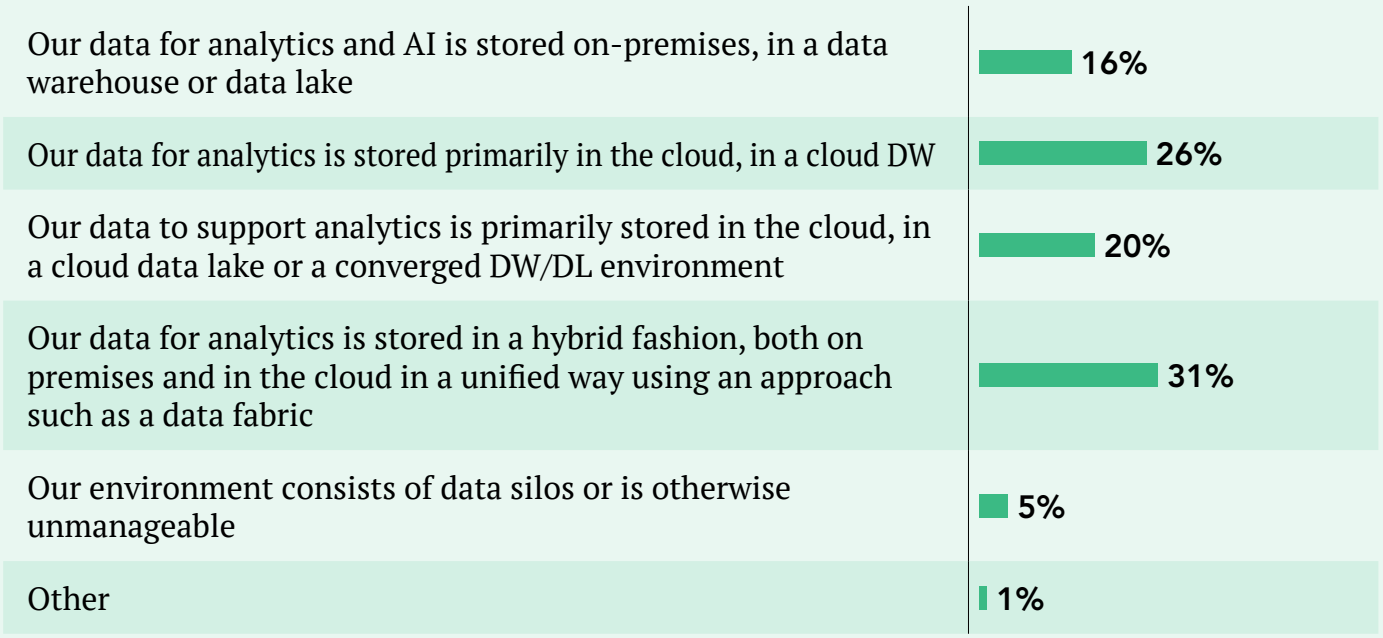

Based on 197 respondents.

their data. Less than 10% of respondents said that their data consists of silos or is otherwise unmanageable. While there are challenges, as we will see later in this report, this reflects the hard work that we've seen organizations putting into their data foundations as they move towards implementing AI and other innovations.

In terms of the relationship between maturity or impact and the kind of data foundation used, a chi-square test suggested no significant relationship between maturity and kind of platform, even when combining similar platform configurations together. Likewise, a similar test shows no significant difference between the extent of AI's business impact and the type of platform used for data management. In other words, the jury is still out as to whether physical or logical unification provides a better foundation for improved chances of AI success. We have seen this in previous TDWI research as well.

> *The next wave of AI is AI applications.*

Likewise, in a separate question, we found there appears to be no significant difference between those who are mature versus not mature with AI in terms of citing the data foundation as insufficient for AI in terms of scalability/performance. This most likely confirms that less mature organizations are not yet engaging in AI initiatives that require robust data foundations, although they may be building towards it. This could also

mean they haven't encountered scalability or performance challenges typical of more advanced AI deployments. Less mature organizations might also lack the expertise to recognize data foundation issues as an impediment.

AI Applications

AI is not simply used for insights; the next wave of AI is AI applications. These may be traditional applications such as recommendation engines or applications that predict when maintenance is needed, or they may be the beginning of generative AI applications such as chatbots that utilize company and customer data. They may even be agentic applications where individual agents accomplish different tasks that come together in an orchestrated manner to achieve a more complex goal, such as in an AI travel planner.

Building AI applications requires skills as well as new technology components. We will examine the skills later in this report. In terms of technologies, new generative AI applications will require new technologies. Adoption and plans for these are outlined in Figure 6. Some of these numbers seem slightly high, perhaps reflecting that many organizations are experimenting with these technologies. These include:

Foundation models. Foundation models are large-scale, pre-trained machine learning models that serve as a base for diverse downstream tasks. These models are trained on vast data sets—often internet-scale—and can be fine-tuned or adapted for specific applications. Different foundation models might be suited to different tasks. For instance, encoder models (e.g., BERT) may be best suited for tasks like sentence classification and summarization. Decoder models (e.g., GPT) may be specialized for text generation. Encoder-decoder models might be useful for language translation. Organizations typically start with off-the-shelf large language models (LLMs) such as GPT, BERT, Gemini, Claude, or Llama for generative AI applications. These models form the backbone of many AI systems, providing capabilities including text generation, summarization, and contextual understanding. Thirty-four percent claim to have these LLMs in place.

Vector embeddings and vector databases. Generative AI applications often require input data to be transformed into a format that the model can interpret—commonly vector embeddings. These are numerical representations of data that capture its semantic meaning. Efficient handling of these high-dimensional vectors may require a vector database designed to store, retrieve, and manage vector representations of data. They enable quick access to relevant information by comparing vector similarities. Twenty-seven percent claim to have this already in place.

Development frameworks. Frameworks provide the structure and tools necessary for building sophisticated generative AI applications. Popular frameworks include Langchain (a Python library for connecting LLMs with company-specific data and creating RAG models). Frameworks can help streamline development by offering reusable templates and workflows, reducing the complexity of building generative AI systems. Retrieval-augmented generation (RAG) frameworks provide real-time data access for generative AI systems. RAG works by retrieving data from external sources (often vectorized databases) and combining retrieved data with user prompts to generate accurate, natural language responses. In other words, RAG helps provide context. Twenty-five percent claim to be using RAG today.

Figure 6

AI technologies currently used and planned by respondents.

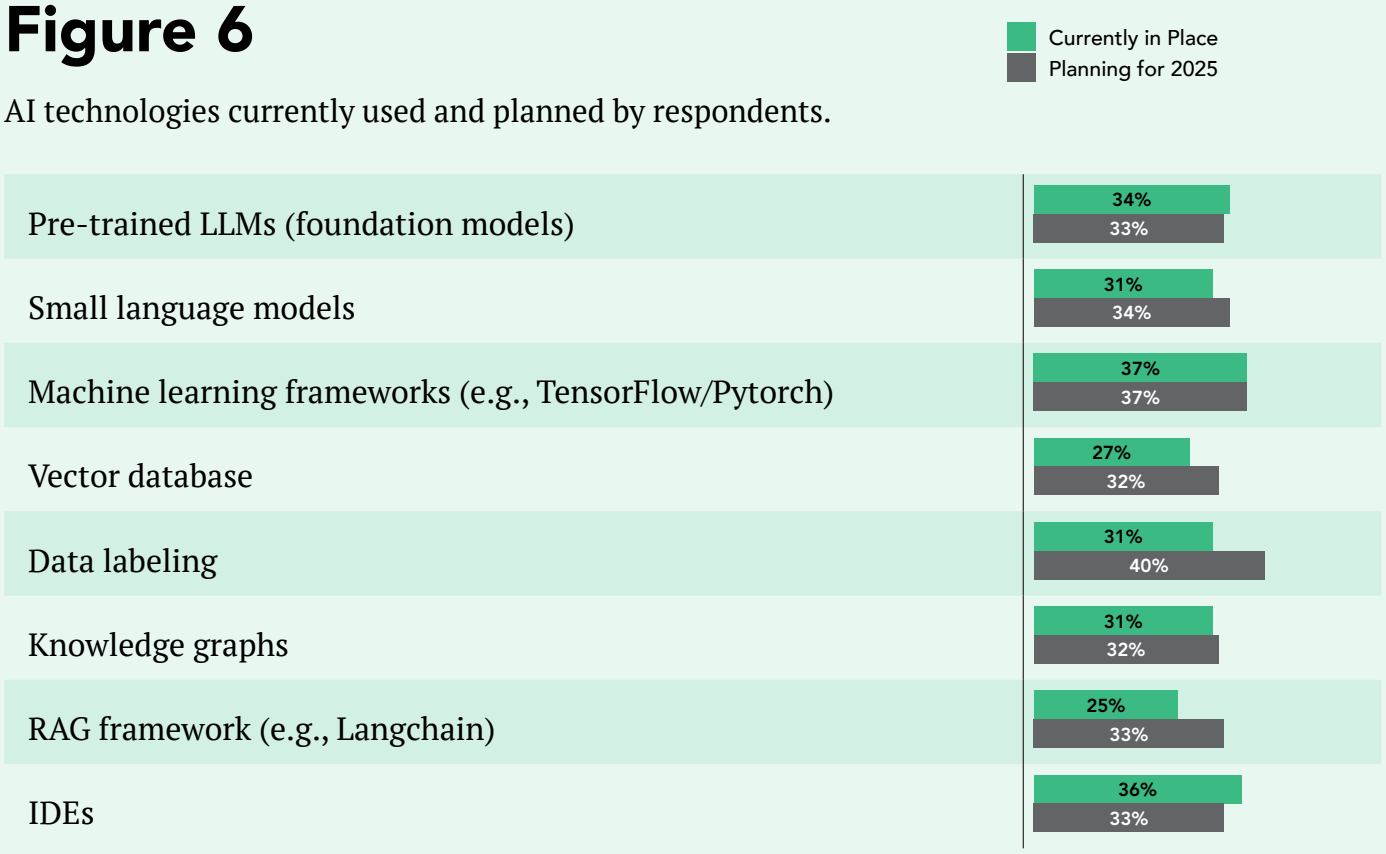

Based on 197 respondents.

Integrated development environments (IDEs). IDEs play an important role in the development of AI applications by providing comprehensive tools for writing, debugging, and managing code. Key IDEs for AI development include Jupyter notebooks (for model prototyping) and GitHub co- pilot (provides AI assisted features that enhance productivity). Thirty-six percent have IDEs in place currently.

Again, the good news is that organizations are not jumping the gun when it comes to building AI applications. An analysis of the data indicates that respondents whose organization's use of AI is "mature" or "very mature" are significantly more likely to have RAG frameworks currently in place compared to those who are "somewhat mature" or "not mature."

Challenges with AI

Even in analytically mature organizations, there can be challenges with AI. We asked respondents about their biggest problems with AI or getting AI off the ground (Figure 7). Respondents could select a maximum of three responses. As illustrated in the figure, challenges often have to do with age-old data quality problems as well as skills. The top three challenges were:

Poor data quality. Twenty-eight percent of respondents cited poor data quality as a problem they have had or expect to have with AI. In TDWI research, we often see that data quality is a challenge. In this survey, the problem was universal across those who were mature as well

as not mature. However, there may be different reasons as to why it is a problem. For instance, for those who are only somewhat mature or not mature in their analytics, they may lack robust processes for data governance, leading them to expect data quality problems when initiating AI projects. For those that are mature, AI uses diverse data types including non-traditional data. While organizations have been collecting this non-traditional data for some time, it is unclear whether they've been analyzing it. Generative AI is giving them the ability to analyze these new data types, and the definition of data quality may be different. On the other hand, the scale and diversity of data needed may magnify quality issues, even for advanced organizations.

> *AI challenges often have to do with age-old data quality problems as well as skills.*

AI skills inadequate. Another top challenge cited by respondents was that skills were insufficient for model building (25%). When asked about what skills are needed for AI in an open-ended question, responses included critical thinking, which highlights the importance of analytical problem-solving in AI initiatives. Other responses emphasized technical capabilities; this underscores the role of specialized expertise in machine learning algorithms and other algorithms used to build AI models, understanding how to put models into production, data engineering for AI models, and building AI applications (which we've already seen may require using new technologies).

AI literacy lacking. AI literacy refers to the knowledge and understanding required to effectively engage with, interpret, and apply AI technologies and concepts. It includes the ability to understand basic AI principles, evaluate AI's implications, and integrate AI tools into decision-making processes responsibly. Skills related to AI literacy include critical thinking, which enables individuals to analyze AI outputs, and ethical considerations, as well as having a basic understanding of how AI technologies work and how to use AI tools. AI literacy also includes understanding the importance of data quality and governance, as well as the ability to frame business problems that can benefit from AI. Twenty-four percent of respondents cite this as a challenge.

Other challenges were cited as well. Number four on the list was insufficient skills to build AI applications. Number five on the list was difficulty with AI data governance. Data governance can be particularly challenging with AI due to the sheer volume and diversity of data often required to train and operate AI models, making it difficult to ensure data quality, consistency, and compliance. Additionally, AI systems often require access to sensitive or proprietary information, raising concerns about privacy, security, and ethical use, which must be addressed within governance frameworks. The dynamic nature of AI, including model updates and evolving regulations, adds complexity, requiring organizations to continuously adapt their governance policies to stay compliant and operationally efficient. We will see later in the report that other aspects of governance, such as AI model governance, are also important.

In Their Own Words: Overcoming the Challenges

So, what are organizations doing to overcome some of these challenges?

Figure 7

What are the biggest challenges your organization has either had or is having getting AI off the ground? Please select a maximum of three responses.



Based on 197 respondents.

For the data quality challenge, many respondents noted generically that they were "working on getting better data quality." Based on previous TDWI research, data quality is a journey; it is not a one-time fix. Some respondents described utilizing the data quality tools in new data platforms and other third-party solutions they were implementing. TDWI has seen organizations starting to utilize automated and augmented data quality tools to address increasing data volumes that cannot be handled manually. In this survey, approximately a third claim to be making use of automated data cleansing tools (not shown). We've seen roughly the same percentages in other TDWI research.[10] Some (20%) are using tools that can identify sensitive information (not shown).

> *Organizations are implementing a range of programs to try to overcome AI challenges.*

Respondents mentioned that they were researching new algorithms to improve data quality, perhaps in the context of newer data types such as unstructured text data or for new applications. For instance, some newer tools on the market can help test for different aspects of AI, from bias and security to toxicity. In terms of addressing issues such as bias, many respondents were not yet thinking about it. For those who were, they mentioned a variety of approaches including implementing structured policies, practices, and frameworks specifically to identify and mitigate bias in AI models. However, a number of respondents mentioned handling bias on a case-by-case basis, addressing issues for each model individually without a standardized approach. This is most likely the case for many organizations.

In terms of data governance, respondents noted that they were trying to "create a better overall structure and ensure we involve department stakeholders." Some mentioned collaborating and building business partnerships. Others mentioned keeping up with the latest trends in data governance and AI governance and keeping up with new data governance methods. There was mention of consultants to assess the data environment. One person noted that their team was trying to "gain executive buy-in for applying resources to data governance." Many mentioned that they were still struggling. Governance is a journey. Many are working on standards to support AI development. At TDWI, we've seen organizations crafting their own governance frameworks using popular AI frameworks from organizations such as NIST, OECD, and big four consulting firms. We will talk more about AI governance later in this report.

For AI literacy, some respondents mentioned that they are conducting skills gap analyses. They are building targeted "tactical" training programs and investing in learning and development. Some of this involves collaborating with external providers (such as TDWI) to offer structured learning opportunities. One respondent mentioned that they are "leveraging what they have done in the past for executive training around data literacy and data storytelling." Others have noted that they are offering "office hours to train employees on the AI tools that have been deployed. Another noted, "We intend to better educate leaders to ensure they have more experience and confidence." Given that some generative AI tools are being used, regardless of literacy, one respondent noted that they were providing safe internal tools to all employees for experimentation.

[10] Unpublished 2024 TDWI Data and Analytics Survey.

In terms of AI skills, many respondents spoke about recruiting new talent or bringing in experienced consultants to help with some of their AI initiatives. Many understood the importance of trying to retain talent. For instance, one respondent noted, "Compete for top talent and create environments that attract and retain it." They discussed performing a skills gap analysis and providing training and developing an AI/ML pilot project to help flesh out the gaps and gain organizational knowledge.

We will talk more about practices for success later in this report.

Intelligent Tools and Platforms

Tools and platforms that are infused with AI can also help to overcome some challenges. These tools, such as modern business intelligence (BI) platforms, CRM and ERP systems, and B2B database platforms, are leveraging artificial intelligence and natural language processing (NLP) to enable more intuitive and efficient data interactions. By integrating NLP interfaces, users can query data conversationally in plain language, making insights accessible to non-technical and technical users. For instance, some tools can identify potential sales opportunities and prioritize leads based on their likelihood to convert. They can provide intent analysis, which identifies buying signals based on online behavior and trends. BI tools infused with AI can surface insights in the data. These kinds of tools help to democratize AI and help overcome some of the skills challenges mentioned earlier. However, AI literacy is still required to use these tools.

In previous TDWI surveys, we've seen that a majority of respondents are making use of SaaS platforms with infused AI capabilities.[11] In the survey for this report, we asked respondents about the challenges of some of these automated and AI tools. The top challenges were around the skills needed to make sure these tools are working properly (46%) as well as the staff to help to train users on these tools (26%). Funding to obtain these tools was also an issue among respondents (23%) (all not shown).

Organizing to Execute

It isn't just technology that is needed for AI success. AI will require new leadership structures, new roles and responsibilities, and new governance demands.

New Leadership Structures for AI

Traditionally, IT has led the data and analytics function. However, that has been changing. About five years ago, organizations started to hire chief data officers and chief analytics officers. The idea is that this person leads a team that drives *business value* in the organization, whether that be in data (CDO), analytics (CAO), or both (CDAO). Often, this person leads the data and analytics strategy, including coordinating the governance strategy. Their team may be part of a data office (sometimes called a data hub or a center of excellence). Titles for the position vary. For instance, the role may be called a general manager or managing director.

In the 2022 *TDWI Best Practices Report: Modernizing the Organization to Support Data*

[11] Unpublished 2025 TDWI Data and Analytics survey.

and Analytics, about a third of respondents had a CDO/CAO leading various efforts including data management, analytics, data governance, and data science.[12] That percentage has not changed significantly in the past few years. Likewise, in 2025, about 25% of respondents state that these programs are led by the VP/director of IT. In 2022, 40% of respondents cited that their data and analytics leaders were part of the C-suite. In this survey, that number has increased to 58% (not shown). This difference is statistically significant, indicating that significantly more respondents to this survey cited that their AI leaders are part of the C-suite.

CTOs and CIOs are more likely to be part of the C-suite than CDOs/CAOs; however, CDOs/CAOs are more likely to be part of the C-suite than VPs/directors of IT. These results are statistically significant. There may be a shift happening where CDOs/CAOs are becoming part of the C-suite. That is good news. In previous TDWI research, those respondents who measured value were more likely to have a CDO/CAO who is part of the C-suite. The same holds true today.

CDOs and CAOs are typically tasked with leveraging data as a core business asset, ensuring that analytics and AI efforts are tied directly to measurable business outcomes such as revenue growth, cost savings, and customer satisfaction. Their presence in the C-suite enables a stronger emphasis on integrating data-driven strategies across the organization, fostering accountability for ROI, and aligning initiatives with broader business goals. Consequently, organizations that prioritize and measure value may elevate these roles to the C-suite to ensure strategic oversight and alignment at the highest level. In the survey for this report, those who have a CDO/CAO in place for AI appear to be more satisfied with their leadership structure than those who have a VP/director of IT leading the effort.

The upshot: the right leadership reporting into the C-suite can make a difference for organizing for AI success. We will come back to this point later in the report.

Important New Roles for AI

In addition to the CDO and the CAO, there are other new roles that organizations are adopting to help with AI. In addition to the data scientist, these include the citizen data scientist, data engineers/DataOps, MLOps, AI application developers, and those who deal in AI governance (see Figure 8).

Data scientists and citizen data scientists. We've been talking about the data scientist and the citizen data scientist for years (if not the past decade or more). Data science is an interdisciplinary field that combines advanced statistics and computer engineering skills to acquire, prepare, and analyze information across a variety of sources. Data scientists have the technical skills to work with and analyze all kinds of data, from structured to unstructured. They build models, develop algorithms, and innovate on data.

Citizen data scientists, on the other hand, are typically business professionals who use analytics to create insights and solve problems without formal training in data science. They often rely on self-service analytics platforms, often augmented with AI or automation, to perform tasks such as building predictive models or analyzing trends. Now they will also utilize generative AI. Citizen

[12] See the 2022 *TDWI Best Practices Report: Modernizing the Organization to Support Data and Analytics*, online at https://tdwi.org/Research/2022/06/PPM-ALL-Best-Practices-Report-Modernizing-the-Organization-Support-Data-Analytics.aspx.

data scientists play a crucial role in democratizing analytics by bridging the gap between technical data experts and decision-makers. In other words, while data scientists typically possess advanced expertise in programming, statistics, and machine learning, citizen data scientists focus on applying prebuilt tools and easy-to-use systems to address specific business needs with less technical depth. In this survey, 58% of respondents had data scientists in place, and an additional 27% were planning to put them in place in 2025. Thirty-three percent have citizen data scientists/modern analysts in place now. An additional 41% were planning to put them in place in 2025.

DataOps is the process by which data is integrated, transformed, and prepared for deployment into business intelligence, reporting, and ad hoc analytics applications. It includes a wide range of functional components for handling and managing data throughout its life cycle. These include pipeline services for discovering, acquiring, extracting, transforming, profiling, cleansing, augmenting, staging, loading, replicating, delivering, indexing, searching, and protecting data. Data engineers are often part of the DataOps team. In this survey, the majority of respondents already had a DataOps team in place (not shown).

MLOps engineers. The MLOps team is responsible for deploying, monitoring, retraining, optimizing, and managing advanced analytics models such as machine learning or deep learning that power many traditional AI and other intelligent applications. Although the data scientist might build the model, MLOps engineers ensure that a model will work in production; they keep it working and update it when it gets stale. They are often highly technical and work with open-source software as well as commercial products. In this survey, 16% of respondents had an MLOps team in place. Fifty-seven percent plan to put one in place. We will see later in this report that this is a key role.

> *Forty percent of respondents claimed to have data product managers already in place.*

AI application developers. The AI application developer plays another important role in creating and deploying AI applications. While the data scientist focuses primarily on data analysis, model selection, and training, the AI application developer is responsible for integrating these models into an application environment. Additionally, they may be developing some of the more innovative applications that utilize AI, particularly generative and agentic AI applications. AI application developers can work closely with data scientists and may build and deliver AI models into scalable, user-facing applications such as chatbots, recommendation systems, and other AI use cases. Their responsibilities include integrating AI models into existing systems, optimizing performance, and ensuring deployment. In some cases (especially with newer model types), there is overlap with MLOps where AI application developers may also monitor and maintain AI systems, retrain models as needed, and ensure compliance with ethical and regulatory standards. In this survey, 42% of respondents already said that they have AI developers in place. This is in line with the group already performing AI.

Data product managers. TDWI defines data products as derived assets from data. These data products can be as simple as data sets or more complex, for example, including scores or industry-specific specialized data. They are shared within and sometimes outside of the organization. Each product might have a data product owner or

Figure 8

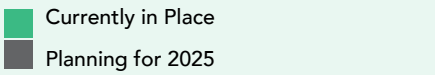

Important roles for AI currently in place and planned by respondents.



Based on 197 respondents.

manager, although the data engineer is important for ensuring that data for the data product is accurately collected, processed, and made accessible in a way that drives business value. Data products are inherently a helpful way to offer up data, because if you're offering up a data product, you want your users to be satisfied, and that means that the product must meet quality guidelines. Here, 40% of respondents said they already had a data product manager in place.

Organizational Models

We also asked respondents how they were organized to execute for analytics and AI. Recently, there have been three organizational models:

- In the *centralized* model, all data functions, ownership, and activities reside within a central office. That office is responsible for data strategy, the data platform, governance, policies, and standards. Thirty-seven percent of respondents are using this model (not shown).
- In the *decentralized/distributed* model, data activities are distributed into different business units. In other words, the business unit owns its data as well as what it wants to do with the data. The data and analytics functions sit in the business units. This was the most popular model, with 34% of respondents citing it (not shown).
- In the *hybrid/hub-and-spoke* model, a data office provides strategies, policies, processes, and standards. It supports the architectural design and delivery of an environment that can be used to support analytics, such as machine learning, so the data can be more easily used for value-generating activities. The data office may also be responsible for coordinating governance activities. The business units own their data but work closely with the data office to ensure that the data is high quality and available for analytics use. Analytics might be owned and occur in the business units, although the data office

may recommend certain tools that work with the infrastructure. In this survey, 29% of respondents are using this model (not shown).

In 2022, when we asked the same question, more respondents were utilizing a decentralized or distributed model (43%) and less (18%) were using a hybrid/hub-and-spoke model. These numbers are close to significantly different than current percentages, perhaps suggesting a move away from the decentralized model to the hub-and-spoke model. The choice between these models depends on the organization's strategic goals, culture, and operational requirements.

Larger companies (e.g., $10B or higher) tend to balance between all three models, while smaller companies ($100M to $499M) tend to lean towards centralized models, possibly due to simpler operations. Mid-sized ($500 million to $999 million) show a stronger preference for distributed models (not shown).

AI Governance

AI governance involves the responsible use of the AI data and technology. It is a proactive approach to identifying and mitigating business, legal, and ethical risks to create trust and deliver tangible business value. For organizations, this means not only adhering to existing legal and compliance frameworks but taking a holistic approach to foster trust and confidence in AI among users and stakeholders. In this survey, about a third of respondents say they have AI governance in place (not shown).

> *Organizations reporting a higher AI impact are more likely to have a dedicated AI governance team.*

AI governance includes governing the data associated with AI as well as the AI models themselves. AI data governance was already noted earlier as a top challenge in this survey, and it is worth calling out here. It includes traditional data governance objectives such as trustworthy data as well as newer concerns such as the ethical use of data. Effective data governance for AI requires data to be trustworthy, ensuring it is complete, accurate, timely, relevant, and robust enough to solve intended problems. It must also be transparent, with clear accountability, explainability, and protections for individuals against loss.

Ethical considerations are important, ensuring data is fair, unbiased, and used with beneficent intent. Data must be secure, with authorized access, encryption, masking, and protection against unauthorized actions like data leakage. Private data must remain private. Additionally, it must be safe, incorporating safeguards against risks like hallucinations, jailbreaks, and irrelevant outputs, with emerging tools and metrics aiding in detecting such issues. Finally, the data must be performant, monitored through robust observability tools to ensure its health and functionality across complex landscapes.

AI governance also includes *governing AI models*. AI model governance involves establishing policies, procedures, and tools to ensure that AI models are developed, deployed, and maintained responsibly and ethically. This includes practices such as monitoring model performance, ensuring compliance with regulatory standards, and addressing issues such as bias, fairness, and explainability. Effective model governance also involves regular audits, version control, and tracking to maintain transparency and accountability throughout the AI life cycle.

This is still an early-stage effort for most respondents to the survey. However, organizations will need to start preparing for this if they want to succeed; an analysis of the data suggests that organizations that have achieved higher impact with AI are more likely to have dedicated AI governance teams (not shown).

Cultural Considerations

Another important consideration is building the culture around AI. Building the culture for AI means fostering an organizational environment that supports the integration and utilization of AI technologies effectively. It involves enhancing AI literacy among employees, ensuring alignment of AI initiatives with business goals, and creating a cross-functional collaboration framework. Additionally, it emphasizes implementing processes such as change management and addressing challenges including resistance to adoption or lack of skills. We asked how organizations are building or planning to build a culture for AI that may require change (Figure 9).

As illustrated by the figure, there are numerous approaches organizations are taking to build the culture for AI. No one approach was utilized by one impact group over the other, although respondents who have achieved high impact

Figure 9

How is your organization building up the culture for AI? Select all that apply.



Based on 197 respondents.

with AI were less likely to use lunch and learns than moderate or low-impact respondents. This suggests that there is no silver bullet using common cultural practices to succeed with AI. For example, an AI incubator program might be effective in one organization but irrelevant in another that lacks foundational data practices. However, previous TDWI research has suggested that leadership and communication are critical cultural components to making AI successful. Leaders provide the vision, alignment, and momentum needed to integrate AI into an organization effectively. Communication is necessary to assuage concerns and to get employees on board with new initiatives.

Considerations for AI Success

As stated earlier in this report, one goal of this research was to examine the AI practices that successful organizations deploy. There may be practices these organizations are implementing that can be important for others as they move through their AI journey. To do this, we examined the relationships between the maturity groups and impact groups mentioned earlier, and the practices cited in the survey around organizational structures, technologies, and processes. Below is a summary of some of these important practices—some have already been mentioned in the report, and are explored here for more clarity, others are additional practices found in the survey data.

Data Foundations to Support AI Are Critical

The majority of organizations that we survey in TDWI research are building towards a data foundation for AI, and that is good news. As mentioned, many of them have already moved to the cloud and are using modern cloud platforms—such as cloud data warehouses, cloud data lakes, and cloud data lakehouses—to try to unify their data. A cloud data lakehouse is a unified data architecture that integrates the scalability and cost-efficiency of a cloud-based data lake with the performance, governance, and structured data management capabilities of a data warehouse. We have seen adoption of this kind of environment growing rapidly over the past few years.

Unifying data on the cloud platform is one way to support AI. Many organizations, however, have a hybrid environment with multiple platforms across cloud environments and on premises. As stated in this report, we are seeing more organizations thinking about a data fabric approach that provides an abstraction layer across multiple platforms—so to the user, it looks like one layer. As mentioned, the jury is out in terms of whether the physical or logical approach is more successful; however, the important point is that organizations should have some sort of data foundation that can support diverse, high-volume, high-velocity workloads. According to the results of this survey, **a foundation built in the cloud is more likely to result in high impact than those that are on premises or siloed**. That doesn't answer the question of platform versus fabric, as both can be in the cloud, but it does provide evidence of the value of cloud platforms.

CDO and C-Suite Presence Are Key

As we saw earlier in this report, about a third of respondents have a CDO or CAO already in place in their organization. The analysis of the survey data indicates that respondents with a

chief data/analytics officer (CDO/CAO) as the AI leader have the highest proportion of satisfaction (34%) among all leadership types. This is higher compared to other leadership types, such as CIOs, CTOs, and VPs/directors of IT (27%) (not shown). While not statistically significant in this data set, these findings suggest that organizations led by a CDO for AI efforts are more likely to be satisfied with their leadership model than those led by other roles. Time will tell if this relationship becomes significant.

Additionally, the survey results presented earlier in this report indicate that, in those cases where the data leader is part of the C-suite, there is more likely to be a higher measured impact for AI. **In other words, organizations with clear leadership for AI report higher impacts.** While not every organization has the luxury of ensuring that their CDO is part of the C-suite, this is a significant finding and something organizations should consider as a best practice.

No One Organizational Model Wins

In this report, we examined three different organizational models—a centralized model, a distributed model, and a hub-and-spoke model. While no one model came out ahead as the best model to use, we are seeing more organizations moving towards the hub-and-spoke model. The hub-and-spoke model can be best for organizations with varied lines of business that have the resources and skills to perform data and analytics activities.

Additionally, the hub-and-spoke model can support the data product concept by providing a balance between centralized governance coordination and decentralized data ownership. The hub offers shared services such as data infrastructure, governance, and metadata management, ensuring consistency and scalability across the organization. Meanwhile, the spokes—individual domains or business units—take ownership of creating and managing data products tailored to their specific needs.

Plan to Upskill and Hire New Talent

Organizations deploying AI will continue to need to upskill. AI is a journey that will require new roles and the use of new technologies. For instance, the way that companies should think about working on newer AI applications (such as generative AI) should include software engineering principles in addition to addressing the complexities of AI (especially in traditional AI that requires feature engineering and other tasks).

> *Working on newer AI applications may be closer to software engineering than ML engineering.*

We saw that many respondents to the survey already have DataOps staff in place to address data engineering. That is good news as the data engineering process may become more complicated with AI. Data products will also add to the mix. Many respondents also have data scientists in place. However, to build and operationalize AI applications, they will also need to add AI developers as well as MLOps talent. **In this survey, greater AI impact is strongly associated with the presence of AI developers/engineers and MLOps staff** (not shown).

MLOps is essential for managing the life cycle of traditional machine learning models and also

has a role in newer AI models. MLOps teams support automated deployment, monitoring, retraining, and governance of machine learning models, treating them like software applications to optimize their life cycle and maintain performance. This helps ensure models are production-ready, scalable, and compliant with governance standards, especially in regulated industries. For generative AI, MLOps still requires version control, performance monitoring, and governance. Issues including model drift, toxicity, bias, hallucinations, and performance metrics will need to be monitored.

Talent and skill are also important in that data scientists, AI developers, and others will need to use new tools that require new skills for generative and agentic AI. For example, in the case of generative AI, 57% of respondents who indicated that they are using company data for their applications also stated they are currently using retrieval-augmented generation (RAG). They will need the knowledge and skills to implement this. AI developers will need to follow a system development life cycle approach. Collaboration will be critical. Depending on the application, AI developers may need data science help and vice versa. They may both need stakeholder input. For instance, while AI developers may be able to build agentic AI applications using pre-built models and tools that can be integrated into an AI application, data scientists may be needed to develop, fine-tune, and optimize more advanced AI agents.

A newer area that is starting to gain popularity are knowledge graphs. These knowledge graphs are another area that will require skill. A knowledge graph is a structured representation of information where the entities and their relationships are stored in graph format. The knowledge graph helps put the data into context. For example, a knowledge graph can be used to represent employees, their roles, skills, reviews, salaries, and departments. These can be important for generative AI for structuring complex relationships in data as well as helping to enhance data quality and provide insights across domains. Sixty percent of those respondents measuring high impact from AI are currently using knowledge graphs for AI. While representing a small number of respondents overall, this was statistically significant, indicating that organizations should start looking at newer tools such as knowledge graphs and planning to train for these and other new technologies.

Literacy Enablement Teams Are Also Important

The concept of data literacy is not new. Data literacy involves awareness and recognition of the value of data, how well people understand and interact with data and analytics, and their ability to communicate data-driven insights to impact behavior and achieve business goals. It includes understanding the business, framing analytics, data elements, critical interpretation, and communication skills. To become more data literate, some forward-looking organizations have created data literacy enablement teams (the name may vary by organization). The team is responsible for planning and delivering on data literacy. These teams are often part of the data office or a Center of Excellence. Data literacy enablement teams determine what roles need what training, devise the strategy, and execute on it. They help build the culture around data.

With the introduction of AI, some organizations are creating AI literacy enablement teams. The idea is the same, except the focus is on how to understand and interact with AI systems. An AI

literacy enablement team is a group dedicated to educating and empowering employees across an organization to understand, adopt, and effectively leverage AI technologies in their roles. This team typically focuses on building foundational knowledge, addressing ethical considerations, and fostering skills for working with AI-driven tools and systems.

In this survey, 35% of respondents offer AI literacy programs; 54% of organizations with higher AI impact levels offer AI literacy programs, compared to only 20% of organizations with lower AI impact levels. While the numbers are small, the data seems to suggest that **organizations achieving higher levels of AI impact are more likely to offer AI literacy programs.** AI literacy programs can be important for value creation as they equip teams with the skills and understanding needed to effectively leverage AI, driving higher organizational impact and better alignment with strategic goals.

Measure the Impact

Measuring the impact of AI is important because it ensures that AI initiatives deliver tangible value, align with business goals, and provide insights for continuous improvement and governance. In this survey, those respondents that stated that they are mature/very mature are significantly more likely to use metrics to measure success in analytics and AI than less mature organizations. This is also the case for those who have measured high impact. In this survey, we specifically asked about ROI associated with analytics and AI efforts. **There was a highly significant relationship between those respondents who are very mature and the use of ROI metrics.** Other metrics that were significant included metrics that tie to the use of analytics and AI as well as those that tie to business goals.

A metric might be something like average time taken from order placement to shipment. Metrics might include worker efficiency, process efficiency, and cost savings. These metrics provide actionable insights, helping organizations optimize AI systems and justify their investments to stakeholders. Adopting measurable success criteria also fosters accountability and continuous improvement, making AI deployment more effective and aligned with strategic priorities.

AI Governance Will Become Critical

We've talked about the importance of both data governance and AI governance. A well-structured data governance team provides better data quality, compliance, and strategy alignment, which are important for effectively leveraging AI. Although the numbers are small, the results of this survey suggest that **organizations with high impact are far more likely to have dedicated AI governance teams** with fewer relying solely on data governance teams.

This group is using various approaches to AI governance. Some seem to be putting AI governance responsibility into a single team or making it part of an AI Center of Excellence. Others cited cross-functional collaboration in a collaborative approach. While AI governance is not yet widespread, it is important to start planning for it now.

Make Use of AI-Infused Tools and Platforms

In addition to the considerations already explored, another important one is to make use of AI-infused tools and platforms. These can help organizations democratize AI. For instance, a B2B

platform infused with generative AI can help a salesperson ask a question of their B2B database platform about their company's history with a customer. AI-infused BI tools can surface insights and even build certain predictive models.

Recommendations

This report has detailed many best practices for succeeding with AI. In closing, we summarize the report by listing the top best practices for successful AI, along with comments about why each is important. Think of the best practices as recommendations that can guide your organization on a successful AI journey. Remember, however, that it is a journey, and roles and processes will evolve.

Tailor your AI adoption paths. Organizations should align AI adoption with their business needs and their business strategy. This is critical. Not every business problem needs an AI solution. From a technology perspective, choose to buy, bolt-on, or build solutions depending on the context. For some companies it may make sense to start with simpler, targeted, and practical high-value use cases and gradually progress to more complex applications like agentic AI.

Focus on measurable outcomes. Treat AI as a strategic capability and align it with measurable business outcomes, such as cost reduction, efficiency gains, or enhanced customer experiences. Set clear expectations and measure ROI.

Invest in platforms that offer strong integration and support for diverse data types. These platforms are often cloud based. Consider both physical and logical architectures and determine what will work best for your organization. Remember the importance that unstructured data will play in new AI use cases.

Pay attention to data quality. Data quality issues continue to plague many organizations. Data quality is a journey; and as new data types are used for both traditional AI and generative AI, organizations will need to pay attention to quality. Consider some of the new tools that are on the market to help with quality for both structured and unstructured data.

Consider new leadership models. Organizations with CDOs who are part of the C-suite appear to be more successful with AI than those whose leaders are not part of the C-suite. Make this recommendation if possible.

Ensure governance and trust. Strong governance processes are critical, particularly for unstructured data. Organizations must establish robust governance frameworks, provide transparency about AI limitations, and develop best practices for implementation. Monitoring processes should evaluate AI outputs for accuracy, bias, and trustworthiness, while robust access controls align with back-end governance tools to prevent security risks and foster trust in AI deployments. Clear communication of AI processes can help manage user expectations and build trust while ensuring secure and scalable AI adoption.

Invest in frameworks and disciplines. This includes processes such as MLOps, as we saw in this report. It will also include processes related to software development as organizations start to build and implement newer kinds of AI applications.

Invest in workforce enablement. AI literacy should extend beyond data scientists to include developers and business users, fostering a culture

where all stakeholders understand and use AI effectively. Educating users on interpreting AI outputs and integrating AI into workflows ensures broader adoption and more effective collaboration across teams.

Make sure new tools as well as AI-infused tools are on your road map. Vendors are offering a host of new tools that infuse AI. These can be very helpful for organizations seeking to gain value with AI and empower their staff.

Communicate. People across the organization need to understand the strategy and plans around AI. They can't invest emotionally in new innovations if they don't know what is happening. Likewise, if their job is to be impacted by AI, they need to know. Make sure that communication is part of the strategy.

Think about the success cycle. We have seen that success begets success. If you are successful with one project, do not stop there. Build on that success to become more sophisticated as an organization.

From Our Sponsor

zoominfo

Go-to-Market Intelligence: The GTM AI Foundation

Today's most innovative companies share a few common growth goals: acquiring more customers, upselling and retaining existing accounts, and improving the productivity of their frontline teams with AI tools.

All these goals demand reliable, actionable data—yet most businesses are still relying on the often-outdated information in their customer relationship management (CRM) platforms.

ZoomInfo's Data-as-a-Service (DaaS) solves these problems by offering innovative teams access to our unparalleled Go-to-Market Intelligence: trustworthy, constantly updated B2B data, and high-velocity buying signals that can fuel their GTM engines.

Some of the world's biggest brands—including 60% of the *Fortune* 100—rely on ZoomInfo's **Go-to-Market Intelligence Platform** to power their growth. And that same intelligence is also available whenever and however the largest enterprise businesses need it: through an API, delivered to cloud data warehouses, or even via flat file.

The data universe that powers ZoomInfo's GTM Intelligence Platform is the most accurate, complete, and actionable business data in the world, encompassing more than 100 million company profiles, over 500 million individual contact records, and billions of high-value buying signals.



Intelligence Built for the AI Era

Even the most advanced AI models can't drive bottom-line results if your underlying data isn't trustworthy. According to a ZoomInfo survey of more than 11,000 GTM professionals, 43% of companies struggle to connect buyer insights across multiple systems, resulting in missed opportunities and lost revenue.

ZoomInfo is already helping GTM leaders bridge these gaps and seize the opportunities their competitors are missing. In fact, ZoomInfo customers were able to increase their TAM by an average of 40% in 2024, thanks to the comprehensive visibility offered by our GTM Intelligence Platform and its underlying data universe.

Biotechnology firm Thermo Fisher approached ZoomInfo after sales leaders suspected there was undetected whitespace in their market. After analyzing more than 32 million job vacancies, we discovered untapped sub-industries and geographies for Thermo Fisher's business. As a result, the company's TAM **is over $20 billion larger** than previously estimated.

We also helped Thermo Fisher build an enhanced definition of its ICP—based on historical data about closed-won new business—and provided the company with intent, news, and technographic data, which they used to create AI-powered propensity models. These insights allowed the company to reallocate frontline sales teams and prioritize marketing spend to the highest-propensity accounts—opportunities that were 80% more likely to convert and three times more likely to close.

At ZoomInfo, we understand the need for actionable, insightful GTM Intelligence that is ready for both people and processes. ZoomInfo's DaaS offerings are fully agent-ready and can be deployed across customers' internal systems to provide unique, bespoke data solutions to businesses' most complex, urgent problems.

We've worked with Amazon Web Services, Google, Snowflake, and many more leading data companies to develop unique workflows, processes, and customized deployments using our world-class data universe.

Fuel Your AI Strategy with Global Impact

ZoomInfo's award-winning GTM Intelligence Platform isn't just accurate and AI-ready—it's also fully compliant with various international data regulations, including Europe's GDPR, to ensure business leaders can act with confidence that their GTM motions are in line with established and emerging data standards around the world.

In 2024, ZoomInfo was the first data company to earn the TRUSTe Responsible AI Certification from TrustArc, a leader in data privacy management and governance solutions, recognizing ZoomInfo's robust commitment to data privacy, security, and ethical oversight.

Year-over-year, ZoomInfo's coverage of company data in North America has increased by 30%. International contact data has expanded by more than 250%, and international company data coverage has increased by more than 55%, giving today's global businesses unprecedented insights into global markets and the immense opportunities within them.

ZoomInfo has been helping businesses in every vertical and sector reach new markets for almost 20 years. Learn more about how ZoomInfo's DaaS can help yours.

TDWI Research provides research and advice for data professionals worldwide. TDWI Research focuses exclusively on data management and analytics issues and teams up with industry thought leaders and practitioners to deliver both broad and deep understanding of the business and technical challenges surrounding the deployment and use of data management and analytics solutions. TDWI Research offers in-depth research reports, commentary, inquiry services, and topical conferences as well as strategic planning services to user and vendor organizations.

tdwi TRANSFORMING DATA WITH INTELLIGENCE™

A Division of 1105 Media
6300 Canoga Avenue, Suite 1150
Woodland Hills, CA 91367

E info@tdwi.org

tdwi.org