The Critical Role of AI Governance for AI Success

Jason Stamper
Research Manager, Software, Data, & Analytics, IDC Europe

Neil Ward-Dutton
Vice President Automation, Analytics and AI, IDC Europe

Table of Contents

CLICK BELOW TO NAVIGATE TO EACH SECTION IN THIS DOCUMENT.

SECTION 1:

AI Use Cases and the Future of GenAI

Interest in Generative Artificial Intelligence Is Astounding — But Will It Be Governed?

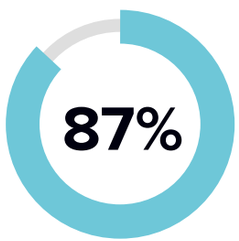

say GenAI **will have an impact** on their business.

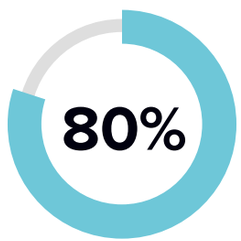

of organizations are **exploring or investing** in generative AI.

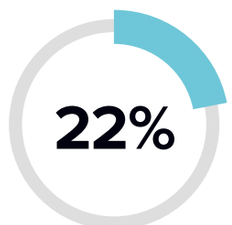

of organizations have clear investment plans for the **next 18 months.**

say GenAI has **already disrupted** their business.

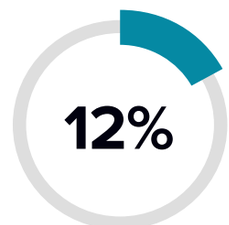

is the mean share of **new IT project budget** allocated to GenAI.

Generative artificial intelligence (GenAI) technologies appeared seemingly from nowhere at the end of 2022 with the release of ChatGPT a major milestone. Within a year, organizations across all industries and geographies moved incredibly quickly to start exploring the possibilities.

Experimentation with AI is a natural way to start understanding the possibilities, but to gain long-term value from and manage the risks of AI investments, organizations must develop and implement AI governance capabilities.

Note: Data weighted by IT spend. n = 1,363; Source: IDC's *GenAI Awareness, Readiness and Commitment Survey,* August 2023

C-suites Want Quick Results from GenAI Implementations

GenAI interest from senior leaders is broad, and they want results fast — but speed must be paired with control.

Top Areas for GenAI Integration

40%	Customer-facing applications
31%	Financial/operational decision support applications
31%	Employee applications (skills/talent/training)
28%	Digital marketing tools
26%	App developer tools
23%	Product life-cycle management and product design applications
21%	Request for proposal/contract process applications

Implementation Timeframe for Major Areas

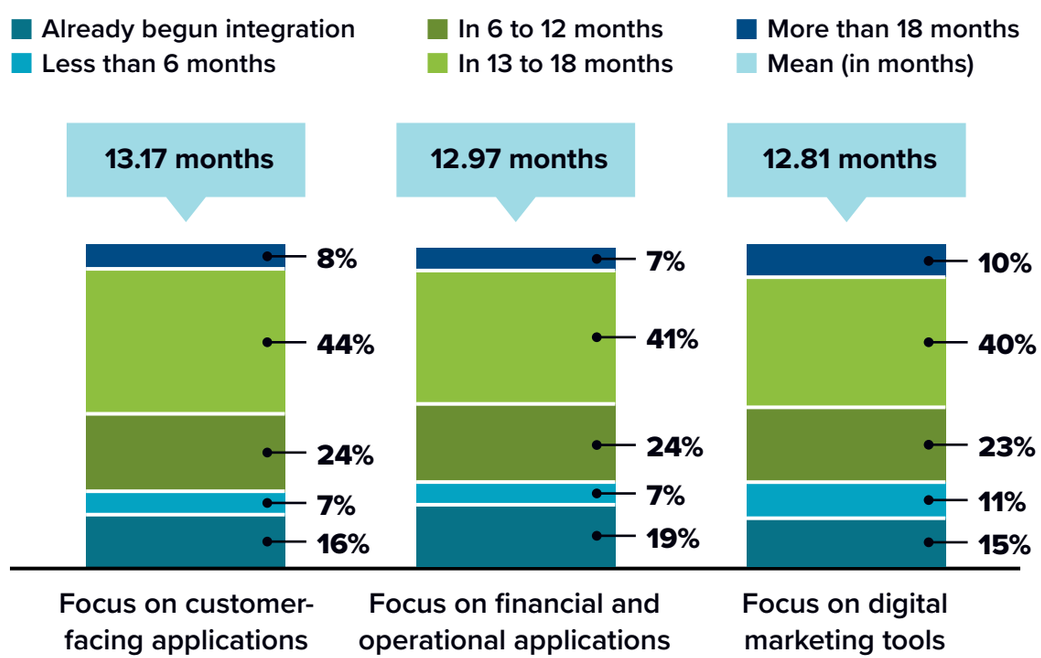

Note: Data weighted by IT spend. n = 1,363; Source: IDC's *GenAI Awareness, Readiness and Commitment Survey*, August 2023 | For an accessible version of the data on this page, see Supplemental Data in the Appendix

Where Are We With AI Today?

Even disregarding the impact of nascent GenAI spending, budgets for AI technology and AI applications, in general, are already highly significant in many organizations.

Companies using AI today are focused on how it can help them improve the most fundamental business objectives: customer experience, operational efficiency, and employee productivity.

Organizations are desperate to move quickly with GenAI, but in the rush to implement, they must take care not to ignore the risks associated with innovative technologies.

Large organizations currently spend, on average, **$134 million** annually on AI.

IDC forecasts that worldwide spending on AI software will exceed **$250 billion** by 2027.

Top 3 business objectives driving AI spending are:

- Customer experience
- Operational efficiency
- Employee productivity

The fastest growing business objective is:

- Improved sustainability

Source: IDC's *Worldwide Artificial Intelligence Spending Guide*, August 2023; IDC's *AI StrategiesView 2022*, May 2022

Common AI Use Cases Across Industries

Investment in AI is not just confined to financial services organizations, high technology firms, or telecommunications service providers.

IDC tracks spending on dozens of use cases across 20 industries. The applications attracting the highest spending are diverse, from helping experts with investigations to making predictions and recommendations, automating customer interactions, automating IT infrastructure operations, and helping prevent equipment failures.

Healthcare — Top 5 use cases

1. Augmented diagnosis and treatment systems
2. Image processing: classification and recognition
3. Sales process recommendation and augmentation
4. Digital assistants
5. Program advisors and recommendation systems

Financial services — Top 5 use cases

1. Augmented fraud analysis and investigation
2. Program advisors and recommendation systems
3. Augmented threat intelligence and prevention systems
4. Augmented customer service agents
5. Augmented claims processing

Telecoms — Top 5 use cases

1. Augmented customer service agents
2. AI infrastructure provisioning
3. Augmented threat intelligence and prevention systems
4. Program advisors and recommendation systems
5. Smart networking

Manufacturing — Top 5 use cases

1. Augmented quality management investigation and recommendation systems
2. Automated preventative maintenance
3. Augmented customer service agents
4. IT optimization
5. Digital assistants

Retail — Top 5 use cases

1. Augmented customer service agents
2. Expert shopping advisors and product recommendations
3. Sales process recommendation and augmentation
4. Augmented merchandising for omni-channel operations
5. Infrastructure provisioning

Government — Top 5 use cases

1. Augmented defense, terrorism, investigation, and government intelligence systems
2. Augmented threat intelligence and prevention systems
3. Augmented public safety and emergency response
4. Program advisors and recommendation systems
5. Automated preventative maintenance

Source: IDC's *Worldwide Artificial Intelligence Spending Guide V1 2023 Forecast*, August 2023

Financial services include banking, securities and investment services, and insurance. **Manufacturing** includes discrete manufacturing and process manufacturing. **Government** includes federal/central, state, and local governments.

Established AI Case Study: Financial Services

One global banking firm with a strong presence in Spain, Mexico, South America, and more **is building advanced analytics and data services in the cloud as part of its data and AI transformation process.** It is harnessing analytics and machine learning (ML) to transform its internal processes, improve risk management, drive growth, and provide innovative solutions for its customers.

The bank is creating a new data platform that will be deployed globally, delivering a secure repository of its operations and customer data. **This platform will provide internal business stakeholders with automated business and market insights, increase operational efficiencies, and attract new customers.**

The company is, however, mindful of data governance and compliance while spurring collaboration and information sharing across key business areas. **By combining data lakes and data warehouses, the platform will provide business units with a unified view of their data and access to more efficient data processing, analysis, and insights, built in the cloud.**

Established AI Case Study: Healthcare

A global healthcare company has grown through multiple mergers and acquisitions over several decades. **Focused on respiratory care products, it has over 20,000 products and over 300,000 customers, all around the world.**

Managing the data the business generates is extremely complex. The company's employees use more than 13 enterprise resource planning and customer relationship management systems, which, combined with increasing use of advanced analytics and machine learning, create a large and complicated data ecosystem. It had reached a point where managing data was getting in the way of doing business.

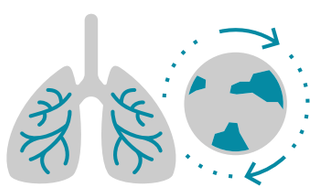

In a highly regulated industry, the company also needed to ensure that, while integrating third-party data through APIs, **it was still able to maintain trust in its data and comply with regulations such as the Health Insurance Portability and Accountability Act (HIPAA) and General Data Protection Regulation (GDPR).**

The company now uses both machine learning and AI, ingesting and analyzing external data for better analytics. The company is able to find, subscribe to, and analyze multiple sources of data from partners and healthcare providers in the cloud.

Established AI Case Study: Retail

As a well-known **British retailer with over 850 stores across the United Kingdom and over 90 websites serving customers both in the United Kingdom and abroad,** this company generates and is trusted with a vast amount of data.

Its customer loyalty program has more than 15 million members in the United Kingdom and Ireland. It uses machine learning and data science to gain customer insights and build brand loyalty by recommending products based on past purchases, sending loyalty emails to rewards members, and creating targeted campaigns.

The previous data science solution relied on many manual processes, and the firm wanted a faster, less labor-intensive way to get models into production with more trusted data quality and accuracy.

The company used machine learning to reduce manual processes and improve the continuous integration/continuous delivery (CI/CD) process. This helped minimize errors and failures in machine learning model production and should ultimately lead to better customer experiences and improved profitability.

The Dawn of the AI-Everywhere Era

Rapid adoption is moving AI from an emerging software segment to a critical technology at the center of a platform transition and a new era of AI everywhere.

The mainstream arrival of GenAI at the end of 2022 with the launch of ChatGPT triggered the opening of this new era because its use of pretrained "foundation models" shifted the economics of AI. By leveraging foundation models, **organizations could implement AI innovations much more quickly and cheaply across a wide range of use cases associated with automation and intelligence.**

The AI-everywhere era will change customer and employee experiences, functionality of products and services, and the flow of business processes and decisions. It will also change how organizations think about protecting and delivering the value of their data.

Digital Business:
Innovation at Scale

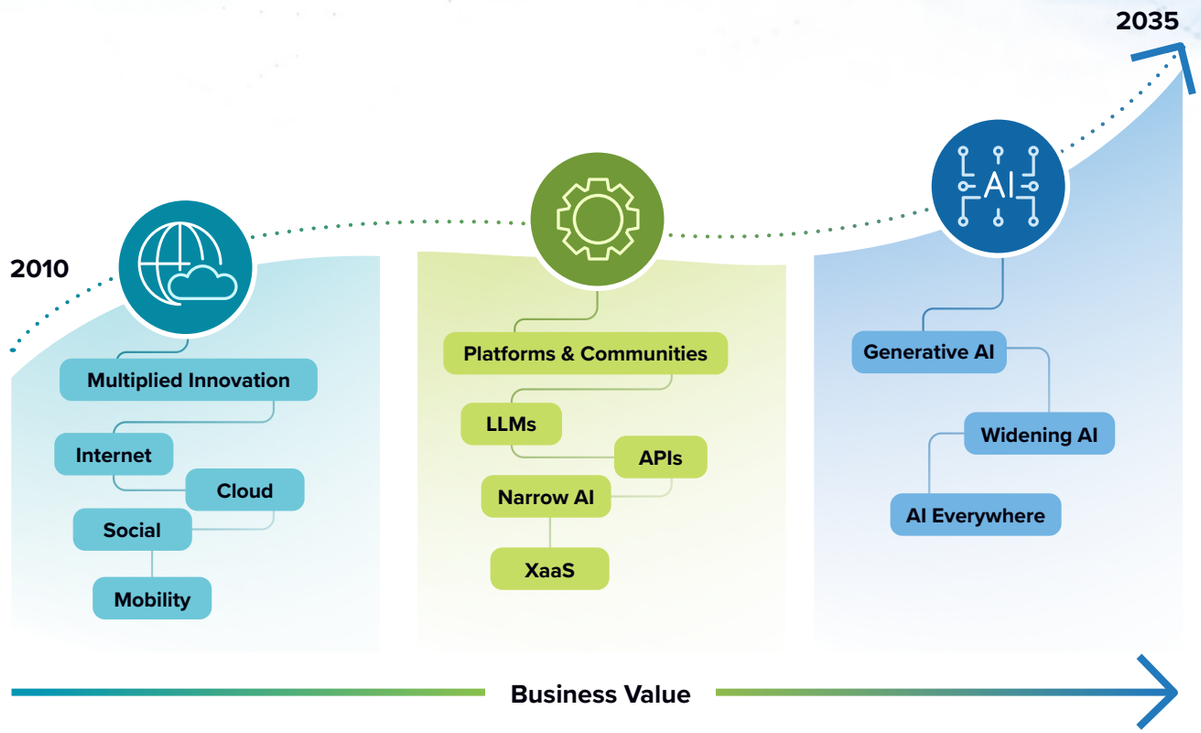

SECTION 2:

AI Risks

An Overview of AI Risks and Challenges

Implementing AI brings a wide spread of issues that organizations must manage carefully, although not all bring the same risks or are at the same level. Ultimately, if these issues are poorly managed, they create material business risks. These can be legal (judgments against corporations or executives), reputational (damage to brand from scandal, etc.), and financial (fines, undetected fraud or errors, etc.). Current and emerging industry regulations are increasing the pressure.

AI Implementation Process

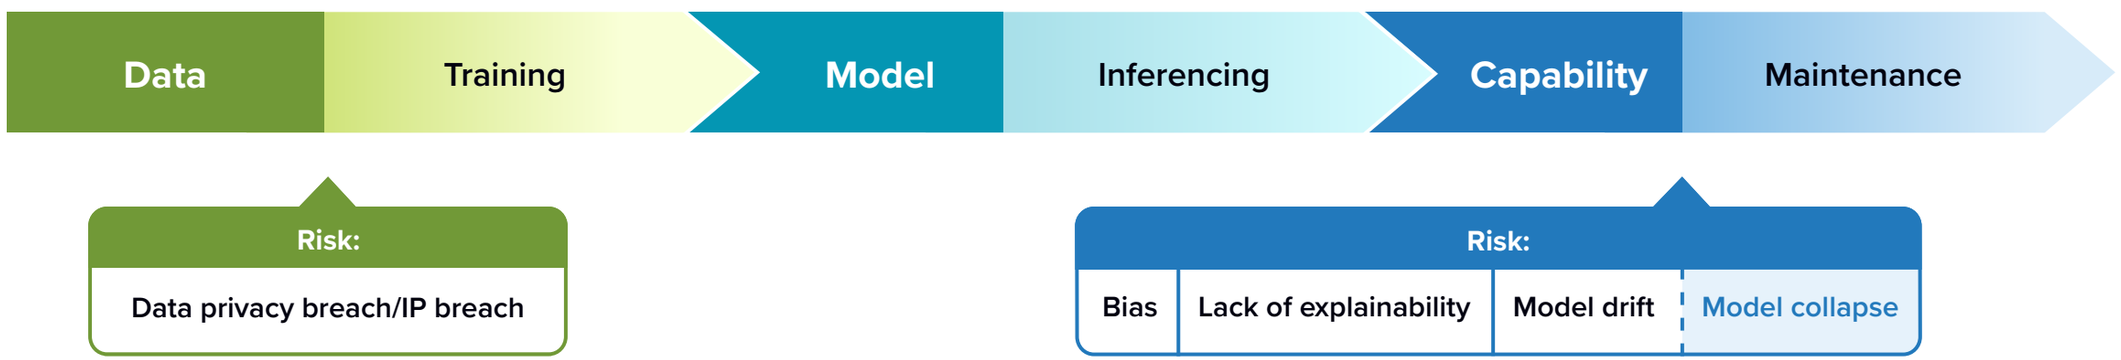

Managing these risks effectively through AI governance is not only increasingly desirable, it is required. Current and planned industry regulations are pressuring organizations worldwide to take these risks seriously across the whole of the end-to-end AI implementation process.

Risk Example: Data Privacy and Intellectual Property Breaches

For years, data privacy has been a major issue for organizations implementing AI. It is particularly critical when organizations want to leverage personal information (as defined in the European Union GDPR, various U.S. states' privacy laws, HIPAA, and data protection legislation in Asia including the Philippines, Japan, and South Korea). The penalties for organizations breaching the data privacy rights of employees, citizens, or customers can be severe, on top of the reputational damage. **When it comes to GenAI, organizations must work to minimize the risks associated with using model prompts that include sensitive information. IDC research shows that many organizations exploring GenAI are already taking these risks seriously.**

Intellectual property (IP) risks also come to the foreground with GenAI. These risks can flow both ways: **Organizations may lose control of their own IP, or they can inadvertently use IP that other organizations have claim over.**

In January 2023, Getty Images sued Stability.ai (the company behind the Stable Diffusion image generator model), claiming that Stability.ai illegally scraped content from Getty Images' library to train its image generation models, without consent.

In April 2023, Samsung suffered from a serious IP breach when engineers pasted confidential software code written by other Samsung teams into ChatGPT in order to understand it, and OpenAI's system then used that code as part of its training data set.

of organizations are experimenting with **private versions of GenAI** models rather than public versions.

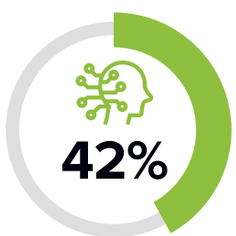

are concerned that **GenAI jeopardizes control of data** and intellectual property assets.

Notes: Data weighted by IT spend; n = 1,363. Source: IDC's *GenAI Awareness, Readiness and Commitment Survey*, August 2023

Risk Example: Bias

When AI systems are trained on data sets that do not accurately reflect the populations they seek to represent, the results of predictions or generations may reflect bias. For example, samples might be skewed toward data with particular characteristics, or a subset of training data clustered around a particular characteristic (or set of characteristics) might be of poorer quality than the rest. Biases can appear in training data sets of multiple kinds, from structured data describing customers' interests, to images of people used to train a recognition algorithm or to generate marketing content. **Bias can be problematic for GenAI, just as it can be for predictive AI.**

In 2019, in the United States, **an AI algorithm used to assess the likelihood of defendants reoffending was found to make an incorrect assessment far more frequently for black defendants than for white defendants.** In contrast, the system was far more likely to incorrectly flag white defendants as low risk.

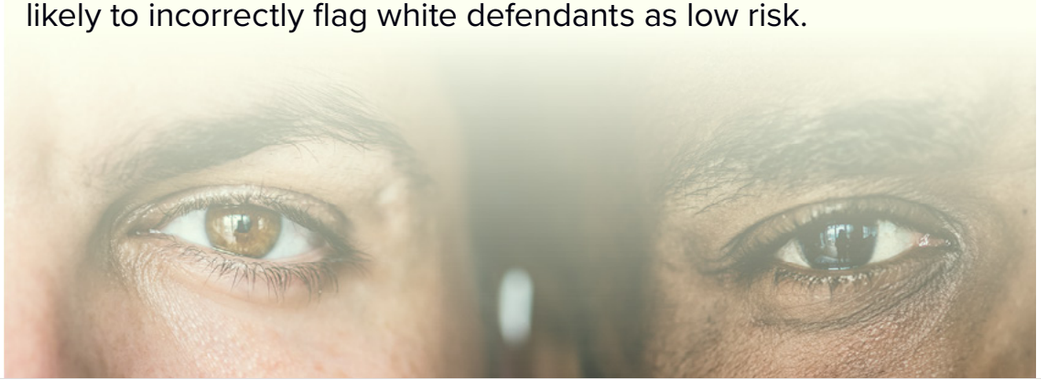

35%

of organizations are concerned about bias in outputs of GenAI models.

The EU's AI Act, scheduled to be introduced in 2024 and fully implemented by 2026, **will require the supplier of any "high risk" AI system to demonstrate that it is taking action to minimize bias.** A high-risk system is one that provides a safety component of another system, operates in areas such as biometric identification or law enforcement, or is likely to have a material impact on the rights of a person (for example, one that decides loan eligibility).

Notes: Data weighted by IT spend; n = 1,363. Source: IDC's *GenAI Awareness, Readiness and Commitment Survey*, August 2023

Risk Example: Lack of Explainability

Explainability of decision-making algorithms is straightforward when algorithms are based on rules or common statistical modeling approaches. It is tougher when algorithms are created through machine-learning processes. Explainability is particularly challenging when "deep learning" techniques based on very large neural networks are used. **Model explainability techniques are now finding their way into popular AI/ML development platforms, but they can be computationally expensive and require skilled data scientists.**

For issues that involve customers, citizens, suppliers, and employees (such as determining product or service eligibility or pricing), and particularly in regulated industries, the explainability of decisions and the systems that underpin them is absolutely essential.

Article 22 of the EU GDPR specifies the **"right to an explanation"** of how a decision was reached if that decision was based on personal information and automated processing.

A key principle underlying the EU AI Act (scheduled to come fully into force in 2026) is the requirement for high-risk AI systems to be transparent and explainable. **This means companies must be able to explain how their AI systems make decisions.**

Risk Example: Model Drift

Just as the requirements shaping traditional software algorithms developed by humans can change over time, the predictions created by AI models can also "drift," meaning predictions become less and less accurate.

There are two main types of model drift:

- Concept drift occurs when the nature of what constitutes an accurate prediction changes. For example, when patterns of fraud change, a model trained on old patterns will become less reliable.
- Data drift occurs when historical data used to train the model ages, and the statistical properties of the data change.

57%

of organizations rebuild and redeploy AI models at least weekly.

n = 1,403; Source: IDC's *AI StategiesView 2022*, May 2022

If a predictive AI use case needs to deliver value consistently over time (rather than being a one-off exploration of an issue or opportunity), **it must be built and maintained using processes and tools that can detect and correct model drift.**

Most Organizations Are Not Systematic in Their AI Governance Approach

When asked about their level of maturity in implementing AI initiatives, the majority of C-suite leaders describe their approach as "unstructured" or "isolated." **Only 3% of organizations really consider and implement AI as a strategic tool.** Less than half of organizations say their level of maturity reflects some degree of consistency in AI governance — though IDC expects this to improve as organizations continue on their AI implementation journeys.

Unstructured

There is no formal AI strategy or coordination across the organization. AI technologies may be used in silos by select individuals on very specific use cases, but insights are not visible to anyone in the C-suite. No generative AI use cases are being explored.

Isolated

An AI strategy and road map for the organization are being developed. AI technologies are used for isolated projects. Data readiness, governance, skills, and technology selection are limited to specific initiatives. A handful of generative AI use cases are being explored.

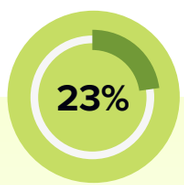

Consistent

AI technologies are used for multiple projects. Data readiness, governance, skills, and technology selection are repeated across those initiatives. An AI strategy and road map are in early stages of development. Some generative AI use cases are being tested.

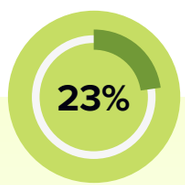

Advanced

An enterprise-wide AI strategy and road map are aligned to business goals. Data readiness, governance, skills, and technology selection are consistent. A significant number of generative AI use cases are being implemented.

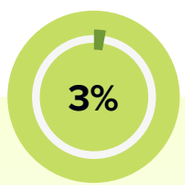

Strategic

An enterprise-wide AI strategy and road map are used to create new business value internally and externally. Data readiness, governance, skills, and technology selection adapt to support the broader strategy. Generative AI use cases are delivering expected business outcomes.

n = 895; Source: IDC's *Worldwide C-Suite Tech Survey,* August 2023

SECTION 3:

The Critical Role of AI Governance

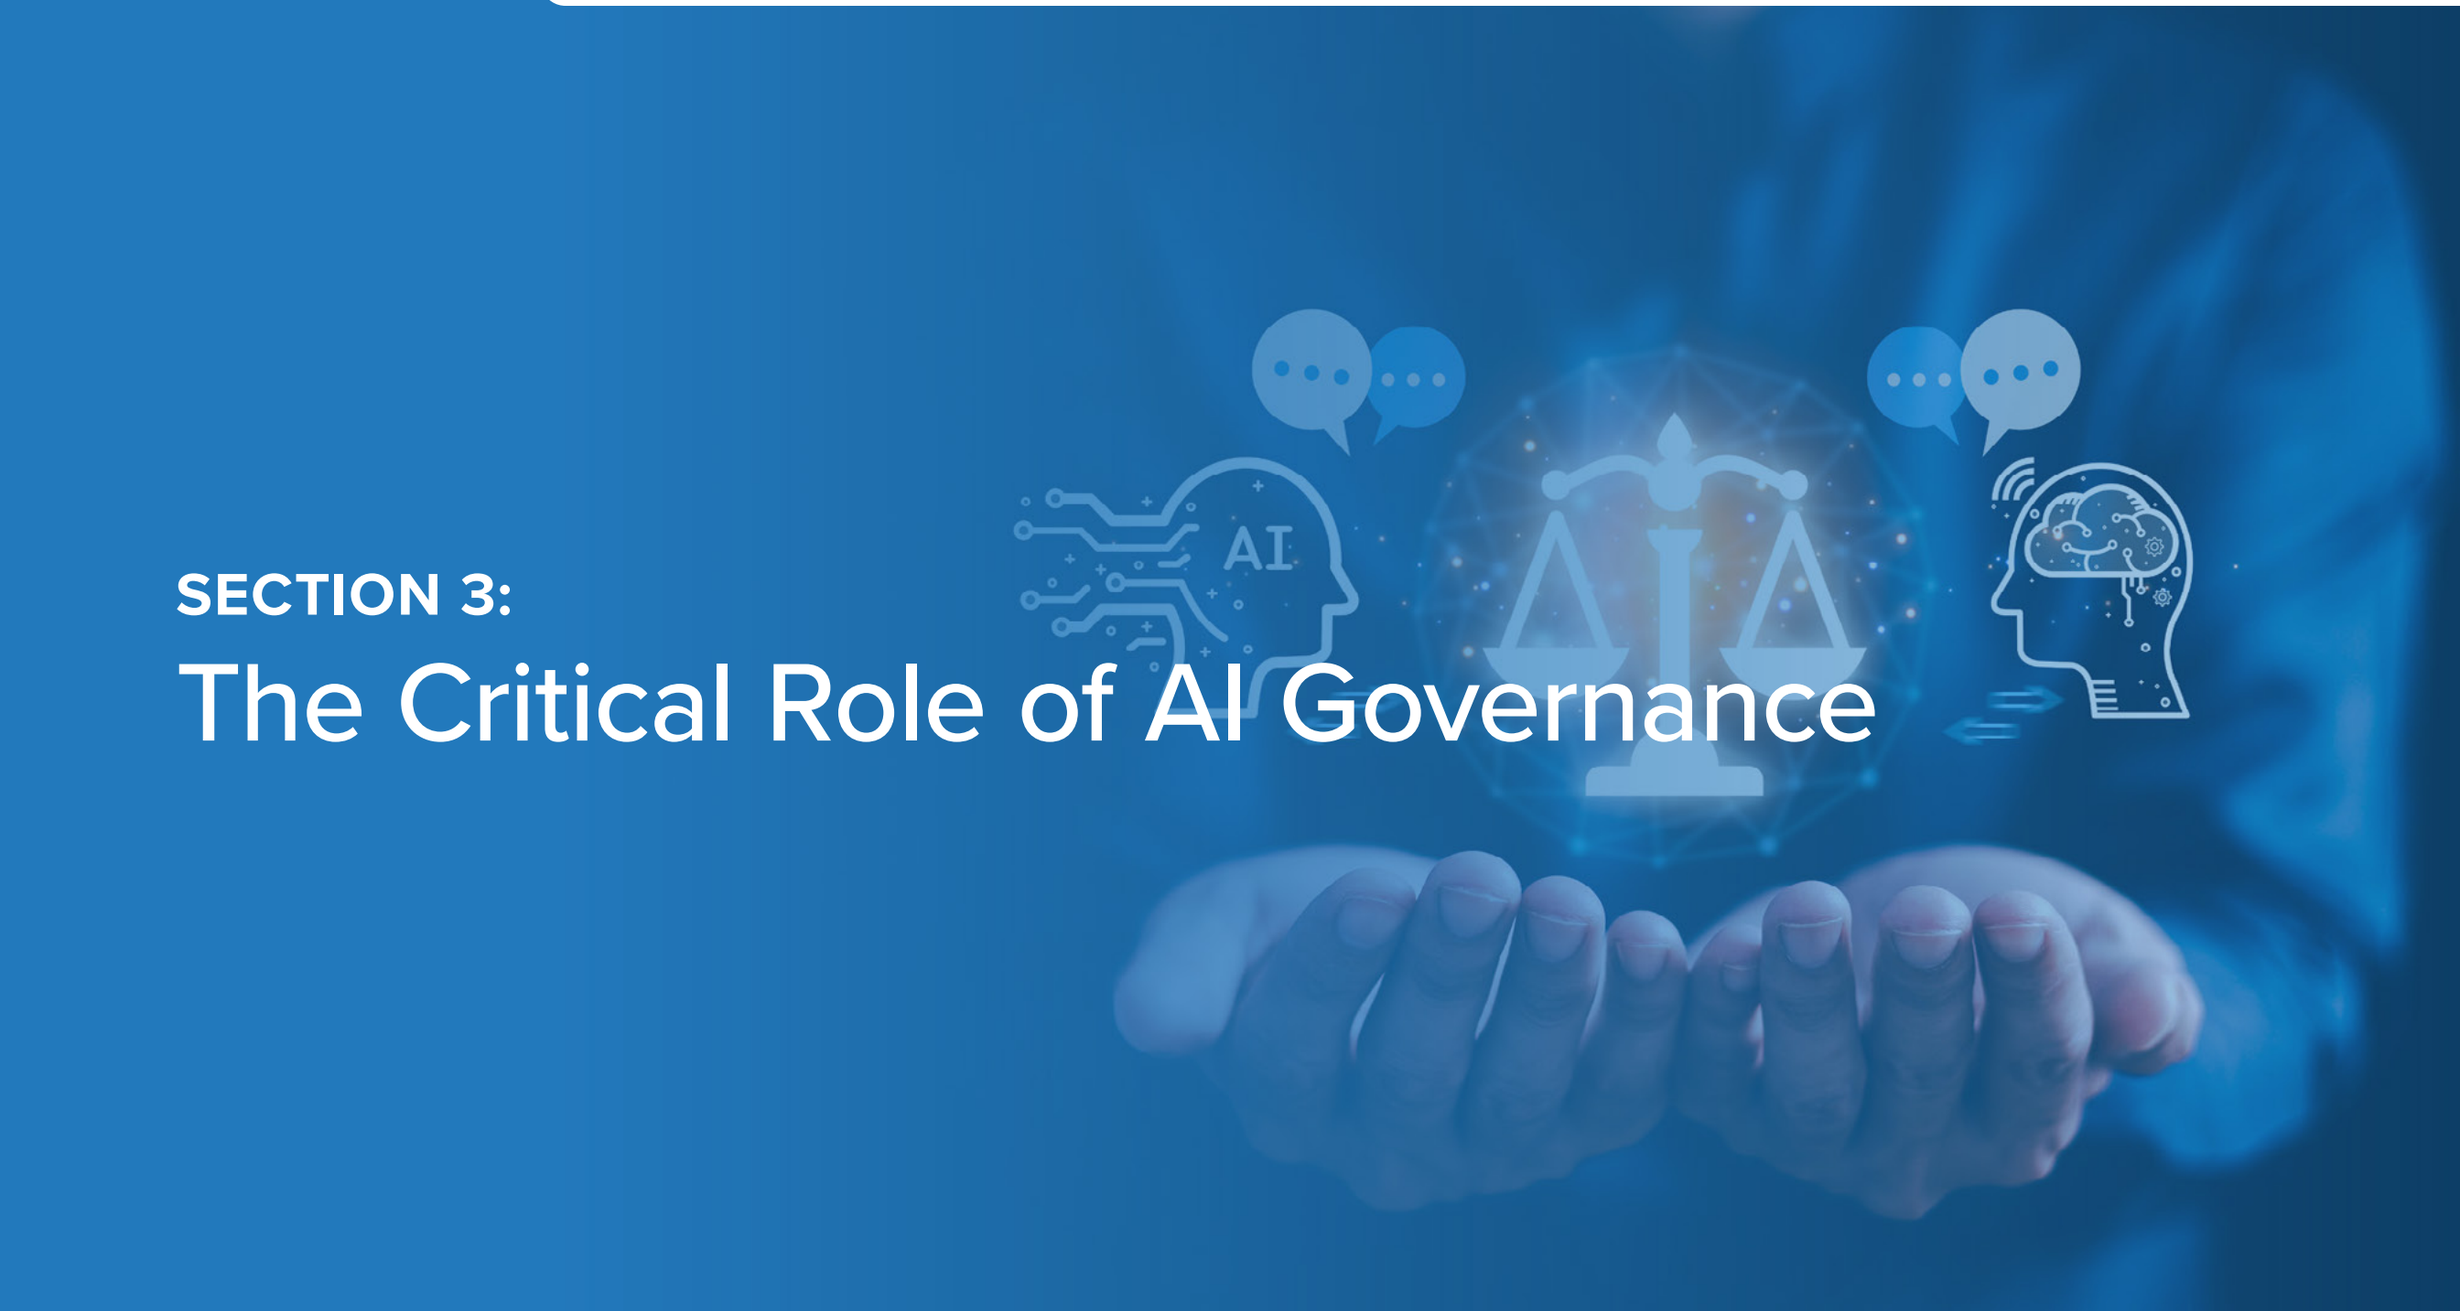

Data, Tools, and Trust Inhibit the Value of AI/ML

What are the biggest inhibitors to realizing value from AI/ML?
Respondents could select multiple answers.

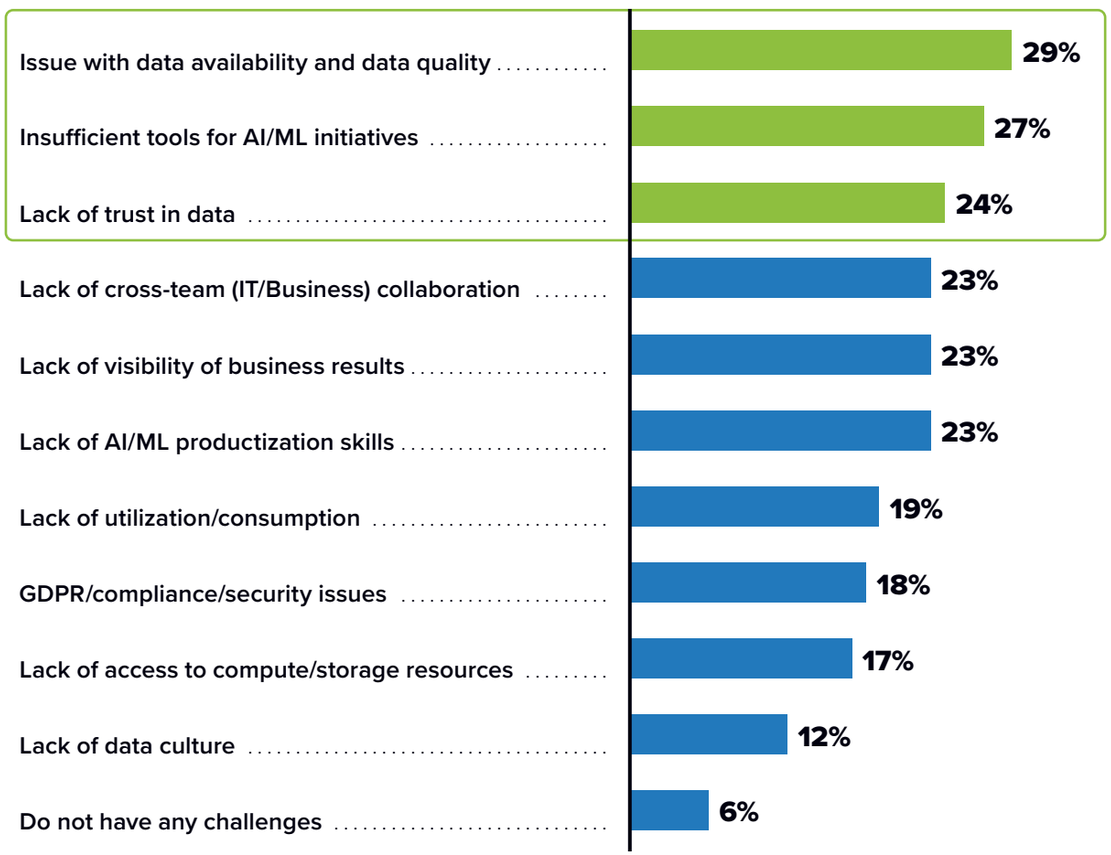

n = 952. Source: IDC's *Future Enterprise Resiliency and Spending Survey, Wave 2*, March 2023

The backbone of AI's capabilities is data, and enterprises rely heavily on vast data sets to train and deploy AI systems. However, leveraging AI effectively within an enterprise comes with data-related challenges that must be addressed to realize its full potential. Effective AI/ML deployment relies on equally good data governance tools, processes, and culture.

Garbage In, Garbage Out

AI without trusted, high-quality data is a disaster waiting to happen. Organizations need to ascertain that the data feeding their AI is...

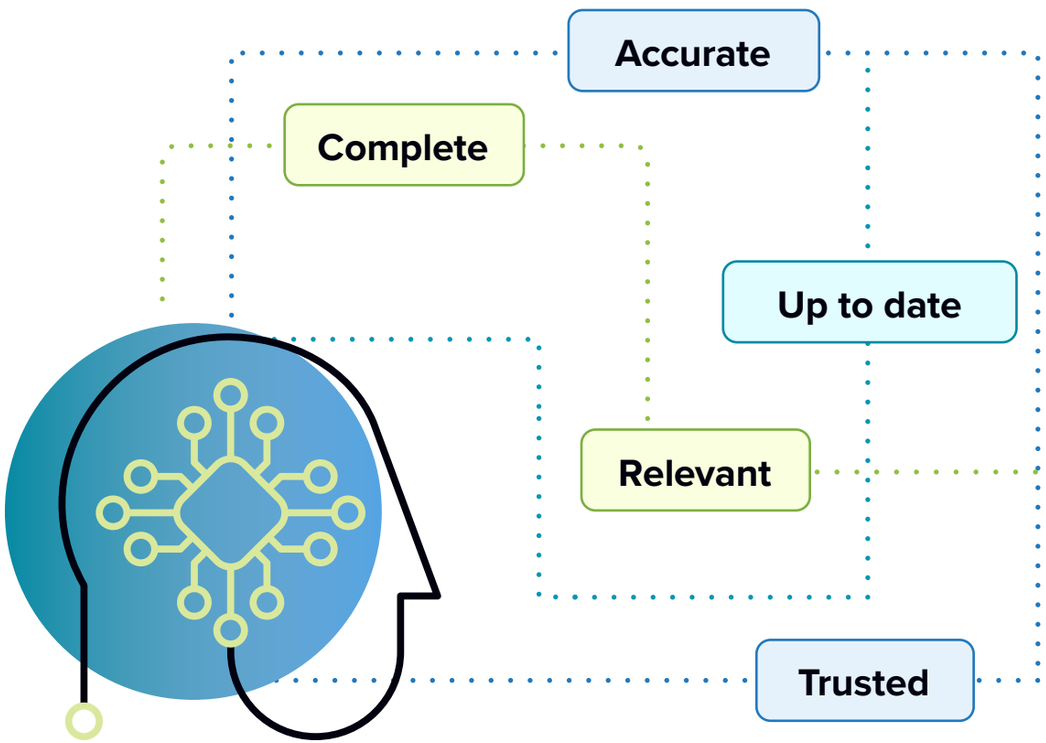

One of the most critical challenges in AI integration within an enterprise is the quality and consistency of data. **AI models require clean, structured, and accurate data to deliver reliable results.** Many organizations struggle with siloed data sources, inconsistent data formats, and data that is outdated or incomplete. **These issues hinder AI's ability to provide meaningful insights or make accurate predictions.**

The Detrimental Business Impacts of Dubious AI

Inaccurate or difficult-to-trace AI and data models carry many risks for organizations.

Enterprises that invest in robust data governance, security measures, and ethical AI practices will be better positioned to navigate the evolving landscape of AI technologies.

Success will require tools, processes, and cultural behaviors that transform data into value.

What are your top 2 potential negative business impacts if AI/ML is not mitigated or implemented responsibly?

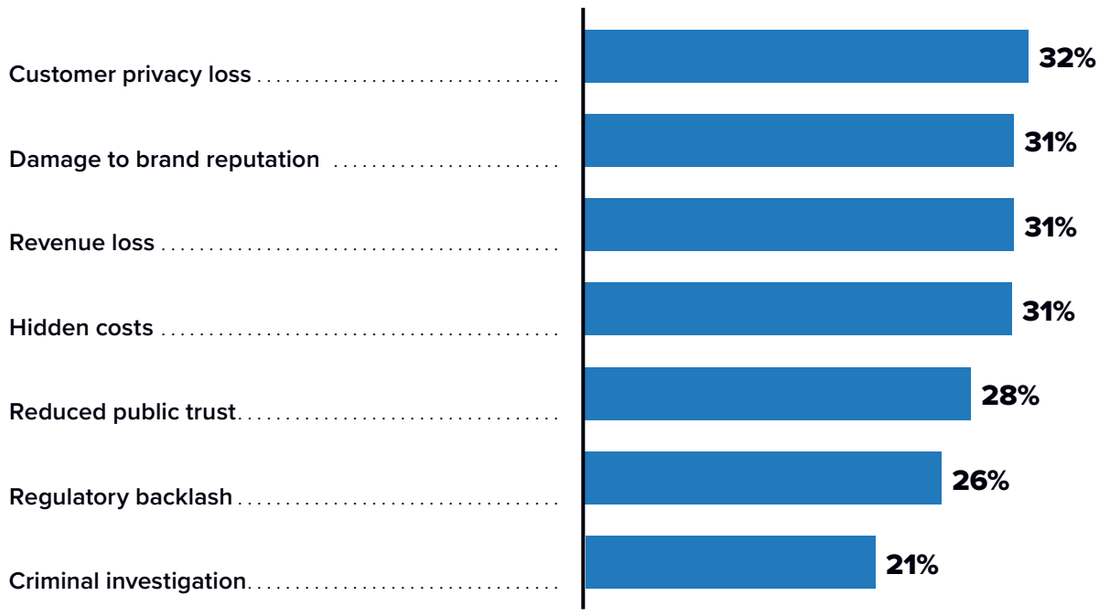

n = 2,053. Source: IDC's *StrategiesView,* May 2022

The Challenge of Trusting AI Data

On a scale of 1 to 5, please rate how challenging a lack of fairness, explainability, transparency, and data lineage tools are holding back AI/ML projects within your organization.

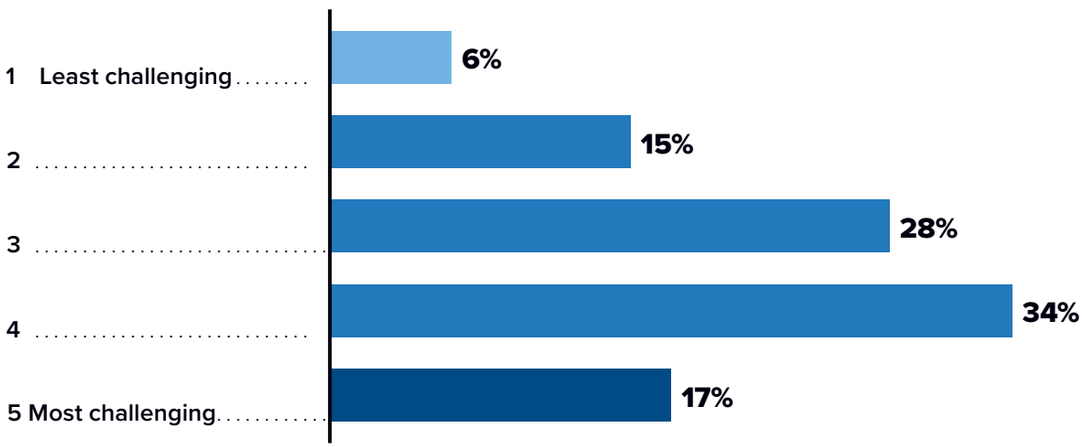

The effective integration of AI is contingent upon addressing the data challenges that come with such projects. Overcoming issues related to data quality and lineage, privacy, scalability, integration, and bias is essential for enterprises to realize the full benefits of AI.

Knowing that data is fair, explainable, and transparent is key to an organization trusting its AI/ML results.

The term **"AI governance"** has been used to describe the application of rules, processes, and responsibilities to drive maximum value from automated data products by ensuring applicable, streamlined, and ethical AI practices that mitigate risk and protect privacy.

n = 2,053. Source: IDC's *StrategiesView,* May 2022

Four Stages of AI Governance

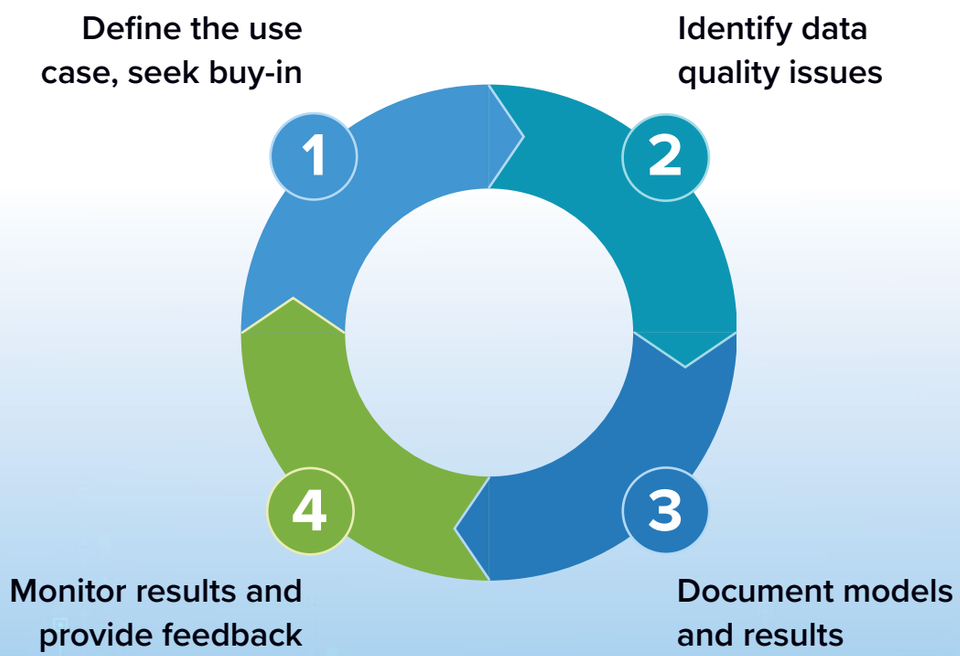

To address AI challenges, enterprises must establish AI governance frameworks, invest in data cleansing and normalization, and ensure quality assurance throughout the AI model life cycle (which also drives a need for data quality assurance). **Regular audits, lineage tracking, and data stewardship programs can help maintain high-quality and consistent outcomes.**

Invest in AI Governance to Deliver Sustainable Value From AI

AI governance is crucial so that organizations can ensure their AI work delivers trusted results and manage the risks that are inherent in AI implementation. **An AI governance framework brings five key benefits to organizations wanting to drive sustainable value from AI at scale:**

Quickly find, understand, and trust data

With a well-structured AI governance framework, data professionals can easily locate the data they need, understand its context, and have confidence in its quality, including when feeding AI models.

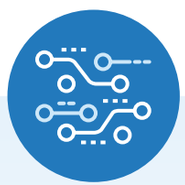

Drive a common language

AI governance facilitates the creation of a common language around the lifecycles and usage of AI models — and the data that feeds those models — within an organization.

Leverage automation

By integrating AI governance with automation tools, organizations can minimize the time and resources required to manage AI implementation lifecycles and keep pace with the rapidly evolving AI landscape.

Mobilize workforce collaboration

An effective AI governance framework encourages collaboration among various stakeholders, including data/AI professionals, business users, and IT teams.

Ensure compliance and mitigate risks

AI governance helps organizations meet regulatory requirements, such as data privacy regulations (e.g., GDPR) and industry-specific standards.

Key Takeaways

1

Generative AI is Ushering in an Era of "AI Everywhere"

AI use in business is already widespread — but the incredible momentum of GenAI interest and adoption will push its use more deeply into every business function, across industries.

2

AI Risks Cannot Be Ignored

As AI begins to deliver strategic business value, and as AI technologies become more complex, the legal, reputational, and financial risks that AI can bring become more urgent to address.

3

AI Governance Must Center on Data Governance

Effective and trustworthy AI implementations can only be built from foundations of well-governed enterprise data — where data, and data quality, are tracked and controlled through the whole data lifecycle.

Appendix: Supplemental Data

The table in this appendix provides an accessible version of the data for the complex figure in this document. Click "Return to original figure" below this table to get back to the original data figure.

SUPPLEMENTAL DATA FROM PAGE 5

Implementation Timeframe for Major Areas

	Focus on customer facing applications	Focus on financial and operational applications	Focus on digital marketing tools
Already begun integration	16%	19%	15%
Less than 6 months	7%	7%	11%
In 6 to 12 months	24%	24%	23%
In 13 to 18 months	44%	41%	40%
More than 18 months	8%	7%	10%
Mean (in months)	**13.17 months**	**12.97 months**	**12.81 months**

Note: Data weighted by IT spend. n = 1,363; Source: IDC's *GenAI Awareness, Readiness and Commitment Survey*, August 2023

Return to original figure

About the IDC Analysts

Jason Stamper
Research Manager, Software,
Data, & Analytics, IDC Europe

Jason Stamper is a research manager working on data management, integration, and analytics topics for IDC in Europe. Over his 30 years in IT, Jason has also met regularly with the most influential people in the industry. He has interviewed the CEOs of hundreds of technology companies, including the likes of IBM, Oracle, HP, EMC, SAP, SAS, and Salesforce. Jason made his transition to the analyst world when he joined the 451 Group in 2013, where he covered data and analytics platforms, with an emphasis on IoT.

More about Jason Stamper

Neil Ward-Dutton
Vice President Automation, Analytics
and AI, IDC Europe

Neil Ward-Dutton is Vice President of Automation, Analytics, and AI at IDC Europe. In this role, he guides IDC's research agendas and helps enterprise and technology vendor clients alike make sense of the opportunities and challenges across these very fast-moving and complicated technology markets. In a 28-year career as a technology industry analyst, Neil has researched a wide range of enterprise software technologies, authored hundreds of reports, and regularly appeared on TV and in print media.

More about Neil Ward-Dutton

IDC Custom Solutions

IDC Research, Inc.
140 Kendrick Street, Building B, Needham, MA 02494, USA
T +1 508 872 8200

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets. With more than 1,300 analysts worldwide, IDC offers global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries. IDC's analysis and insight helps IT professionals, business executives, and the investment community to make fact-based technology decisions and to achieve their key business objectives.