

Ethics can't be delegated

Generative AI is unlike any technology that has come before. It's swiftly disrupting business and society, forcing leaders to rethink their assumptions, plans, and strategies in real time.

To help CEOs stay on top of the fast-shifting changes, the IBM Institute for Business Value (IBM IBV) is releasing a series of targeted, research-backed guides to generative AI, on topics from data cybersecurity to tech investment strategy to customer experience.

This is part eight: Responsible AI & ethics.



Human values are at the heart of responsible AI

As companies race to discover all the incredible new things generative AI *can* do, CEOs must lead the conversation about what it *should* do.

Each use case comes with its own ethical dilemmas and compliance concerns: How can companies protect sensitive data? How can they use AI in a manner that respects copyrights? Are AI outputs biased, discriminatory—or just plain wrong?

While answering these questions takes the entire team, CEOs must set the organization's moral compass. The course they chart will define how the business will balance cutting-edge innovation with the age-old principles of integrity and trust.

CEOs must implement policies and processes that provide transparency and accountability across the board, offering clarity on how and where technology is being used, as well as the source of the data sets or underlying foundation models. This work will be ongoing, as the organization will need to continuously monitor and evaluate their AI portfolios to ensure they remain in line as policies and processes evolve.

Leaders must foster a culture focused on AI ethics, which aims to optimize AI's beneficial impact while reducing risks and adverse outcomes for all stakeholders in a way that prioritizes human agency and well-being, as well as environmental sustainability. This will be a socio-technical challenge that can't be solved with technology alone. Ongoing investments in organizational culture, workflows, and frameworks are necessary to be successful at scale.

The court of public opinion will judge whether companies are behaving ethically—and in line with consumer values. In this way, fairness and appropriateness will be gauged subjectively. But compliance will not.

The IBM Institute for Business Value has identified three things every leader needs to know:

1. CEOs can't pass the buck on AI ethics.



2. Customers are judging every decision you make. Don't jeopardize their trust.



3. Some companies freeze in the headlights of regulatory ambiguity.



And three things every leader needs to do right now:

1. Give ethics teams a seat at the table—not an unfunded mandate.



2. Earn trust by aligning with customer expectations.



3. Bake in ethics and regulatory preparedness for all AI and data investments.



1. Strategy + Generative AI

What you need
to know



CEOs can't pass the buck on AI ethics

Generative AI is like the wild west. The rush to riches has outpaced rules and regulations—and early prospectors have the chance to strike it rich.

But at what cost? Organizations that push forward without considering the intricacies of AI ethics and data integrity risk damaging their reputation for short-term gains.

Executives understand what's at stake: 58% believe that major ethical risks abound with the adoption of generative AI, which would be very difficult to manage without new, or at least more mature, governance structures. Yet, many are struggling to turn principles into practice. While 79% of executives say AI ethics is important to their enterprise-wide AI approach, less than 25% have operationalized common principles of AI ethics.

That's why CEOs must take the reins and blaze a trail for others to follow. Roughly 3x more executives look directly to CEOs for guidance on AI ethics than the Board of

Directors, General Counsel, Privacy Officers, or Risk and Compliance Officers. And 80% of executives say business leaders—not technology leaders—should be primarily accountable for AI ethics.

That accountability extends beyond decision-making. CEOs must also hold themselves responsible for educating other leaders on emerging ethics issues. By elevating conversations about trustworthy AI to the rest of the C-suite and the Board of Directors, CEOs can ensure these key stakeholders aren't sidelined. Taking a proactive, inclusive approach helps ensure everyone understands the risks—and the clear action plan for managing them. This lets the organization move faster while keeping leaders in lockstep.

While **79% of executives** say AI ethics is important to their enterprise-wide AI approach,

less than 25% have operationalized common principles of AI ethics.



1. Spending + Generative AI

What you need
to do



Give ethics teams a seat at the table— not an unfunded mandate

Roll up your sleeves to close the gap between intentions and actions. Champion ethics teams, policies, and monitoring. Report progress to the Board of Directors and externally as appropriate.

Take charge, even if it's outside your comfort zone. Consider appointing a Chief AI Ethics Officer or another leader to run point on enterprise-wide efforts, and make accountability clear among current executive roles. Align executives to common AI ethics goals across business units and functions. Make sure the right people come to the table—including your risk and information security execs.

Create effective human + technology collaborations. Set the tone and strike the right balance between automation and augmentation. Recommend that a design guide for AI is created and adopted, and a specific section on algorithmic accountability be incorporated into the company's code of business ethics. Promote the training, AI and data literacy, and change management agenda. Treat impacted employees with dignity and respect.

Establish “ethical interoperability.” Augment your innovation ecosystem by identifying and engaging key AI-focused technology partners, academics, startups, and other business partners. Affirm values as part of your corporate identity and culture—and make sure the values of all your partners are aligned.

2. Trust + Generative AI

What you need
to know



Customers are judging every decision you make. Don't jeopardize their trust.

It takes decades to build a blue-chip brand—and only days to destroy it. In an era of data breaches and distrust, consumers, employees, and partners are unforgiving of companies that act without integrity.

More than half (57%) of consumers say they are uncomfortable with how companies use their personal or business information and 37% have switched brands to protect their privacy. And consumers rank companies in many traditional industries, including retail, insurance, and utilities, lowest in responsible use of technology.

Partners, investors, and boards of directors are also watching companies closely, though they appear inclined to support responsible AI advancement. CEOs say they feel over 6x more pressure from their boards to accelerate generative AI adoption rather than to slow it down.

Employees are eager to work for companies that share their values, as well. 69% of workers say they would be more willing to accept a job offer from an organization they consider to be socially responsible and 45% say they would be more willing to accept a lower salary to do work at such an organization.

Taken together, these perspectives showcase why companies with stronger data practices create more value. Our *2023 Chief Data Officer (CDO) Study* found that roughly 8 out of 10 CDOs from these companies say their organization outperforms in data ethics, organizational transparency and accountability, or cybersecurity.

45%

of employees say they would be more willing to accept a lower salary to do work at a **socially responsible organization**.



2. Trust + Generative AI

What you need
to do



Earn trust by aligning with customer expectations

Ensure AI ethics and data integrity are an organization-wide priority from the top down. Achieve greater focus by building a collaborative culture of trust from the bottom up. Make ethics everyone's responsibility—and governance a collective noun.

Stay ahead of customers' ethical expectations. Recognize that your customers experience ethical failures in every part of their lives every day. Build trust by defining your ethical values clearly. Communicate them widely and transparently. Then communicate them again. And again.

Put people first. Re-skill your employee base to understand AI and the proper and improper use of it. Build AI ethics and bias identification training programs for employees and partners to reinforce the importance of trustworthy AI. Clarify when to get help from experts. Empower your teams to be stewards of ethics across and beyond your organization to cultivate customer trust.

Hold everyone accountable. Take personal responsibility with an expectation that executives and other employees will take personal responsibility as well. Ask business and AI leaders to sign their names and put their individual reputations on the line—starting with yourself. Prioritize technology ethics as a key part of procurement's ethical sourcing criteria. Make these promises public.

3. Compliance + Generative AI

What you need to know



Some companies freeze in the headlights of regulatory ambiguity

The EU's AI Act is imminent. China is moving ahead with robust regulations and guidelines. And business leaders around the world are feeling pressure to prepare—but they're not exactly sure what they're preparing for.

Globally, fewer than 60% of executives think their organizations are prepared for AI regulation—and 69% expect a regulatory fine due to generative AI adoption. In the face of this uncertainty and risk, CEOs are pumping the brakes. More than half (56%) are delaying major investments until they have clarity on AI standards and regulations. Overall, 72% of organizations will forgo generative AI benefits due to ethical concerns.


While some remain stuck in this regulatory quagmire, companies that proactively address ethical concerns can press on with confidence. Good data and AI governance is

necessary no matter how regulations evolve—and implementing responsible and trustworthy AI from the start will help them achieve compliance when the time comes. Plus, those with strong ethics and governance capabilities have a chance to stand out from the crowd, with three in four executives citing ethics as a source of competitive differentiation.

Prioritizing ethics can help CEOs act decisively in the face of regulatory ambiguity and embrace the early benefits of generative AI without compromising values. This may be why investments in AI ethics are on the rise—tripling from 3% of AI spend in 2018 to nearly 9% in 2025.

72%

of executives say their organizations will forgo generative AI benefits due to ethical concerns.



3. Compliance + Generative AI

What you need
to do



Bake in ethics and regulatory preparedness for all AI and data investments

The broad strokes of emerging regulations are clear enough to guide CEOs' hands. You can course-correct and recalibrate once the details are finalized. Regardless of how it plays out, stay focused on trustworthy AI. Good governance is essential.

Communicate, communicate, communicate. Advocate for regulation that makes sense. Make sure use cases are easily explainable, that AI-generated artifacts are clearly identified, and that AI training is transparent and open to continual critique.

Document everything—and then some. To help manage risk, develop a culture of documentation of AI use in the organization and the current governance around it. Ensure AI-generated assets can be traced back to the foundation model, data set, prompt, or other inputs by creating an inventory of every instance where AI is being used. Seed this source information into digital asset management and other systems.

Steer the ship while it's moving. Be prepared to make adjustments as the regulatory winds shift—or new ones blow.

Responsible AI & ethics

The statistics informing the insights on this page are sourced from proprietary IBM Institute for Business Value surveys conducted in collaboration with Oxford Economics and MomentiveAI/SurveyMonkey, published IBM Institute for Business Value studies, and one external reference from Cisco. The first survey was asked to 200 US-based executives regarding generative AI and AI ethics and was fielded in August–September 2023. A second survey was asked to 300 US-based executives regarding the intersection of generative AI and talent and was fielded in May 2023. A third survey was asked to 16,349 consumers in 10 countries regarding opinions on social responsibility and sustainability and was fielded in 2021. Published IBM Institute for Business Value studies referenced include: the *2023 CEO Study: Decision-making in the age of AI*, the *2023 CDO Study: Turning data into value*, and *AI ethics in action* (2022).



© Copyright IBM Corporation 2023

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America | October 2023

IBM, the IBM logo, ibm.com and Watson are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at: ibm.com/legal/copytrade.shtml.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

This report is intended for general guidance only. It is not intended to be a substitute for detailed research or the exercise of professional judgment. IBM shall not be responsible for any loss whatsoever sustained by any organization or person who relies on this publication.

The data used in this report may be derived from third-party sources and IBM does not independently verify, validate or audit such data. The results from the use of such data are provided on an “as is” basis and IBM makes no representations or warranties, express or implied.

6JAM3VB7-USEN-00