



# Gartner Business Quarterly

Third Quarter 2023

# AI

HERE & NOW



» Balance the Environmental Perils  
and Promises of Generative AI

» 4 Decisions to Make When  
Creating a Generative AI Policy



# AI Here and Now

## Letter From the Editor

Enterprise leaders cannot ignore generative AI because their rivals won't. As the hype swirls — will it save untold hours of work? will it end life as we know it? — only one thing is clear: It won't go away.

Nearly half of more than 2,500 executives we polled said they have been planning to spend more on AI of all kinds since ChatGPT rolled out in late 2022. And separately, dozens of executives have told us in conversations that they want to move faster than they did on past AI initiatives and get something into production within a year.

Their biggest question: What steps must we take over the next 12 months so we don't have to play catchup for years to come — while protecting against pitfalls known and unknown?

This issue of *Gartner Business Quarterly* will help you act now. It's packed with use cases along with guidance for considering investments, enticing customers, and rethinking your talent strategy. Just as importantly, you'll find analysis of very real dangers ahead, some that have hit the headlines and some that have not.

Our article on balancing the environmental impact of GenAI encapsulates our practical perspective. Dream big: use these applications to suggest sustainable materials, markets and design. Minimize the negative: Consider so-called composite AI, which uses knowledge graphs and causal networks for optimal efficiency — tamping down the models' voracious appetite for energy. And buy clean power where you can.

Company tactics come from ABB and HP Inc., while important context emerges in our Q&A interviews with EU lawmaker Dragos Tudorache and decision-making experts Daniel Kahneman and Olivier Sibony.

GBQ helps you and your team align with others and reach peak effectiveness, so your enterprise can achieve its goals, be bold and principled, and bring employees, investors and the public along for the ride.

Our standing departments keep you up to speed — Cutting Edge is a look at provocative new data; Briefs offer short takes about smarter spending and planning, talent and culture, growth and innovation, and data and technology.

We welcome your feedback. Please contact me at [judy.pasternak@gartner.com](mailto:judy.pasternak@gartner.com).

— Judy Pasternak

# Credits

**Research & Advisory** Val Sribar, Executive VP

**Executive Sponsor** Scott Christofferson, Group VP

**Global Chief of Research** Chris Howard,  
Distinguished VP Analyst

**Editor in Chief** Judy Pasternak

**Managing Editor** Richard Eames

**Associate Editor** Kelly Fiveash

**Contributing Editors** Dian Zhang, Jonah Shepp

**Project Manager** Carolina Abalos

## R&A Studio

**Creative Manager** Steve Hennessey

**Senior Designers** Nicole Daniels, Divya Malkani

**Graphic Designer** Tal Yaari

**Associate Graphic Designer** Akansha Garg

**Senior Editor** Caroline Baker

## Contributors

Amisha Ajay

Gene Alvarez

Vaughan Archer

Van Baker

Alexandra Bellis

Kathon Betterton

Erick Brethenoux

Frank Buytendijk

Sharon Cantor Ceurvorst

Karen Carter

Uma Challa

Arun Chandrasekaran

Jian Chen

Laura Cohn

Jeff Cribbs

James Crocker

Meg Day

Michelle DeClue

Stephen Emmott

Pieter den Hamer

Jeremy D’Hoinne

Ellen Eichhorn

Dan Essig

Mike Fang

Peter Firstbrook

Ben Fisher

Ron Friedmann

Tsuneo Fujiwara

Rupeesha Galhotra

Dennis P. Gannon

Dan Gottlieb

Dan Gutter

Sonu Gupta

Irina Guseva

Ilona Hansen

Robert Hetu

Allysa Johnson

Roma Kaur

Rajesh Kandaswamy

Frances Karamouzis

Lauren Kornutick

Angela Kreiter

Eliza Krigman

Mike Lashinsky

Leah Reidy Leachman

Avivah Litan

Kaelyn Lowmaster

Abbott Martin

Kate McLaren-Poole

Simon Mingay

Kristin Moyer

Anthony Mullen

Grant Faulkner Nelson

Kanwarpreet Singh Oberoi

Cian Ó Móráin

Pedro Pacheco

Iga Pilewska

Helen Poitevin

Amrita Puniani

Mark Raskino

Gabriele Rigon

Eser Rizaoglu

Jose Rosario

Melissa Ruiz Hernandez

Moutusi Sau

Don Scheibenreif

Tom Scholtz

Steve Shapiro

Tony Sheehan

Sandy Shen

Matthew Shinkman

Jasleen Kaur Sindhu

Gurbani Singh

Shivendra Singh

Lauren Smith

Marco Steecker

Stuart Strome

Anushri Subramaniam

Noha Tohamy

Bettina Tratz-Ryan

Kayla Velnoskey

Sarah Watt

Lauren Wheatley

Madison Wilson

Shubhangi Yadav

# Contents

## » Feature Articles

**10** Balance the Environmental Perils and Promises of Generative AI

**17** The Whiteboard: Big Questions About How to Invest in Generative AI

**26** 4 Decisions to Make When Creating a Generative AI Policy

**31** Prepare for AI Regulation by Addressing 4 Critical Areas

**38** AI Systems 'Need a Moral Compass'  
A Q&A With Dragos Tudorache,  
Co-Architect of the EU's AI Act

**43** Generative AI Will Affect Information Security (and CISOs)

**48** How Generative AI Can Help Meet Customer Experience Expectations

**53** Learn to Serve Your AI-Powered Customers Before They 'Walk Away'

**59** What Generative AI Means for Your Talent Strategy

**64** Can AI Overcome Flaws in Human Decision Making?  
A Q&A With Daniel Kahneman and Olivier Sibony

**68** ABB Electrification Lights Up ML Models for Financial Planning

## » Departments

**05** The Cutting Edge 3Q23: Cool New Data Points

**72** Briefs: Quick Takes on Fresh Research

» Department

# The Cutting Edge 3Q23

Cool New Data Points

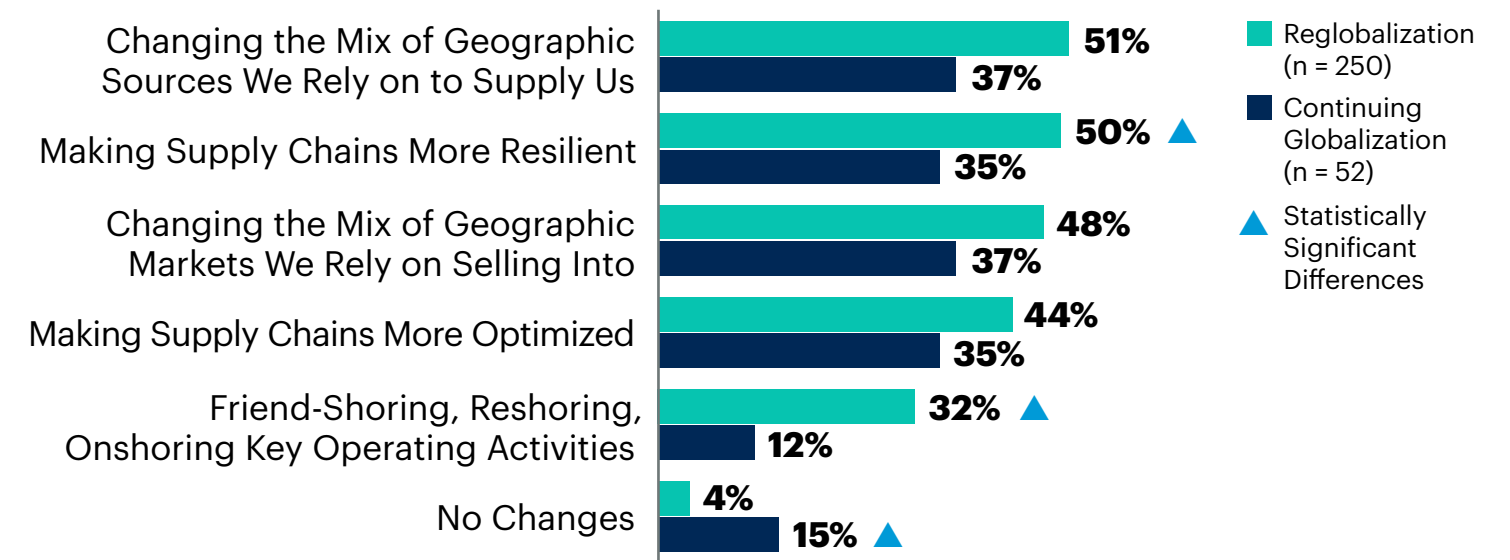
Compiled by Laura Cohn and Dian Zhang

## Power Shifts Among Nations Prompt CEOs to Adapt Strategic Priorities

The actions they are taking depend on whether they believe global trade is continuing much as it has or whether they see a reconfiguration.

### Response Strategies by CEOs' Views of Globalization Scenarios

Percentage of Respondents (Multiple Responses Allowed)



n varies

Q. What kind of changes, if any, are you making in response to continuing globalization/reglobalization/slowbalization or deglobalization?

Source: 2023 Gartner CEO and Senior Business Executive Survey

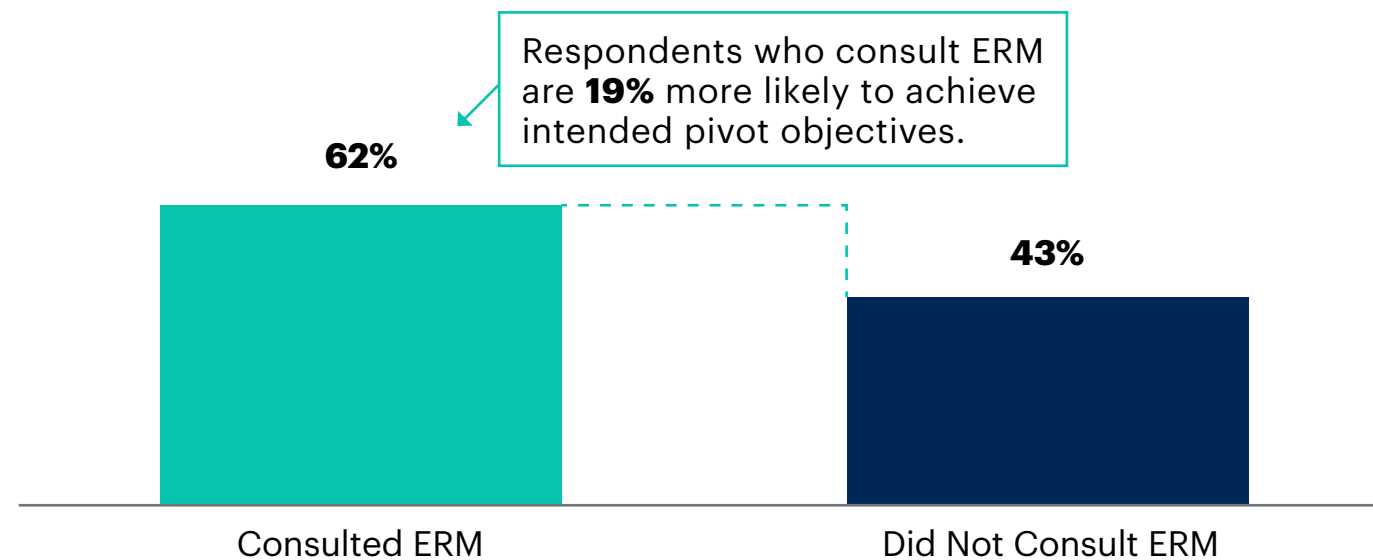
Note: Respondents could choose from the following scenarios on globalization: globalization is continuing but changing, as power shifts between major nations (reglobalization); globalization is slowing down (slowbalization); globalization is continuing much the same as it has for decades; and globalization is reversing (deglobalization).

**When Making a Strategic Pivot, Involving ERM Boosts the Chance of Success**

Enterprise risk management teams can discuss relevant risk information, run scenario planning workshops, and map out and review the disruption response plans.

**Impact of Consulting ERM on Strategic Pivot Outcomes**

Percentage of Respondents Selecting



n = 250 respondents

Q: What actions did you take that involved the risk and assurance functions: I consulted my organization's ERM team. State your agreement with the following statements about the pivot, in retrospect: The pivot fully achieved its intended objectives

Source: 2021 Gartner Resilience Survey


Note: Correlation is significant at the 0.01 level

**Despite Reports of ESG Backlash, Business Leaders Have Increased Commitments**

Almost three-quarters (72%) say they are building out their programs.

**Business Leaders Who Report Increased ESG Commitments Across the Past 12 Months**

By Percentage



n = 206

Source: Gartner 2023 ESG Goal-Setting Survey

**GC-Cited Examples of Increasing ESG Commitments at their Organizations**

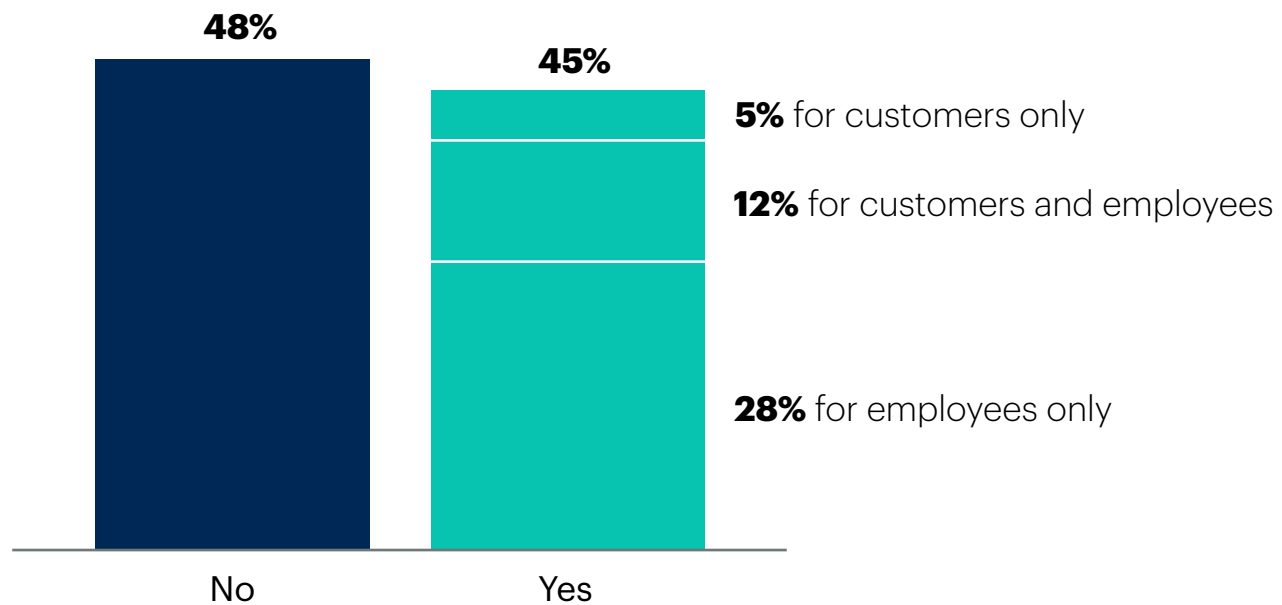
- Bringing on dedicated FTEs in ESG-specific roles
- Increased ESG funding and/or establishing a dedicated ESG budget
- Establishing ESG committees and/or working groups
- Issuing the organization's first ESG report
- Investing in Board ESG education
- Creating internal ESG intranet sites to communicate and solicit feedback

Source: Gartner

### Nearly Half of Organizations Use Biometric Verification

Companies are turning to biometric verification just as more jurisdictions plan to regulate it. To manage this risk, legal and privacy should work with information security and the functions using biometrics to check how data will be stored and protected, and how consent will be obtained.

#### Does Your Organization Use Biometric Verification?



n = 58

Source: Poll from Privacy Strategy and Operations: Improving Privacy Metrics Webinar, 16 February 2023

Note: 7% Don't Know


### Even Top Performers Struggle With Feelings of Futility About Doing 'Enough'

Your top people need performance confirmation, not just conversation. Make sure you reward and recognize great achievements in real time to keep employees motivated.

#### Prevalence of High Futility by Recent Performance Rating

Percentage of Employees Experiencing High Futility

- Futility**
- "I feel anxious about whether I am performing"
  - "Reaching high performance feels hopeless"
  - "Someone will always do more than me"



n = 2,118 remote-capable employees who received a recent performance rating

Source: 2023 Gartner Achieving High Performance Employee Survey

**Only Half of Boards Are Effective — But the GC Can Help**

While boards face a host of challenges — whether it’s increased regulatory action, heightened cybersecurity risk, or rising geopolitical tensions — there’s clearly room for improvement.

**Three Steps to Improve Board Operations**

When general counsel take steps to make the board more effective, they play a big role in helping companies navigate an ever-shifting web of economic, social and geopolitical pressures.

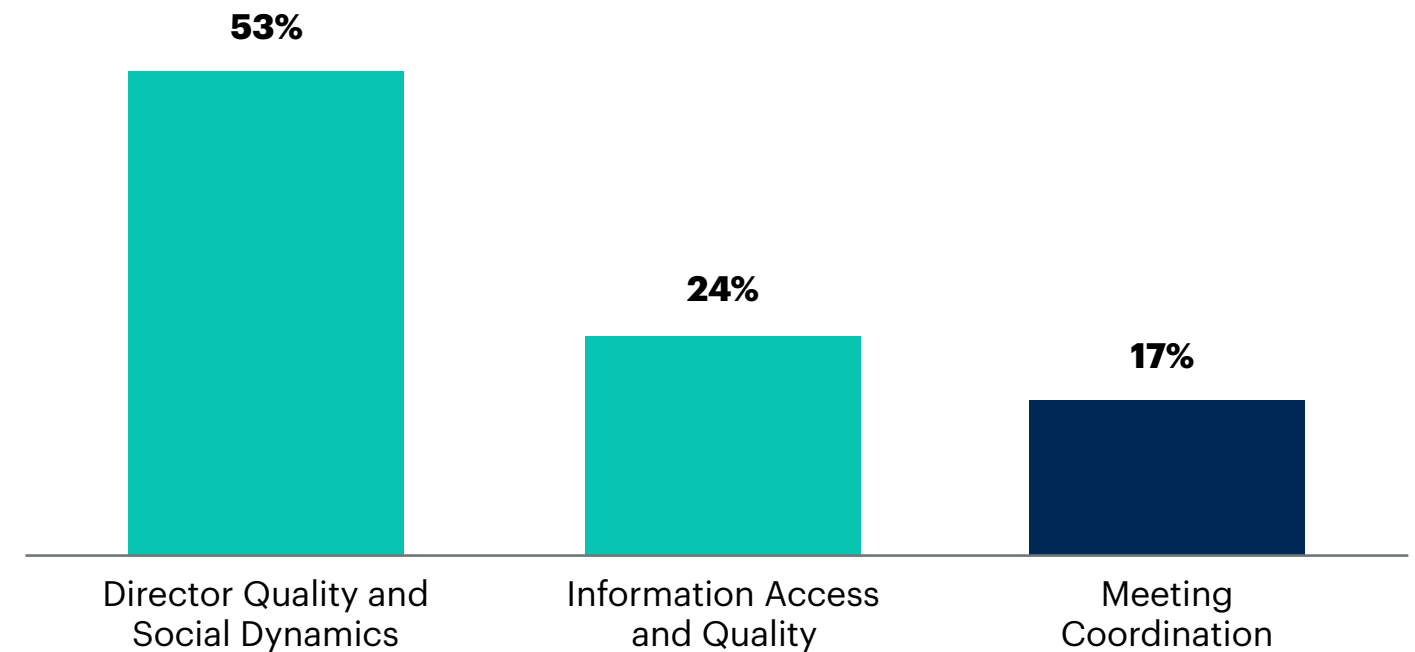
**Percent of Organizations With Effective Boards**

Percent of Respondents Scoring an Average ≥6 on the 7-Point Board Effectiveness Index<sup>a</sup>



n = 92  
 Source: 2023 Gartner Corporate Governance and Board Management Benchmark Survey  
<sup>a</sup>Board Effectiveness Index combines Executive Accountability and Board Oversight of Strategy and Risk.

**Maximum Impact of Each Factor on Board Effectiveness**



n = 92  
 R<sup>2</sup> = 55%  
 Source: 2023 Gartner Corporate Governance and Board Management Diagnostic Survey



# Build a strategic plan that works — even in volatile times!

Only 29% of strategists say their organizations change plans fast enough to respond to disruption.

In an increasingly volatile and uncertain world, strategy and its execution can quickly go off course. Use the Gartner one-page Strategic Planning template to:



**Set clear, strategic objectives** for your organization.



**Combat seven mistakes** common to strategic planning.



**Capture and communicate** your plans with an **exclusive one-page template**.

**Download the Strategic Planning guide for your function.**


Source: Gartner  
© 2023 Gartner, Inc. and/or its affiliates. All rights reserved. CM\_GBS\_2412262

**Gartner**

# Balance the Environmental Perils and Promises of Generative AI

by Kristin Moyer, Sarah Watt and Pieter den Hamer

Generative AI consumes a lot of electricity and water. The technology can accelerate positive sustainability and financial outcomes, but executive leaders need to be selective about use cases to make sure they are doing more good than harm.



The rapid adoption of ChatGPT and similar general-purpose tools has elevated the negative environmental impacts of generative AI from an insider topic to an immediate concern for executive leaders. Plenty of seemingly sustainable use cases powered by this emerging technology will do more harm than good in terms of greenhouse gas (GHG) emissions and electricity and water consumption. But generative AI is not automatically or entirely bad news for sustainability.

For example, Orbital Materials, a UK-based startup, is seeking to apply it to accelerate the development of clean energy technologies.<sup>1</sup>

*If used in the right way and with human oversight,* generative AI can potentially help companies mitigate sustainability risk, optimize costs and drive growth. Executive leaders should balance the perils and promises of this technology by simultaneously:

- Recognizing and reducing generative AI's energy footprint
- Identifying, evaluating and prioritizing its environmental sustainability use cases

## Recognize That Tech Has a Consumption Problem

Generative AI relies on huge models trained with massive amounts of data, making it thirsty for cooling water and hungry for electricity. In some cases, the technology can consume vast quantities of both. Electricity-related GHG emissions will decline in the long term, as the generation capacity of renewable energy grows. However, more powerful generative AI models will likely require increased computing capabilities.

The problem of tech-related electricity and water consumption goes well beyond generative AI. Information and communications technology (ICT) consumes less than 1% of global electricity today, but it is forecast to use 6.4% of total output by the end of the decade.<sup>2,3,4</sup> If new power-hungry technologies scale quickly, ICT energy consumption could be even higher.

By 2030, 75% of CIOs will experience electricity constraints.<sup>5</sup> For example, Singapore “implicitly imposed” a moratorium on building new data centers in 2019 (but lifted it in 2022).<sup>6,7</sup> The Dutch government tightened rules for hyperscale data centers in 2022.<sup>8</sup> As the needs of technology and society compete more intensely with each other, CIOs do not want to be caught in the position of vying with local communities for finite resources.

## Reduce Generative AI's Energy Footprint

Executive leaders can make their organization's generative AI more environmentally friendly by leading efforts to:

- **Make it as efficient as the human brain.** One of the reasons the brain is so energy-efficient is that it organizes knowledge in network structures. Consider adopting so-called composite AI, which uses similar techniques such as knowledge graphs, causal networks and other symbolic representations to complement the current brute force, deep learning method.
- **Put it on an electricity and water diet.** Monitor energy consumption during machine learning. Improve awareness across the company about AI's environmental footprint. Stop training AI as soon as improvements flatten out and no longer justify the costs of continuing. Keep data for model training local, but share improvements at a central level; this federated machine learning reduces electricity consumption and bolsters data privacy. Reuse models that have already been trained, contextualizing them if needed. Use more energy-efficient hardware and networking equipment. Balance “follow the sun” data center workloads (which are better for clean energy production) with “unfollow the sun” measures (which are better for water efficiency).



- **Run it in the right place and at the right time.** When and where a generative AI workload happens matters. The carbon intensity of the local energy supply varies by country, generating authority, time of day, weather conditions, transfer agreements, fuel supply and other factors. Best practice is to use energy-aware job scheduling for generative AI, along with carbon tracking and forecasting services to reduce related emissions.
- **Aim to buy new clean power where you plan to consume it.** Not all renewable electricity purchases are equal. Buying renewable energy certificates (RECs), which reduce or offset GHG emissions, does not guarantee that the clean energy was actually consumed or that any new renewable power generation capacity was built to match increasing demand. Power purchase agreements (PPAs) can help but are not always available. The Greenhouse Gas Protocol, a global standard-setting body, is considering requiring companies to provide more detailed analysis of clean power by location, time of day or both.<sup>9</sup>

Until then, executives should procure PPAs when possible, or source RECs from recently commissioned projects (thus adding new renewable energy to the grid) from the region where their company will consume the electricity. For example, IBM,<sup>10</sup> Walmart<sup>11</sup> and Salesforce<sup>12</sup> have policies to purchase renewable energy generated in the same grid region as they consume it and/or to add new clean power to the system.

## Identify Potential Environmental Sustainability Use Cases

We picked out valuable examples of generative AI use cases that will accelerate environmental sustainability by mitigating risks, optimizing costs or driving growth. We then developed descriptions of each that incorporated some suggestions from ChatGPT. We found that the tool frequently overestimated its abilities but that we underestimated its potential.

### Mitigate Risks

- **Comply with regulations.** Generative AI can identify relevant sustainability laws, standards, directives and reporting requirements, including updates over time. It can interpret complex regulatory language and provide explanations in simpler terms. Generative AI can develop an action plan to achieve compliance, including insights on how to establish internal controls and auditing. It can develop training materials to educate employees on specific regulations.
- **Avoid stranded assets.** Generative AI can identify trends, regulations and technology disruption that increase the risk of stranded assets.
- **Develop organizational sustainability policies.** Generative AI can help employees identify sustainability policy best practices; develop a framework for such policies; and draft statements, guidelines, procedures and protocols.
- **Develop sustainability KPIs.** Generative AI can help identify key performance indicators to evaluate sustainability efforts and progress.

## Optimize Costs

- **Support decision making.** Generative AI can analyze internal sustainability data and identify patterns, trends, areas for improvement, feasibility, risks and benchmarks that support decision making. It can provide insight into how organizational decisions will impact sustainability and forecast likely future performance. Enterprises can therefore plan and select optimal pathways to reach GHG emissions reduction goals.
- **Improve corporate sustainability communications.** Generative AI can generate press releases, website content, articles, social media posts and campaigns, blogs, and external and internal communications on sustainability. It can suggest key messages and best practices for conveying information. ChatGPT can develop frequently asked questions and answers.
- **Engage supply chain partners.** Generative AI can develop criteria for supplier assessment, create a code of conduct, suggest engagement strategies, develop training content, provide suggestions for incentives and corrective actions, and answer questions.
- **Train employees.** Generative AI can develop employee training materials, create sustainability awareness communications, provide insight on sustainable work practices, explain sustainability concepts, answer employee questions and support learning by recommending materials and courses.
- **Embed generative AI in products to make them more sustainable.** Generative AI can suggest sustainable product uses to customers by automating reply generation, using the right response tone and gamifying the customer experience through rewards.

## Drive Growth

- **Discover alternative resources and materials.** Generative AI can provide suggestions for sustainable substitutes for conventional inputs; insights on technological innovation such as nanomaterials; and information on availability, performance and environmental impact.
- **Support sustainable product development.** Generative AI can help with new, sustainable product ideation and insights on its viability; market research on trends; suggestions on prototyping; sustainable design principles; information on patents, trademarks or copyrights; competitive analysis; and pricing advice.
- **Leverage generative design for buildings, parts and products.** Generative AI can suggest alternative design approaches, and parameter and constraint options, as well as provide insight on design strengths and weaknesses.
- **Accelerate sustainability innovation.** Generative AI can support ideation, opportunity assessment, risk and compliance mitigation, and business case development.

## Evaluate the Positive and Negative Impacts

Executive leaders weighing generative AI use cases for sustainability should analyze their:

- **Business value** — Financial and sustainability benefits, as well as feasibility (the level of difficulty and cost to implement the use case)
- **Negative environmental impact** — As measured by GHG emissions and electricity and water consumption

We evaluated sustainability use cases, employing general-purpose generative AI tools (such as ChatGPT) based on a combination of primary and secondary research, proprietary data and analyst experience (see Figure 1).

» Figure 1. Analysis of Generative AI Use-Case Examples

Higher ● ◐ ◑ ◒ ◓ Lower

		Business Value			Negative Environmental Impact
		Financial Value	Sustainability Value	Feasibility	
Mitigate Risks	Comply with regulations	◐	◑	◑	◑
	Avoid stranded assets	◐	◑	◑	◑
	Develop organizational sustainability policies	◐	◑	◑	◓
	Develop sustainability KPIs	◐	◑	●	◓
Optimize Costs	Support decision making	◑	◑	◑	◑
	Improve corporate sustainability communications	◓	◓	◓	◓
	Engage supply chain partners	◑	●	◑	◑
	Train employees	◐	◑	◑	◑
Drive Growth	Embed GenAI in products to make them more sustainable	◑	◑	◑	◑
	Discover alternative resources and materials	●	●	◑	◑
	Support sustainable product development	◑	●	◓	◑
	Leverage generative design (buildings, parts, products)	●	●	◑	◑
	Accelerate sustainability innovation	◑	◑	◑	◑

Source: Gartner



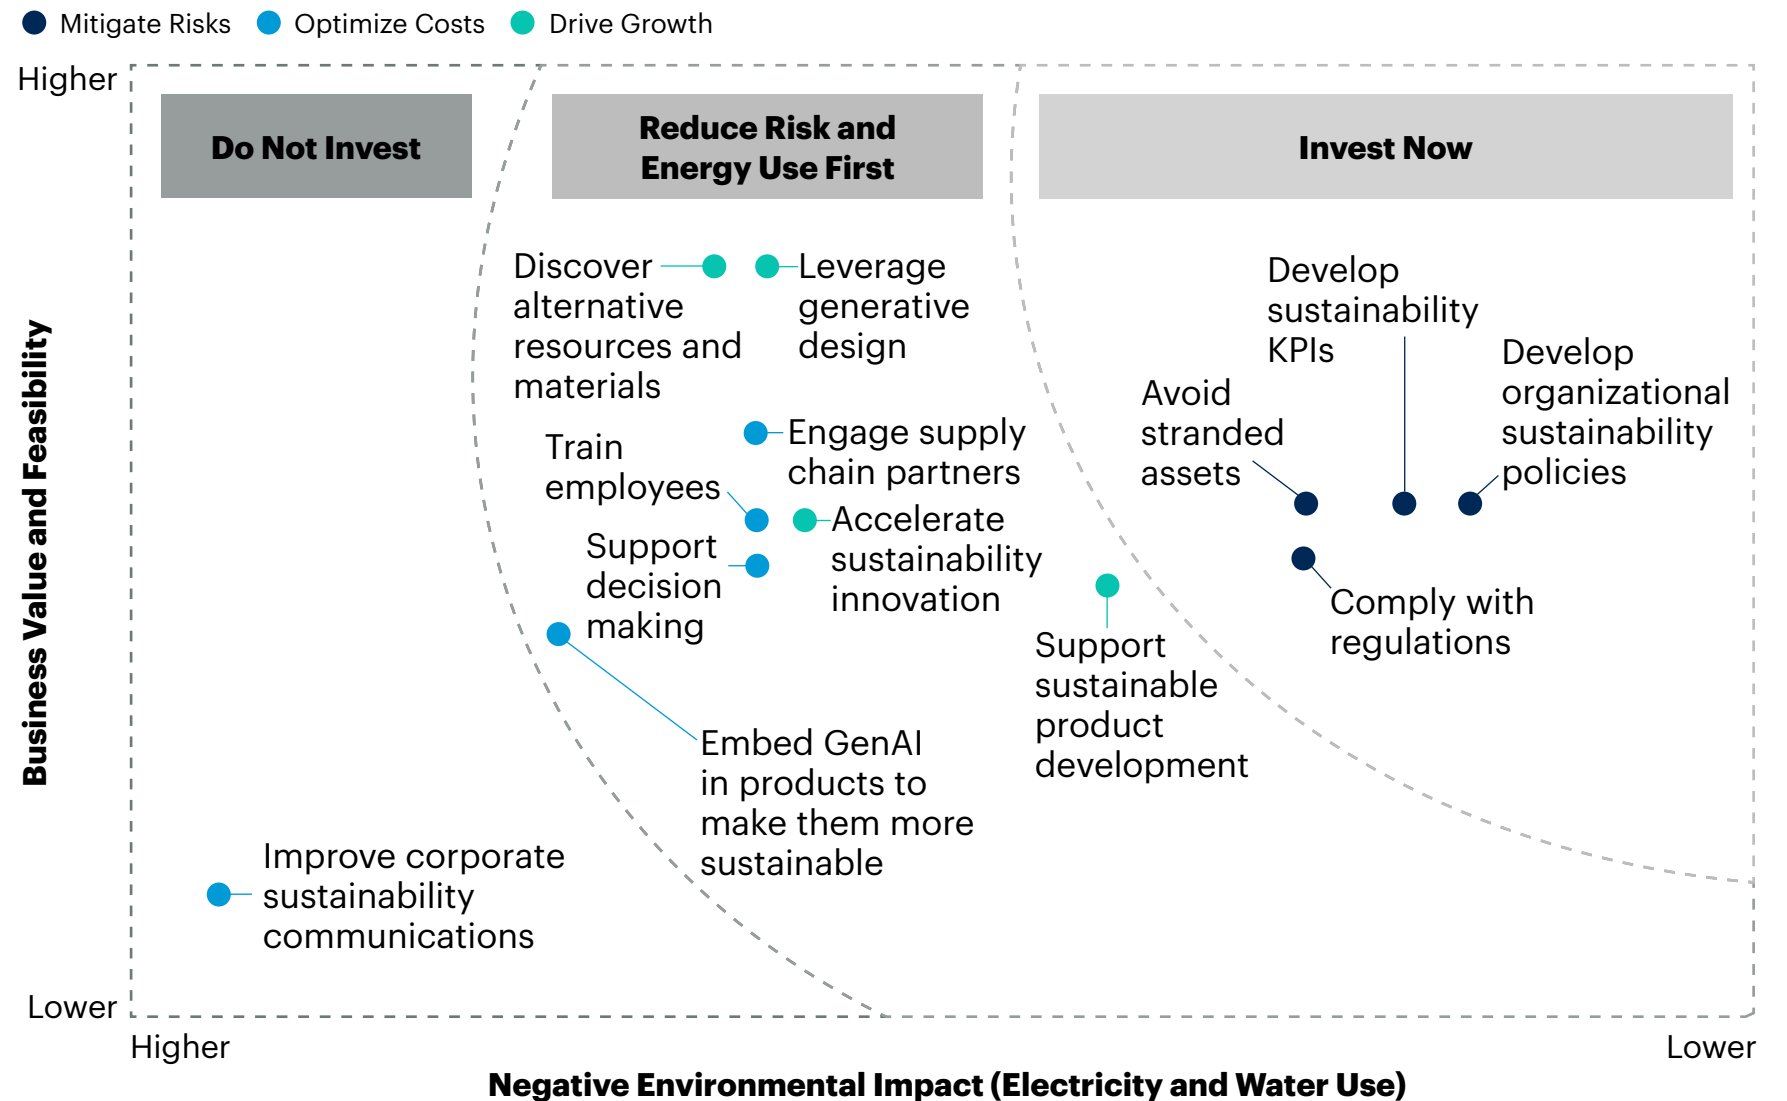
## Prioritize Your Investments

Based on this assessment, executive leaders should categorize generative AI use cases into three levels of priority (see Figure 2):

- **Invest now** — Act immediately to advance use cases that will drive business value with a relatively low negative environmental impact.
- **Reduce risk and energy use first** — Lower intellectual property and proprietary data risk and/or improve the energy efficiency of generative AI before proceeding.
- **Do not invest** — Avoid use cases that could damage business value or the environment.

Generative AI can help organizations accelerate positive sustainability outcomes, provided they leverage only those use cases that create more value than they destroy.

» Figure 2. Prioritization of Generative AI Use-Case Examples



Source: Gartner

<sup>1</sup> DeepMind Alum Wants to Use AI to Speed the Development of Green Materials, Bloomberg.  
<sup>2</sup> The Decadal Plan for Semiconductors, Semiconductor Research Corporation.  
<sup>3</sup> World Energy Outlook 2022 — Stated Policies Scenario, International Energy Agency.  
<sup>4</sup> Maverick Research: Net Zero Will Stall Tech Growth and Innovation.  
<sup>5</sup> Gartner prediction based on Semiconductor Research Corporation's Decadal Plan Update and the International Energy Agency's World Energy Outlook 2022 — Stated Policies Scenario.  
<sup>6</sup> Singapore Hits Pause on Building New Data Centres; Short-Term Rents Up, The Business Times.  
<sup>7</sup> Singapore Authorities Invite Applications for New Data Centers, DCD.  
<sup>8</sup> Dutch Call a Halt to New Massive Data Centres, While Rules Are Worked Out, DutchNews.  
<sup>9</sup> Survey on Need for GHG Protocol Corporate Standards and Guidance Updates, Greenhouse Gas Protocol.  
<sup>10</sup> Energy and Climate, IBM.  
<sup>11</sup> Climate Change, Walmart.  
<sup>12</sup> Inside Salesforce's Clean Energy and Carbon Programs, Salesforce.

Gartner®

# The Future of AI: Reshaping Society

**Can the autonomous decisions made by AI for us, be compared to human agency?**

With recent developments in generative AI, such as ChatGPT, AI is no longer merely a technological or commercial instrument. The age of AI has well and truly arrived, and it is about to transform society as a whole.

**Download the research for more on:**



**What is up with AI:**  
Introduction to some  
fundamental issue




**What is next with AI:**  
Future AI scenarios  
through 2033



**What else with AI:**  
More insights on  
the impacts of AI

[\*\*Download Research ↗\*\*](#)



# The Whiteboard: Big Questions About How to Invest in Generative AI

by Van Baker, Erick Brethenoux and Arun Chandrasekaran

Contributions by Jeff Cribbs, Meg Day, Ellen Eichhorn, Dennis P. Gannon, Dan Gottlieb, Robert Hetu, Grant Faulkner Nelson, Pedro Pacheco, Eser Rizaoglu, Moutusi Sau, Steve Shapiro, Tony Sheehan, Jasleen Kaur Sindhu, Lauren Smith, Marco Steecker, Noha Tohamy and Shubhangi Yadav



Suddenly, executives recognize how easily they can take advantage of the most disruptive set of technologies and capabilities to hit the global market in decades.

But don't get carried away by the hype surrounding generative AI. Concentrate on practicalities such as use cases, cost and vendors.

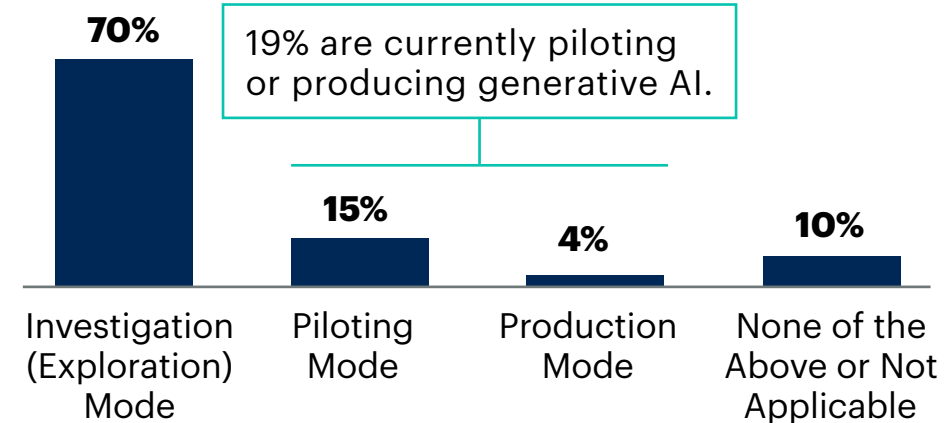
**Let's head to the whiteboard and sketch out how organizations can make smart moves as they pilot and scale their initiatives.** →

## How are other companies advancing with generative AI?

One in five organizations say they're already piloting or producing it **and 70% are in the investigation stage!**

### » Current Phase of Generative AI

Percentage of Respondents



n = 2,544

Source: Beyond the Hype: Enterprise Impact of ChatGPT and Generative AI Webinar Polls; 30 March and 21 April 2023





## What new questions must we ask about possible investments?

### Impact

**Vision:** How might the capabilities created by GenAI impact our strategy?

**Benefits:** What new opportunities and capabilities will it present?

**Objective:** What are we trying to gain by applying it to this problem?

**Adoption:** How will staff react to its use?

### Process & Infrastructure

**Stakeholders:** Who is likeliest to be impacted?

**Processes:** Which ones will we need to transform or create?

**System architecture:** How will we train GenAI to leverage our infrastructure?

**Data:** What data is available to feed into GenAI? How much is there? Does it have sufficient depth (i.e., metadata)?

**Skills:** What skills are required to iterate and collaborate with this tool?

### Governance

#### Protection:

- How will we safeguard proprietary and sensitive data?
- What disclosures will we need to make?

#### Prevention:

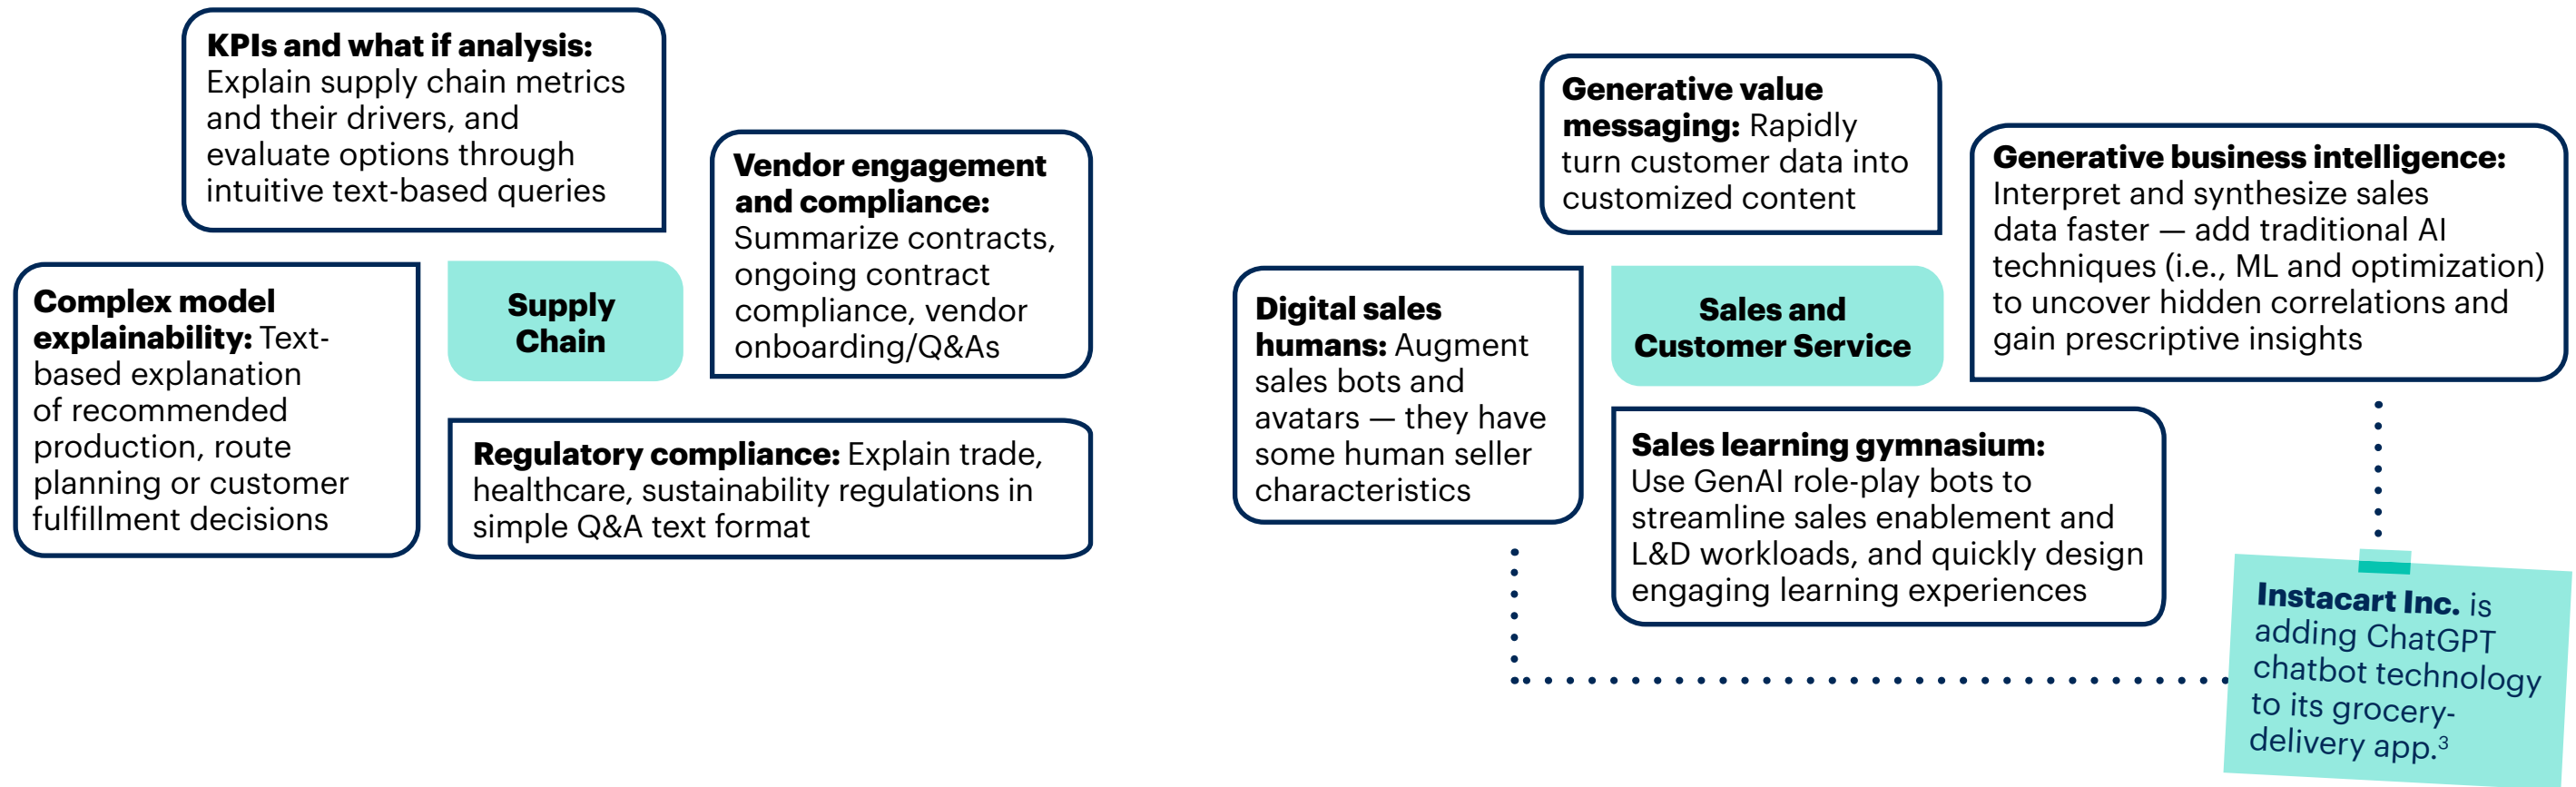
- Who will be accountable for GenAI's outputs, and to what degree?
- What governance structures will we need to detect and respond when it acts in unexpected ways?



Image generated using DALL-E, <https://openai.com/dall-e-2>.



## What are some emerging functional use cases?<sup>2</sup>







## What are some emerging functional use cases?<sup>2</sup>



### HR operations:

- Augment and improve staff productivity in classifying, summarizing and responding to employee queries
- Automated drafting of policies and document generation

### Tailored job descriptions:

Quickly create more compelling and inclusive text to better attract diverse talent with in-demand skills

### HR

**HR chatbots:** Use these tools, with enhanced NLP capabilities from generative AI, to

- handle routine employee requests using human-like text
- provide a better user experience

### Company reporting:

Create initial drafts of management discussion and analysis, financial footnotes, etc., for quarterly/annual filings

### Finance

### Coding assistance:

Translate code from older coding languages, like COBOL, into modern programming languages, like SQL, KnowledgeSQL, and Python

### Management updates:

Explain forecast and budget variances for FP&A to use in business reviews; summarize trends and insights for the board and executives

### Contract and document review:

Scan contracts for errors and specific terms; summarize and categorize documents for sorting, review, retrieval

**Zoom Video Communications** is using generative AI to create first drafts of tax documentation.<sup>4</sup>

### Competitive intelligence:

Summarize public financial disclosures (e.g., competitors' quarterly earnings transcripts) and industry publications to compile reports faster

### Strategy

### Assist in curating agendas and frameworks:

e.g., for annual planning offsites or defining industry ecosystems

*Users should have appropriate security and privacy policies in place for these types of use cases*

### Internal knowledge management:

Ingest unstructured/nonintegrated datasets and artifacts (e.g., business cases, project plans) to organize and easily query expertise

### Planning assistance:

Pre-fill templates (e.g., auto-generate baseline Five Forces and SWOT analysis) to accelerate BU leaders' planning efforts



## What are some possible industry use cases?<sup>2</sup>



### Financial Services

- **AI frontline co-pilot:** Chat interface helps client-facing employees get important information faster
- **Morgan Stanley** is training GPT-4 to help its financial advisors.<sup>5</sup>
- **Compliance and regulatory monitoring:** Assist in verifying communications with clients against internal codes and rules
- **Personalized customer support:** Recommendations for contact center agents and relationship managers based on customer profile, needs and expectations
- **Claims management:** Individualized suggestions/explanations on claims coverage and applicant-friendly reasons for denials

### Healthcare and Life Sciences

- **Conversational patient self — triage and checking symptoms:** Chatbot makes suggestions and guides patients regarding acute symptoms, chronic condition management, health and wellness activities, or behavioral health needs
- **Auto-composition of clinical messages:** Automatic replies based on content and tone of patient message, accessible clinical data, and clinician's tone and preferences
- **Mass General Brigham**, a health care system in the U.S., is testing generative AI for patient portal messages and clinical notes.<sup>6</sup>
- **Scientific literature discovery:** LLMs help scientists identify relevant research, extract insights, aggregate findings and generate new hypotheses
- **Coding assistant for mainframe support:** Helps software developers generate, test, debug code snippets in languages common to mainframe technologies, like COBOL — often used in U.S. healthcare payers' claims processing systems
- **Consultative population health analytics:** Users ask plain language questions of a report or dashboard in areas like population health, costs and care activities

### Education

- **Student tutors:** Conversational UI to support personalized learning
- **Language training:** AI reading and speaking companion
- **Faculty assistant:** Accelerate authoring of quizzes, tests, presentation materials, curricula, lesson plans, feedback, student referral letters
- **Virtual student assistant:** Chat interface to integrated student data
- **Student recruitment/enrollment/persistence:** Including nudging students toward course completion



## What are some possible industry use cases?<sup>2</sup>



### Retail

**Tesco** is using GenAI and other technologies to enhance customer experience, predict demand, analyze consumer behavior and prevent fraud.<sup>7</sup>

- **Enhanced search and upselling:** Improve customers' abilities to find what they are looking for, and encourage more expensive purchases or add-ons
- **Social media customer sentiment:** Quickly monitor customer and influencer social media content, spot trends and sentiments, predict outcomes and inform future decisions

- **Supply chain optimization:** Improve predictions for sourcing and procurement, logistics, transportation, and collaboration with suppliers
- **Conversational chat interface:** Interact with customers and associates, which may include facilitating a transaction -- enable human customers to converse via their platform of choice
- **Associate hiring, onboarding:** Enhance recruiting and training through interactive individual experiences

### Manufacturing

- **Education and training:** Direct an employee with or without relevant technical knowledge to verify a factory-floor machine in their chosen language(s)
- **Product innovation:** Suggest alternative ingredients and packaging based on user sentiment and aggregated trends/shopping patterns
- **Digital product interaction:** Download new behaviors/capabilities to digital products based on aggregated voice feedback
- **Product servicing:** Help humans and AI agents in continuously diagnosing issues; order parts, complete programmable maintenance or schedule recommended servicing needs. (Goal: reduce unplanned downtime)

### Transportation


- **Customer interaction:** Use of LLM chatbots
- **Maersk** is using ChatGPT on its website to auto-generate FAQs and improve search accuracy.<sup>8</sup>
- **Vehicle damage estimation for insurance claims:** Help a smartphone camera recognize damage more precisely even where visibility and contrast are poor
- **Estimation of vehicle resale value:** Use GenAI on computer vision to enable a smartphone camera to assess value more accurately
- **Assessment of mechanical condition:** Enable more precise evaluations





## How much will generative AI cost?

From almost nothing to many millions of dollars, depending on:



**Small and midsize enterprises ...**  
 ... could use the free versions of public, **openly hosted applications such as OpenAI** or **pay low subscription fees**

**GenAI capabilities will increasingly be built into everyday software products**

- Effectively a “free,” or close to free, tier with productivity benefits
- Vendors will pass costs on to customers as part of bundled incremental price increases

**Larger organizations ...**

- ... or enterprises that want to analyze or use their own data more
- ... or are looking to invest in significant process or market disruption

**will need higher levels of security, IP and privacy protections ...**  
 ... requiring a range of custom services.

Do we **build our own licensed, customizable and proprietary models** with foundation model platforms that will require working with vendors and partners?

**Gartner Generative AI Framework** *The more you customize, the more it costs.*

Desired Content Type	Non-Sensitive Text	PII/Enterprise IP Included	Enterprise Data and Model Instructions Needed	Model Fine Tuning to Improve Use/Performance	Custom Model is Created for Unique Use Case
Enterprise Control Level	No Controls Needed	LLM With Privacy	LLM With Privacy, Policies and Data	LLM w/Privacy, Policies, Data and Added Model Layers	Enterprise Hosted Custom Differentiated Model
Solution Needed	ChatGPT	LLM API Accessed via Application Frame	LLM API with Standard Privacy, Policy and Per Incident Data Injection	Modified LLM With Transfer Learning/Added Layers to Adjust Model Output	Custom Build LLM Using Enterprise Accessible Data
Required/Recommended Technologies	Open AI Hosted Application	Cloud Instance With LLM APIs	Cloud Instance With LLM APIs, Prompt Engineering, Custom Policies, Indexed Database	Licensed Customizable Model/ Proprietary Model, Data, ML Platform	Custom Build Proprietary Model, Data, ML Platform
Cost	Negligible	\$	\$\$\$\$	\$\$\$\$\$\$\$	\$\$\$\$\$\$\$\$\$\$\$\$



## What does the vendor landscape look like?

### Generative AI Apps:

Horizontal or vertical – available from both incumbent software vendors and new SaaS start-ups

### Generative AI Engineering Tools:

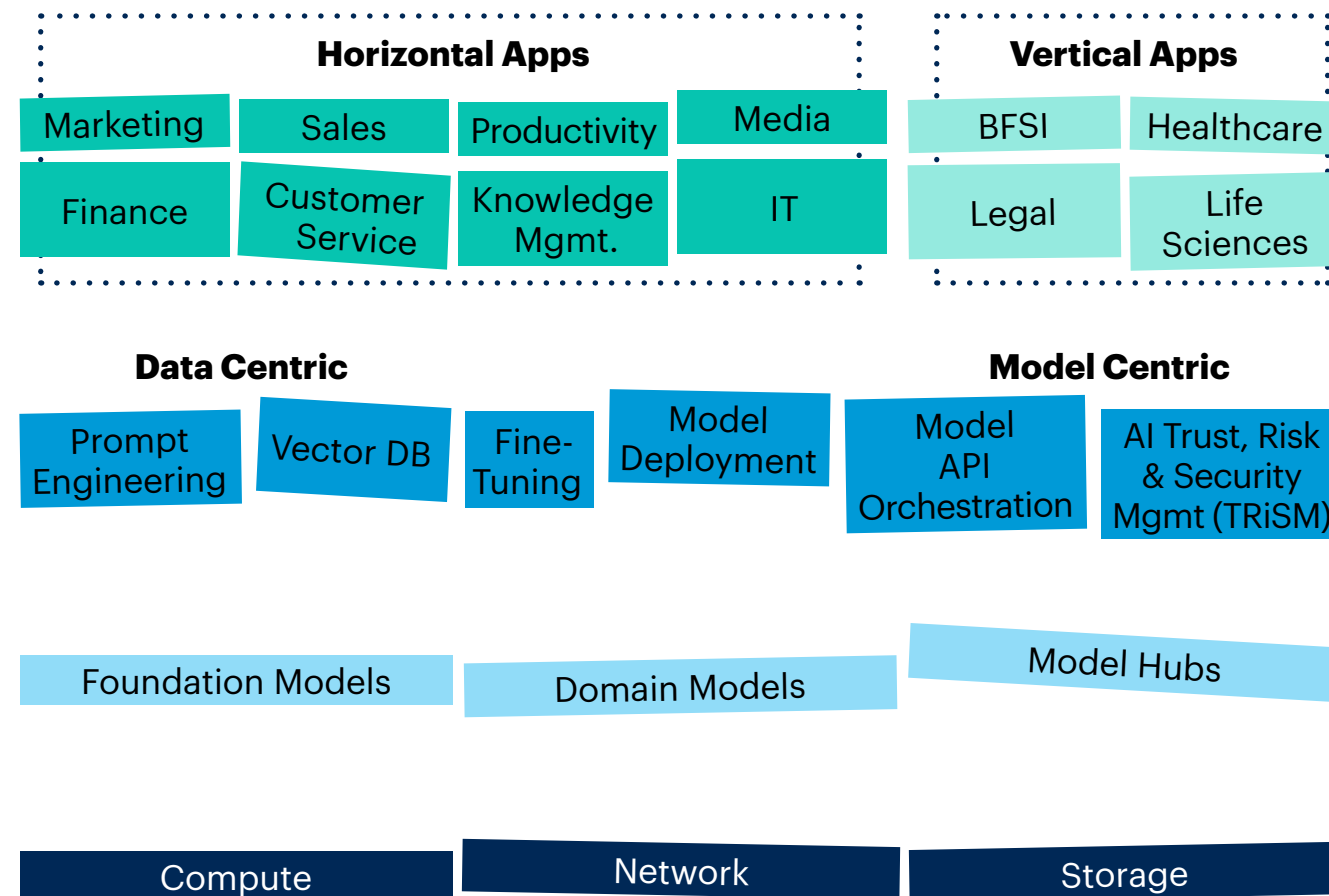
to operationalize and run generative AI models safely and efficiently in production

### Generative AI Models:

Versatile ML systems that can create human-like artifacts

### Infrastructure:

Infrastructure training and running generative AI models



**Vendors include:** Jasper AI, Harvey, Rytr, Mem, Notion, Glean, Algolia, GitHub

**Vendors include:** Pinecone, Weaviate, Arize, Credo.AI, Snorkel AI, Scale AI

**Vendors include:** OpenAI, Microsoft, Google, Cohere, Anthropic, AI21 labs, AWS, IBM, Alibaba, Tencent, Baidu, Hugging Face, EleutherAI, Stability AI

**Vendors include:** Nvidia, AMD, Intel, AWS, Microsoft, Google Cloud

*Note: Vendor lists are non-exhaustive and do not imply any Gartner endorsement or ranking of the providers mentioned.*

<sup>1</sup> 2023 Gartner Beyond the Hype: Enterprise Impact of ChatGPT and Generative AI Webinar Polls. This webinar was held on 30 March and 21 April 2023 with 1,465 and 1,079 respondents to the polling, respectively, for a total of 2,544 responses. Results of these polls should not be taken to represent all executives, as the survey responses come from a population that had expressed interest in AI by attending a Gartner webinar on the subject.

<sup>2</sup> Sourced from news articles and conversations with executives.

<sup>3</sup> [Instacart Joins ChatGPT Frenzy, Adding Chatbot To Grocery Shopping App](#), The Wall Street Journal.

<sup>4</sup> [Businesses Aim to Harness Generative AI to Shake Up Accounting, Finance](#), The Wall Street Journal.

<sup>5</sup> [How Morgan Stanley Is Training GPT To Help Financial Advisors](#), Forbes.

<sup>6</sup> [Mass General Brigham Creates Video With Generative AI; Testing It to Draft EHR Messages, Clinical Notes](#), Becker's Hospital Review.

<sup>7</sup> [How This Company is Redefining Retail Experience with Generative AI](#), Analytics India Magazine.

<sup>8</sup> [The AI chatbot ChatGPT is a game-changer!](#) maersk.com

# 4 Decisions to Make When Creating a Generative AI Policy

by Laura Cohn, Stuart Strome, Anthony Mullen, Avivah Litan and Lauren Kornutick

Corporate leaders have been grappling in 2023 with how to manage workplace use of generative AI applications such as ChatGPT and Bard.

Though some of the most prominent voices in generative AI development have warned that these products need close oversight and regulation, that task for now falls to the general counsel and C-suite colleagues who will lead implementation or use of the new technologies.<sup>1,2</sup>

Government requirements do loom on the horizon — the European Union, China, the U.S., the U.K. and Canada are all working on rules and guidance now. Having policies in place will also prepare your enterprise for whatever measures authorities enact in the future.

Based on our analysis of AI policies already instituted by companies and city governments, the GC should direct organizations to consider the following questions before establishing a policy:

1. What is our risk tolerance for use of generative AI?
2. What restrictions should we put in place to mitigate risks — and how will those differ for publicly available applications (such as OpenAI's ChatGPT) versus generative AI models tailored to our business needs?
3. Who has the authority to make decisions on generative AI use?
4. What information do we have to share, and with whom?



## What Is Our Risk Tolerance for Use of Generative AI?

To get started on determining the company's risk tolerance, legal leaders should borrow a practice from enterprise risk management and guide a discussion with senior management on "must-avoid outcomes." Such scenarios are defined as severe but avoidable events that will prevent a company from achieving its strategic goals. The C-suite should use these sink-the-enterprise possibilities as the basis for establishing guardrails.

First, you'll need to know the potential applications of generative AI models within the business. Once you have that list, consider the potential results and sort them into must-avoid outcomes and those that carry acceptable risk given the likely benefit of AI.

If your company makes consumer products, for instance, and a generative AI-powered application gives legally, financially or morally unsound advice, consumer trust in

the brand might falter or even evaporate — a result that might constitute a must-avoid outcome. What if, however, an LLM-powered chatbot provides potential ideas for uses of your product, and some of those, while not dangerous, are simply not optimal? That may fall into the "acceptable risk" category.

## What Restrictions Should We Put in Place?

Now it's time to dig a little deeper. Reach out to departments that may adopt generative AI models such as marketing, data and analytics, sales, and customer service, to name a few. Ask them about timelines for adoption of generative AI models to help understand how urgent it is to address those use cases in your policy. And make sure to ask whether they have witnessed, are aware of, or have any planned uses in their part of the business for generative AI applications that train on external data rather than customized information from your own business.

Compile a list of these more detailed use cases and organize them according to perceived risk — both the likelihood and severity of the threat. If the organization already has a team managing AI use cases, suggest the group extend its existing work to include policy and risk considerations. Some potential risks to look out for include unreliable outputs, outputs that violate copyright or IP protections, and regulatory compliance. Tools from a trusted vendor (e.g., Microsoft, Google, Amazon) may pose lower risk than those from publicly available applications.

For higher-risk situations (e.g., producing content for customer consumption), consider applying more comprehensive controls, such as requiring approval from a manager or AI committee. And in the highest-risk cases, you may want to think about outright prohibition. For lower-risk use cases (e.g., coming up with ideas for a fun activity during a staff off-site, or translating jargon or regulations written in a foreign language), consider applying basic safeguards such as requiring human review. Sample categories, examples and approval requirements can be found in Figure 1.

Legal leaders should apply different restrictions to publicly available applications than they do for more business-specific applications from a trusted third-party provider. Irrespective of provider, legal leaders should consider some universal safeguards. For instance, since all generative AI models may produce incorrect information or exhibit bias, be sure to include close monitoring for “hallucinations” — a common term of art meaning erroneous claims — and checking for skewed viewpoints.

An example of this in practice comes from the City of Boston’s AI policy, which warns employees who use AI applications to have a human being scrutinize any output (see Figure 2).<sup>3</sup>

Restrictions on publicly available applications will need to be more comprehensive. For example, they often use information entered by users to improve the model and that means sensitive material must stay out of the mix.

» Figure 1. Guidance on Use Cases and Restrictions of Generative AI (Illustrative)

Category	Examples	Who to Notify
Approved	Brainstorming, translations, trusted software within the enterprise’s security perimeter.	e.g., manager, HR
Requires Approval	Purchasing new software, using algorithms in HR or another department.	e.g., legal, procurement
Prohibited	Putting confidential company information, client data or personally identifiable information into AI platforms that don’t protect it.	Follow misconduct reporting procedure

Source: Gartner

» Figure 2. City of Boston’s Guidance on Human Review of Generative AI

**Fact Check and review all content generated by AI, especially if it will be used in public communication or decision making.**

- **Why:** While Generative AI can rapidly produce clear prose, the information and content might be inaccurate, outdated, or simply made up. It is your responsibility to verify that the information is accurate by independently researching claims made by the AI.
- **What to look for:**
  - Inaccurate information including links and references to events or facts.
  - Bias in the positions or information. We want to make sure that vulnerable populations are not harmed by these technologies. Think about how racial and ethnic minorities, women, non-binary, people with disabilities or others could be portrayed or impacted by the content.

Source: The City of Boston

When drafting policies, call out employee obligations in simple, unambiguous language. Consider including a set of “Dos and Don’ts” such as:

- Don’t input any personally identifiable information.
- Don’t input any sensitive information.
- Don’t input any company IP without checking with the team charged with reviewing use (e.g., the security team).
- Do turn off history if using external tools that enable that choice.
- Do closely monitor outputs, which could contain factual errors and biased or inappropriate statements.

Finally, don’t be overly restrictive. Banning use of these applications outright, or applying hard controls (e.g., restricting access to websites), may result in employees simply using them on their personal devices.

Consider defining low risk, acceptable use cases directly in your policy (e.g., producing a high-level summary of secondary research on a non-sensitive topic) to provide more clarity.

## Who Has the Authority to Make Decisions?

Work with functional, business and senior leadership stakeholders to:

- Align on who is responsible for mitigating a specific risk
- Review duties (e.g., whose responsibility it is to assess a novel use case or the output of an approved use case)
- Understand decision rights (e.g., who should approve new use cases or exceptions), as well as any other related activities, such as assessing risks, monitoring, imposing disciplinary actions or fielding questions

Be clear if there are uses that do not need approval, specify what they are directly in the policy, and provide examples — but also make sure employees know where to take their questions. Communicate which uses do need signoff to the workforce, clearly document the role that can provide approval, and list the relevant contact information.

Common points of contact for policies include:

- Legal
- Information security
- IT
- The company AI committee (this might include the CEO, chief legal officer, chief information officer and other business leaders)

## How Much Information Should We Disclose?

A critical tenet common of proposals in a variety of jurisdictions, including the standard-setting EU, will be that companies should be transparent about their use of AI. And consumers want to know if you use generative AI applications to craft corporate messages, whether the information appears on a public website, social channel or app.<sup>4</sup>

Therefore, it’s sensible to include a policy of disclosing the use and monitoring of generative AI technologies to both internal and external stakeholders. The GC should help companies consider:

- Whom to disclose to (customers, employees, third parties or all three)
- What they need to disclose
- What they need to be open about beyond use (e.g., employee monitoring for non-approved use of AI)



In practice, that means requiring employees to make sure the output is recognizable as machine generated by clearly labeling text. One option to consider: watermarks in AI-generated images to the extent technically feasible.

Here, again, Boston’s guidance provides an example. The city mandates agencies and departments not only reveal the use of generative AI, but also name the application used. Boston’s guidelines include an explanation of why disclosure is important — as well as suggested language to let users know how the content was created (see Figure 3).<sup>3</sup>

## A Final Tip: Build in Flexibility

Organizations that have established guidelines for employees typically make sure to note that they will be updated from time to time as:

- Generative AI develops
- Regulatory changes occur
- Company risk appetite shifts

### » Figure 3. City of Boston’s Generative AI Policy’s Disclosure Rules

Disclose that you have used AI to generate the content. You should also include the version and type of model you used (e.g, [Open AI's GPT 3.5](#) vs [Google's Bard](#)). You should include a reference as a footer to the fact that you used generative AI:

- **Why:** even when you use AI minimally, disclosure builds trust through transparency and it might help others catch errors.
- **Suggestions:** document how you used the model, the prompts you used etc. it could be helpful to you and your colleagues to better understand how you can use these technologies better and more safely.
- **Sample credit line:** “This description was generated by ChatGPT 3.5 and edited by Santiago Garces”
- **Sample credit line:** “This text was summarized using Google Bard”

Source: The City of Boston

<sup>1</sup> “The Godfather of AI” Warns of AI Possibly Outperforming Humans,” npr.com.

<sup>2</sup> OpenAI CEO in “historic” move calls for regulation before Congress, axios.com.

<sup>3</sup> Boston Isn’t Afraid of Generative AI, wired.com.

<sup>4</sup> 2023 Gartner Consumer Community Survey. Gartner Consumer Community (n =320, 10 to 17 February 2023). While the Gartner Consumer Community (n = 500) resembles the U.S. general population, the data cited is based on the responses of community members who chose to take each activity. These samples may not be representative of the general population and the data should only be used for directional insights.

# Prepare for AI Regulation by Addressing 4 Critical Areas

by Laura Cohn

with contributions from Dan Essig, Mike Fang, Melissa Ruiz Hernandez, Lauren Kornutick, Anthony Mullen, Sandy Shen and Stuart Strome

Even before large language models (LLMs) took the world by storm, the number of enterprises deploying AI was on the rise.<sup>1</sup> But the public launch of ChatGPT and other LLM applications catalyzed both AI investment and global regulatory efforts.

The general counsel (GC) can prepare senior leaders and the board even as the rules come into focus. A close reading of proposed and new regulations and guidance — in the EU, Canada, China, the U.S. and the U.K. — reveals shared underlying principles. The GC can help organizations develop corporate AI strategy by prioritizing them.

Although lawyers specializing in privacy, cybersecurity and AI told us the rules in the EU and Canada may not take effect until at least 2025, considering what is under discussion now will prevent companies from running into trouble with regulators later. China announced in July that interim regulations will take effect on 15 August, 2023.<sup>2</sup>

Two of those proposals contain fines for noncompliance; maximum penalties in the EU's proposed AI law exceed the attention-grabbing fines in the bloc's General Data Protection Regulation (GDPR). Even if authorities don't impose the highest possible fines, companies need defensible AI oversight and risk management that enable innovation and experimentation.

## EU Regulation May Influence Others

The EU's proposed AI law, first put forward in April 2021 and still under discussion, builds on the GDPR. Like that regulation, lawyers and business leaders expect it to become a roadmap for other countries to follow.<sup>3,4,5,6</sup> Elsewhere, Canada is weighing the Artificial Intelligence and Data Act (AIDA), and China has interim rules for generative AI.<sup>7,8,9</sup> Both the U.S. and the U.K., meanwhile, have issued guidance in the form of AI white papers (see Figure 1).<sup>10,11</sup>

» Figure 1. Global AI Regulation and Guidance at a Glance

	<b>EU's AI Act</b>	<b>Canada's Draft Artificial Intelligence and Data Act</b>	<b>White House Blueprint for AI Bill of Rights*</b>	<b>U.K.'s AI White Paper*</b>	<b>China's Interim Measures for the Administration of Generative Artificial Intelligence Services</b>
<b>Release Date</b>	April 2021	June 2022	October 2022	March 2023	July 2023
<b>Scope</b>	"... providers placing on the market or putting into service AI systems** in the Union."	Companies that "design or develop a high-impact AI system" or "make a high-impact AI system available for use" or "manage the operations of an AI system."***	Guidelines apply to the public and private sectors.	"We will not assign rules or risk levels to entire sectors or technologies. Instead, we will regulate based on the outcomes AI is likely to generate in particular applications."	Only applicable to services accessible to the general public within China
<b>Maximum Penalties</b>	€30 million or up to 6% of total worldwide annual turnover <sup>5</sup>	CA\$25 million or up to 5% of global revenue <sup>6</sup>	Not applicable	Not applicable	Not applicable
<b>Effective Date</b>	To be determined	To be determined	Not applicable	Not applicable	15 August 2023

\* Nonbinding government guidelines

\*\* The EU AI Act defines an AI system as "software that is developed with one or more of the techniques and approaches listed in Annex I and can, for a given set of human-defined objectives, generate outputs such as content, predictions, recommendations or decisions influencing the environments they interact with."<sup>6</sup>

\*\*\* Canada's draft act states that "high-impact AI systems" will be defined by forthcoming regulations to account for what's laid out in the EU's AI Act and technological advances. Among the potential considerations: whether the system will impact human rights and whether opting out is not possible.<sup>7</sup>

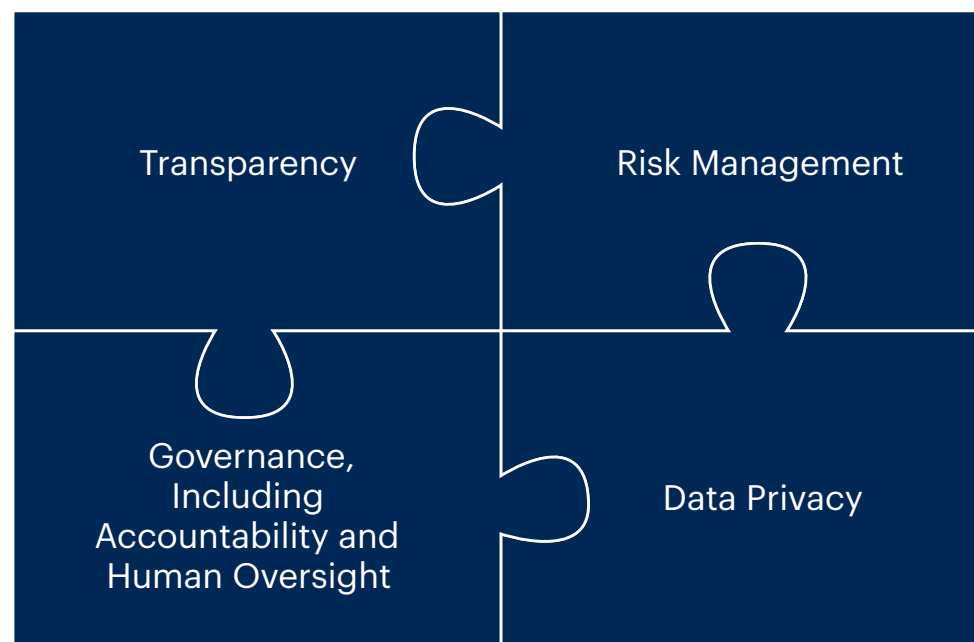
Source: Gartner



According to our analysis of the proposals, our interviews and our best-practice research, GC can get ahead by directing their organizations to prioritize the following (see Figure 2):

- Transparency
- Risk management
- Governance, including accountability and human oversight
- Data privacy

» Figure 2. Common Principles in Global AI Regulations and Guidance



Source: Gartner

## Be Transparent About AI Use

*Providers shall ensure that AI systems ... are designed and developed in such a way that natural persons are informed that they are interacting with an AI system...<sup>6</sup>*

### **EU's proposed AI law**

A critical tenet common across jurisdictions is that companies should be open about using AI.

One way that lawyers we interviewed suggest organizations should prepare: Disclose AI use in marketing content and the hiring process. At a bare minimum, legal leaders can help by updating the privacy notices and terms and conditions on their company's website to reflect AI use. But it's better to develop a separate section on the organization's online "Trust Center." Or post a point-in-time notice when collecting data that specifically discusses the ways the organization uses AI, assures individuals that their privacy rights won't be negatively impacted and, just as important, makes clear that such use will deliver value to the customer or individual.

Assess whether your company needs stand-alone AI guidelines. Organizations such as IBM make it easy for clients and customers to understand how they are using AI by linking their stand-alone AI policy to the code of conduct, IT security policy and privacy policy.<sup>12</sup>

Let departments working on AI initiatives know they can help the business avoid privacy risk by being transparent from the start. Otherwise, legal and privacy leaders risk finding out about such projects only when they are finalized.

## Risk Management Must Be “Continuous”

*The risk management system shall consist of a continuous iterative process run throughout the entire life cycle of a high-risk AI system, requiring regular systematic updating.*<sup>6,13</sup>

### **EU’s proposed AI law**

Like the proposed EU legislation, China also requires ongoing monitoring. China’s interim measure states, “Before using generative AI products to provide services to the public, a security assessment must be submitted to the state cyberspace and information department [i.e., the Cyberspace Administration of China].” Then organizations must register the algorithm on the official government website.<sup>2</sup>

The emerging practice of an algorithmic impact assessment (AIA) can document decision making, demonstrate due diligence, and reduce present and future regulatory risk and other liability. Creating an AIA must be a cross-functional endeavor. Besides legal, GC should involve information security, data management, data science, privacy, compliance and the relevant business units to get a fuller picture of risk. Since legal leaders typically don’t own the business process they recommend controls for, consulting the relevant business units is vital. The White House blueprint calls for organizations to make such assessments “public whenever possible.”<sup>10</sup>

Canada’s existing Directive on Automated Decision Making, which requires Canadian government organizations to conduct AIAs, includes a tool that guides organizations through the process — starting with an assessment of risk areas such as the reasons for automation as well as the source and type of data used.<sup>14,15,16</sup>

## Governance Must Include Human Oversight and Board Updates

*Human oversight shall aim at preventing or minimizing the risks to health, safety or fundamental rights that may emerge when a high-risk AI system is used.*<sup>6</sup>

### **EU’s proposed AI law**

Keep humans in the loop of AI development; regulatory measures call for it. That means establishing controls for humans to view, explore and calibrate AI system behavior. Where possible, these systems should be able to spell out why a particular result was achieved — sometimes called explainable AI — so users understand AI decisions. Legal leaders should mandate that third-party solution procurement and internal AI initiative development include the use of human-in-the-loop tactics to provide explainability and decrease risk.

AI-mature organizations are much more likely to involve legal teams in the AI development process, specifically in coming up with ideas for AI use cases.<sup>17</sup>

The GC could also establish a digital ethics advisory board of legal, operations, IT, marketing and outside experts to help project teams manage ethical issues.<sup>18</sup> The White House blueprint notes that independent ethics committees can both review initiatives in advance and monitor them to check whether “any use of sensitive data” infringes on consumer rights.<sup>10</sup>

Human oversight includes the board of directors. With 89% of boards considering digital innovation an implicit part of their growth strategy,<sup>19</sup> decisions to roll out LLMs within the organization should start with a well-informed discussion with directors. Legal leaders should make sure board members maintain a current understanding of the rapidly evolving LLM landscape. For a discussion guide, see Figure 3.

» **Figure 3. Discussion Guide on the Opportunities and Risks of LLMs**

### Opportunities for Using LLMs

- Are we targeting the right LLM business opportunities?
- How do we identify appropriate LLM use cases for our business while balancing risk with reward?
- What are our competitors doing?

### Risk Mitigation Plan

- Are there any risks we should not accept regarding LLMs?
- Are we taking an appropriate amount of risk?
- Do we need any additional controls to mitigate risk?
- How much oversight of LLMs do directors want?

Source: Gartner

## Guard Against Data Privacy Risks

*Instill “appropriate safeguards for the fundamental rights and freedoms of natural persons.”<sup>6,20,21</sup>*  
**EU’s proposed AI law**

The advent of ChatGPT has also led national regulators within the EU to act. In May 2023, France’s data protection authority, the Commission Nationale de l’Informatique et des Libertés (CNIL), announced it will create a new department focusing on AI.<sup>22</sup> The CNIL wants AI development to respect personal data rights such as individuals’ rights of access, rectification and opposition.

Legal and compliance leaders should manage privacy risk by applying privacy-by-design principles to AI initiatives. For example, require privacy impact assessments early in the project or assign privacy team members at the start to assess relevant risks. Better still, work with business partners such as the project management office to build privacy impact assessment requirements directly into project life cycles.

With public versions of LLM tools, organizations should alert the workforce that any information they enter may become part of the training dataset. That means sensitive or proprietary information used in prompts could find its way into responses for users outside the business. Leaders must establish guidelines, inform staff of the risks involved and provide direction on how to safely deploy such tools.



Privacy liaisons, who act as a link between the privacy team and a business unit, can help highlight the importance of privacy impact assessments and the risks to the workforce. One tactic is to use liaisons to detect noncompliance or overly risky uses of AI as a first line of defense. Recruit them from parts of the business with suspected exposure to AI risk, and provide extra training and guidance to help them identify and assess potential problems.

## Be Aware of These Differences

Despite the commonalities, jurisdictions have their own unique AI philosophies and requirements. These two stand out:

- The U.K.'s guidance notes that “rushed attempts to regulate AI too early would risk stifling innovation.” The document adds: “We recognize the need to build a stronger evidence base before making decisions on statutory interventions.”<sup>11</sup>
- China’s interim measure states: “Content generated through the use of generative AI shall reflect the Socialist Core Values, and may not contain: subversion of state power; overturning of the socialist system; incitement of separatism; harm to national unity; propagation of terrorism or extremism; propagation of ethnic hatred or ethnic discrimination; violent, obscene, or sexual information; false information; as well as content that may upset economic order or social order.”<sup>2</sup>

With organizations under pressure to use AI innovatively, the GC and other legal executives must watch the regulatory landscape closely. Doing this can help establish appropriate guardrails to manage potential risks — and make sure their companies don’t fall behind competitors in the race to harness new technology.

---

<sup>1</sup> Gartner CIO and Technology Executive Agenda Surveys, 2019 through 2023. Question: What are your enterprise's plans in terms of the following digital technologies and trends? In 2023, 32% of respondents said they plan to deploy AI, up from 14% in 2019. n = 2,882 (2019), 1,063 (2020), 1,825 (2021), 2,363 (2022), 2,186 (2023).

<sup>2</sup> [Interim Measures for the Management of Generative Artificial Intelligence Services](#), Cyberspace Administration of China. Direct quotations came from an earlier draft for comment translated into English by the Stanford Cyber Policy Center of Stanford University and remain in the new interim measure.

<sup>3</sup> [Regulating AI in EU: Three Things That You Need to Know, and Three Reasons Why You Have to Know Them!](#), Dentons.

<sup>4</sup> [Advancing Cooperative AI Governance at the 2023 G7 Summit](#), CSIS.

<sup>5</sup> [EU Closes In on AI Act With Last-Minute ChatGPT-Related Adjustments](#), Computerworld.

<sup>6</sup> [Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence \(Artificial Intelligence Act\)](#), EUR-Lex.

<sup>7</sup> [The Artificial Intelligence and Data Act \(AIDA\)](#), Government of Canada.

<sup>8</sup> [Bill C-27](#), Parliament of Canada. This includes the Consumer Privacy Protection Act, the Personal Information and Data Protection Tribunal Act and the Artificial Intelligence and Data Act.

<sup>9</sup> [China Releases Rules for Generative AI Like ChatGPT After Alibaba](#), Baidu Launch Services, CNBC.

<sup>10</sup> [Blueprint for an AI Bill of Rights](#), The White House.

<sup>11</sup> [A Pro-Innovation Approach to AI Regulation](#), gov.uk.

<sup>12</sup> [IBM Trust Center](#), IBM.

<sup>13</sup> [United States: Things We Learned at the 2023 IAPP Global Privacy Summit](#), Mondaq. The EU's AI Act defines a high-risk AI system as one that poses "significant risks to the health and safety or fundamental rights of persons." This could include AI systems found in medical devices and autonomous vehicles.

<sup>14</sup> [Algorithmic Impact Assessment Tool](#), Government of Canada.

<sup>15</sup> [Bill C-27: Canada's First Artificial Intelligence Legislation Has Arrived](#), Norton Rose Fulbright.

<sup>16</sup> [Directive on Automated Decision-Making](#), Government of Canada.

<sup>17</sup> **2022 Gartner AI Use Case ROI Survey.** This survey sought to understand where organizations have been most successful in deploying AI use cases and figure out the most efficient indicators that they have established to measure those successes. The research was conducted online from 31 October through 19 December 2022 among 622 respondents from organizations in the U.S. (n = 304), France (n = 113), the U.K. (n = 106) and Germany (n = 99). Quotas were established for company sizes and for industries to ensure a good representation across the sample. Organizations were required to have developed AI to participate. Respondents were required to be in a manager role or above and have a high level of involvement with the measuring stage and at least one stage of the life cycle from ideating to testing AI use cases. *Disclaimer: The results of this survey do not represent global findings or the market as a whole, but reflect the sentiments of the respondents and companies surveyed.*

<sup>18</sup> [Holding AI Accountable: Three Considerations for the Tech and the Minds Behind It](#), Fast Company.

<sup>19</sup> **2023 Gartner Board of Directors Survey on Business Strategy in an Uncertain World.** This survey was conducted to understand the new approaches adopted by nonexecutive boards of directors (BoDs) to drive growth in a rapidly changing business environment. The survey also sought to understand the BoDs' focus on investments in digital acceleration; sustainability; and diversity, equity and inclusion.

The survey was conducted online from June through July 2022 among 281 respondents from North America, Latin America, Europe and Asia/Pacific. Respondents came from all industries, except governments, nonprofits, charities and NGOs, and from organizations with \$50 million or more in annual revenue. Respondents were required to be a board director or a member of a corporate BoD. If respondents served on multiple boards, they answered for the largest company, defined by its annual revenue, for which they are a board member.

*Disclaimer: The results of this survey do not represent global findings or the market as a whole, but reflect the sentiments of the respondents and companies surveyed.*

<sup>20</sup> [EU Lawmakers Back Transparency and Safety Rules for Generative AI](#), TechCrunch.

<sup>21</sup> [Data Protection](#), European Data Protection Supervisor. Privacy is considered a fundamental right in the EU.

<sup>22</sup> [Artificial Intelligence: The Action Plan of the CNIL](#), CNIL.



# AI Systems ‘Need a Moral Compass’

A Q&A With Dragos Tudorache, Co-Architect of the EU’s AI Act

by Laura Cohn

 Table of Contents



*Photo courtesy of the office of MEP Dragos Tudorache*

**Dragos Tudorache** co-leads the European Parliament’s work on the AI Act. A lawyer who began his career as a judge in his native Romania, Tudorache was elected to the European Parliament in 2019. Along with artificial intelligence and new technologies, his legislative work includes security and defense, transatlantic issues and internal affairs. Prior to becoming an MEP, he headed legal departments at the Organization for Security and Co-operation in Europe and the United Nations Mission in Kosovo.



In a video interview from Brussels, Dragos Tudorache discussed what business should home in on to prepare for AI regulation, lessons from the GDPR, and his thoughts about warnings from tech leaders that AI could lead to human extinction. This interview has been edited for length and clarity.

#### **What do organizations need to do now to prepare for the EU's AI Act?**

Well, if you are a company that uses AI in industrial processes, in optimizing flows in a factory, in things that have nothing to do with humans, let's say, then you shouldn't even bother asking your lawyers to read the text.

If you are developing AI in areas that have to do with us humans — whether it's recruitment, whether it's medical, whether it's AI that is optimizing decision making — then it is likely that your AI is influencing environments that we believe need to be protected, and therefore you will need to ask your lawyers to start looking at this text.

#### **When it comes to AI regulation, a critical tenet across jurisdictions globally is transparency. In your view, what does good transparency look like? What do companies need to reveal?**

How your algorithm works, how you trained it, how you developed it, how you instructed it to function, and to actually reach its decisions or recommendations, or whatever it is that it generates as content. Is it something that induces risk?

#### **How can organizations explain their algorithms — and where should they do that?**

The regulation will provide templates for how you can comply with the technical data that you have to provide. There will be what is called an EU-wide AI registry, which is going to be a public database where high-risk applications of AI will have to publish all of this data that they have to provide as part of their transparency obligations.

So if I want to understand how the algorithm works, and the kind of data sets used, and what kind of instructions it received, then I open that, and look inside and say, 'Aha.'

That's how it was trained. That's how it was instructed. That's how it was tweaked to work. And then I can see whether there is anything wrong in the way it was done or not.

#### **What transparency obligations do you see for producers of large language model applications?**

These systems need a moral compass. And that's exactly what we're saying with this text to developers of large language models. They are great and fascinating, they are beautiful products, no doubt about it — but they need a moral compass. They need some rules and you have a responsibility to proactively, in the design and development, introduce safeguards that the content that is going to come out of your machine is not harmful and against the law. That's one obligation and the second has to do with copyrighted material.

#### **What's the obligation for companies there?**

Large language models are also absorbing a lot of copyrighted material, whether it's scientific articles, songs by ABBA or works of art.



Developers have an obligation to be transparent about the copyright material that they use. So that if I'm Drake and the Weeknd, I need to know that ChatGPT actually learned my songs. So in case it produces a song that very much resembles mine, then I can go and knock on the door and say, 'Fella, you need to give me some dollars for that.'<sup>1</sup>

**Let's talk about human oversight. Do companies need to assign an AI point person, team or committee to review any outputs from AI, for example?**

We have deliberately tried not to be overly prescriptive. This is in order to leave room for companies to achieve compliance in the best way they see fit.

**But can you give business leaders some examples of how companies can provide effective human oversight?**

I can't tell companies exactly what to do. But you have to have a human in the loop. Which means what? That I can't let an algorithm be the sole deciding factor without having, at some point, a human who takes a look. So I need to have someone, a person, who can have that final human touch to validate its content, its product.

**For the EU's AI Act, how important are the biometric verification safeguards?**

This has been, and remains, the most ideological point in the text. Many misunderstand what we're doing by saying that we're going to inhibit biometric technology. No, we're not. Biometric technology will remain. The one thing that cannot be done — and that has less to do with companies developing it, but more to do with law enforcement agencies in Europe — is, you cannot put an algorithm to run live, 24/7, in public spaces. We don't want this in Europe.

**Why is this so critical?**

We don't want to give a quasi-invitation to law enforcement to run these systems all the time on the grounds that maybe a criminal or missing child might pass by on the street.

How do you prevent abuse? And also, how do you protect privacy? Because basically what it means is that I will know that I will be walking in the street and there's always someone who will biometrically identify me.

That's something that's one step too far, according to how we understand privacy here in Europe.

**In terms of future enforcement of the AI Act, have you learned any lessons from the EU's General Data Protection Regulation (GDPR)?**

We're certainly trying to learn some lessons. First of all, in terms of avoiding the silo effect of different national regulators without much coordination, we are introducing a stronger mechanism of ensuring coherence in the implementation of the law across the EU. That's number one. Number two, we put some teeth in this law.

Regulators will have the possibility to investigate algorithms, to request information, if it's not clear from the data that companies have provided as part of the transparency exercise. So regulators can go in, knock on the door, and investigate the algorithm further. And if they find infringements, basically the options can go from shutting down the AI system, to withdrawing it from the market, to requesting changes, to very significant fines.

**You were in Washington recently, talking to U.S. officials. What's the outlook for international cooperation on AI regulation?**

Unlike the GDPR, where major jurisdictions were not ready to accept the same level of protection as we were introducing in Europe, my feeling from talking to literally almost every democratic jurisdiction out there is that with AI, there is already a much better starting point. There is quite a lot of convergence in terms of understanding the challenges.

Will all of them adopt an act like we do? Most likely not, but my narrative has always been that this matters less. As long as we agree on the big principles and on the big political objectives and what we want to achieve, then you can accept diversity in terms of the type of legislation in place.

**What was your take on the message from OpenAI's Sam Altman and other tech leaders that AI could lead to human extinction, and that mitigating the risks "should be a global priority?"<sup>2</sup>**

It's interesting that someone like Sam Altman tells you that there is a risk of extinction. Compared to any other piece of digital legislation that we have worked on until now, this time around I've heard businesses big and small say, 'Listen, we think it's time to have some rules in place.' It is good that they understand the responsibility they have. Also, as tech leaders, it is good that they engage on this and put pressure on lawmakers to start making decisions.

They know better than anyone else what's behind this technology. If they sound alarm bells, that means that our instincts in putting forward rules were the right instincts.

So for me, their plea confirms that we have done the right thing in writing up some rules.

---

<sup>1</sup> [An A.I. Hit of Fake 'Drake' and 'The Weeknd' Rattles the Music World, nytimes.com](#)

<sup>2</sup> [A.I. Poses 'Risk of Extinction,' Industry Leaders Warn, nytimes.co](#)



# Find out how generative AI will reshape your marketing function

One thing remains constant about generative AI tools, is that we all have questions about them. For marketing leaders, the questions include how it will change the function, what it means for how their organization operates, and what impact it will have on their interactions with customers.

## But how do marketing leaders determine what opportunities to pursue?

Join this complimentary marketing webinar, as we explore what changes generative AI could create for your organization, your enterprise and your industry.



Find out how generative AI tools could reshape marketing.



Explore the impact of generative AI.



Determine the next steps for CMOs.

**Watch now to find out what CMOs can and should do to take advantage of generative AI.**



# Generative AI Will Affect Information Security (and CISOs)

by Jeremy D’Hoinne, Avivah Litan and Peter Firstbrook

The cyber security impacts of a flood of experiments with ChatGPT and large language models (LLMs) will require chief information security officers (CISOs) to move quickly — adapting to rapid change while helping business units safely build and consume the new technology.

That means preparing in four areas:

1

## **Generative cybersecurity AI.**

Exploit opportunities to better manage risks, optimize resources, defend against emerging attack methods or reduce costs.

2

## **Safe consumption of generative AI.**

ChatGPT was the first example; embedded generative AI assistants in existing applications will be the next. These applications all have unique security requirements that are not fulfilled by legacy controls.

3

## **Development of new generative AI applications for the enterprise.**


These techniques have an expanded attack surface and require adjustments to application security practices.

4

## **The evolving threat landscape.**

Malicious actors will refine their tactics and leverage generative AI where it is a good fit.





## Generative cybersecurity AI

Start with AI assistants for developing secure applications and security operations chatbots. The former can be included in code generation tools such as GitHub Copilot, Tabnine and Amazon CodeWhisperer. The latter have the potential to improve the average security analyst's productivity and to lower the bar for skill requirements.

Examples and first demos of generative cybersecurity AI products and features appeal to many security professionals. However, our conversations with CISOs have highlighted questions and challenges related to cost and quality. One example of a concern: building new workflows and dependencies relying on immature technology that would require more work in a few years.

### **Recommendations:**

- Prepare and train security and risk management teams to deal with direct (privacy, IP, AI application security) and indirect impacts (from other teams) of generative AI uses in the enterprise.
- Run experiments with new features from existing security providers, starting with targeted, narrow use cases in operations and application security.
- Establish or extend existing metrics to benchmark generative cybersecurity AI against other techniques, and to measure expected productivity improvements.
- Determine your corporate position on providing feedback to the applications and improve their efficacy in the long run.
- Identify changes in data processing and dependencies in the supply chain for security tools and require your security providers to be transparent about data usage.
- Adapt documentation and traceability processes to augment internal knowledge and avoid feeding the tools with only your insights.
- Monitor the release of fine-tuned or specialized models that align with the relevant security use case or for more advanced teams.
- Remember that using generative AI base models “as is” might not be adequate for advanced security use cases.

# 2

## Safe consumption of generative AI

Consumption of generative AI applications such as LLMs through business experiments and unmanaged, ad hoc employee adoption creates new attack surfaces and risks related to individual privacy, sensitive data and organizational intellectual property (IP). CISOs should work with IT, legal and compliance and directly affected business lines to minimize unsanctioned uses of generative AI.

### Recommendations:

- Define a governance entity and workflow. Generative AI applications might require a specific approval workflow and continual usage monitoring where possible.
- Monitor and block. When ChatGPT became public, many organizations blocked access to OpenAI domains or applied some data leakage prevention. Now, they are shifting to a mix of “disclaimer pages,” which include links to user policy and relevant contacts, and more stringent blocking measures.
- Quickly communicate a short acceptable use policy. Often, a one- or two-page policy document can be used to share the internal contact for approval, highlight the risks of the applications, forbid the use of client data, and request documentation and traceability of generative AI outputs.
- Investigate prompt engineering and API integrations. Intercepting inputs and outputs with a custom-made prompt might improve results but also enable more security controls. Larger organizations might investigate prompt augmentation techniques, such as retrieval augmented generation (RAG).
- Favor private hosting options when available, which affords extra security and privacy controls.

# 3

## Development of new generative AI applications for the enterprise

In addition to using third-party generative AI applications and services, many organizations will build their own versions that will pose new risks, such as attacks on ML models and data poisoning. That in turn creates new requirements for AI application security. CISOs should adjust security controls and implement our AI trust, risk and security management (AI TRiSM) framework — governance practices establishing security policies that are specific to AI.

Implementing controls for GenAI applications will depend heavily on what kind of AI architecture an organization is using. If an organization builds and trains an in-house model, security and development teams share responsibility for protecting against the entire AI attack surface. Using synthetic data might improve security but also reduce the accuracy of the applications. These techniques are not available when hosting a pretrained out-of-the-box model.

With third-party models, companies must still address adversarial direct and indirect “prompt injections,” or the insertion of hidden instructions or context in the application’s conversational or system prompt interfaces, or in the generated outputs. There are also digital supply chain management challenges related to the model itself. Then, even when hosting an out-of-the-box third-party application or model, CISOs and their teams will be responsible for infrastructure and data security, but will also have to be involved in managing third-party risks and vulnerabilities.

### Recommendations:

- Adapt your control strategy to hybrid development models, such as front-end wrapping (e.g., prompt engineering), private hosting of third-party models and custom GenAI applications with in-house model design and fine-tuning.
- Team up with the data science team to evaluate the data security options when training and fine-tuning models, notably the potential impact on accuracy or additional costs. Weigh this consideration against requirements in modern privacy laws, which often allow individuals to request that organizations delete their data.
- Train your security champions as soon as training on secure GenAI coding is available.
- Apply AI TRiSM principles and update security best practices to include AI applications requirements
- Add requirements for testing against adversarial prompts and prompt injections.
- Monitor model operation tool improvements.

# 4

## The evolving threat landscape

As with any innovation (consider the rise of social media and cryptocurrency as two recent examples), malicious actors will seek creative ways to exploit immature security practices and limited awareness. Perhaps the biggest risk will be generative AI's potential to rapidly create believable fake content.

CISOs and their teams should anticipate the following:

- **Generative AI as a lure:** Expect counterfeit GenAI apps, browser plug-ins, applications and websites promising “the best Gen AI” but delivering malware-infected applications.
- **Digital supply chain tampering:** Embedded third-party GenAI components (libraries, models) could be affected by adversarial tactics, such as training data manipulation.
- **Adversarial prompting:** This threat includes prompt injection.

Generative AI tools will enable cyberattackers to increase the quantity and quality of attacks at low cost. Attackers have already started to exploit the technology's capabilities to craft better malicious and fraudulent content at scale.

Uncertainty about how successful they will be will require more flexible cybersecurity roadmaps. Security leaders need to reinforce their investments in resilience and in reducing exposure to categories of threats, rather than individual known attacks.

### Recommendations:

- Remember that the right order for any security investment is people, process and, only then, technology. Address threats that are tied closely to human interpretation of generated content for which there are no existing technical controls.
- Monitor industry statistics to measure the impact of generative AI on the attacker landscape.
- Make sure you can measure drift in detection rate from existing controls.
- Elevate requirements for more adaptive behavioral and ML defenses in your existing security controls.
- Challenge all existing and prospective security infrastructure vendors to outline how their product and research will evolve to address threats and tactics now possible due to generative AI. Beware of overstated claims.
- Evaluate the business dependency of key digital supply chain software and develop playbooks for “zero-day vulnerability” attacks, where malicious actors discover and exploit weaknesses before they are publicly known.
- Reduce the number of “blind spots” — assets, transactions and business processes that you cannot monitor for anomalies.
- Address the potential increased risk of fraudulent content and influence operations on corporate brands.
- Add generative AI content to security awareness training and phishing simulations. Make business workflows more resistant to realistic phishing campaigns and voice and video impersonations.




# How Generative AI Can Help Meet Customer Experience Expectations

by Gene Alvarez, Uma Challa and Leah Leachman

Customer retention and growth feature heavily in the plans of executives interested in generative AI. Thirty-eight percent of leaders see improving customer experience and retention as the primary purpose of initiatives to deploy applications trained on large language models, while 26% highlight revenue growth. Only 17% cite cost optimization (see Figure 1).<sup>1</sup>

Generative AI could help enterprises achieve all three goals by addressing rapidly changing consumer expectations. CEOs report price sensitivity as the top shift in customer behavior as inflation begins to bite, and 21% of them regard AI as the leading disruptive technology.<sup>2</sup> Executive leaders should therefore pilot promising customer-oriented use cases of tools such as ChatGPT and other AI techniques while guarding against the biggest risks.

» Figure 1. Primary Focus of Generative AI Initiatives



n = 2,544

Source: 2023 Gartner Beyond the Hype: Enterprise Impact of ChatGPT and Generative AI Webinar Polls  
Note: Results of these polls should not be taken to represent all executives as the survey responses come from a population that had expressed interest in AI by attending a Gartner webinar on the subject.

For example, snack company Mondelez International created an AI avatar campaign to coincide with the Diwali festival in India.<sup>3</sup> Wendy's is using Google Cloud's generative AI technology to enhance its drive-thru ordering service.<sup>4</sup> And U.K.-based Octopus Energy says AI is doing the work of 250 people at the company by replying to customer emails.<sup>5</sup>

## Make Customer Experience and Retention Your Core Aim

Customers expect effortless access to content, quick resolution of issues and seamless product experiences. Churn rates can be high if:

- Information provided by the company is inconsistent or even contradictory.
- Resolving problems takes a lot of customer time.
- The value of the product is not intuitive or doesn't meet expectations.

Smoother encounters with your organization are the prize to pursue — not just cost cutting. Replacing service people with bots that annoy or frustrate those seeking help will not boost retention.

Generative AI, however, has great potential to improve customer experience — and customer relationships — even though some of the use cases are complex and require significant investments.

First, the technology can **enhance value for customers** by:

- Educating them on how to use products and services better
- Advising them on new uses
- Validating their purchase decisions
- Anticipating their needs
- Helping them achieve a goal

For example, consider using generative AI to:

- **Onboard new customers or educate existing ones.** Generative AI solutions

integrated into self-service channels such as virtual assistants can interact in human-like fashion to provide detailed information on products and how to use them.

- **Identify value enhancement opportunities for customers in real time** and automate or assist agents to give contextualized information and advice.

Second, generative AI can **reduce customer effort**. Companies could integrate tools such as ChatGPT or other AI techniques to:

- **Better understand customer intent and sentiment** (through virtual assistants) and help resolve issues efficiently — or route the customer to the best channel to solve them.
- **Summarize customer interactions.** Generative AI models can do after-call work such as producing call notes, providing updated summaries during the interaction on both voice and text channels, and placing information into

enterprise repositories and search engines for continuous improvement. The technology can therefore increase accuracy and significantly reduce the time service agents spend compiling call notes, thus improving their experience and productivity.

- **Close the loop on customer issues.** Service organizations can generate emails and texts to customers instantly and efficiently, and enterprise communication platforms can personalize follow-up messages.
- **Assist customer support agents** by monitoring live interactions to suggest the next best sentence or action. A study by Stanford University and MIT found that agents who used AI tools that created conversational scripts increased their productivity, measured as issues resolved per hour, by 14% on average.<sup>6</sup>

## Customer-Centric Content Is Vital for Revenue Growth

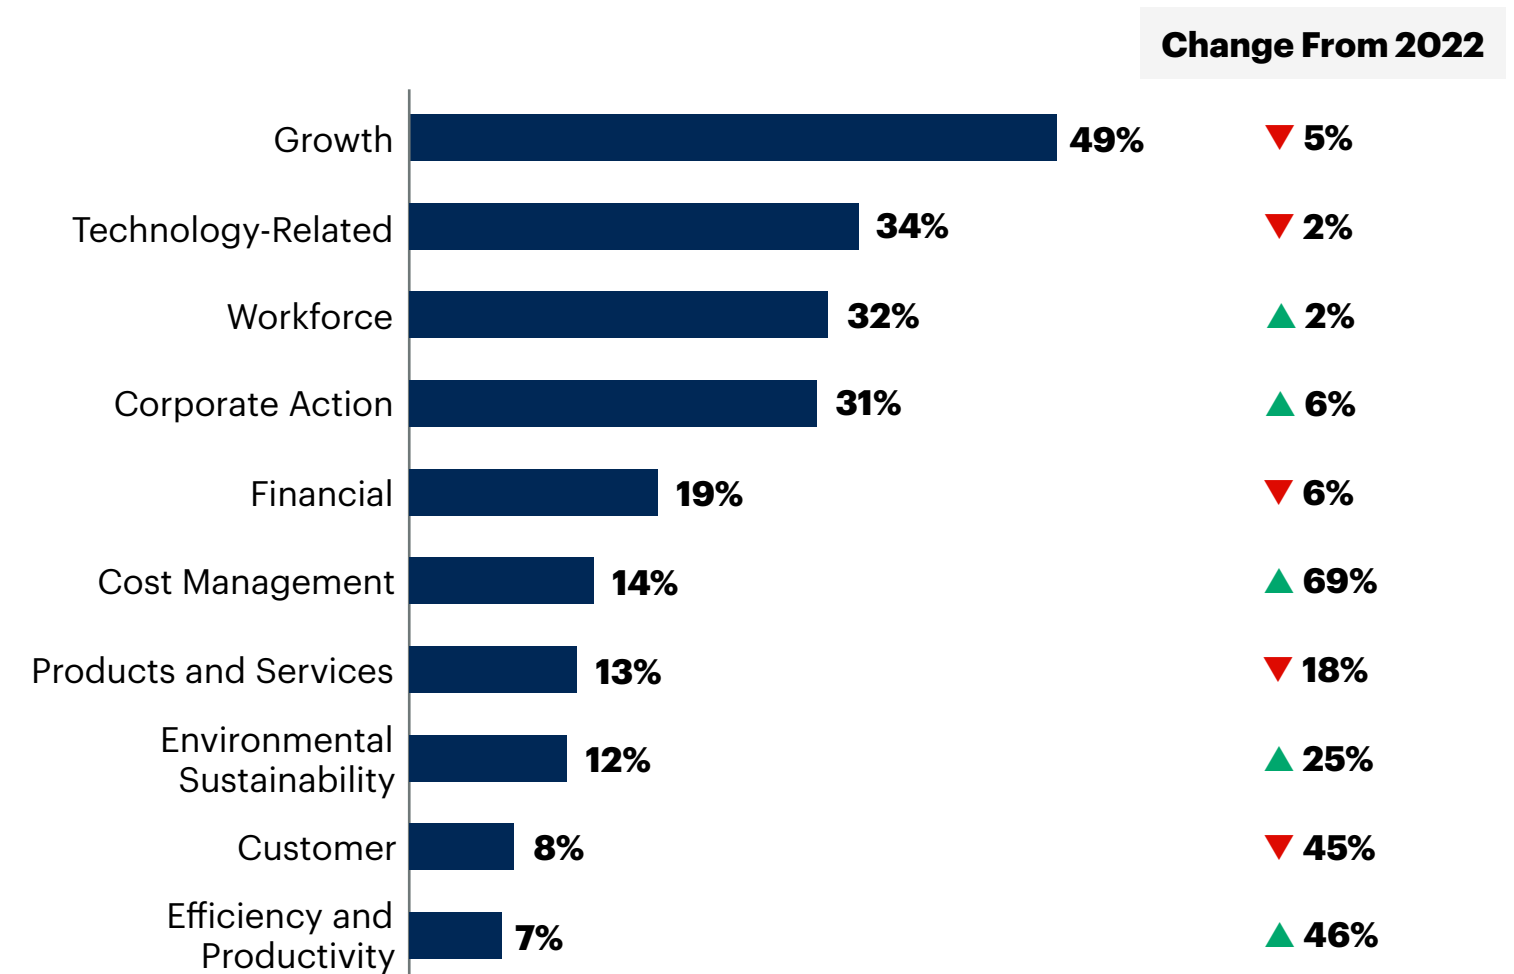
Growth is CEOs' No. 1 business priority (see Figure 2).<sup>2</sup> Economic uncertainty means efficiency is imperative, and generative AI's ability to automate the rapid production of more personalized creative content makes it a "must have" for marketing, sales and digital commerce initiatives. More than half of marketing teams (64%) are already leveraging AI heavily to create, execute and optimize campaigns.<sup>7</sup>

Although generative AI promises increased productivity, simply doing more does not guarantee growth. While the volume of content, ads and notifications that customers receive daily is proliferating, the value they attach to the content is strikingly low. Only 5% of consumers recall experiencing valuable digital interactions with a brand when considering a purchase.<sup>8</sup>

Executive leaders should look for points of overlap where generative AI can enhance both content production and personalization. Don't leave this mission on autopilot. Guide it with clear use cases, a renewed emphasis on customer research and human oversight to avoid biases and inappropriate uses. Get this step right, and:

- Marketing will feel empowered to shift focus from more tactical tasks, such as design production, to customer research, strategy and content curation. More resonant and engaging customer experiences will result. Regard generative AI as a complement to human expertise, not a substitute.
- The speed and range of content produced will enable your organization to reach customers worldwide more efficiently and allow your brands to manage complex customer journeys more effectively.

» Figure 2. CEO's Top 10 Strategic Business Priorities for 2023 and 2024



n = 422, all respondents

Q. Tell us about your organization's top five strategic business priorities for the next two years (2023 and 2024).

Source: 2023 Gartner CEO and Senior Business Executive Survey

- Optimizing product choices and potentially influencing customers' decisions — by nudging them toward different options — can accelerate adoption of new products or services.

But be cautious with recent advances in generative AI. These tools create derivatives of what already exists, which include hidden biases and the possible violation of ethical, copyright and regulatory standards.

## Explore How Generative AI Can Boost Operational Efficiency

In 2023, CEOs rank talent shortages and inflation as higher concerns than a recession or rising interest rates.<sup>2</sup> Organizations are therefore seeking productivity gains to offset labor constraints and increasing cost pressures. Generative AI innovations can help raise efficiency, including through:

- **Summarization.** ChatGPT and other generative AI techniques can give customers concise instructions to solve a problem — freeing time for customer-facing employees to handle more requests.
- **Content creation.** Generative AI can help salespeople create personalized RFP responses with industry terms and processes aligned to the products or services being sold. The technology will base these drafts on all the RFPs the organization has previously replied to and will use best practices. This tactic can accelerate responses to customer requests and reduce overall deal cycle times.

Marketing can use generative AI's capabilities to efficiently create:

- More personalized and in-context customer messaging, increasing the number of customers that purchase the product or service.

- Better webpages and merchandise offerings, boosting sales volume.
- Virtual influencers and brand avatars to mitigate friction in the customer journey, alleviating some of the burden on service teams.

## Experiment With Pilots and Mitigate Risks

Executive leaders should:

- Try small generative AI pilots aimed specifically at improving customer retention, revenue growth or operational efficiency.
- Locate and carefully curate data sources that will train the models for these initiatives.
- Question whether generative AI would best serve your use case or whether other more traditional techniques would better address both customer and commercial objectives.

Remember this emerging technology is not risk-free. Consider the following factors to plan and implement mitigating steps:

- Nonsense answers (which nonetheless sound authoritative) are a pervasive problem.
- In a publicly shared model, data privacy is an issue.
- As with every machine learning technique, models may provide biased responses, depending on the training data.
- Copyright violations are a further risk.
- Governments may impose or interpret regulations that could add risk to current implementations in the future.



<sup>1</sup> **2023 Gartner Beyond the Hype: Enterprise Impact of ChatGPT and Generative AI Webinar Polls.**

This webinar was held on 30 March and 21 April 2023 with 1,465 and 1,079 respondents to the polling, respectively, for a total of 2,544 responses. Results of these polls should not be taken to represent all executives, as the survey responses come from a population that had expressed interest in AI by attending a Gartner webinar on the subject.

<sup>2</sup> **2023 Gartner CEO and Senior Business Executive Survey.**

This survey was conducted to examine CEO and senior business executive views on current business issues, as well as some areas of technology agenda impact. The survey was fielded from July 2022 through December 2022, with questions about the period from 2022 through 2024. One-quarter of the survey sample was collected in July and August 2022, and three-quarters was collected from October through December 2022. In total, 422 actively employed CEOs and other senior executive business leaders qualified and participated. The research was collected via 382 online surveys and 40 telephone interviews. The sample mix by role was CEOs (n = 277); CFOs (n = 95); COOs or other C-level executives (n = 19); and chairs, presidents or board directors (n = 31). The sample mix by location was North America (n = 169), Europe (n = 105), Asia/Pacific (n = 102), Latin America (n = 29), the Middle East (n = 11) and South Africa (n = 6). The sample mix by size was \$10 million in revenue to less than \$50 million (n = 3), \$50 million to less than \$250 million (n = 51), \$250 million to less than \$1 billion (n = 102), \$1 billion to less than \$10 billion (n = 190) and \$10 billion or more (n = 76).

*Disclaimer: Results of this survey do not represent global findings or the market as a whole, but reflect the sentiments of the respondents and companies surveyed.*

<sup>3</sup> [2022 One Show - Creative Use of Data](#), Shah Rukh Khan-My-Ad, The One Club for Creativity.

<sup>4</sup> [Wendy's Is Working With Google to Test an AI-powered Drive-thru](#), Restaurant Business.

<sup>5</sup> [AI Is Doing the Work of 250 People at an Energy Company and Satisfying Customers Better Than Trained Workers](#), CEO Says, Business Insider India.

<sup>6</sup> [Stanford and MIT Study: AI Boosted Worker Productivity by 14% — Those Who Use It 'Will Replace Those Who Don't'](#), CNBC.

<sup>7</sup> **2023 Gartner Multichannel Marketing Survey.** This survey was conducted to determine best practices for maximizing multichannel marketing investments in response to evolving customer journeys in a fluid marketing environment. It was conducted online from November through December 2022. In total, 397 respondents were surveyed in their native languages across North America (n = 201), Western Europe (n = 161) and the Nordics (n = 35). Qualifying organizations reported enterprisewide annual revenue for FY 2021 of at least \$100 million, with 83% of the respondents coming from organizations with \$1 billion or more in annual revenue.

The respondents came from a variety of industries: financial services (n = 61), manufacturing (n = 58), consumer products (n = 46), retail (n = 45), travel and hospitality (n = 45), healthcare (n = 44), pharmaceuticals (n = 35), media (n = 33), and IT and business services (n = 30). All respondents were senior leaders who manage multiple marketing channels or were responsible for the execution of their organizations' multichannel marketing strategies. Eighty-five percent of respondents were aligned to the marketing function, 9% to brand management and 7% to sales or other business units.

*Disclaimer: The results of this survey do not represent global findings or the market as a whole, but reflect the sentiments of the respondents and companies surveyed.*

<sup>8</sup> **2022 Gartner Digital Customer Engagement Survey.** This survey was conducted to explore consumers' and B2B buyers' experiences in interacting with brands using digital channels and their perceptions of the value of

their interactions with brands. The research was conducted online from November through December 2022 among 1,076 consumers from North America (n = 717), Western Europe (n = 136) and Asia (n = 223) and among 509 B2B buyers from North America (n = 327), Western Europe (n = 78) and Asia (n = 104). Consumer respondents were required to have considered a purchase within the last month, used a brand's digital information sources and recalled a specific digital interaction with a focal brand. B2B buyer respondents were required to be involved in a significant purchase on behalf of their organization in the last six months, used a supplier's digital information sources and recalled a specific digital interaction with a focal supplier or department (n = 205).


Forty-seven percent of B2B respondents came from organizations with \$2 billion USD in annual revenue. B2B respondents came from a wide variety of industries, including technology products (n = 78), financial services (n = 76), manufacturing and natural resources (n = 71), retail (n = 69), IT and business services (n = 41), construction (n = 30), healthcare (n = 24), wholesale (n = 19), transportation (n = 17), consumer products (n = 16), real estate (n = 11), government (n = 10), media (n = 10), pharmaceuticals (n = 10), travel and hospitality (n = 9), education provider (n = 7), agriculture (n = 3) and other (n = 8).

*Disclaimer: The results of this survey do not represent global findings or the market as a whole, but reflect the sentiments of the respondents and companies surveyed.*

# Learn to Serve Your AI-Powered Customers Before They 'Walk Away'

by Mark Raskino and Don Scheibenreif

AI will accelerate the replacement of human customers by machine customers, which will act very differently. By adopting smart tactics now, executive leaders can make sure their organization benefits from this emerging megatrend.



The popular obsession with AI's effect on employment is understandable, but it deflects attention from a greater business challenge. AI will accelerate the replacement of human customers by machine customers. This megatrend will disrupt markets, often by stealth.

While AI bot replacements for employees might do as you say, AI-enabled machine customers — nonhuman economic actors that obtain goods and services in exchange for payment — will do as *they* see fit. If you serve them well, they can grow your addressable market. If you ignore their needs, you may not even see them walk away.

The new generation of deep neural network AI systems will change the role of the customer — in many cases, for the better. While some shopping for handbags, rare whiskeys or a new car might count as a leisure activity, a great deal of customer work is a laborious chore. For example, few buyers enjoy shopping for toilet paper, life insurance or a large technology purchase for their organization. Likewise, a recent study of a four-day workweek experiment in the U.K. found employees devoted their additional time off largely to “life admin” such as chore shopping.<sup>1</sup> AI can and will help chip away at that work too.

## Let Go of the Idea That All Customers Are Human

Machines are increasingly becoming customers on our behalf, and AI is speeding up that trend.

- Walmart's AI-powered negotiation software with a text-based interface (i.e., a chatbot) has so far closed agreements with 68% of suppliers approached, with each side gaining something it values.<sup>2</sup>
- Ticketmaster admitted that a “staggering number of bot attacks” disrupted its site selling tickets for a recent Taylor Swift concert.<sup>3</sup>
- LexCheck, a contract review platform, uses AI to help study legal, procurement and sales proposals.<sup>4</sup>
- Nike adjusted its terms of sale to curb the effect of “sneaker bots” that resellers sometimes use to automate the buying of scarce inventory.<sup>5</sup>

In 2023, large language model (LLM) generative AI has added an extra twist.

- Expedia announced in April 2023 that it's adding ChatGPT features to its mobile app to provide an interactive experience for basic travel details. The AI service cannot plan an entire trip yet, but that is in the works.<sup>6</sup>

- DoNotPay's ChatGPT bot renegotiated a customer's broadband provider bill with Comcast.<sup>7</sup>
- Greg Brockman, co-founder of OpenAI, demonstrated in a recent TED Talk how ChatGPT could help him plan a dinner party menu, send the menu to DALL-E to create an invitation, make a shopping list and send that list to Instacart.<sup>8</sup>
- Lazada, an e-commerce company in Singapore, created LizzieChat, a ChatGPT-powered bot that provides users with a more personalized shopping experience.<sup>9</sup>

AI is accelerating the trend of machines becoming your customers. That is not a contrived or exaggerated perspective: Machines are gradually taking over all the tasks of being a customer. They can determine needs, they can shop and compare, and they can negotiate, agree and pay. The argument that the “real customer” is always the human or organization behind the bot is irrelevant to your ability to sell. Once the machine assumes the role of purchaser, it is the machine you must convince to buy.

## The Automotive Industry May Show the Road Ahead

Cars will be one of the most important breakthrough categories of machines that will become customers for everything from tires and recharging, to passenger snacks and cleaning services. The demise of autonomous vehicles has been exaggerated: They are already operating on the streets of San Francisco, Phoenix and Los Angeles,<sup>10</sup> and in Chongqing, Wuhan and Beijing.<sup>11,12</sup>

What's at stake is so profound that big car companies are radically reorganizing themselves to accelerate transformation and insure against disruption. In 2022, Ford split itself into two divisions: Ford Blue to maintain the heritage combustion engine business, and Ford Model e to aggressively develop electric autonomous vehicles.<sup>13</sup> Also that year, General Motors CEO Mary Barra declared that the company's Cruise subsidiary could generate \$50 billion of annual revenue by 2030 from autonomous vehicle services and technology.<sup>14</sup>

Meanwhile, Elon Musk has repeatedly asserted that Tesla cars will earn money in the future by operating as autonomous taxis.<sup>15</sup>

These developments point to a world in which driverless cars roam around earning money *and spending it* directly from their own payment wallets. But machine customers do not have to be physical things like cars or home appliances. Software agents also will be important players, as embedded generative AI massively boosts the abilities of intelligent assistants such as Amazon's Alexa. Once we all have access to powerful shopping robots in our work and home lives, what won't we be prepared to have them buy for us?

### See the Opportunity for Market Growth

In many B2C and B2B settings, human consumers and business professionals will start delegating machines to do their customer work for them. Often, that work is the "self-service" effort that companies foisted on their customers in the name of convenience, but really

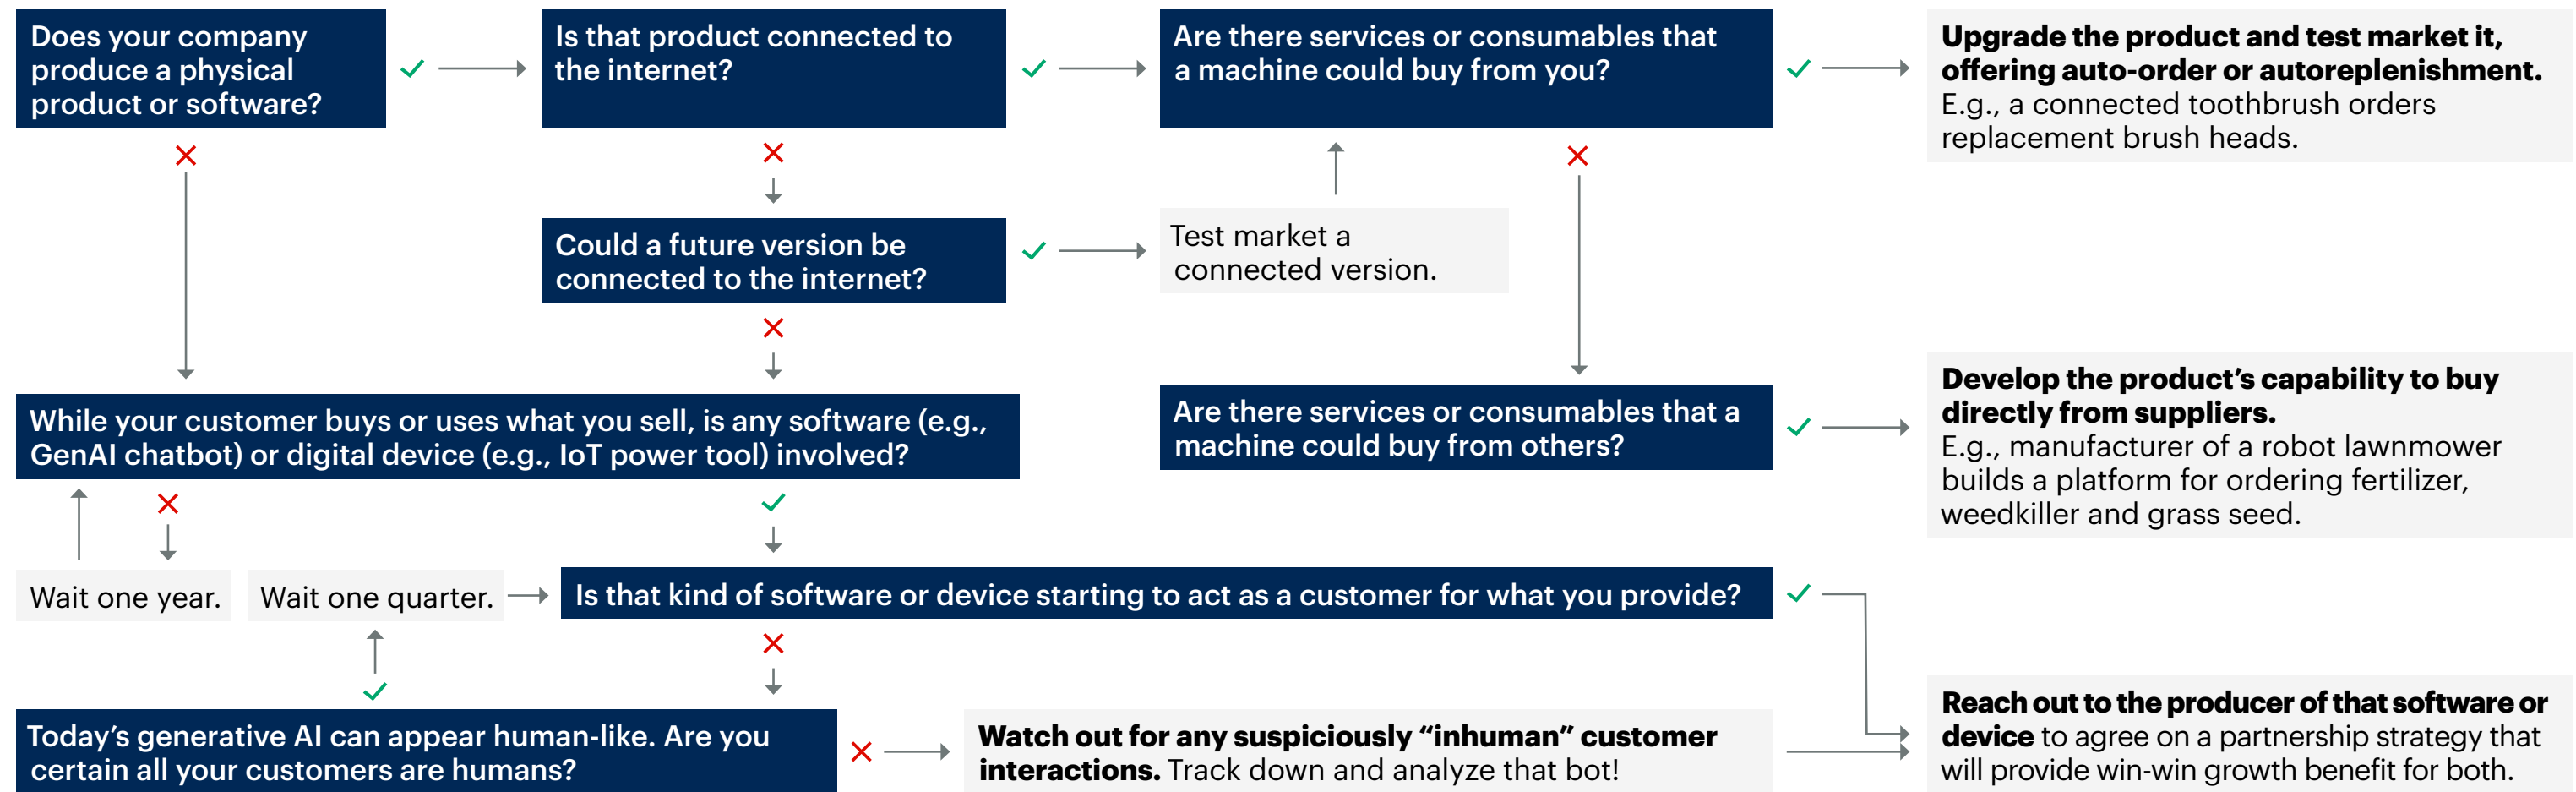
to reduce internal operating costs, in a prior technology wave. Be ready for this monumental transition as it unfolds throughout this decade. For example, organizations that produce physical products or software have some vital decisions to make (see Figure 1).

The spread of machine customers might shrink some markets, but for many or most companies, it will be an opportunity for total market growth. That's because humans tend to be inefficient customers who frequently dither, forget and fail to complete specific actions. Did you ever miss the moment to buy that event ticket, or open your toolbox and find you still had not got around to replacing that worn out drill bit?

Little wonder, perhaps, that HP Inc. now has over 11 million customers for its HP Instant Ink service,<sup>16</sup> where people delegate to their printer the task of remembering to buy ink. That's a lot of people who get to print pages when they otherwise might not have — expanding the total market. In 2022, HP Inc. extended that service into an adjacent revenue opportunity: Your printer can now order paper for you too.<sup>17</sup>



» Figure 1. Machine Customer Decision Path



Source: Gartner

## Here's What You Should Do

Executive leaders can position their organizations to benefit from the AI-fueled machine customer trend by taking the following steps:

- **Make all your product and service information easily accessible to machine customers.** They may be searching on 100 different variables, and you'll need to provide data for all of those, depending on where they are in the purchase process. Provide and encourage API access, and make sure CAPTCHA and other bot-thwarting tools are not shutting out legitimate machine customer revenue.
- **Add machine customers to your core digital commerce and sales strategy.** Strive to be exceptional at digital commerce. It will be the first place machine customers go when they want to buy from you. Consider the impact of these buyers on the way you sell and provide information.
- **Develop a strong commercial partnership between sales, marketing, supply chain, IT and analytics.** This team should create one to three scenarios that explore what happens when AI-enabled machine customers start buying from you. The supply chain must be agile enough to respond to unexpected demand patterns.

- **Start hiring and training sales and service staff to work with AI agents.** They will need to understand and possibly crack the algorithms that drive a machine customer's purchase behavior or after-sale service demands. These employees will need a basic understanding of how the technology works; some positions may require a data science background.
- **Alert and train customer-facing staff to spot machine customers.** Machines (generative AI LLM-powered bots) posing as humans may already be trying to negotiate, book and buy from you through text-chat-based customer service lines, and perhaps even via telephone calls using high-quality voice synthesis.

- <sup>1</sup> [Would You Prefer a Four-Day Working Week?](#) University of Cambridge.
- <sup>2</sup> [How Walmart Automated Supplier Negotiations](#), Harvard Business Review.
- <sup>3</sup> [Taylor Swift | The Eras Tour Onsale Explained](#), Ticketmaster.
- <sup>4</sup> [LexCheck](#), LexCheck.
- <sup>5</sup> [Nike Moves to Curb Sneaker-Buying Bots and Resale Market With Penalties](#), CNBC.
- <sup>6</sup> [Expedia App Integrates ChatGPT](#), Forbes.
- <sup>7</sup> [Watch DoNotPay's AI Chatbot Renegotiate a Comcast Bill to Be \\$120 Lower](#), PCMag.
- <sup>8</sup> [The Inside Story of ChatGPT's Astonishing Potential](#), TED.
- <sup>9</sup> [Lazada Unveils New eCommerce AI Chatbot LizzieChat](#), Marketing-Interactive.
- <sup>10</sup> [Waymo Is Starting Driverless Taxi Tests in Los Angeles](#), Engadget.
- <sup>11</sup> [Baidu Launches China's First Driverless Taxi Services in Chongqing and Wuhan in Landmark Moment for Autonomous Motoring](#), South China Morning Post.
- <sup>12</sup> [Baidu, Pony.ai Win Permits to Offer Driverless Robotaxi Services in Beijing](#), Reuters.
- <sup>13</sup> [Announcing Ford Blue™ and Ford Model e™](#), Ford.
- <sup>14</sup> [GM's Cruise Robotaxi Unit Drives Deeper Into the Red](#), Reuters.
- <sup>15</sup> [Watch Elon Musk's Full Interview With CNBC's David Faber on Twitter, Tesla and A.I. Advances](#), CNBC.
- <sup>16</sup> D. Scheibenreif and M. Raskino, "When Machines Become Customers," Gartner, 2023.
- <sup>17</sup> [Never Run Out of Paper Again!](#) HP Inc.



# What Generative AI Means for Business

Business leaders face 3 new sets of expectations — and generative AI can help



**Investors** will expect new sources of growth and better margins.

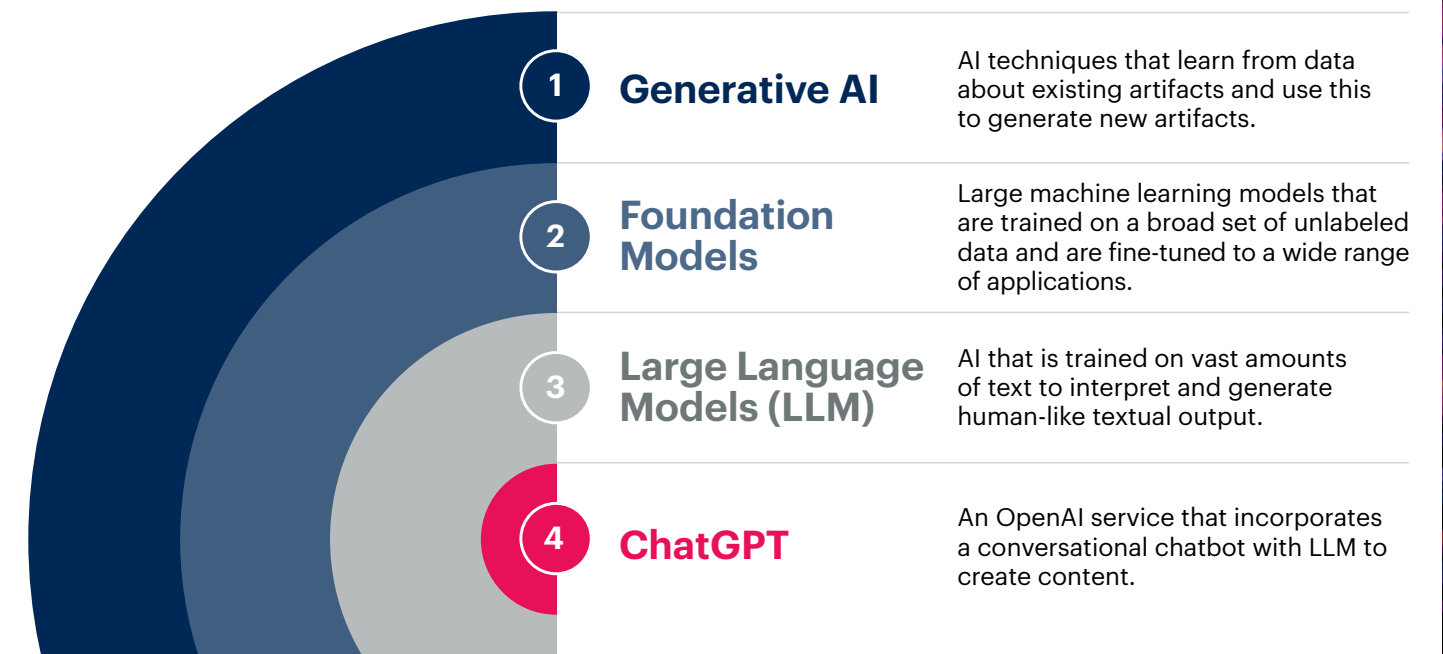


**Customers** will leverage GenAI in their daily lives — and expect you to do the same.



**Employees** will leave organizations where humans are doing work that GenAI could handle.

## What Is Generative AI?



**Explore Gartner resources to get started and move forward with GenAI in your enterprise.**

# What Generative AI Means for Your Talent Strategy

by Helen Poitevin and Pieter den Hamer

The accelerated investment in generative AI has unsurprisingly led to concerns about how this technology will affect jobs, including those once thought impervious to automation. Some roles will cease to exist, while many others will change radically, encompassing new tasks and requiring new skills. Faced with a potentially historic disruption, executive leaders should shape their talent strategies using a new framework for understanding generative AI's impact on the workforce.

AI applications have been affecting workers for years, but this time feels different. During the late 2010s, few executive leaders would claim publicly that layoffs were in any way due to investments in AI or automation. They would usually say they identified impacted staff, retrained them and shifted them to other roles. Today, however, as use of generative AI tools such as ChatGPT spreads rapidly, leaders are more willing to explicitly call out AI as one of the reasons positions will disappear — whether the technology is truly the culprit or not.<sup>1,2</sup>



## Assess Impact Within Your Business Context

Generative AI will affect each organization differently. Executive leaders must consider their enterprise's unique situation as they anticipate the effects of this technology (and AI more broadly) on the workforce. Most available models for measuring automation risk from AI analyze only existing roles and estimate the potential for AI to replace current employees. While these models have merit, they omit two important factors in any executive leader's talent decisions:

- **Demand drivers** show whether your organization should scale up products and services, scale them down or hold them steady.
  - Talent impact: Future demand for specific capabilities within your organization, industry or broader market will determine where you invest in talent and technology. Demand also affects your ability to meet your talent needs: Sought-after workers are harder to attract, and low-demand jobs become less desirable and more challenging to fill.

- **Technology drivers** relate to how your organization implements and uses technology. They affect how you organize work within and across teams.
  - Talent impact: Organizations adopting generative AI will design roles and workflows differently depending on what they are using it for. Executive leaders must also account for the combinations of skills available within the organization and the wider labor market.

Generative AI will directly alter tasks such as content creation, question answering and discovery, translation, document summarization and software coding. But how it will transform individual jobs is more challenging to predict. For example, knowing generative AI will disrupt copywriting and customer support doesn't tell you what it means for the people providing those services in your organization. Your business context will determine whether you use the technology to make these employees more productive, phase out jobs that can be done by machines or reconfigure these roles to provide new types of services.

The following steps will help executive leaders identify the impact of AI on their organization's workforce and adapt their business, talent and technology strategies accordingly.

## Step 1: Forecast Demand

First, examine how you expect demand for a given product or service to change, based on forecast business conditions and company strategy. Apply this exercise at the enterprise, business unit, department or team level, and for both internal- and external-facing services. This helps clarify how much a given team will need to deliver in the future.

Demand for a product or service could be:

- **Increasing/High** — An offering's success is driving a marked rise in demand, with high growth targets. This demand may also be for a particular capability — such as AI, software or other digital technologies — that needs to scale to an increasing number of domains.
  - Examples: Medical diagnostics, software engineering, AI development, prompt engineers
- **Stable/Core** — Volume and future evolution are predictable, possibly with some variability linked to seasonality or similar factors. There are no plans for either high growth or phasing out the service.
  - Examples: Payroll services, customer service supporting stable markets, internal IT services, insurance claims processing

- **Decreasing/Low** — The service or product targets a niche market or is being phased out. Ongoing support is required, but there are no plans for investing further to try scaling up to a bigger offering. This includes instances where a company continues to serve existing customers for a product or service but no longer offers it to new ones.
  - Examples: Legacy products no longer sold but with multidecade contracts that need to be honored; niche products such as parametric insurance

## Step 2: Envision AI Technology Applications

After forecasting demand, formulate your future vision for how you will apply generative AI and AI more broadly. This helps identify the degree of impact on a given team or role. Will people in certain jobs be replaced? Will they need new skills? How critical will it be for a worker's success and competitive positioning that they use generative AI effectively?

Consider what your planned applications of AI mean for existing professions and industries.

You could:

- **Stay within their boundaries** — Many emerging use cases for generative AI are far from transformative. For example, a recruiter could use embedded generative AI functionality in a talent sourcing tool to automatically apply search criteria based on natural language input, without having to manually apply filters. This change simply makes an existing work pattern easier and faster.
- **Go beyond their boundaries** — Generative AI, alongside other AI techniques, can change how people access products and services. For example, AI can help create webpages, videos, apps or other content quickly, without requiring any technical knowledge. This capability allows for a fully personalized customer experience.
- **Create new boundaries** — Autonomous business, in which AI manages or carries out the majority of operations, has already started to emerge. Machine customers, augmented managers and autonomous operations are central features of these new business models. We anticipate the labor ratio to substantially change, with enterprises needing fewer people to generate the same amount of revenue.

## Consider a Matrix of AI Futures

Combining these demand and technology factors creates a range of situations (see Figure 1). Within each organization, and even within each team, multiple cases are likely to apply. For each situation, specific investments in talent development and workforce planning will help executives lead their teams through an AI-driven wave of transformation.

We expect generative AI to have two broad effects on the workforce, each with significant implications for executive leaders.

» Figure 1. Talent Impact Matrix for Generative AI and AI More Broadly

Low AI Impact  High AI Impact

		How AI Is Applied in Industries and Professions			
		Within the Boundaries:	Pushing Boundaries:	Breaking Boundaries:	
Business Context and Demand Volumes	Increasing/High:	Scale up and grow to meet demand.	<b>Scale Up:</b> Productivity goes up with AI and humans accomplishing more work with less effort. Still hiring to meet targets.	<b>New Impact Level:</b> High performance expectations in many newly configured roles, transformed through AI and human partnership.	<b>Symbiosis:</b> Machine economy with few employees compared to impact, reach and revenue. Top design teams required.
	Stable/Core:	Scale through efficiency and productivity gains.	<b>Shift and Not Replace:</b> Work shifts. Many are not replaced, and most complex work remains. Roles must be reconfigured, often toward assisted multiskilled generalist roles.	<b>Run Smarter:</b> Increased performance expectations per worker in terms of reach or impact. Only those who can work effectively with AI remain competitive.	<b>Synergy:</b> New pockets may emerge for machine customers and autonomous operations. Innovative redesigner teams required.
	Decreasing/Low:	Scale down or maintain low niche-style activity.	<b>Last Ones Standing:</b> Hard to attract talent and likely losing talent. Automation is an opportunity to maintain service despite talent shortage.	<b>Niche Innovations:</b> Hard to attract and find talent in hyperspecialized markets. May open up some new markets that can now be served profitably.	<b>Experiments:</b> Rare, but micro niche investments. Few employees.

Source: Gartner

<sup>a</sup> Demand can be for internal or external services, at the organization, business unit or department level. It drives activity volumes and portfolio prioritization.

## Generative AI Within the Current Boundaries

Using generative AI in this way will result in the shifting of roles over time. Fewer people will be needed to complete the same amount of work. Mass layoffs driven solely by generative AI adoption are unlikely, especially considering that labor markets remain historically tight. However, people who leave jobs affected by the technology are less likely to be replaced. Employees hired into these roles will be expected to accomplish higher work volumes more quickly. It will also become increasingly challenging to find talent willing to take on jobs that automation will likely displace.

Executive leaders implementing generative AI in this context should anticipate headcount reductions over time. They will need to redesign jobs displaced or disrupted by AI into smaller numbers of multiskilled generalist roles that encompass a wider range of capabilities and offer a more compelling employee value proposition.

## Generative AI Beyond the Current Boundaries

Second, using generative AI to push the boundaries of professions and industries, or even break them, will set off a race for performance. Expectations will be higher in newly configured roles, and organizations won't be able to compete without using AI. The question will be not which tasks go to AI and which to humans, but how people can use AI creatively to reach new heights. New, highly specialized jobs will emerge where generative AI and related technologies are used creatively and strategically to transform what teams do and what their clients expect. This requires a rich blend of business and technological acumen that few possess.

For example, organizations will need executive-level business architects who wield both types of expertise in an entrepreneurial way. When it's too hard to find the right fit for these roles, build on experience with cross-disciplinary fusion teams that design and deliver digital products and services. Whether this specialized role is filled by a person with this rare mix of talents or by a fusion team, it will be critical for enterprises that make the leap to create new boundaries in a machine economy.

<sup>1</sup>BT to Cut 55,000 Jobs With Up to a Fifth Replaced by AI, BBC News.  
<sup>2</sup> IBM CEO Among the First Major Executives to Say They'll Replace Jobs With AI, Axios.





# Can AI Overcome Flaws in Human Decision Making?

A Q&A With Daniel Kahneman and Olivier Sibony

by Steve Shapiro

 Table of Contents



Photo courtesy  
Wikimedia Commons

**Daniel Kahneman** is professor of psychology and public affairs emeritus at the Princeton School of Public and International Affairs, the Eugene Higgins professor of psychology emeritus at Princeton University and a fellow of the Center for Rationality at the Hebrew University in Jerusalem. He is the author of *Thinking, Fast and Slow*.



Photo courtesy  
Wikimedia Commons

**Olivier Sibony** is a professor, author and expert in strategic thinking and decision processes. He is professor of strategy at HEC Paris and an associate fellow of Saïd Business School in Oxford University.

When presented with the same facts, people will disagree and even draw different conclusions themselves from one day to the next.

This variability in decision making — what Daniel Kahneman, Olivier Sibony and their co-author Cass Sunstein call “noise” in their book *Noise: A Flaw in Human Judgment* — can create costly problems for organizations. Disagreements don’t cancel out, but add up, they argue.

We should expect such inconsistency in humans — after all, we are not machines. But with machines now making decisions, can we overcome this flaw in human psychology? And what does that mean for organizations?

In separate conversations, Kahneman and Sibony told us why decision making often goes wrong and whether AI can provide answers. The interviews have been edited for length and clarity.

### **What’s the biggest thing that people get wrong about “noisy” decisions?**

**Kahneman:** People are not aware of the amount of noise. And they’re certainly not aware of the damage that it does. One of the motivations for the book was a conversation with somebody who runs a hedge fund and he thought that errors cancel out. With that misperception you’re not going to take noise very seriously. When they have disagreements, it appears to be a one off, but disagreement is actually the rule. People ignore the problem.

### **Has the emergence of new generative AI tools such as ChatGPT changed how you think about noise?**

**Sibony:** One of the main reasons to use models and algorithms instead of humans to make judgments is that they may not be unbiased every time, but at least they are noise free — if you ask the same question twice, you will get the same answer. Now, as we’ve all experienced with ChatGPT or Bard or whatever your favorite large language model is, that’s not entirely true.

If you ask ChatGPT why it’s not giving you the same answer every time, its answer, tellingly, is that in order to look and

feel more natural, meaning more human to the human user, it intentionally introduces an element of randomness to the answers so that it will not feel like a machine. But at a minimum, it is clearly noisy, so if you were hoping that those models would be a substitute for human judgment, that’s going to be a problem.

### **What about algorithms in general? Can they solve the problem of noise if they don’t act like humans?**

**Kahneman:** Moving to algorithms will improve the quality and the accuracy of judgments. Algorithms tend to beat humans, or at least tie with them.

The main reason for human inferiority in that regard is noise. Bias in human judgment tends to be drowned out by the noise. When you have a system that is not noisy, biases stand out very clearly so there is something utterly nonsensical in all this talk about bias in AI. My guess is that it is going to be a rare case where the AI is more biased than the humans that it replaces.

I think some of the resistance to algorithms that we see comes from a deep misunderstanding of what’s going on. Humans are noisy and that’s why you sometimes don’t see how biased they are.

**Sibony:** Using algorithms in the most generic sense to make a judgment is going to be less noisy than using human judgment. That point obviously still stands. ChatGPT is the exception, not the rule, and any simple algorithm is noise free. But since we have a preference for sophisticated and user-friendly algorithms over basic ones that apply simple rules, it's going to be more and more of a problem.

### **Should we use algorithms to try to eliminate noise completely?**

**Kahneman:** I think the moment you start becoming aware of noise you know that you cannot eliminate it. And you don't want to reduce it to zero because you want individuals to exercise their judgment. You do not want to completely turn them into machines.

There are costs to noise, but there are also costs to reducing noise and you want to reduce those costs as much as possible. You want people to feel that they're expressing themselves. People need to see AI and algorithms as noise reduction tools that help them rather than as bureaucratic constraints on the way that they operate. You can't get to zero noise if you're using human judgment, but zero is not the best possible outcome.

### **What about situations where you have competing AI models or algorithms?**

**Kahneman:** It's going to probably produce less noise than humans would and that's how you've got to look at it. You've got to look at the alternative. And of course, now people find it shocking when different models disagree. But if they knew how much people disagree, they would be less shocked.

Of course there will be noise with multiple algorithms; that's bound to happen. But there is something else that happens when there are competing algorithms. Because typically algorithms are based on a lot of data, it's also possible to reconcile them and reduce the noise in a way that is really not possible when it's people generating noise.

**Sibony:** If you've got several models, you will have several points of view, just like when you have several humans. There is an advantage though here, which is that when you have one point of view and it's yours, that judgment seems to you a lot more correct than the judgment of another human. When you're looking at three different algorithms, you will tend to assume that each of them is to be taken with a grain of salt, and that's a better attitude.

If you can move to the mindset where your own judgment is one of the inputs, another human is another input, one AI model is another input and another AI model is yet another input, then you've recognized noise and you're doing something about it. You're not giving precedence to a single point of view, which happens to be yours.

So far it's slightly theoretical because I haven't seen many situations where people are using more than one AI model to solve the same problem and wondering which device to follow. But more and more this is going to be the case.

The more inputs we have access to from various models, the more we recognize that under uncertainty, a multiplicity of points of view is actually a good thing.

---



Gartner®



Over 60,000 attended a Gartner conference last year — Don't miss out this year

“Gartner Conferences offer a unique opportunity to network with your peers, learn new and best-in-class practices and be exposed to new solutions all in one event. Well worth the time to step away from the business for a few days to expand your expertise.” **2023 Gartner Conference Attendee**

- ★ New!
- ◆ Most popular
- Reintroduced!

**Gartner  
IT  
SYMPOSIUM | Xpo. ◆**

**11 – 13 September 2023**  
Gold Coast, Australia  
**16 – 19 October 2023**  
Orlando, FL  
**6 – 9 November 2023**  
Barcelona, Spain



**Gartner CFO & Finance  
Executive Conference**

**18 – 19 September 2023**  
London, U.K. ★

**Gartner Supply Chain  
Planning Summit**

**30 – 31 October 2023**  
London, U.K. ★  
**29 – 30 November 2023**  
Phoenix, AZ ●

**Gartner ReimagineHR  
Conference**

**11 – 12 September 2023**  
London, U.K.  
**23 – 25 October 2023**  
Orlando, FL  
**4 – 5 December 2023**  
Sydney, Australia ★

**Gartner Security & Risk  
Management Summit ◆**

**26 – 28 September 2023**  
London, U.K.

**Why attend Gartner  
conferences ↗**

© 2023 Gartner, Inc. and/or its affiliates.  
All rights reserved. EVT\_M\_NC\_2398966





# ABB Electrification Lights Up ML Models for Financial Planning

by Madison Wilson and Mike Lashinsky

Machine learning (ML) can make financial planning more efficient and accurate, serving as the basis for generating meaningful forecasts that help executive leaders prepare for future disruptions. Yet, when FP&A leaders rush to replace traditional forecasting with this technology, their underdeveloped models — trained on poor historical data — can lead to untested algorithms that frustrate progress.

Alessandro Marchesano, head of FP&A at Switzerland-based ABB Electrification (ABB EL), a business unit of ABB, recognized that making headway with ML in financial planning required his team to take the time to learn how the technology works and understand what it does best. The other vital step was to identify the specific role humans will play in building, training and governing ML models.

ABB EL tested a human-machine learning loop method that empowered the FP&A team to:

1. Integrate complex external drivers into ML models.
2. Test, iterate and refine ML-based drivers.
3. Refine the algorithms regularly to maintain performance.

This ML pilot helped them gather unbiased data results and gain executive leadership buy-in to build on the project's success. As a result, they plan to scale the use of ML to other financial planning activities, such as on-demand scenario analysis and planning assumption updates. The team is also well-positioned to expand ML to other areas of planning and, eventually, other areas of ABB.

---

“Our forecasting transformation triggered deep discussions on key business drivers, is as accurate as our traditional bottom-up process, and is much faster,” said Marchesano.

## 1 Integrate Complex External Drivers Into ML Models

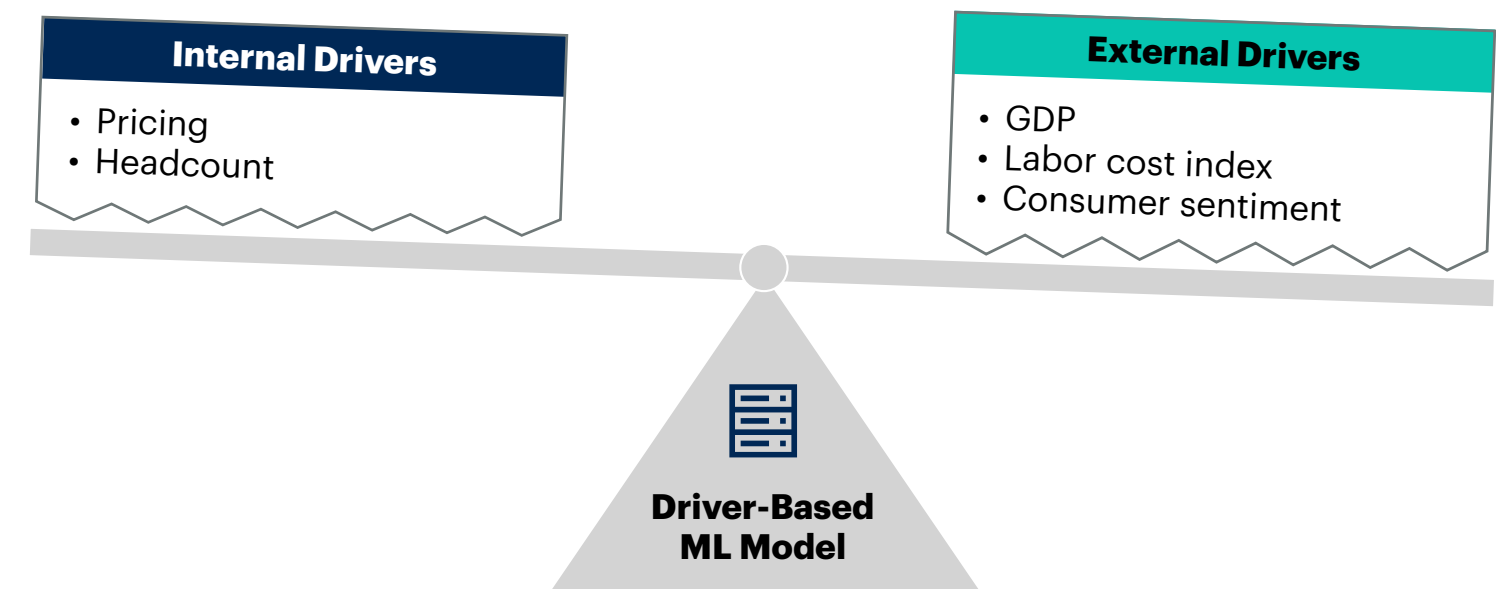
First, ABB EL created driver-based ML models, which are advantageous because they:

- Explain the story behind a forecast.
- Find material relationships between business drivers and financial outcomes.
- Model short-term effects of external market changes or disruptions.
- Adapt plans to market volatility and better map business complexity.
- Highlight the data with the largest influence on predictions.

As part of this strategy, ABB EL strengthened the ML algorithms with more complex external drivers that a business cannot directly control or influence, such as GDP or consumer sentiment.

These external drivers change frequently, which trains the ML model to self-correct and enables a quick, optimized business response to shifting market conditions. Put simply, when one driver changes, multiple algorithms can automatically revise their forecasts with limited human oversight. This technique also provides the company with early-warning signals to model multiple future possibilities (see Figure 1).

» Figure 1. Weighting of Internal and External Drivers as ML Model Inputs  
Illustrative



Source: Adapted From ABB

## 2 Test, Iterate and Refine Your ML-Based Drivers

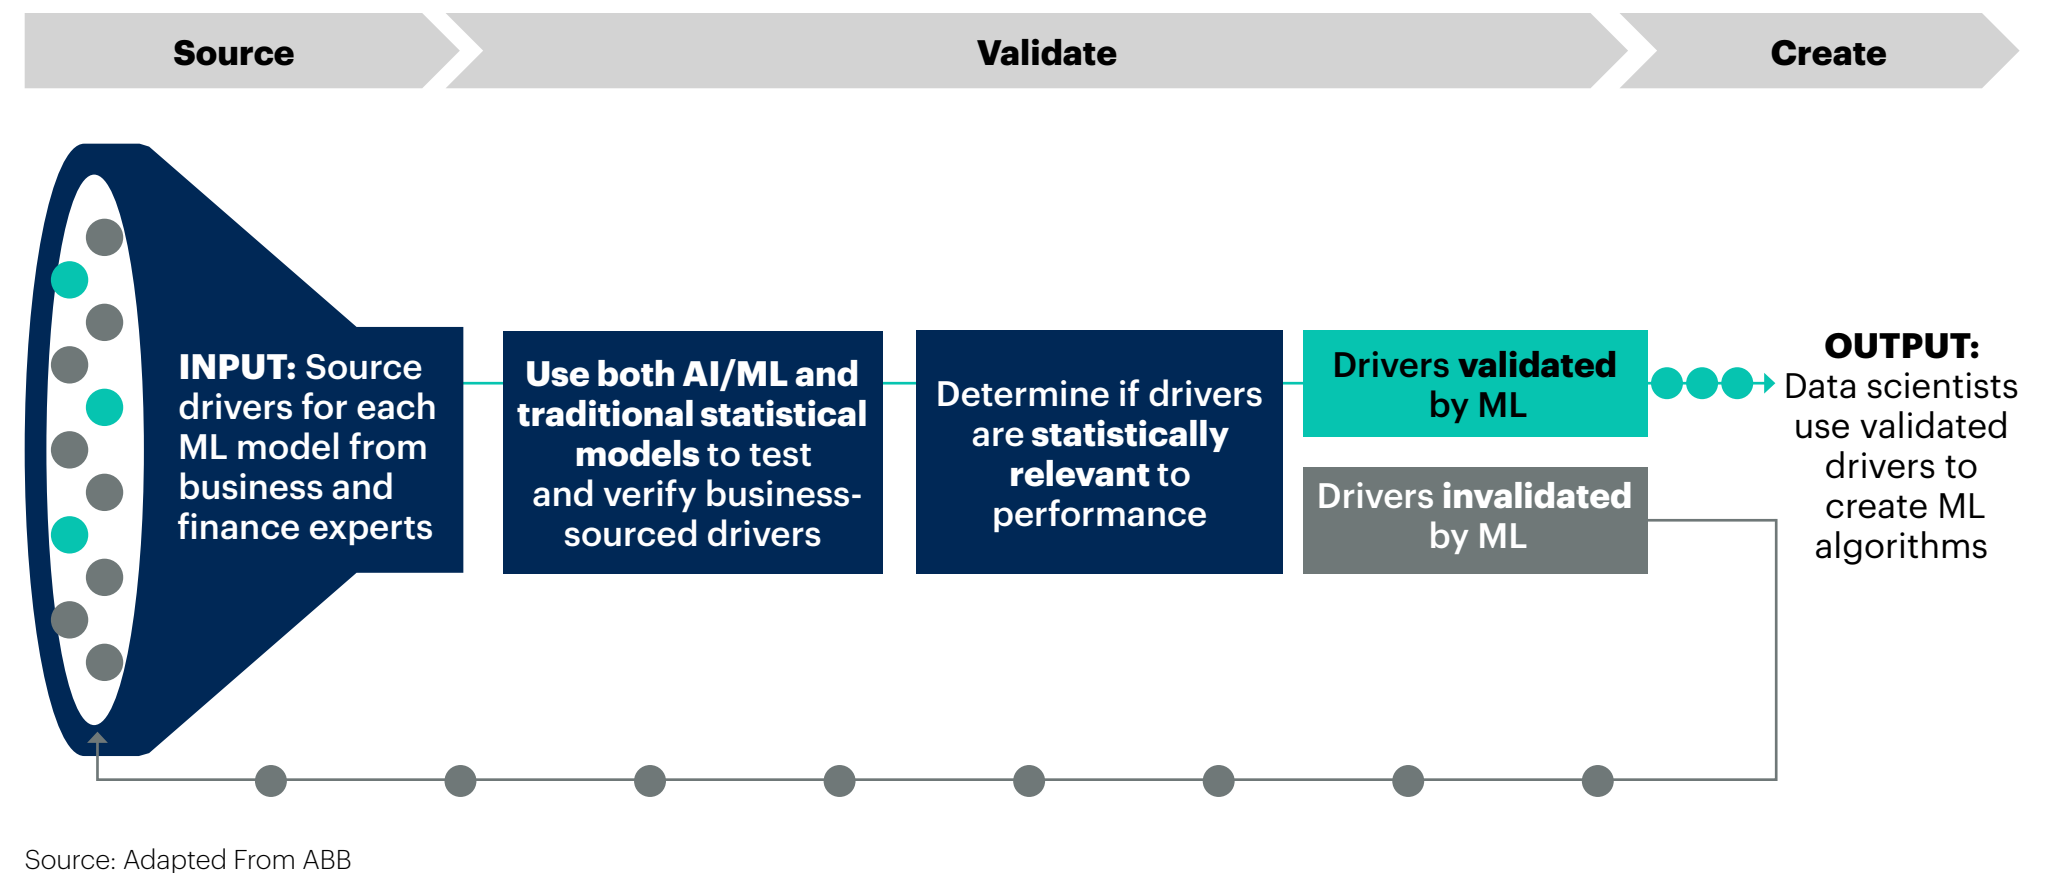
Organizations tend to only refine drivers if there is a material macroeconomic change. But when executive leaders overlook driver validation, they lose the opportunity to better understand what spurs financial performance.

Having recognized that the first set of drivers chosen for each algorithm would always need fine-tuning, ABB EL's FP&A team quickly set about substantiating the ML results to make sure they reflected financial performance. To achieve this goal, FP&A created a three-step process to test the relevance of the drivers and measure their effect on outcomes (see Figure 2).

The team:

1. Sourced external drivers for each complex business area from business and finance experts.
2. Used ML and statistical models to validate those business-sourced findings — and, if the FP&A team hit a dud, business analysts went back to the first step and continued to iterate until they uncovered the most effective drivers.
3. Created algorithms with data science teams using the validated drivers from the second step.

» Figure 2. Validation Process for Human-Sourced Drivers  
Illustrative



### 3 Refine the Algorithms Regularly to Maintain Performance

After creating first drafts for each algorithm, the team members didn't stop there. They knew the first version of an ML algorithm was unlikely to be the best possible one.

They found some initial assumptions were incorrect, drivers could be more specific or ML models didn't accurately reflect the market environment. To combat this issue, they assessed ML performance on a rolling monthly basis, hand-picking algorithms that missed the mark on expected accuracy yet were critical to company performance.

The team now has a framework to repeat this process until they identify and improve all weak algorithms.

FP&A guides this refinement with four questions:

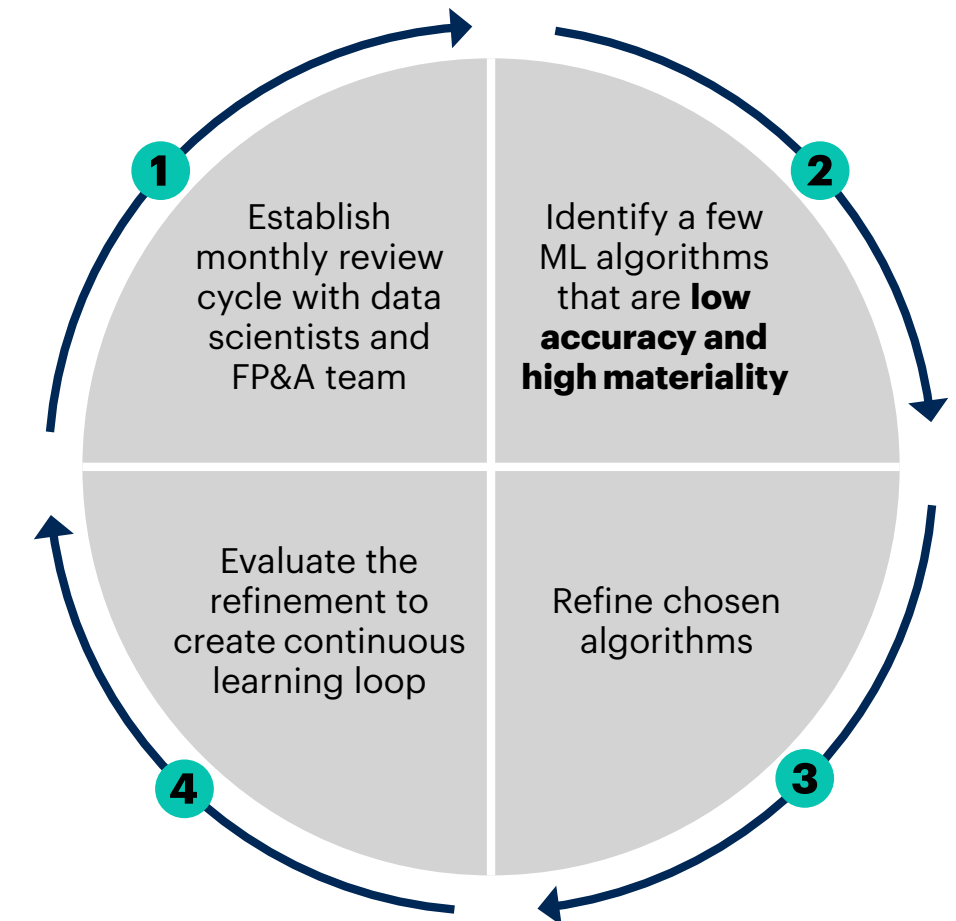
- **Did we choose the most impactful inputs for the algorithm?** This question reminds them that it might be necessary to revisit the drivers put into each algorithm to make sure they are still relevant.
- **Did we write the best algorithm to model this business area?** This question prompts data scientists to check for a better way to create this algorithm and make it a more accurate model.

- **Do we need to refine any underlying trend models?** Many algorithms are based on seasonality or trend models. Answering this question tests if those models are not performing as well as they could.

- **Should we consolidate or break up any business areas to make them easier to model?** Sometimes, a business area may be too large or too small for a single algorithm to best represent it. This question reminds the team that it might be necessary to adjust the size of the business area in their models to make sure the algorithm performs as expected.

Weeding out underperforming algorithms helped the team learn how to spot languishing ML models and correct how they are governed, which establishes the human-machine learning loop. The models then iteratively improve over time, making ABB EL's financial planning speedier and more objective (see Figure 3).

» Figure 3. Learning Loop for Refining ML Algorithms



Source: Adapted From ABB



# Smarter Spending & Planning

## Are Alternatives to Layoffs Better for Shareholder Returns?

Beyond hefty severance payouts, reducing the workforce can carry hidden costs. The short staffing that follows layoffs may lead to the need for expensive contractors. The remaining employees may demand faster promotions. And a vicious cycle may kick off: overstretched people leading to lower morale, increasing turnover. When it's time to get back to hiring, candidates may only accept offers at higher pay.

Small wonder, then, that only 11% of organizations maintain layoff savings for three consecutive years. C-suite executives must collaborate to define the potential financial and non-financial risks of layoffs.

Consider these alternatives:

1. Voluntary reduction in hours — Many employees may willingly work fewer hours at commensurately lower pay.
2. Internal redeployment — Even where headcount must decrease, some may transfer to other parts of the business where their skills are in demand.
3. Sabbatical — An unpaid break can give employees the opportunity to pursue a professional interest that still contributes back to the company.
4. Executive compensation cuts — Lower base pay temporarily while preserving long-term incentives.

None of these options will make everyone happy, but they can keep a company in position to recover faster when better days arrive.

— Vaughan Archer

## A Funnel to Send Cost Cuts to Growth Projects

When CFOs promote cost-cutting measures in response to pressures on profitability, they can end up starving, rather than protecting, their organization's in-progress investments.

But it doesn't have to be this way. One answer lies in "cost savings winbacks," a mechanism that strikes the right balance between these competing priorities.

CFOs have a choice of two types:

1. Direct return — Give a portion of the resource saved to the budget owners to use for projects that will drive long-term revenue increase.
2. Central — Send the money saved to an enterprisewide pool for reinvestments designed to grow revenue based on organizational priorities.

For instance, the finance function at a U.S. medical equipment company offered business leaders concrete benefits based on central winbacks. They could submit proposals for what they would do with the money if they could reclaim half their savings.

The result? Different business units competed to unearth more cost reductions — increasing the amount available to fuel growth.

They can use these funds as a cushion against budget overruns, without delaying other projects. Plus, it will bake the concept of uncovering resources into a "business as usual" practice instead of a one-time action.

— Roma Kaur

# Talent & Culture

## A Four-Day Workweek Is More Feasible Than You Think

The four-day workweek is a valuable tool for attracting and retaining talent, preventing burnout and boosting flexibility. For most candidates, this benefit — even more than a pay bump — is the top perquisite that would draw them to a position. Media buzz about this trend is not just hype: 21% of organizations have already started experimenting with a shorter workweek.

What holds HR leaders back from putting the concept into practice? They worry that leaders will resist and that operational adjustments will prove too challenging.

Finding the model that works best for your organization can mitigate those perceived barriers.

HR leaders have more options than they think. Ask:

1. How many hours should employees work?  
Choose a reduced (e.g., 32 hours in four days) or a condensed schedule (e.g., 40 hours in four days), depending on the nature of employees' work.
2. When should you offer the additional day off?  
Opt for a universal (e.g., Fridays off for all employees) or distributed (different days off for different employees) weekday off, depending on business continuity needs.
3. How widely should the program apply?  
Employees can decide whether to participate or you can mandate it for everyone.

— Shivendra Singh and Amrita Puniani

## Three Cues to Help Employees Do Their Best and Keep It Up

Sixty-five percent of employees whose roles could be remote experience at least one of the following:

1. Devoting too much time to tasks that are necessary but not considered in their formal workload (for instance, checking schedules to plan a meeting).
2. Growing exhausted because they don't rest.
3. Feeling that more effort is futile because performance expectations seem unattainable.

Performance cues from the organization can signal to employees if they are working efficiently, setting a sustainable pace, and making progress.

These require both HR policy changes and management execution:

1. Giving context so employees can work flexibly but waste less time on mundane decisions — One HR team taught managers to facilitate discussing all stakeholders' needs before setting team norms for how they will work together.
2. Building wellness into planning and execution of work — HR executives should consider making the team's well-being part of leaders' performance goals so they see it as essential to success.
3. Recognizing great work in real time — One total rewards group within an HR function created a system for managers to pre-define high-performance actions and reward teams when they deliver.

— Kayla Velnoskey, Iga Pilewska and Kate McLaren-Poole

# Growth & Innovation

## Marketing Must Do Less to Get More and Drive Profitable Growth

Chief marketing officers are investing in more technology, channels, data and personalization — with demonstrably diminishing returns. Few can afford to be stuck in this “cycle of more,” especially when companies face budgetary pressures and are focusing on margins.

The key is discipline: creating fewer, but more meaningful, customer experiences. In both business-to-business and consumer purchase decisions, people who have just one interaction that changes understanding of their own needs and makes them feel more confident moving in a new direction, are twice as likely to engage in commercially productive behaviors — such as paying a premium or referring other customers to the brand.

This shift requires a compatible set of leadership techniques:

1. Clarity — Decide what initiatives you will, and will not, support over the life of your strategy. Give equal thought to those that don’t make the cut, and why.
2. Connections — Success is almost always a collaborative endeavor. Start with a shared understanding of where your function intersects with others to get meaningful work done.
3. Courage — Defend your choice to do less, pushing back on urgent but unimportant requests. All major investments should be able to justify themselves based on measurable contribution to future goals, not those in the past.

— Sharon Cantor Ceurvorst

## Develop the Five Markers of Innovation Potential

Technical expertise is common in the R&D workforce but is not enough to constantly generate new value to the enterprise. What R&D leaders need is a team with high innovation potential that can move beyond solving technical problems — identifying and developing novel solutions for customer challenges and building stakeholder support for big bets.

R&D executives looking for talent to nurture should check for these five behavioral markers of innovators:

1. Result seeking — They leave no stone unturned in achieving goals.
2. Customer empathy — They understand customer needs and expectations, and relate them to relevant technologies.
3. Idea integration — They pinpoint issues and extract key information from large datasets.
4. Influencing — They network proactively.
5. Risk taking — They take calculated risks, and identify opportunities and threats from market information and competitors.

Managers should set performance objectives that develop these capabilities in addition to the usual proficiency oriented goals.

Some examples: Internalize customer perspective while developing ideas and solutions. Lead a topic-specific brainstorming session to generate unique insights. Work with business partners to agree on transformational targets and update them quarterly about whether project teams are taking enough risk in their work.

— Amisha Ajay

# Data & Technology

## Changes in Digital Business Strategy Make Cybersecurity More Collaborative

Business owner accountability for cyber risk is fundamental to effective digital transformation.

Sixty-seven percent of CEOs and senior executives want more technology work done directly within business departments. And 57% of enterprises have made business leaders formally accountable for the cybersecurity risks associated with their resources. That's sensible; cyber risk decisions cannot be made in isolation — organizations must balance such decisions with reputational, competitive and legal impacts.

As a result, those who drive digital change initiatives must join forces with the chief information security officer (CISO) to:

1. Support effective cyber risk assessment with risk acceptance, escalation, exemption procedures and a representative steering committee that includes business unit leaders to enable shared decision making.
2. Make security policy more flexible by allowing choice among technology options that achieve the same objectives. Additionally, rationalize the number of security policies, while co-creating new policies and guidance with the business.
3. Go beyond traditional security awareness training with tools and playbooks that build business technologists' independent cyber judgment.
4. Clearly communicate to the enterprise that the primary role of the CISO is enabling digital innovation — rather than enforcement of security controls.

— Tom Scholtz

## A Shipbuilder's Hackathon to Steer Toward Digital Technology

When organizations introduce a new technology, employees often lack the hands-on experience to leverage its full capabilities. General Dynamics Electric Boat found an engaging and effective solution: an all-day, offsite hackathon. Participants learned how to make the most of a new digital visualization tool by creating deliverables for real-world use.

High adoption was just one outcome. One team's work resulted in a dashboard now in use. Several teams set up new visualization projects. Employees also began experimenting with new technologies beyond the one they'd practiced working with.

Three elements led to success:

1. Using real company datasets for hackathon activities quickly connected visualization to participants' daily responsibilities. It also made the outputs immediately useful for their job.
2. Training high-potential talent as peer guides provided support where facilitators couldn't. The fact that these colleagues had only recently received their own introduction to the tool reduced skepticism in the ranks.
3. Inviting top executives to evaluate the outputs offered employees exposure to senior leadership, and motivated presenters to impress when they discussed their work. Management benefited, too, by witnessing the technology's impact and potential firsthand, acquiring more ideas for innovation, and gaining an understanding of the business case for additional investment.

— Jose Rosario



Access to Gartner Content Caveat

Some content may not be available as part of your current Gartner subscription.

© 2023 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. It consists of the opinions of Gartner's research organization, which should not be construed as statements of fact. While the information contained in this publication has been obtained from sources believed to be reliable, Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner research may address legal and financial issues, Gartner does not provide legal or investment advice and its research should not be construed or used as such. Your access and use of this publication are governed by Gartner's Usage Policy. Gartner prides itself on its reputation for independence and objectivity. Its research is produced independently by its research organization without input or influence from any third party. For further information, see "Guiding Principles on Independence and Objectivity."

Any third-party link herein is provided for your convenience and is not an endorsement by Gartner. We have no control over third-party content and are not responsible for these websites, their content or their availability. By clicking on any third-party link herein, you acknowledge that you have read and understand this disclaimer.

795276

**Gartner**<sup>®</sup>

# Actionable, objective insight

Position your organization for success. Explore these recent editions of the Gartner Business Quarterly:

Already a client?  
Get access to even more resources in your client portal. [Log In](#)

## Journal

Q1 2023 Gartner Business Quarterly:  
Help Your Talent Keep Transformation  
On Track

Achieve digital ambitions to thrive through upheaval

[Download Now](#)



## Journal

Q2 2023 Gartner Business Quarterly:  
Getting From Pilot to Innovation at Scale

Helping C-suite leaders get more experiments into the business for the biggest impact and the best outcomes.

[Download Now](#)



## Webinar

Reduce Information Overload to Boost Your  
Business Strategy and Retain Employees

Identify and mitigate information overload in your organization.

[Watch Now](#)



## Webinar

Don't Own Innovation; Enable It Across  
the Business

Establish competitive advantage by enabling your organization to innovate autonomously.

[Watch Now](#)

