Managing the risks of generative AI

A playbook for risk executives – beginning with governance

May 2023

Table of contents

Introduction 2

What's at stake for business? 4

The new and amplified risks to manage 6

For the chief information security officer 8

For the chief data officer and chief privacy officer 10

For the chief compliance officer 12

For the chief legal officer and general counsel 14

For internal audit leaders 16

For the chief financial officer and controller 18

How to get started 20

Bringing it all together 21

Bottom line 23

Introduction

Something truly revolutionary happened in November 2022. Suddenly, anyone with an internet connection, armed only with the ability to hold a conversation in a chat app, could wield the transformative power of artificial intelligence (AI).

Within a single week, more than a million users, with ChatGPT's help, produced short-form articles, wrote computer code, made art and summarized long sources into pieces perhaps better and more concise than the originals.

Meanwhile, malicious threat actors test-drove ChatGPT to write malware, more believable phishing emails and more convincing fake identities, rapidly and for widespread dissemination – potential harbingers of large-scale fraud, privacy violations, disinformation and cyber attacks.

Mere months after the debut of ChatGPT, generative AI continues to become ever more deeply intertwined into our lives and businesses. We've seen the fastest ramp-up in active consumer users ever. We've seen a leap in capabilities from OpenAI's GPT3 to GPT4 – achievements recorded in coding and mid-level professional writing. In quick succession, tech companies have launched/re-launched competing products; start-ups have released models for bespoke applications; and companies, including PwC, have announced massive investments to create their own "CompanyGPT" for internal use and new service offerings.

But generative AI comes with a warning label. "AI systems with human-competitive intelligence can pose profound risks to society and humanity," concerned citizens, including experts, caution. Even top providers of this technology acknowledge these risks.

Managing them is key to success. If your company wants to launch successful generative AI initiatives and gain a competitive edge, you will need to assess the risks the technology might pose enterprise-wide. For that, you will need a risk management framework that also allows you to embrace opportunity.

A risk-based approach to generative AI will start you on the right digital foot with regulators, consumers and other stakeholders. Earning trust as you deploy generative AI will position you to take full advantage, quickly, of the benefits this game-changing technology offers.

Are companies at risk of trading off trust for speed?

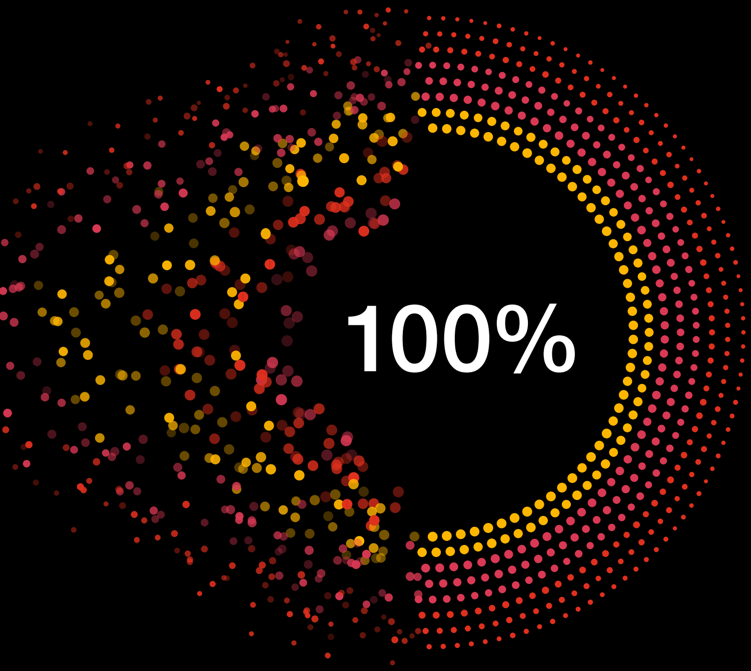

Nearly all business leaders say their company is prioritizing at least one initiative related to AI systems in the near term.*

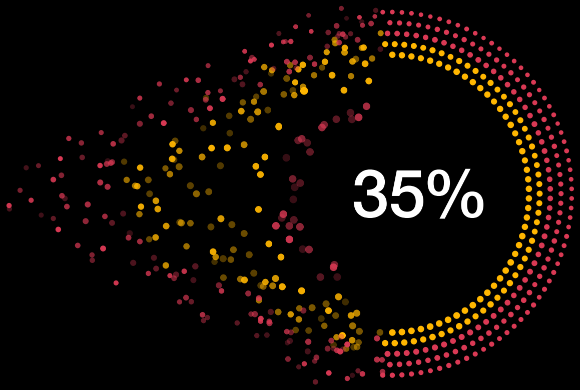

But over the next 12 months, only 35% of executives say their company will focus on improving the governance of AI systems.*

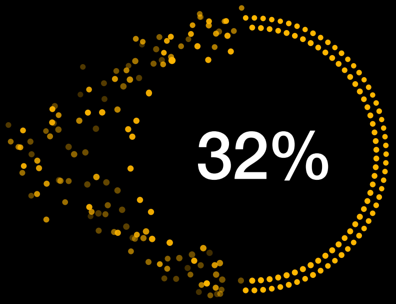

And only 32% of risk professionals say they're now involved in the planning and strategy stage of applications of generative AI. **

*2023 Trust Survey, 500 business executives; ** 2023 US risk survey, 308 risk executives

What's at stake for business?

Generative AI, a powerful subset of all AI technology, is having a truly transformative impact on business. It can automate and enhance aspects of nearly all business operations, including customer service, software development and data analytics.

It might improve how you engage with your customers by personalizing interactions with them. It could automate high-volume tasks, such as processing insurance claims and communications or performing certain software development tasks.

It may make it easier for your teams to understand unstructured data including contracts, invoices, customer feedback, policies, insurance adjuster notes, performance reviews, medical records and more.

Employee productivity could soar. By OpenAI's estimate, approximately 8 of every 10 US workers could use generative AI to automate at least 10% of their work tasks, and this is just the beginning. By automating routine tasks, generative AI tools could free employees to work creatively, innovate and gain a fuller understanding of complex topics and tasks for more advanced critical thinking.

As the demand for this technology continues to grow, so do its capabilities. In four months alone, AI language systems advanced significantly in sophistication and use, and they aren't likely to stop anytime soon.

The key to sustainably riding this growth will be to enlist your risk professionals from the earliest stages. Doing so can help you build confidence in your generative AI projects.

Your risk managers will have to manage new and amplified risks as well as a slew of business, legal and regulatory challenges. One after another, the White House, US Congress, Federal Trade Commission, Cyberspace Administration of China and the European Union (EU) have moved to regulate generative AI. Meanwhile, several nations (Italy, Canada, Spain, France, Germany) started investigations in response to complaints or concerns about generative AI's collection, use and disclosure of personal information without consent, in violation of data protection laws.

Vigilant regulators don't want to waste a hard-earned lesson from previous innovation boom cycles. "The trajectory of the Web 2.0 era was not inevitable – it was instead shaped by a broad range of policy choices," wrote Lina Khan, Chairman of the Federal Trade Commission (FTC).

Your risk professionals can help your company use generative AI safely, securely and resiliently. They can help confirm that it's appropriately private, fair with harmful bias managed, valid and reliable, accountable and transparent, and explainable and interpretable.

In other words, that it's trusted.

The new and amplified risks to manage

We see four broad risks inherent to the technology that organizations need to understand and manage:

Data risks

Error propagation, intellectual property (IP) or contractual issues (due to lack of approvals to use data for such purposes), or misleading and harmful content caused by low-quality data used to train generative AI models.

Model and bias risks

Breach of ethical and responsible AI principles in the language model development, leading to discriminatory or unfair outputs.

Prompt or input risks

Misleading, inaccurate or harmful responses due to unsophisticated prompts or questions being provided to the AI model.

User risks

Unintended consequences due to users becoming unwitting parties to the creation of misinformation and other harmful content. For instance, they might pass off AI-generated hallucinations – erroneous or nonsensical responses – as fact.

You may incur other risks, as well, depending on how your company uses generative AI – particularly if you plan to create proprietary models connected to the foundational models and add proprietary or third-party data.

Your risk professionals are the ones who will activate generative AI toward trusted outcomes, so that trust-by-design, not speed alone, is your value proposition to your customers, investors, business partners, employees and society.

Risk domain specialists should consider the whole host of risks to privacy, cybersecurity, regulatory compliance, third-party management, legal obligations, intellectual property, and collaborate with one another to manage overall *enterprise* risk.

In parallel, you should work with your talent/HR leaders to develop training programs at all levels to familiarize everyone with the risks and rewards of generative AI. Put experienced humans in place to validate “rough draft” generative AI outputs. Monitor human performance to guard against “skills atrophy,” complacency or drop in quality over time.

Established frameworks, such as the NIST AI risk management framework and the ISO framework for AI systems using machine learning, are a good start for designing and deploying trusted AI applications. So too are industry requirements and norms, such as guidance from the Office of the Comptroller of the Currency, Consumer Financial Protection Bureau and the Federal Reserve. And there are more than 800 national AI policies from more than 69 countries, territories and the EU.

Having an effective AI governance strategy will be vital because beyond the risk professionals, many people inside and outside your organization can influence your ability to use generative AI responsibly. They include data scientists, data engineers, data providers, domain experts, socio-cultural analysts, experts in the field of diversity, equity, inclusion and accessibility, affected communities, user experience designers, governance experts, system funders, product managers, third-party entities, evaluators and legal and privacy professionals.

PwC's Responsible AI framework

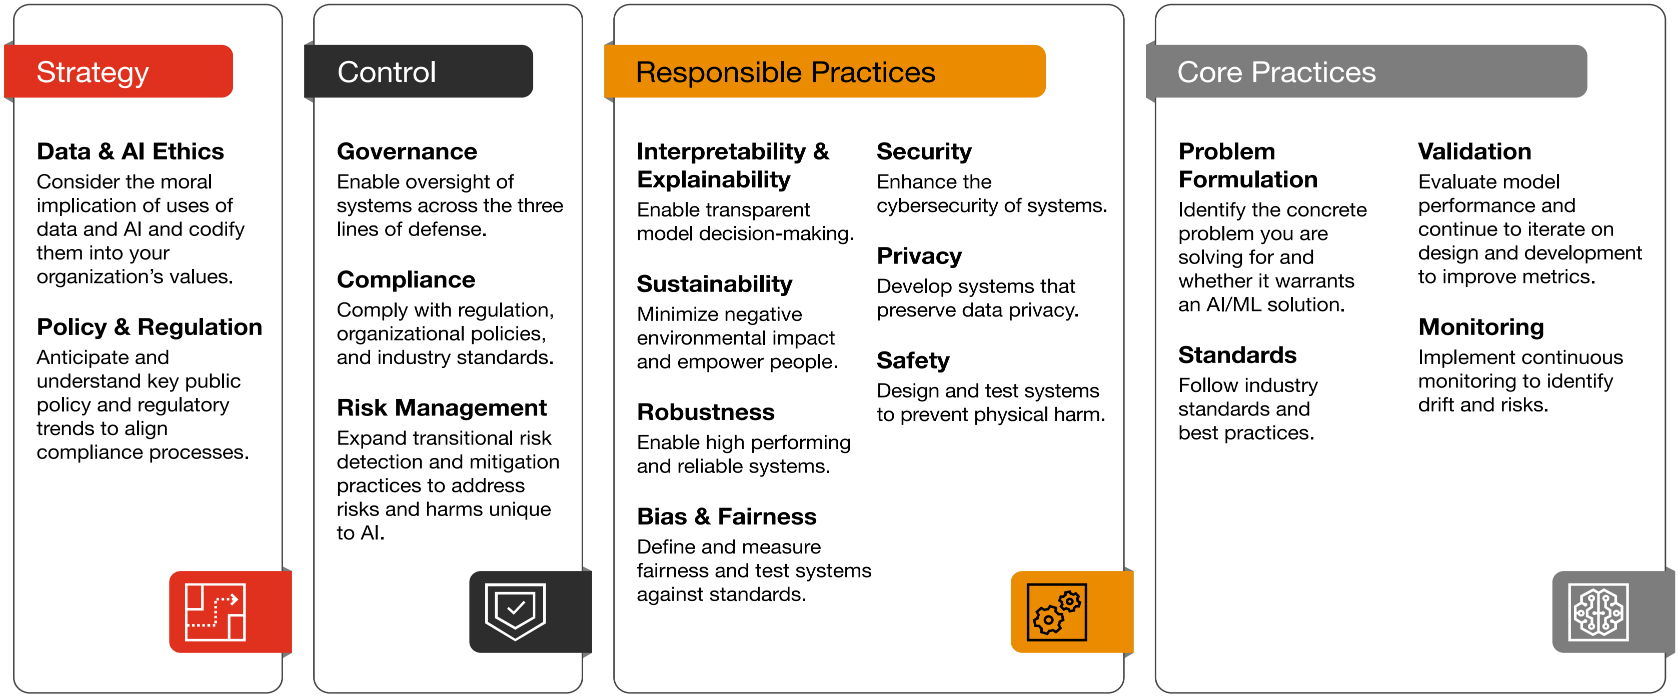

Strategy

Data & AI Ethics
Consider the moral implication of uses of data and AI and codify them into your organization's values.

Policy & Regulation
Anticipate and understand key public policy and regulatory trends to align compliance processes.

Control

Governance
Enable oversight of systems across the three lines of defense.

Compliance
Comply with regulation, organizational policies, and industry standards.

Risk Management
Expand transitional risk detection and mitigation practices to address risks and harms unique to AI.

Responsible Practices

Interpretability & Explainability
Enable transparent model decision-making.

Sustainability
Minimize negative environmental impact and empower people.

Robustness
Enable high performing and reliable systems.

Bias & Fairness
Define and measure fairness and test systems against standards.

Security
Enhance the cybersecurity of systems.

Privacy
Develop systems that preserve data privacy.

Safety
Design and test systems to prevent physical harm.

Core Practices

Problem Formulation
Identify the concrete problem you are solving for and whether it warrants an AI/ML solution.

Standards
Follow industry standards and best practices.

Validation
Evaluate model performance and continue to iterate on design and development to improve metrics.

Monitoring
Implement continuous monitoring to identify drift and risks.

Key risks that GenAI poses and actions that risk executives can take are in the following sections.

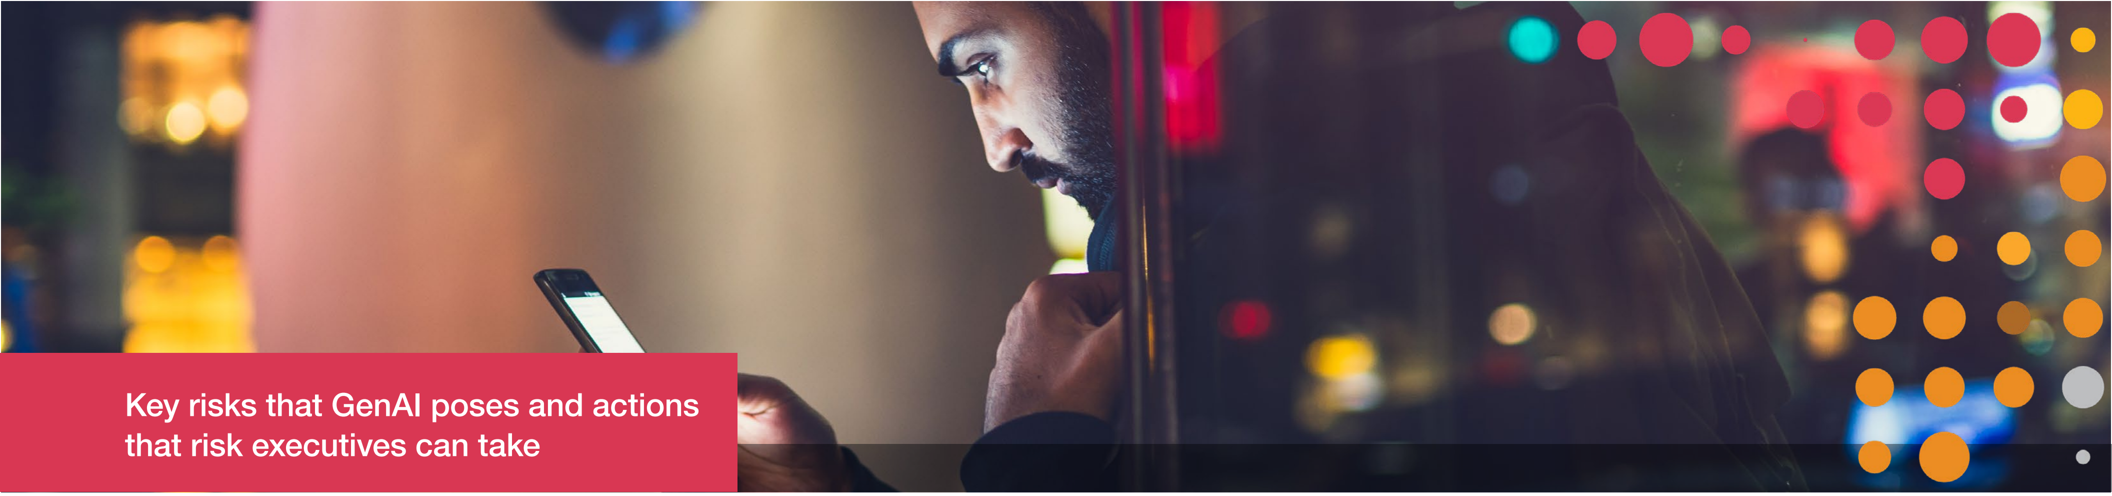

Key risks that GenAI poses and actions that risk executives can take

For the chief information security officer

Generative AI can reduce barriers to entry for threat actors. The most immediate risk to worry about? More sophisticated phishing. More compelling, custom lures used in chats, videos, or live generated "deep fake" video or audio, impersonating someone familiar or in a position of authority. Even before generative AI was launched, researchers mapped 33 offensive AI capabilities to the MITRE ATT&CK framework.

For the CISO, generative AI adds a valuable asset for threat actors to target – and for your organization to manage. For example, they could manipulate AI systems to make incorrect predictions or deny service to customers. Your proprietary language and foundational models, the data and new content will need stronger cyberdefense protections.

1 Automate, update, and upgrade cyber countermeasures.

a. Continuously assess access privileges on a user-by-user basis to identify probable attack vectors and chains powered by generative AI.

b. Ramp up detection countermeasures. Build an endpoint detection and response (EDR) platform using generative AI to detect anomalies with few to no false positives.

c. Continuously evaluate the models' vulnerabilities to adversarial attacks in different domains using emerging evaluation toolkits.

2 Prepare for higher-resolution threat models and insights, predictions and scenarios.

a. Analyze collected vulnerability data and compromise assessment data, and draft assessment reports and remediation plan/activities.

b. Generate executable threat scenarios specific to your company's environment and identify most effective mitigation strategies.

3 Put data loss prevention controls in place.

a. Establish controls to manage public use of your generative AI.

b. Identify potential exfiltration of generative AI-related data.

c. Use generative AI-focused data protection processes to see and safeguard sensitive data in use, in transit and at rest.

d. Identify privacy and other data-risk controls needed to use generative AI in a risk-based manner.

Beyond cyber defense, the CISO will play an important role in the selection of technology providers – the most critical third parties – by determining the kind of model testing results and documentation that vendors need to provide.

4 Protect internal/local generative AI models and associated data.

a. Put controls in place to protect the models against misuse and unauthorized use, in line with the company's legal, privacy, security and ethics policies and procedures.

b. Create internal security controls around generative AI tools to prevent manipulation of data in models or unauthorized use that may cause these tools to deviate from intended parameters.

c. Understand the security posture and controls used by vendors of your internal generative AI instances and related data environments that you use and correct where needed.

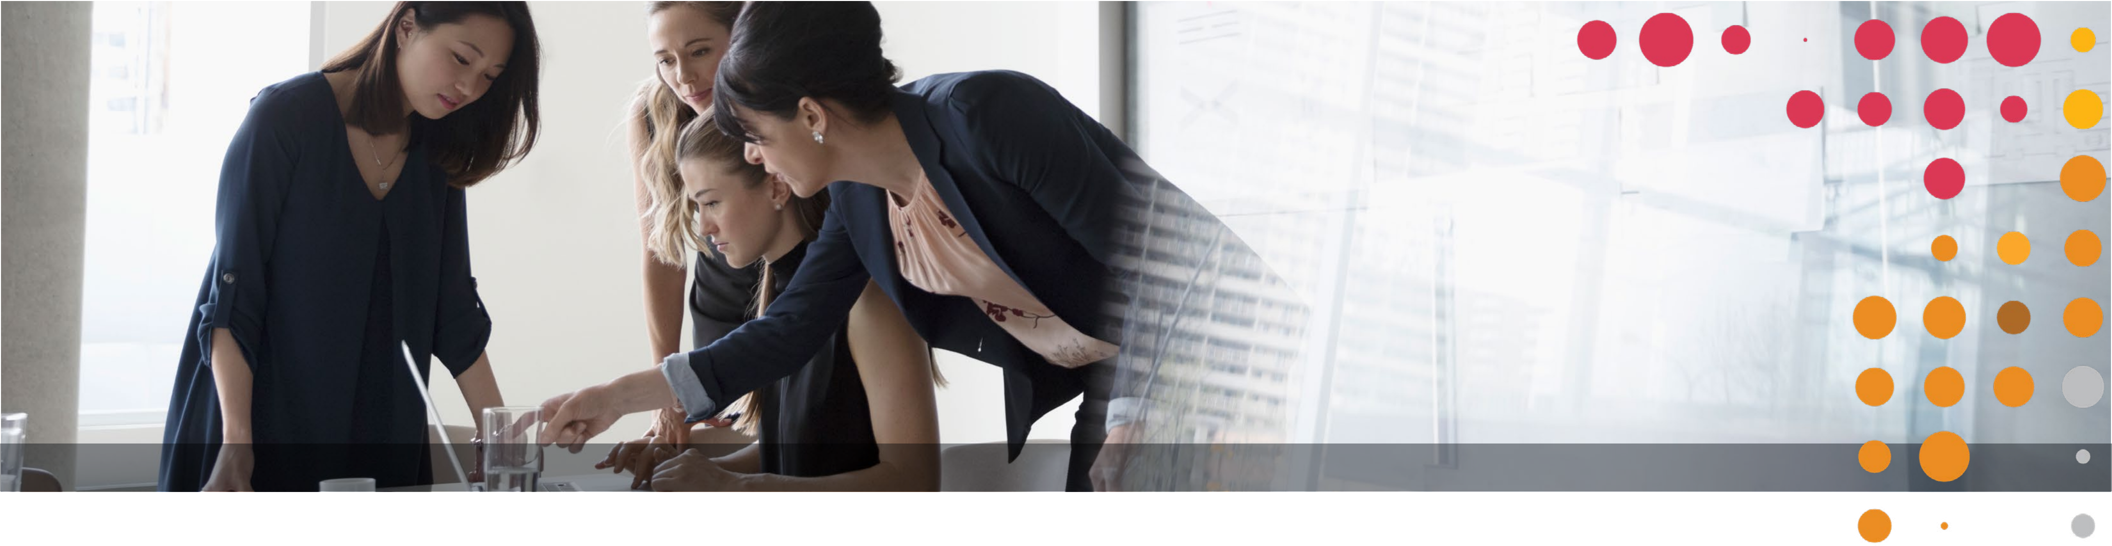

For the chief data officer and chief privacy officer

Generative AI applications could exacerbate data and privacy risks; after all, the promise of large language models is that they use a massive amount of data and create even more new data, which are vulnerable to bias, poor quality, unauthorized access and loss.

Employees entering sensitive data into public generative AI models is already a significant problem for some companies. Generative AI, which may store input information indefinitely and use it to train other models, could contravene privacy regulations that restrict secondary uses of personal data.

Here are specific risk-mitigation actions that CPOs and CDOs should take.

1 Enhance your data governance protocols.

a. Assess data inputs for generative AI, including legal bases, processing purposes and privacy-enhancing technologies (PETs).

b. Craft a strategy for maintaining access to the underlying data needed to operate and improve generative AI.

c. Specify in data controls which data sets can be used in which circumstances. This should go beyond regulated data sets such as personally identifiable information, personal health information and personal consumer information to encompass all sensitive data.

2 Shore up privacy measures.

a. Analyze current and planned use of generative AI, including the use and generation of personal information, against existing privacy laws and regulator guidance.

b. Improve privacy impact assessment processes to account for potential privacy risks and effects of future generative AI use cases. Add bespoke components to address those.

c. Consider whether using data sets to fine tune large language models violates the principles of differential privacy.

d. Mitigate privacy risks upfront by identifying and taking inventory of generative AI uses, datasets, third parties and other related information. Establish controls to limit data use outside of approved cases.

e. Create evergreen inventories of your AI/model use so you can know which areas of your enterprise may need to comply with future regulations.

3 Monitor and protect sharing of organizational data to external generative AI models.

a. Use your cybersecurity protocols to apply data protection controls such as Data Loss Prevention, Cloud Access Security Broker to see and restrict attempts to upload sensitive data to external generative AI services.

b. If you do need to input sensitive data into generative AI, go through your third-party risk management process to do it securely.

For the chief compliance officer

A nimble, collaborative, regulatory-and-response approach is emerging with generative AI, requiring, perhaps, a major adjustment for compliance officers.

1 Keep up with new regulations and stronger enforcement of existing regulations that apply to generative AI.

As companies worldwide plunge into using generative AI in their products and services or race to develop their own models, policymakers are scrambling to set limits and increase accountability.

2 Map your organization's planned use of generative AI applications to existing laws and regulations.

In financial services, for example, applying generative AI in consumer lending will require compliance with an array of rules, such as the Equal Credit Opportunity Act (ECOA), Fair Housing Act (FHA), Home Mortgage Disclosure Act (HMDA), Community Reinvestment Act (CRA) and Truth in Lending Act (TILA).

Federal health care regulators, notably the Food and Drug Administration (FDA), have also been advancing guidelines for AI use. And many states are considering regulating data and AI, possibly affecting all industries.

3 Upgrade your regulatory reporting capabilities.

Prepare for scrutiny from multiple regulatory regimes. Be ready to provide evidence that generative AI applications did not negatively impact your reporting.

4 Monitor how FTC actions may affect your contracts with AI developers.

The agency will assert its legal jurisdiction to handle collusion, monopolization, mergers, price discrimination and other unfair methods of competition that could be brought on by the rapidly developing AI sector.

5 Assess the compliance posture of your generative AI deployments.

Update your standard compliance assessment processes, assessing whether your enterprise uses of generative AI are in compliance with internal policies, applicable laws and regulations.

6 Update your core compliance artifacts, quickly.

Update process maps, controls, control narratives and risk control matrices for processes where generative AI is being deployed. Include governance protocols, detailed policies and procedures and training documentation as part of the update package. You may need to update these artifacts more frequently as this technology quickly changes.

7 Understand how generative AI alters the organization's productivity over time.

Understand also the degree to which it becomes a critical technology in delivery, requiring its own recovery time and recovery point objectives. When implemented at scale, interruptions in availability of generative solutions may affect business resilience and business continuity.

8 Establish strong model governance processes.

Model governance puts guardrails around generative AI use by monitoring the foundational models and company models connected to them. It stipulates auditing and testing to avoid inaccuracy and bias and standardization and transparency that regulators look for. Use a platform that manages, monitors and helps you govern your generative AI models end-to-end.

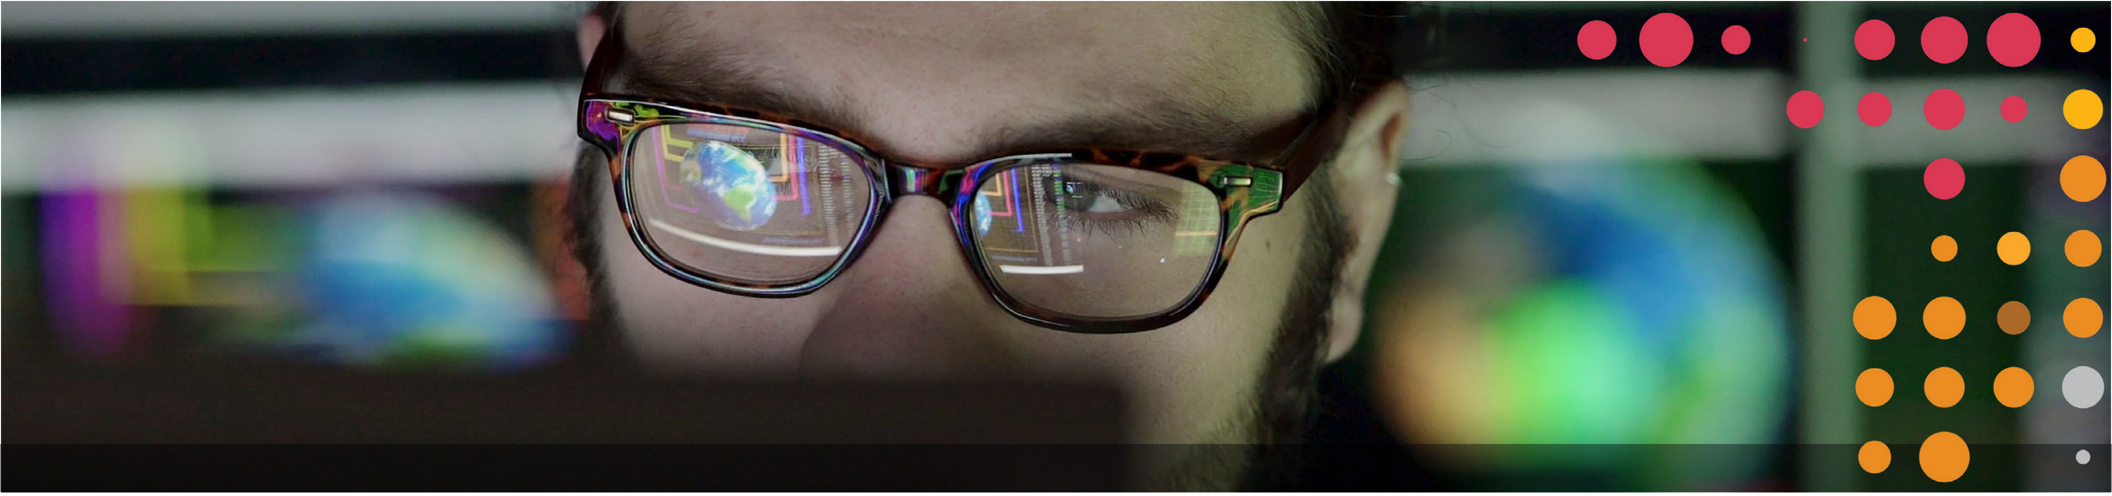

For the chief legal officer and general counsel

Without proper governance and supervision, a company's use of generative AI can create or exacerbate legal risks. Lax data security measures, for example, can publicly expose the company's trade secrets and other proprietary information as well as customer data. And not thoroughly reviewing your generative AI outputs can result in inaccuracies, compliance violations, breach of contract, copyright infringement, erroneous fraud alerts, faulty internal investigations, harmful communications with customers and reputational damage.

Legal departments will need to step up to the challenge. To challenge and defend generative AI-related issues, your legal teams will need deeper technical understanding that lawyers typically don't have. They'll also need to participate in developing generative AI tools. And they will need to collaborate with peers to develop a coordinated response to generative AI's many legal and compliance risks.

Specific risk-mitigation actions that CLOs and general counsel should take include:

1 Limit your IP exposure.

a. Negotiate clauses requiring your generative AI service to segregate trade secrets and other proprietary information, and license protected material in the model's training data.

b. Work with compliance teams to develop policies, procedures and training to mitigate the risk of copyright, patent or trademark infringement liability when using generative AI. Establish an IP review process to screen generative-AI powered processes for potential liability before they are released publicly.

2 Guard against improper secondary uses of data.

a. Identify which data sets lack the necessary consent for use in the generative AI model. Establish protocols to protect that data from unauthorized ingestion.

b. Advise on privacy-disclosure changes to make now to authorize data set use in the future.

c. Review and update policies and procedures for legal risk related to unauthorized secondary uses of data and develop a mitigation plan.

d. Coordinate with other risk and compliance teams to put controls and training in place to restrict unauthorized ingestion of data in the technology.

3 Plan for litigation and investigations.

a. Assess your company's potential exposure to legal claims – e.g., litigation or enforcement actions alleging privacy violations, copyright infringement, bias, harmful communications, regulatory violations, breach of contract – arising from generative AI uses.

b. Assess your vendors, suppliers and other third parties for these same risks.

c. Prepare to defend using generative AI in litigation and investigation processes. Document decisions and uses; keep track of what's happening in other cases.

d. Develop a mitigation plan using your legal exposure assessment. Work with compliance teams to establish a governance framework, policies, training, controls and supervision protocols. Coordinate with internal audit on your generative AI tool's performance and remediation of issues, to support an audit trail.

For internal audit leaders

Auditing will be a key governance mechanism to confirm that AI systems are designed and deployed in line with the company's goals.

But to create a risk-based audit plan specific to generative AI, Internal Audit must design and adopt new audit methodologies, new forms of supervision and new skill sets. Auditing procedures should include elements of both governance and technology audits.

Existing audit procedures may be undermined by four characteristics of large language models: generativity (their open-endedness of applications), emergent abilities (sudden and unexpected appearance of new behaviors), lack of grounding (lack of basis in the real world) and the fact that models are only accessible via application programming interfaces. Practitioners, rising to this challenge, are seeking to evaluate the performance of auditing frameworks for generative AI.

It's difficult and ineffectual to assess the risks that generative AI systems pose independent of the context in which they are deployed. Understanding the problem the company is trying to solve using generative AI is an important starting point.

1 As the third line of defense, Internal Audit should focus on understanding the company's goals and uses, as well as enterprise risk management's view on risks and mitigation plans.

a. Collaborate with key stakeholders as they adopt generative AI for their function. Consider the system and model's design, how it will be used and how those uses fit with company policies. Recommend ways to resolve any issues you find.

b. Adapt your existing internal audit risk assessment process to include generative AI risks.

c. Teach your people about generative AI capabilities and risks.

d. Work with system and model owners to design manual intervention, where possible, for high-risk areas.

e. Engage the audit committee and the board on applications and risk management plans.

2 Create a plan to audit the core data sets used in training, tuning and running the system and models.

a. Follow existing techniques to audit data sets, including elements of data privacy and compliance, data security and data governance.

b. Upgrade capabilities to do audits on large data sets.

3 Design an audit plan around the generative AI systems and models.

a. Learn in detail how the system and models work, the data they're based on, whether that data has previously been audited and how the system and models will be used.

b. Design audit procedures tied to the system or models' objectives while also balancing existing audit objectives such as completeness and accuracy of data.

c. Develop materials to make a generative AI system and its models explainable and interpretable for technical and non-technical stakeholders.

d. Evaluate reliability testing and validation processes.

4 Create an audit plan for output of models.

a. Design procedures to test fairness and bias. Analyze the system and models' robustness and reliability by testing its outputs under various conditions such as changes in input data, model parameters and external factors.

b. Assess outputs by reviewing a representative sample against the defined metrics. Highlight patterns or anomalies.

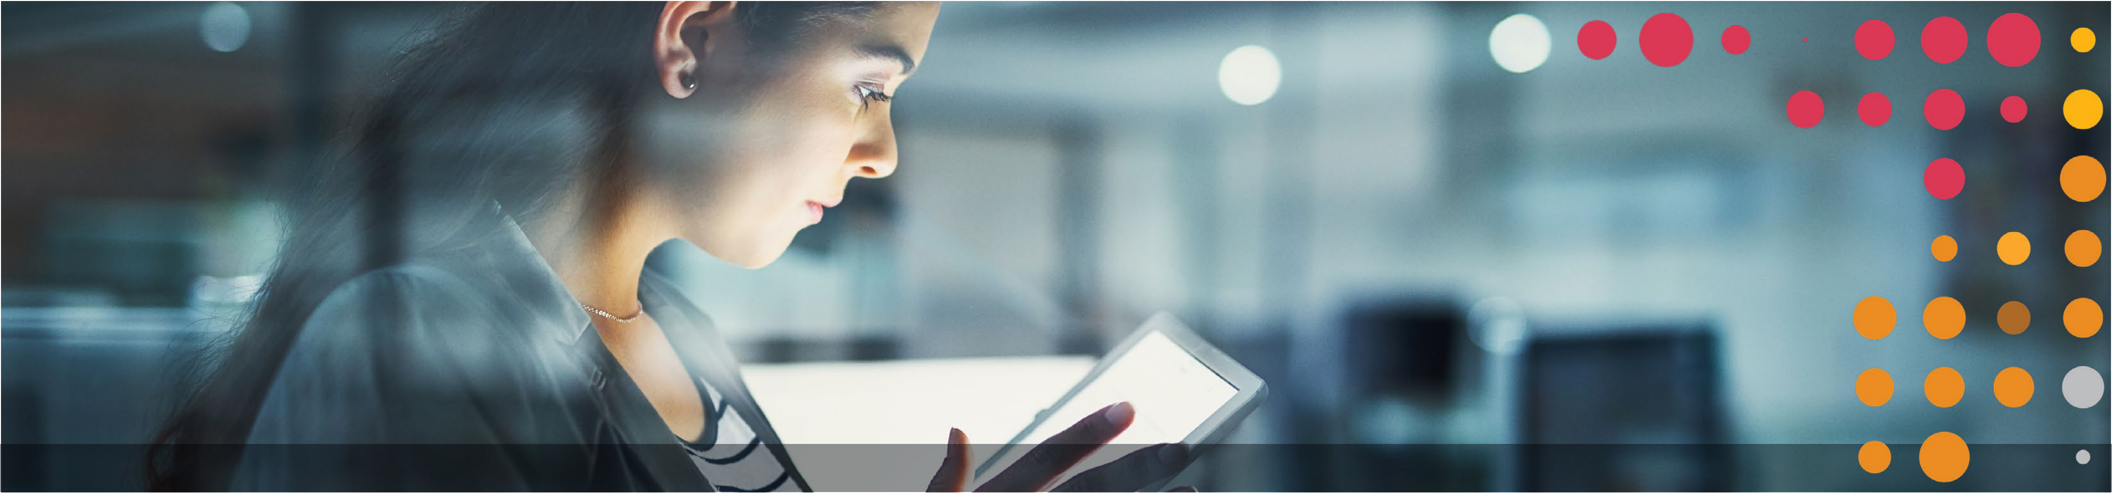

For the chief financial officer and controller

Without proper governance and supervision, a company's use of generative AI can create or exacerbate financial risks. If not used properly, it opens the company to "hallucination" risk on financial facts, errors in reasoning and over-reliance on outputs requiring numerical computation. These are high-consequence risks that CFOs face in the course of their normal duties, often in a regulated environment. Highly-visible, unintended financial reporting errors result in loss of trust with customers, investors, regulators and other stakeholders and have resulted in severe reputational damage that is costly to recover from.

1 Identify internal controls, SEC and Sarbanes-Oxley statutory requirements over financial reporting relevant to generative AI use cases.

a. Create an innovation sandbox for financial and accounting teams to help de-risk generative AI use cases based on a robust innovation and risk management framework.

b. Lead internal policy review in consultation with General Counsel to develop more expansive access to and use of enterprise data of all types, while remaining compliant with applicable statute and regulatory requirements.

c. Work with the Chief Digital and/or Data Officer to curate and prepare the necessary financial and accounting data that should be fed into the generative AI system.

d. Partner with the CISO to confirm the security and confidentiality of financial and accounting data.

e. Co-develop generative AI solutions in collaboration with reputable firms to benchmark and compare current approaches vs. generative AI-enabled outcomes toward completeness, accuracy, timeliness and reliable disclosures (i.e., regardless of how they are prepared).

2 Inventory financial and accounting tasks in the organization.

a. Explore how generative AI might be developed, deployed and scaled throughout the organization to positively impact the preparation of critical financial or, increasingly, non-financial information that is important to internal and external stakeholders (e.g., ESG or climate and diversity, equity and inclusion reporting).

3 Develop and implement a human resources upskilling and reskilling plan

a. Do a task inventory of accounting and finance professionals and walk through an innovation exercise to determine which of them are (i) lower value propositions of human resources and time investment and (ii) promising candidates for augmentation and/or outsourcing entirely to a generative AI solution.

How to get started

This moment calls for an enterprise-wide playbook on responsible generative AI because of the scope, novelty and breadth of risks. That's the reason why the NIST AI risk management framework puts governance at the center of its framework. It's the mechanism that forces the necessary integrated risk management.

To build meaningful governance, it's important to first establish a strategy and vision for governance, supported by process and policy that different groups can follow. Form the governance committee and name its champion.

How do you begin your company's deployment?

1. Establish a governance structure and enterprise-wide generative AI risk management framework.
2. Engage with the leading AI technology providers.
3. Perform legal diligence on your contracts and intellectual property.
4. Engage your employees in identifying use cases in their work and for the company's customers.
5. Evaluate and prioritize use cases based on risk/reward. Look for common "patterns" that apply to the majority of use cases, are reusable and applicable to future use cases. Start with those use cases.
6. Build your generative AI factory, with tooling and enhancements and with the appropriate security and controls.
7. Do trial runs in sprints.
8. Roll out for broader use under a dedicated enterprise program office.
9. Monitor your foundation models and applications for compliance and drift periodically, using a model governance tool such as Model Edge.
10. Adopt robust generative AI systems and model metrics and monitor for "concept drift," toxicity and bias.

Bringing it all together

Let's look at two examples in which executives and their teams might collaborate to manage the risks, and how strong governance can help.

Example 1: The opportunities and risks of building a generative AI-powered medical consultation chatbot

A healthcare provider contemplates using generative AI to offer medical advice in place of tele-health sessions with clinical staff. The provider gathers years of patient data, symptoms, diagnoses and treatments to train the model.

- CDO: Make sure the data is accurate and clean with no overweighting of certain populations, age groups, etc.
- CCO: Determine whether the use of the data meets compliance obligations under state boards of health, HIPAA/HITECH, CMS rules and state Medicaid rules.
- CPO: Collaborate to take a privacy-by-design approach and make it clear to users how their inputs will be used and which data will be retained.
- CTO: Design a dedicated instance for this use case so as to not inadvertently commingle the data with other operational generative AI tools.

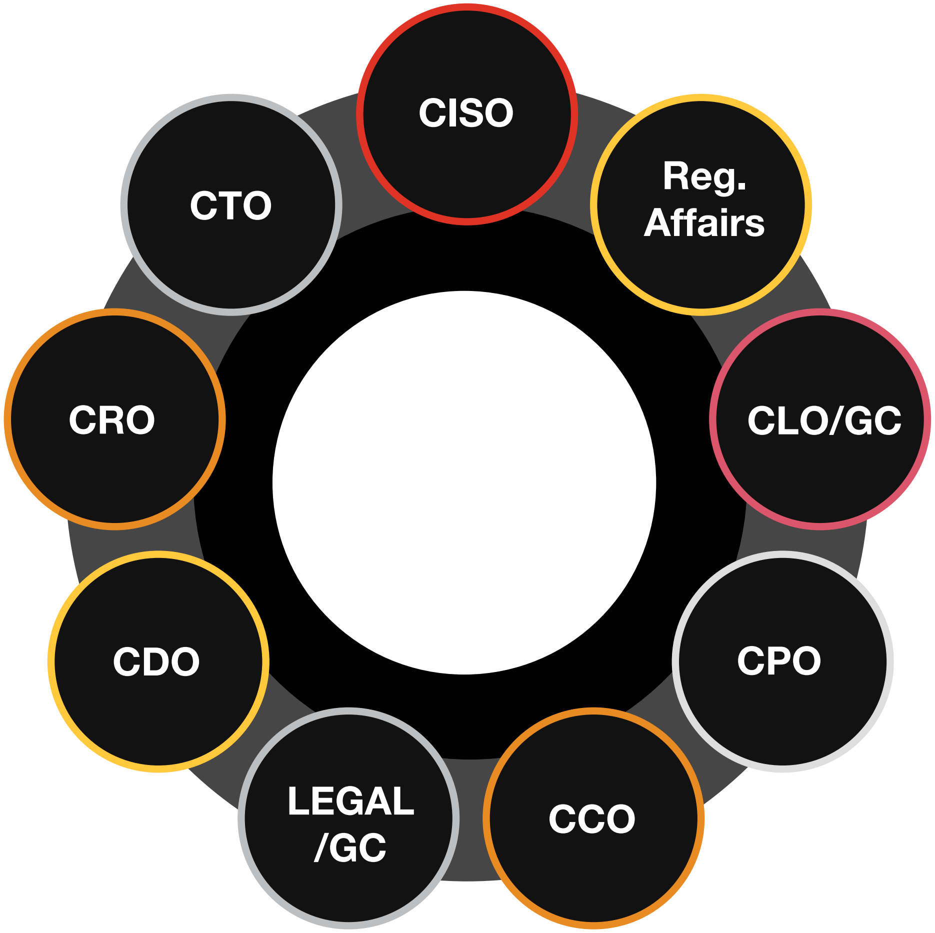

- Legal/GC: Negotiate contractual assurances with the generative AI platform that patient data will remain segregated from the AI model's public instance.

- CISO: Designate this application and data store as a "crown jewel" and provide adequate protections for it based on the most sensitive data classification.

- Internal audit: Develop an audit risk assessment and plan around the proposed system and model – including legal and compliance risks based on HIPAA, HITECH and/or CMS policy and procedures – and assess reliability and performance of system and models.

- CRO: Coordinate with the CCO to establish policies, training, testing and controls to confirm that AI-generated medical advice is accurate and compliant with state medical board standards.

Example 2: Validating credit analysis efficiently and with awareness of the risks

A bank considers using generative AI to automate manual processes for performing annual credit checks on every counterparty documented in counterparty credit evaluations, as well as quarterly checks for high-risk customers based on market events and other triggers.

- CDO: Make sure the data are accurate and clean, and that there is no inherent bias weighting towards certain demographics. Set up dedicated sandbox instances to support the product.

- CCO: Update process maps and compliance artifacts to show how the technology is being used to reach decisions and to demonstrate evidence that it complies with regulatory requirements such as ECOA, FHA, TILA, etc.

- Regulatory affairs: Update reporting protocols.

- CPO: Call for a privacy-by-design approach, making it clear to end users how the data they provide will be used and what will be retained.

- CLO/general counsel: Negotiate contractual assurances from credit agencies and other data vendors to allow use of their data for generative AI, as well as assurances from the generative AI platform that customer data will not be commingled with or used to train other instances.

- CFO/controller: Confirm that SOX frameworks address the implications of AI use in processes that impact financial reporting, and that external auditors are aware and engaged in necessary changes to consider implications for audits.

- CISO: Designate this application and data store a "crown jewel" and protect it based on the most sensitive data classification.

Bottom line

Using generative AI requires constant, swift changes and adaptations – for AI developers, commercial users, investors, policymakers and citizens.

To truly get the most benefits from this groundbreaking technology, you need to manage the wide array of risks it poses in a way that considers the business as a whole. Stakeholders will need to come together as never before to consider all the effects and issues of bringing on board each new generative AI solution. Demonstrating that you're balancing the risks with the rewards of innovation will go a long way toward gaining trust in your company – and in getting a leg up on the competition.

Ultimately, the promise of generative AI rests with your people. Invest in them to know the limits of using the technology as assistant, co-pilot, or tutor, even as they exploit and realize its potential. Empower your people to apply their experience to critically evaluate the outputs of generative AI models – after building your enterprise risk guardrails. Every savvy user can be a steward of trust.

Contact us to learn more

Sean Joyce
Global Cybersecurity & Privacy Leader, US Cyber,
Risk & Regulatory Leader
sean.joyce@pwc.com
202 684 5782

Mir Kashifuddin
Data Risk & Privacy Leader, Cyber Risk & Regulatory
mir.kashifuddin@pwc.com
817 683 8296

Vikas Agarwal
Risk & Regulatory – Financial Services Leader
vikas.k.agarwal@pwc.com
216 789 0314

Jennifer Kosar
Partner, Digital Assurance & Transparency
jennifer.kosar@pwc.com
646 591 2070

Bret Greenstein
Partner, Data & Analytics
bret.greenstein@pwc.com
475 204 1458

Tim Persons
Partner, Digital Assurance & Trust
tim.persons@pwc.com
202 302 2764