Australia

pwc

A playbook for risk executives – beginning with governance

Managing the risks of generative AI

Artificial Intelligence
Accelerate responsibly.

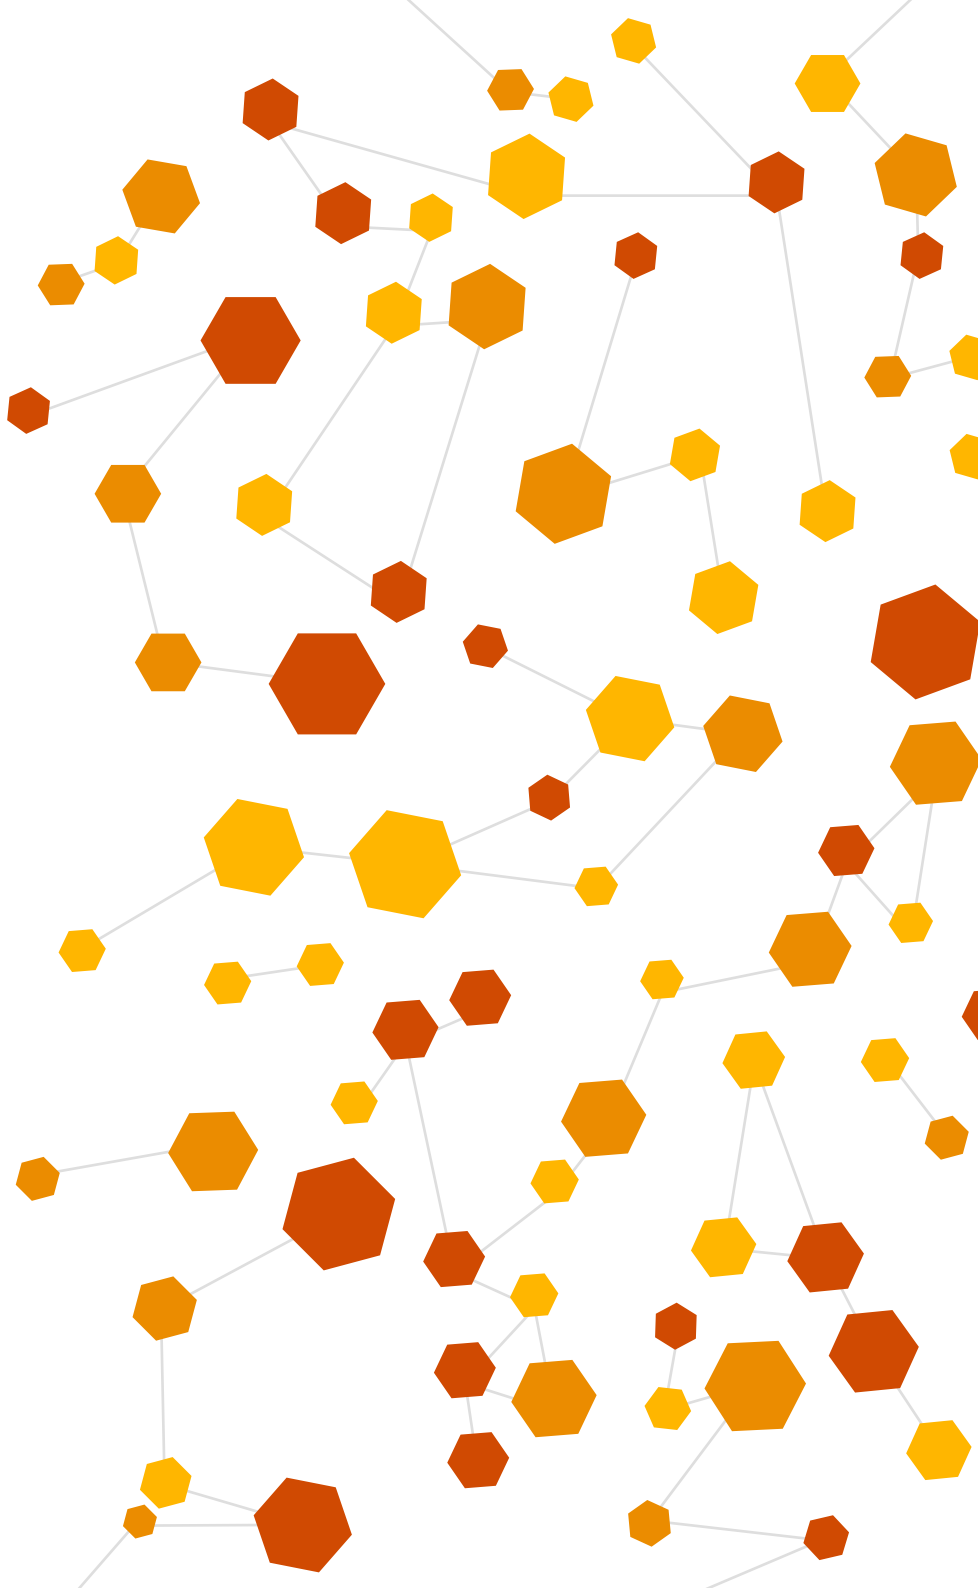

Table of contents

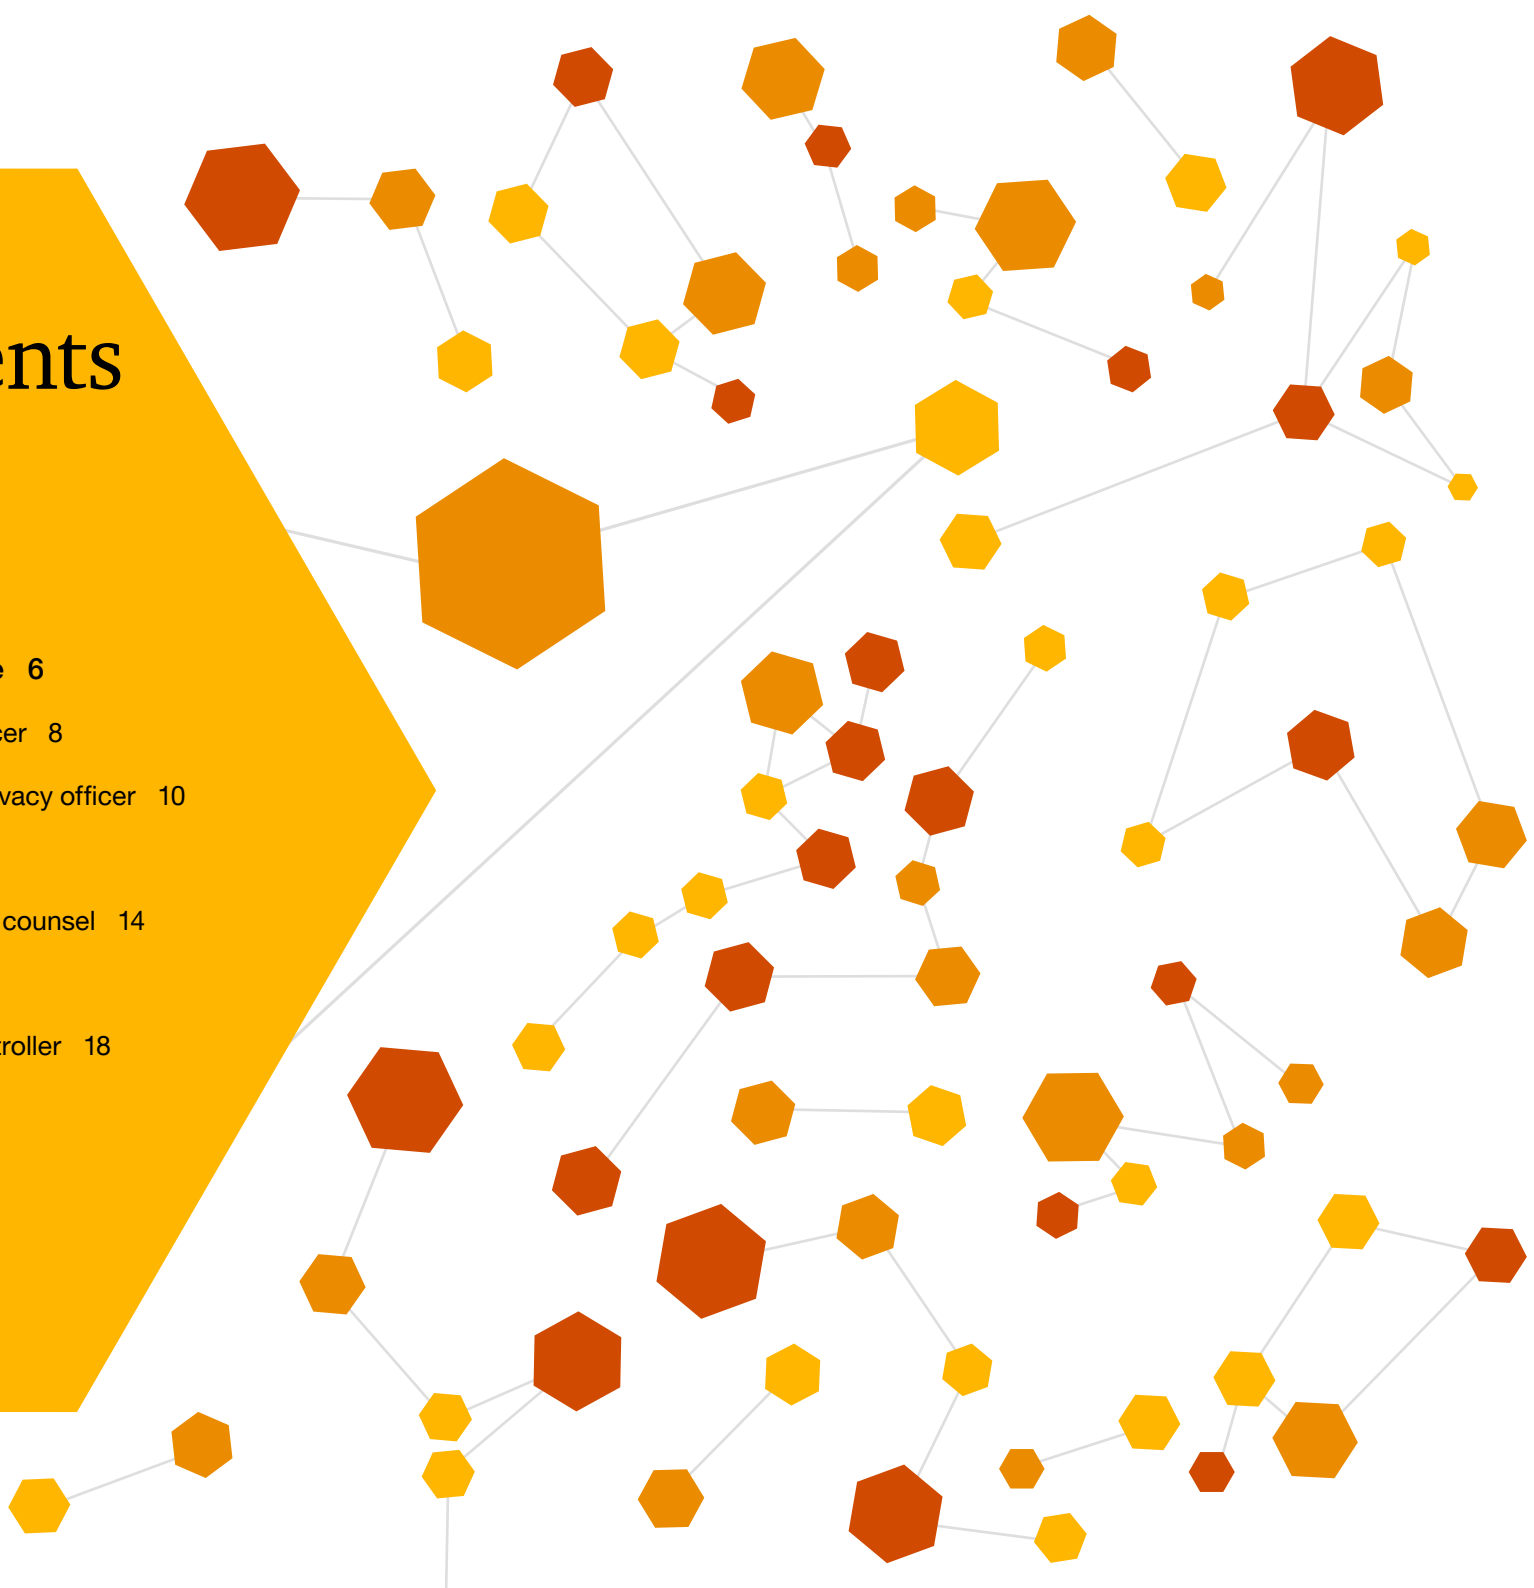

Introduction

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 summarised long sources into pieces perhaps better and more concise than the originals.

Meanwhile, malicious threat actors test-drove generative AI 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?

7 in 10

Australian businesses are investing in new and emerging technologies, including AI. (Source: PwC, 2023)

Less than 1 in 20

Australian businesses rate themselves as 'mature' in navigating the risks of AI. (Source: Fifth Quadrant, 2022)

Only 1 in 2

employees feel positive about the impact that AI will have on their career (Source: PwC, 2023)

What's at stake for business?

Generative AI, a powerful subset of Artificial Intelligence (AI), 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 personalising 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* 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.

In Australia, as recently as June 2023, the Department of Industry, Science and Resources has sought external opinions on the potential risks of AI and strategies to mitigate them, with the aim of promoting safe and responsible use. A discussion paper has been released to incorporate the received feedback into regulatory and policy responses.

"Using AI safely and responsibly is a balancing act the whole world is grappling with at the moment. ... The upside is massive, whether it's fighting superbugs with new AI-developed antibiotics or preventing online fraud. [But] there needs to be appropriate safeguards to ensure the safe and responsible use of AI." - The Hon Ed Husic, Minister for Industry and Science.

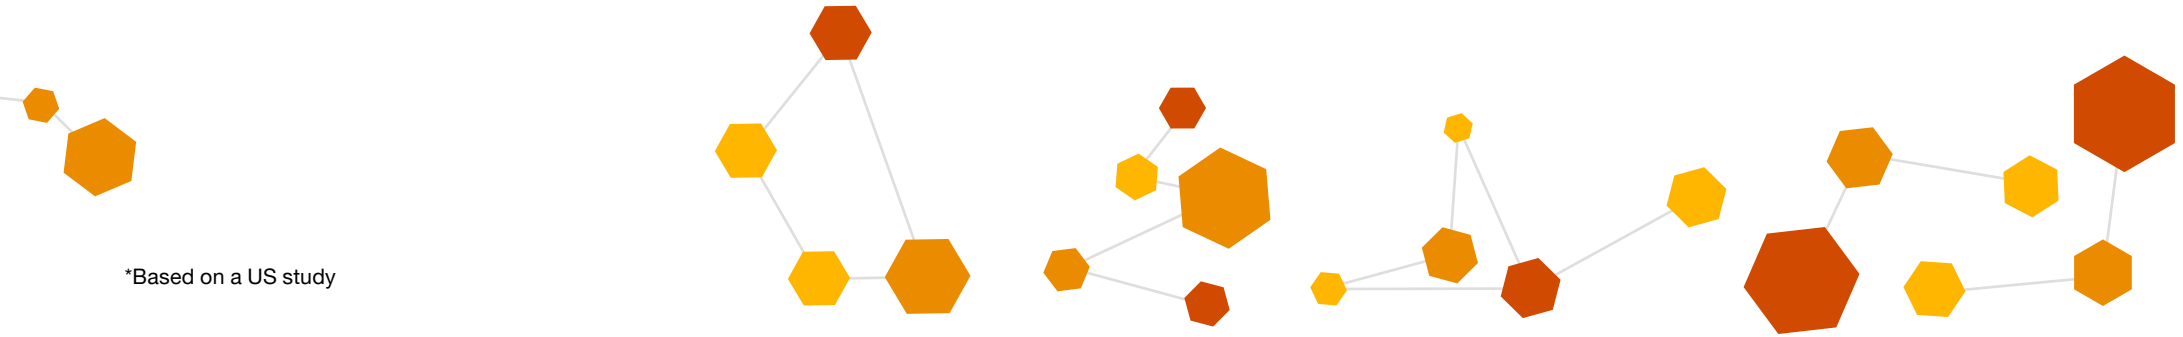

*Based on a US study

Your risk professionals can help your company accelerate responsibly with Generative AI. 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 organisations 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 familiarise 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.

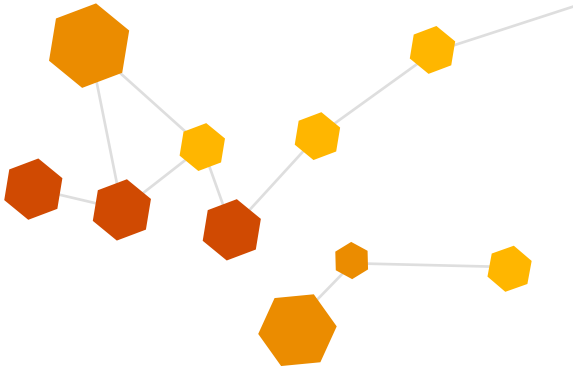

Established frameworks, such as the NIST AI risk management framework, the ISO framework for AI systems using machine learning, and Australia’s AI Ethics Framework are a good start for designing and deploying AI applications. So too are industry requirements and norms, such as guidance from the Office of the Australian Information Commissioner (OAIC), the Commonwealth Ombudsman. And there are more than 800 national AI policies from more than 69 countries, territories and the EU.

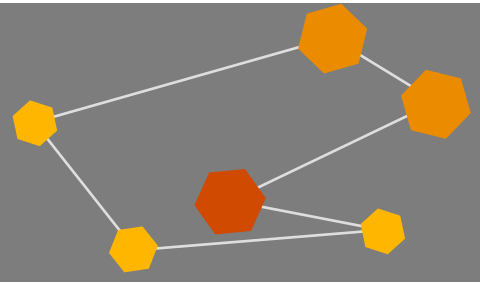

Having an effective AI governance strategy will be vital because beyond the risk professionals, many people inside and outside your organisation can influence your ability to use generative AI responsibly. They include data scientists, data engineers, data providers, domain experts, socio-cultural analysts, 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.

Key AI Governance Considerations for AI

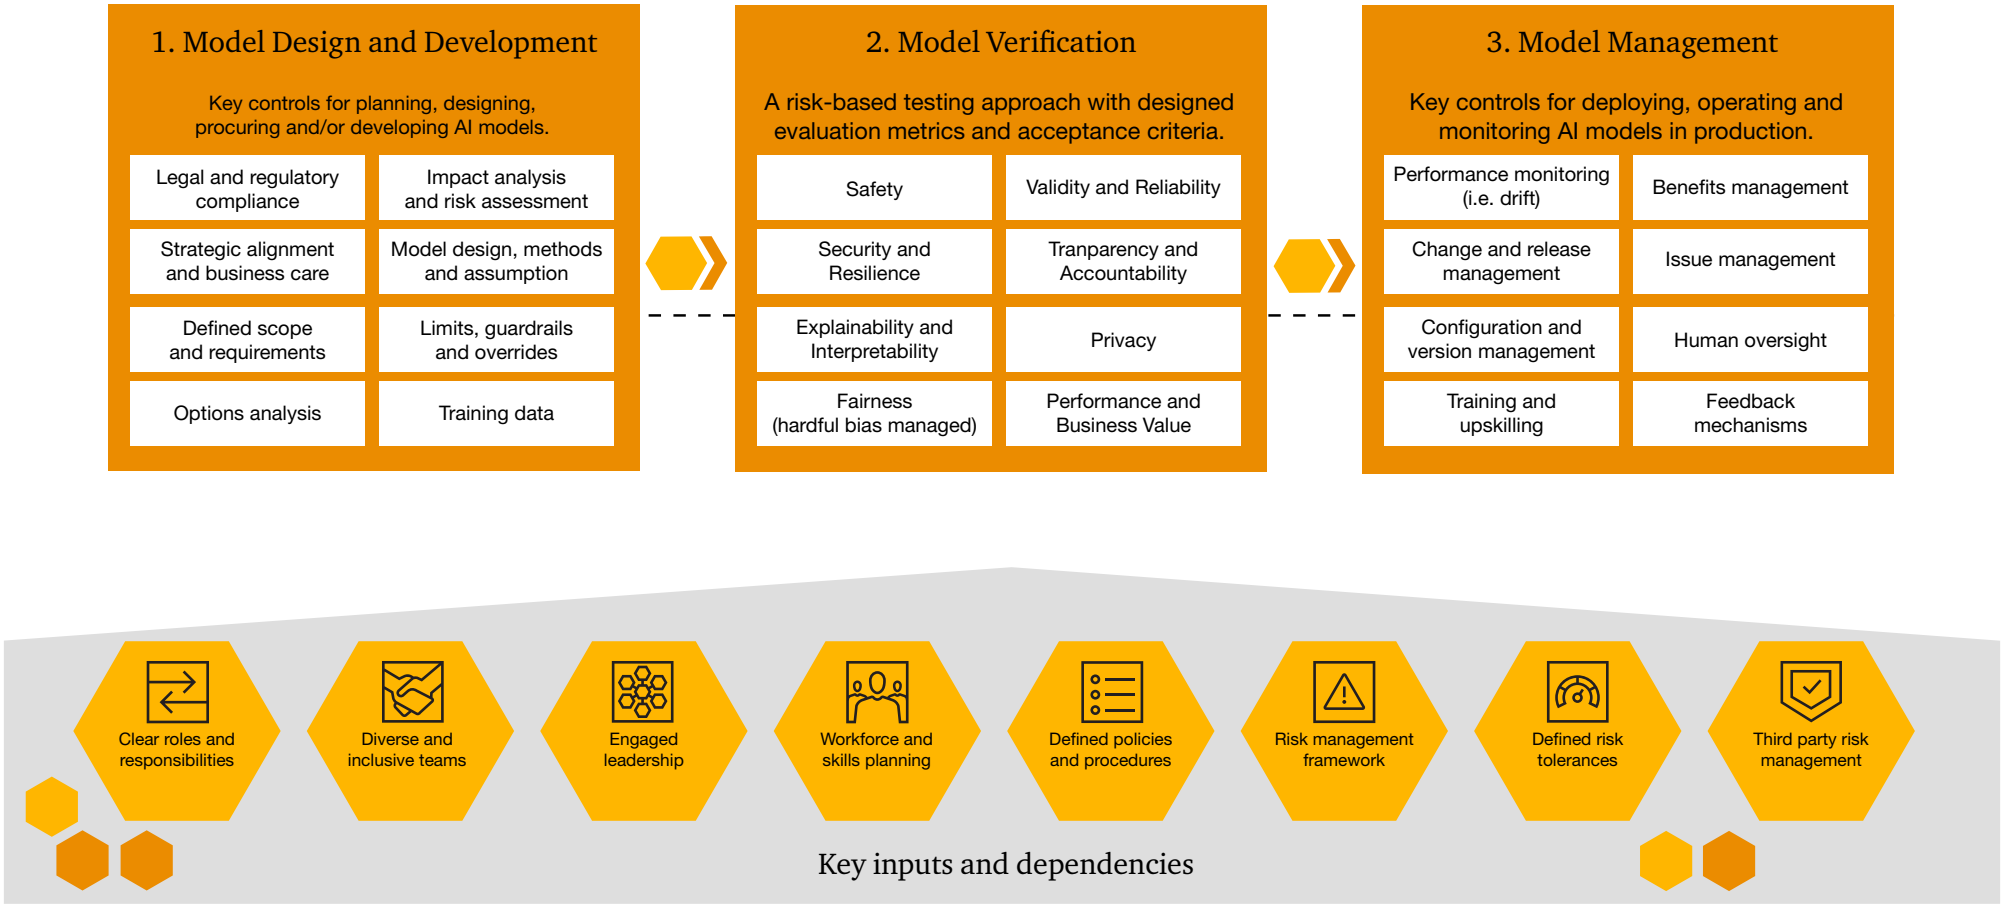

Key risks that generative AI poses and actions that risk executives can take are in the following sections.

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 organisation 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 cyberdefence 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. Analyse 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 defence, 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 unauthorised 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 unauthorised 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, unauthorised 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.

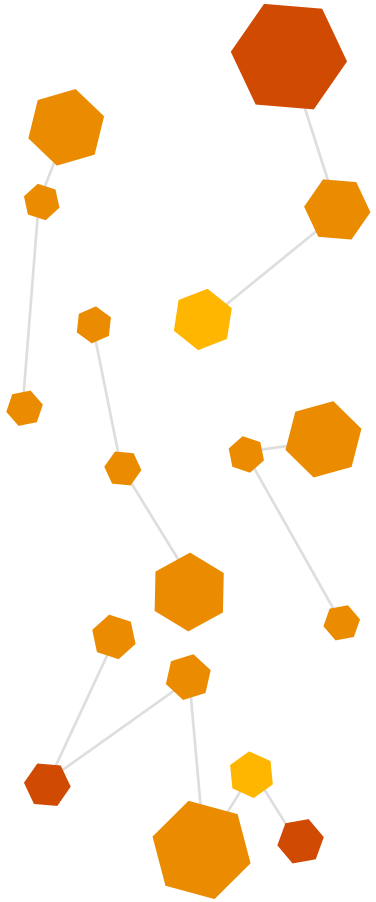

1 Enhance your data governance protocols.

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. Analyse 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.

f. Creation of a Triage process that prioritises proactive assurance for high risk AI models that are planned or already implemented.

g. Creation and implementation of a training and awareness campaign.

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

a. Use your cybersecurity protocols to apply data protection controls such as data loss prevention and cloud access security brokers 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.

c. Use and set up generative AI services within your own compliant and credentialled environment.

d. Continuously monitor and control any generative AI outputs to ensure compliance with policy and governance requirements.

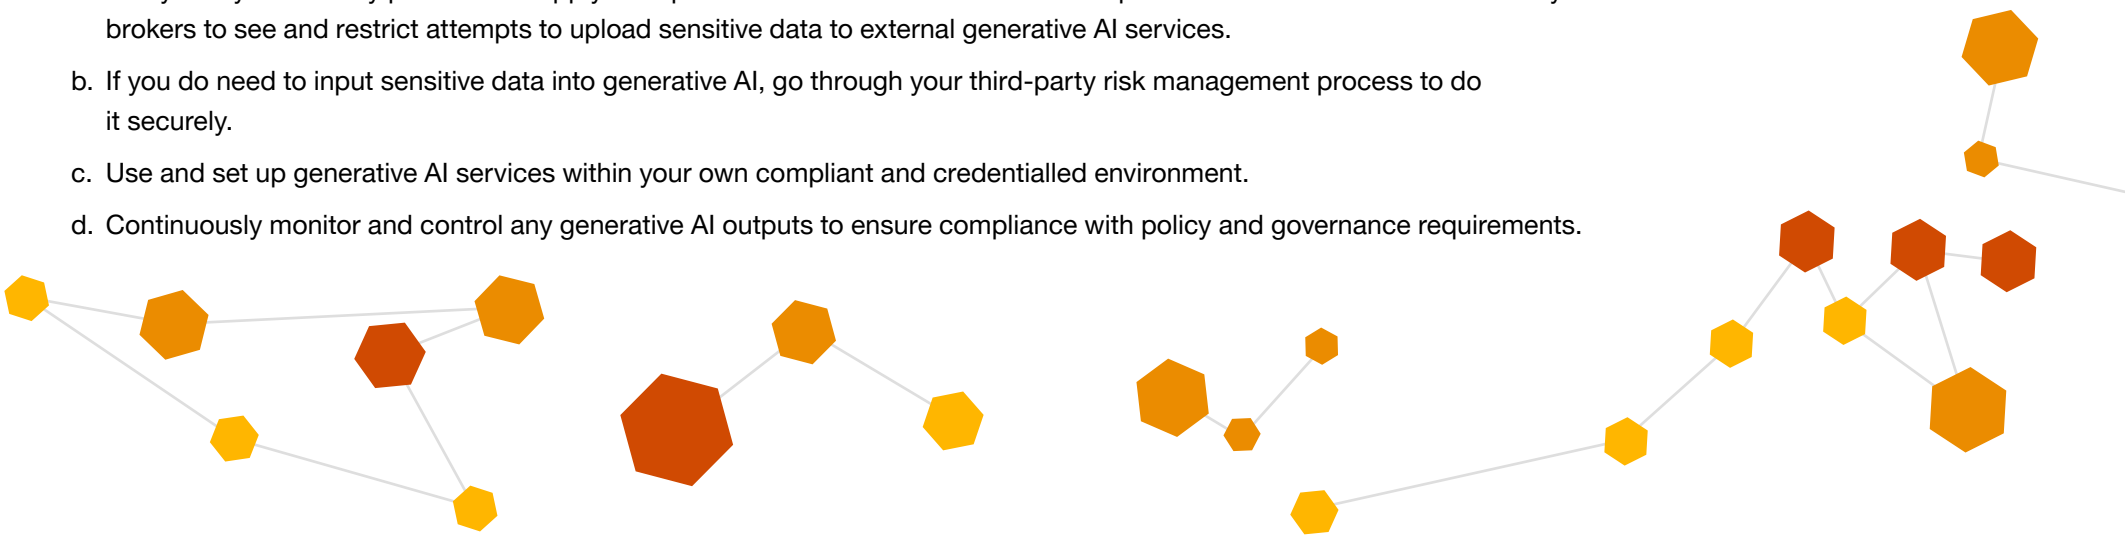

For the chief compliance officer

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 organisation's planned use of generative AI applications to existing laws and regulations.

Depending on the nature of your organisation's proposed used of AI, a range of laws in the existing legislative environment and at common law may apply. For example, your organisations should consider:

- intellectual property laws
- the Privacy Act 1988 (Cth) and other federal and State health information and record legislation
- the Security of Critical Infrastructure Act 2018 (Cth)
- APRA standards for financial services entities
- employment laws (e.g. Fair Work Act 2009 (Cth)) and enterprise bargaining agreements
- federal and state work health and safety laws such as Work Health And Safety Act 2011 (Cth)
- any federal and state anti-discrimination law
- state human rights charters such as Charter of Human Rights and Responsibilities Act 2006 (Vic)
- surveillance laws
- contractual obligations
- negligence laws.

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 ACCC actions may affect your contracts with AI developers.

Keep abreast of how the ACCC's views on AI evolve in the context of collusion, monopolisation, mergers, price discrimination and misuse of market power. These could impact the agreements you enter with AI developers and your use of large datasets as part of LLM development.

5 Assess the compliance posture of your generative AI usage.

Uplift your standard compliance assessment processes to continuously monitor and intervene Generative AI based on usage and ensure compliance with internal policies, applicable laws and regulations. Introduce a virtual risk and compliance team dedicated to understanding, educating and monitoring your organisation's use of Generative-AI.

6 Update your core compliance artefacts, 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 artefacts more frequently as this technology quickly changes.

7 Understand how generative AI alters the organisation'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 standardisation and transparency that regulators look for. Use a platform that manages, monitors and helps you govern your generative AI models end-to-end.

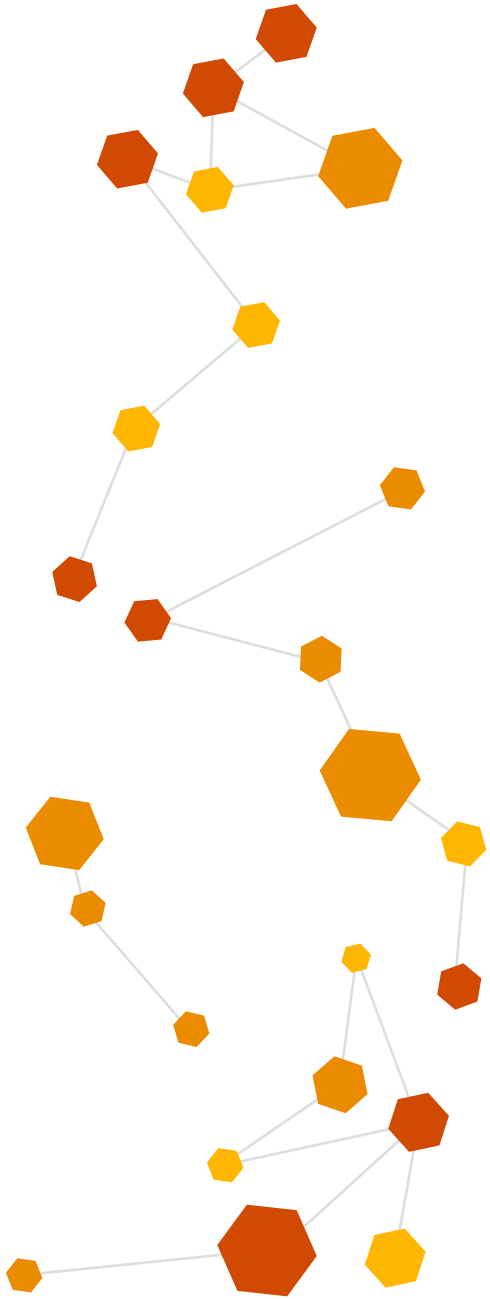

For the chief legal officer and general counsel

Without proper governance and supervision, a company's use of generative AI can create or exacerbate legal risks. For example, a failure to appropriately consider intellectual property laws in relation to AI use could result in infringement claims and an inability to protect key business knowledge and know-how developed by AI solutions. Further, a failure to appropriately secure AI systems could expose personal information of customers, employees and other stakeholders to unauthorised access creating risk of regulatory fines.

Outputs of AI can also create significant risk to an organisation. Inaccuracies, compliance violations, breach of contract, copyright infringement, erroneous fraud alerts, faulty internal investigations, harmful communications with customers, defamation claims and reputational damage could all result if outputs are not reviewed and controlled.

Legal departments will play a key role in paving the way for an organisation to get the most from AI in a responsible manner. 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 Assist the Chief Compliance Officer with identifying applicable laws and regulations for AI.

a. Assist with identifying applicable laws both domestic and international that may apply to the use of AI within the organisation. These laws may relate to inputs of AI (surveillance and privacy laws) and outputs (anti-discrimination and consumer protection laws).

b. Create a plan to ensure ongoing compliance with legislative requirements and legal obligations, including setting out clear interpretation of the applicable laws and how they can be navigated.

2 Manage IP effectively.

a. Consider how IP laws may apply to both the inputs and the outputs of AI and put in place strategies to mitigate the risk of the operation of those laws on the organisations.

b. Where possible, look to negotiate clauses requiring your generative AI service to segregate data provided by your organisation, and license protected material in the model's training data.

c. 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.

3 Guard against unauthorised use of data.

a. Consider what rights and obligations that your organisation has in relation to particular data sets and on what basis your organisations collects, uses and discloses that data.

b. Identify which data sets lack the necessary rights for use in the generative AI model. Establish protocols to protect that data from unauthorised ingestion.

c. Consider changes to your collection notices and contracts to allow the organisation to utilise data sets in AI models in the future.

d. Review and update privacy policies and procedures to manage legal risk related to uses of data in the AI model and develop a mitigation plan to ensure that data is appropriately protected from unauthorised disclosure and use.

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

4 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, defamation – 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.

e. Consider developing processes and procedures for dealing with complaints or claims in relation to generative AI which involves lawyers as early as possible to manage the risk of escalation of claims and to protect your organisation's interests.

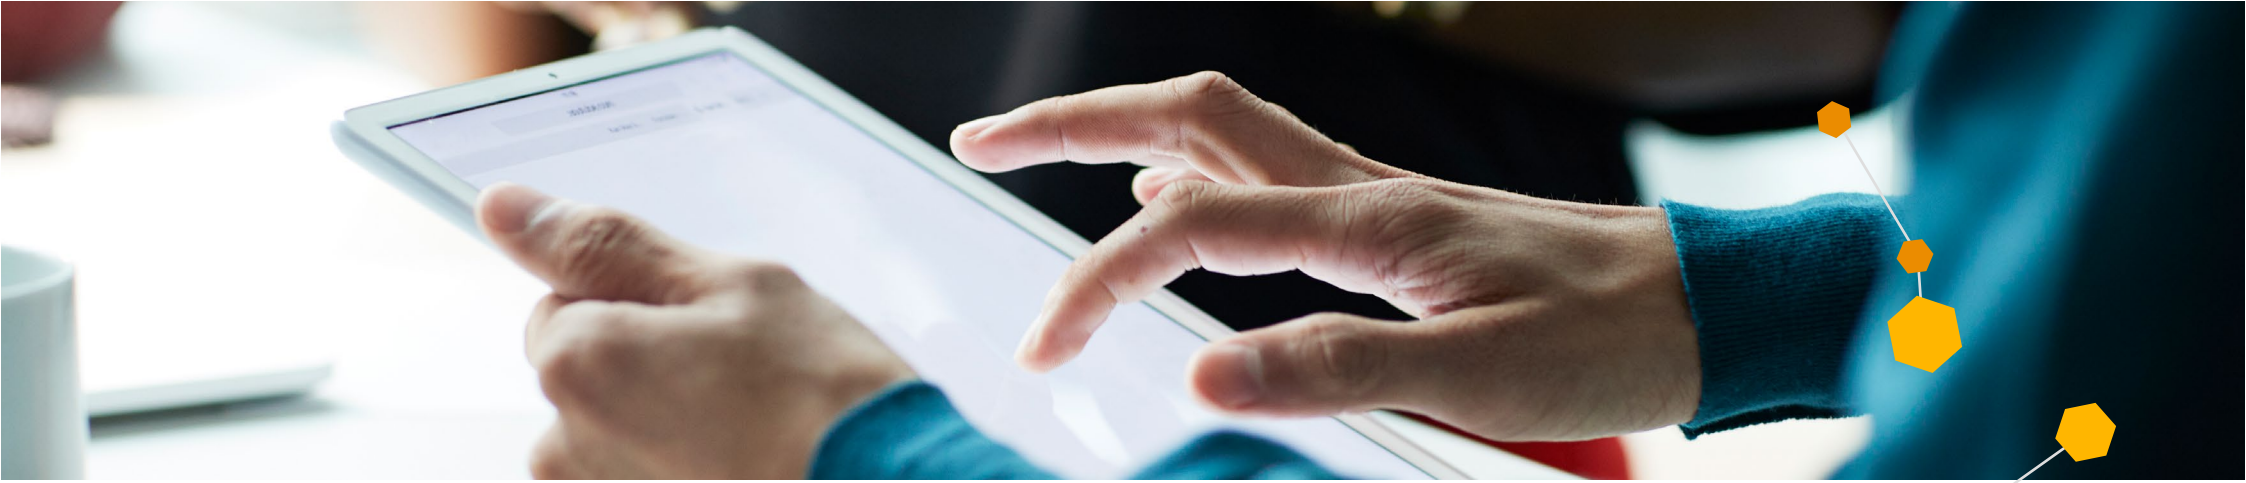

For internal audit leaders

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 data 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 behaviours), 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 defence, 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. Analyse 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.

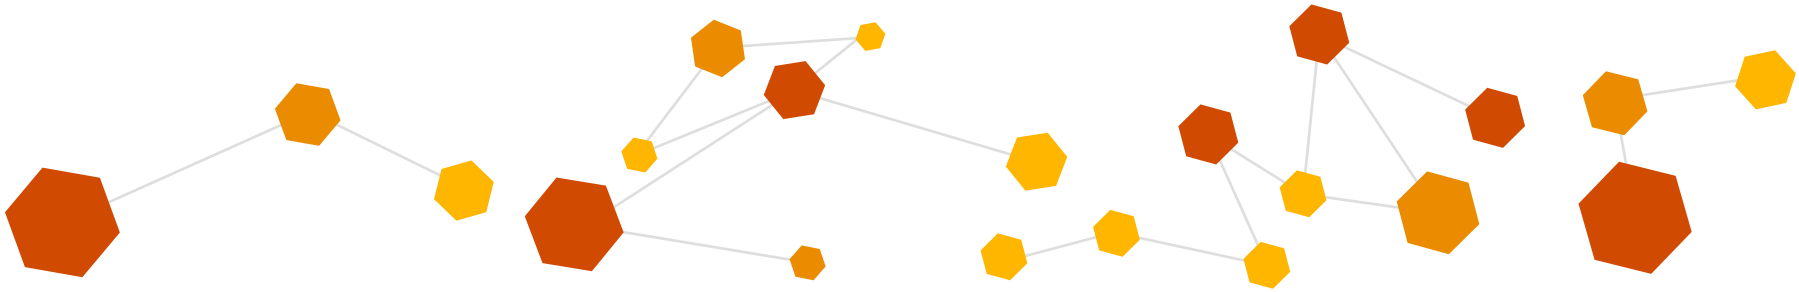

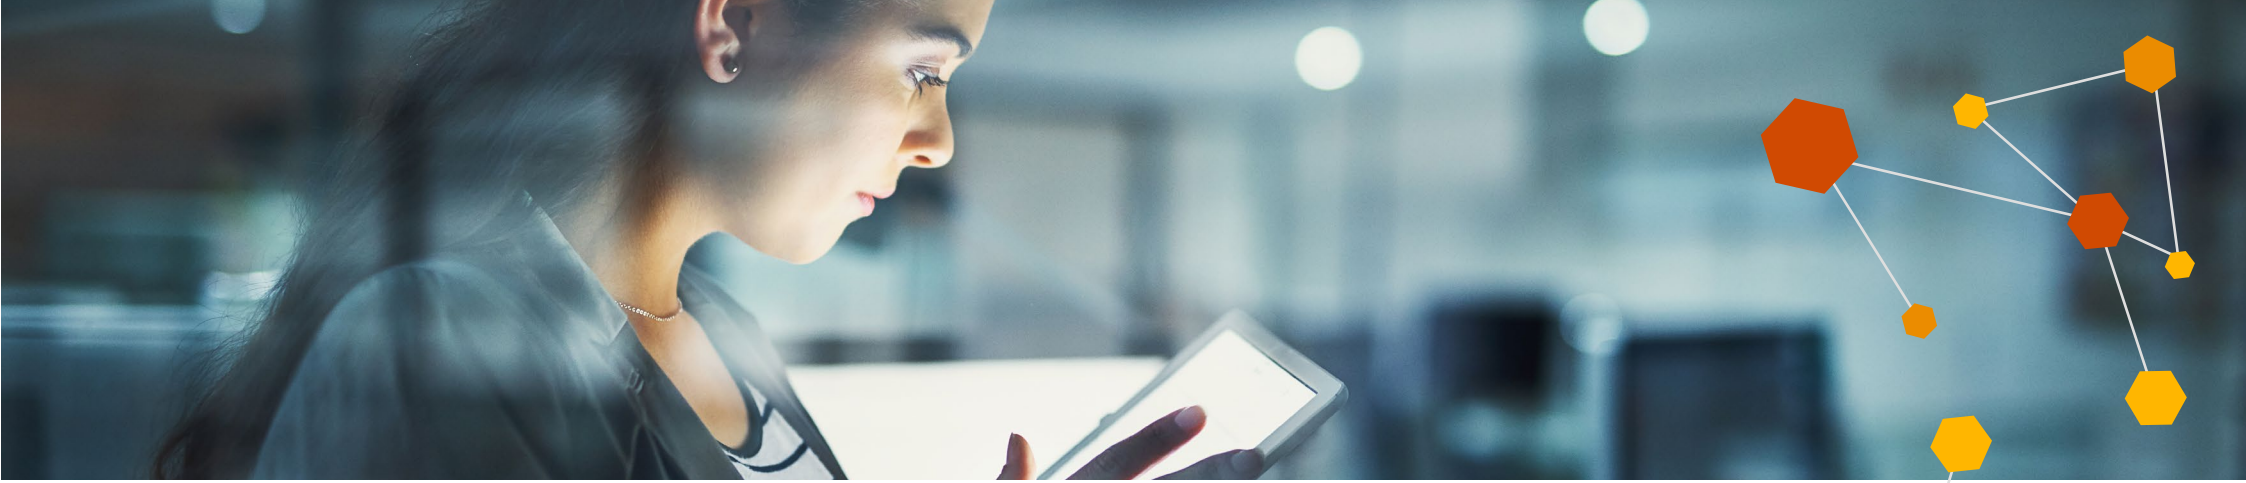

For the chief financial officer and controller

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 Inventory and accounting tasks in the organisation.

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

2 Identify and apply internal controls that are relevant to generative AI use cases.

a. Create an innovation sandbox 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 regulatory requirements.

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

d. Partner with the CISO to confirm the security and confidentiality of the 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).

3 Design an audit plan around the company's generative AI systems and models

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.

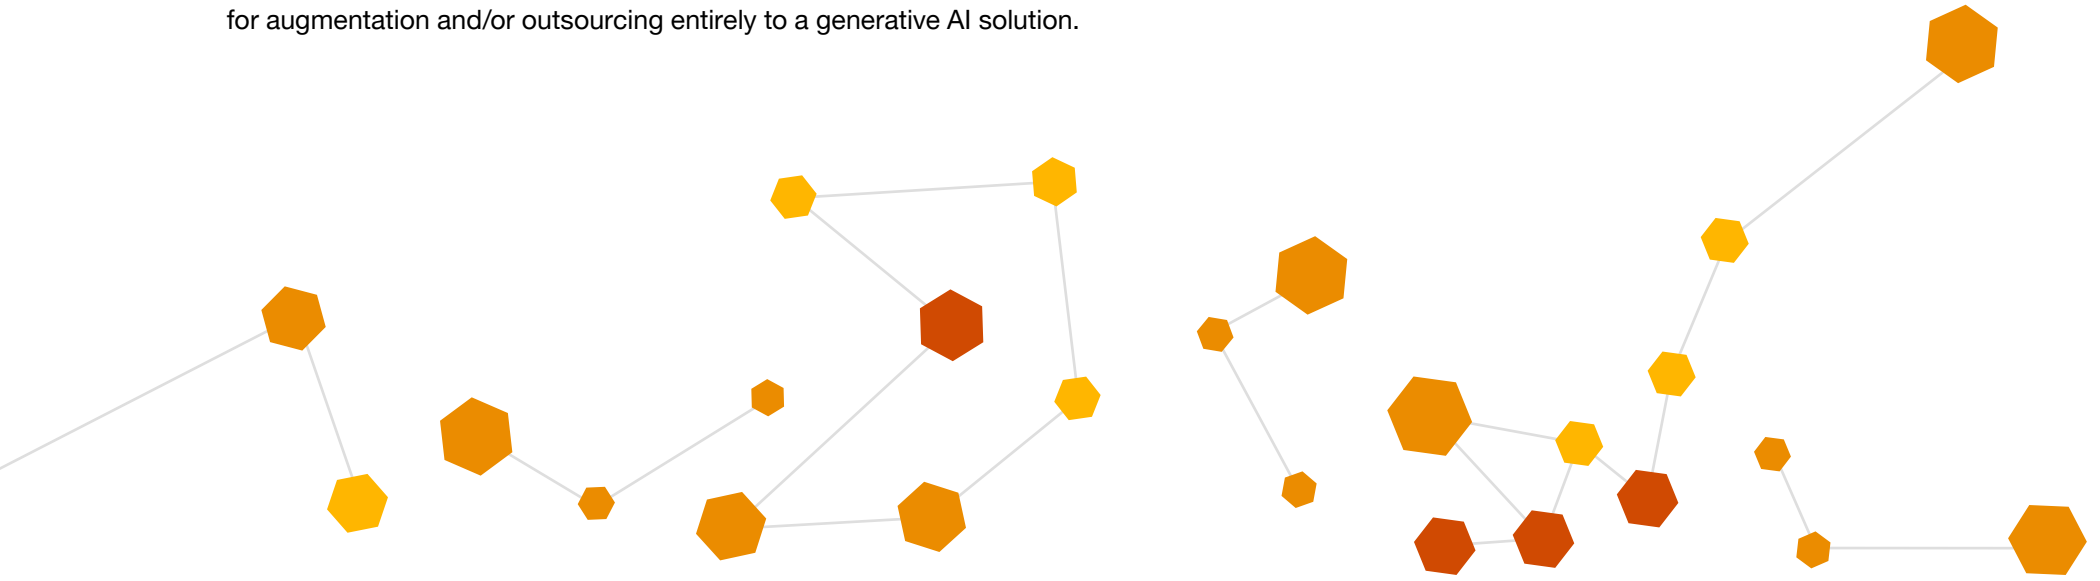

Bringing it all together

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 data meets compliance obligations under state-based health records legislation and The Commonwealth Privacy Act. I.e.: Health Insurance Act, My Health Records Act, Aged Care Act and state based non-health privacy laws for non health related PI if public.
- **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:** Responsibly for compliance, particularity with health and data laws, potential legal analysis on IP and data, risks associated with negligent advice and negotiate contractual assurances with the generative AI platform that patient data will remain segregated from the AI model's public instance.

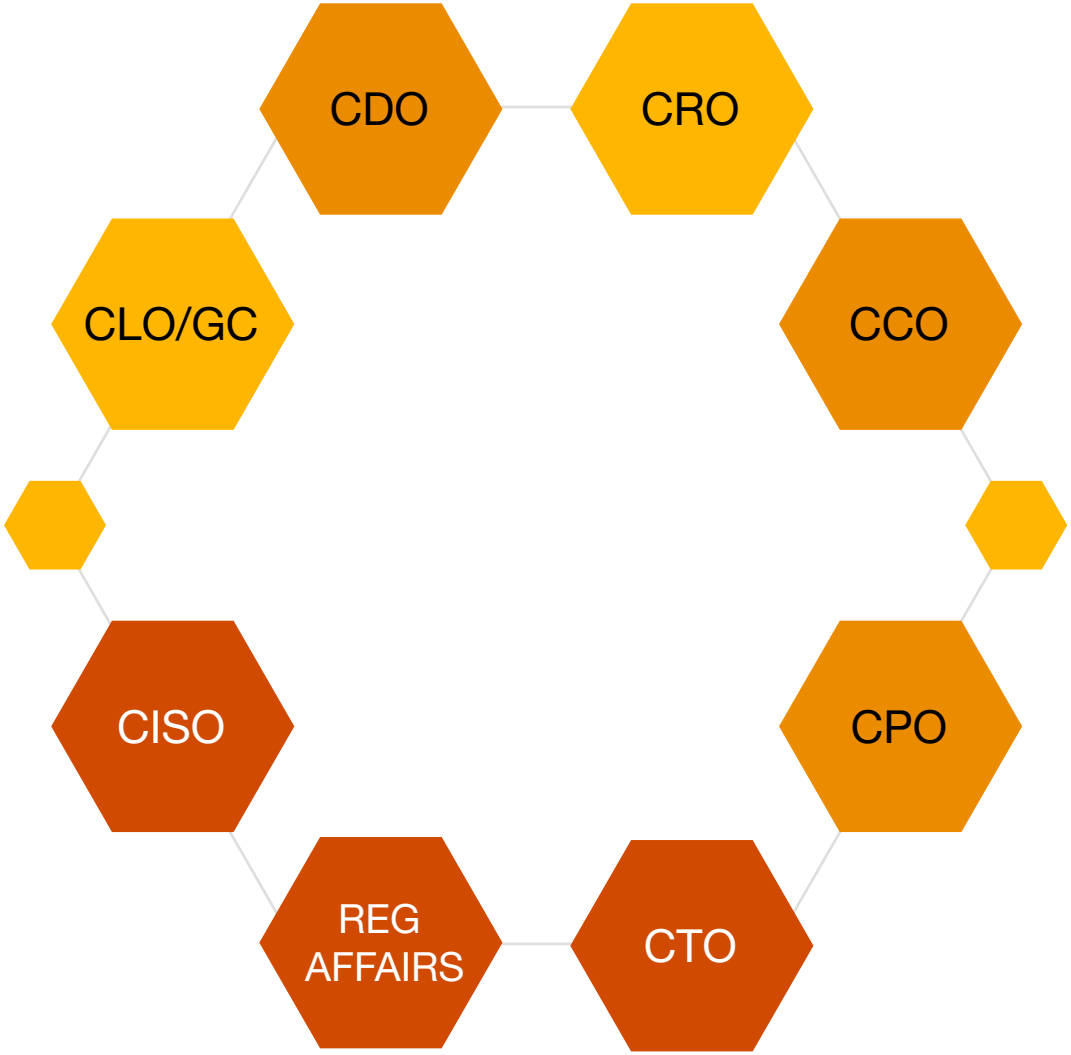

- **CISO:** Designate this application and data store as a “crown jewel” and provide adequate protections for it based on the most sensitive data classification.
- **Internal audit:** Develop an audit risk assessment and plan around the proposed system and model – including legal and compliance risks based on state-based health records legislation and The Commonwealth Privacy Act. I.e.: Health Insurance Act, My Health Records Act, Aged Care Act and state based non-health privacy laws for non health related PI if public – 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 is 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 artefacts to show how the technology is being used to reach decisions and to demonstrate evidence that it complies with Part IIIA of the Privacy Act, the Privacy (Credit Reporting) Code, The Competition and Consumer Act (CCA) and Discrimination laws.
- **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:** Ensure that the internal controls environment and relevant risk and controls frameworks is adequate in addressing the potential implications of AI adoption, by having the confidence on the effectiveness of the design and operation of the controls on the integrity of financial reporting process and meeting regulatory and legislative requirements such as IFRS or SOX 404.
- **CISO:** Designate this application and data store a “crown jewel” and protect it based on the most sensitive data classification.

Move through complexity with clarity

Our AI practice brings together a team of experts in data science, engineering, data ethics, digital law, risk and governance - along with experts in your industry - to help you accelerate responsibly and take the next step on your organisation's journey towards AI.

Understand, Establish, Accelerate and Assure is our formula for AI in the enterprise, and it's underpinned by 25 years of experience as a leader in technology and data governance, our award-winning AI consulting services and our global ecosystem of technology alliances.

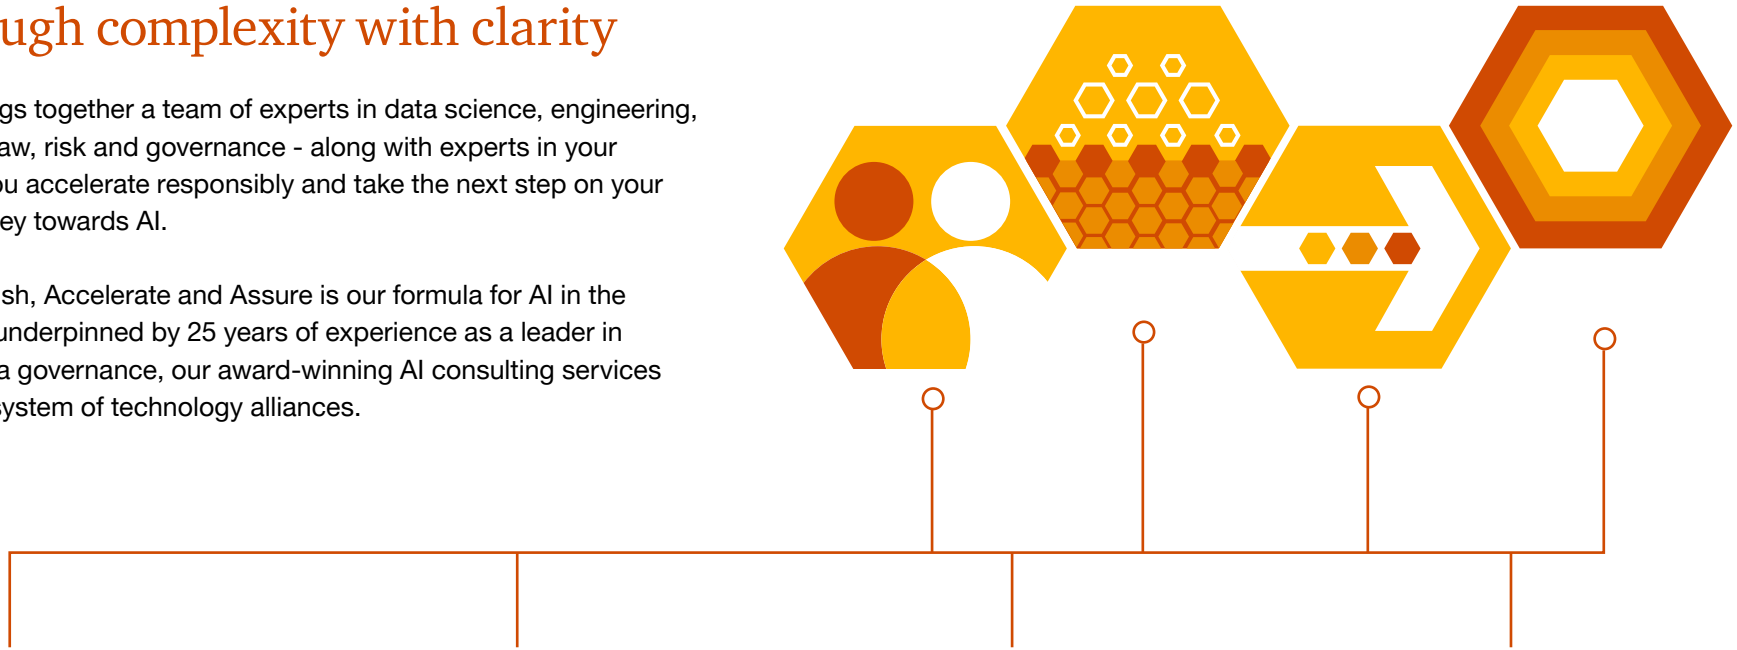

Understand

Understanding your baseline and readiness to respond to the opportunities and risks of artificial intelligence.

Establish

Establishing the foundations that organisations should have in place for embracing AI swiftly yet safely.

Accelerate

Accelerating your ideation, exploration and prototyping. Scaling successful ideas and monitoring their efficacy and safety.

Assure

Reviewing your existing AI frameworks, platforms, models and implementations to help build and maintain trust with your stakeholders.

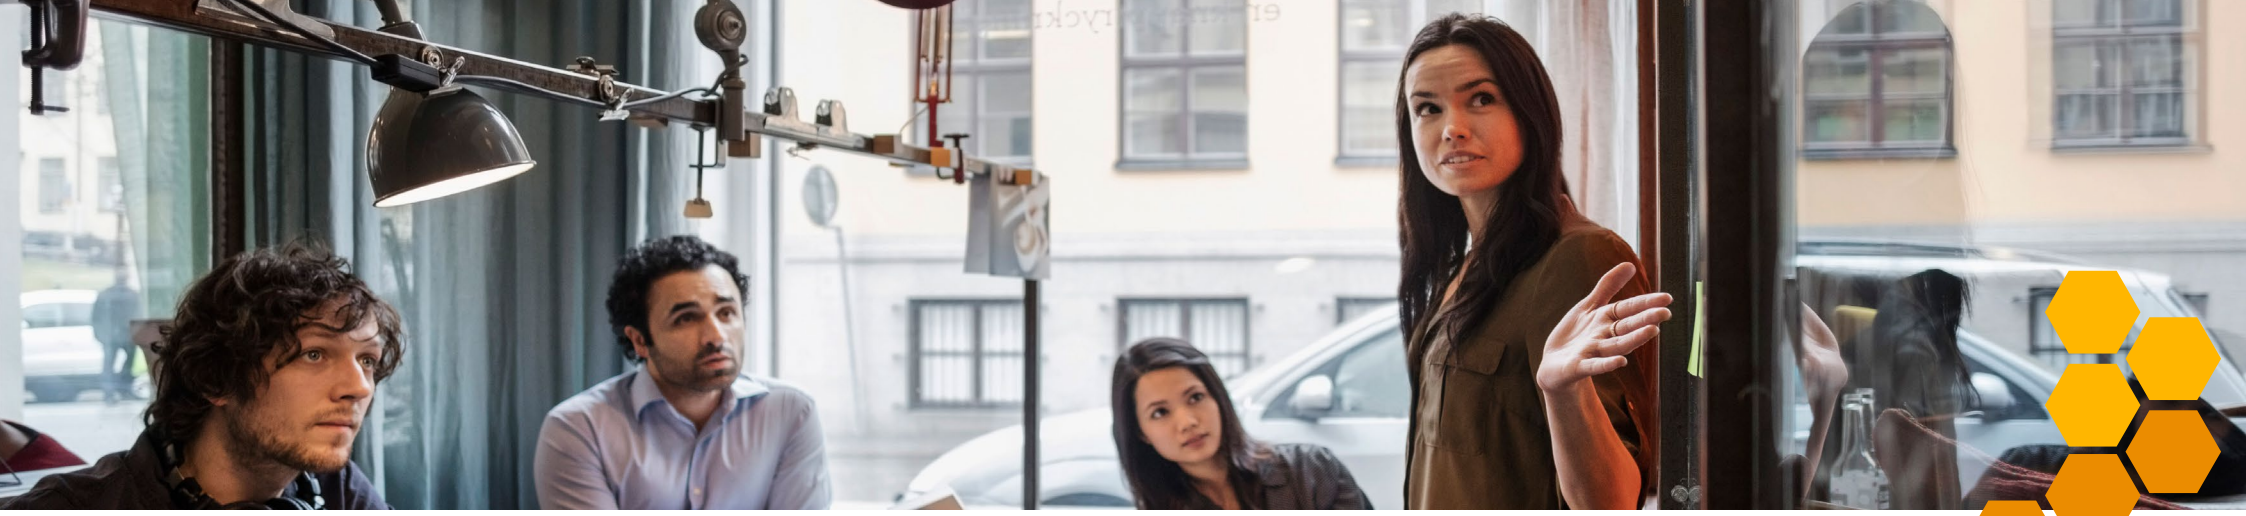

Bottom line

For the organisations that apply it wisely, generative AI has the potential to save time and money, improve products and services and even strengthen reputations. But the approach should be human-led and tech-powered, rather than the other way around.

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 realise 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.

The organisations that have a strong understanding of generative AI risks and how trustworthy AI systems can be designed, measured and managed can afford to move faster on AI transformation and are more ready to unlock high value use cases than those who don't.

Contact us to learn more

Tom Pagram
Lead Partner, AI
tom.pagram@au.pwc.com
+61 7 3257 5440

Adrian Chotar
Partner, Head of Digital, Cyber and Technology Legal
adrian.chotar@au.pwc.com
+61 2 8266 1320

John Studley
Partner, Chief Analytics Officer
john.studley@au.pwc.com
+61 417 220 198

Jacqui Visch
Partner, Chief Digital & Information Officer
jacqui.visch@au.pwc.com
+61 466 415 851

Jon Benson
Partner, Data Trust & Privacy
jon.benson@au.pwc.com
+61 3 8603 1669

Carlos Aggio
Partner, Cloud & Digital
carlos.aggio@au.pwc.com

WLTD0553529