Navigating the generative AI and cybersecurity journey: from perils to profits

Research insights from 600 business leaders and IT/security professionals

Bell

Contents

1. Overview

1.1. Executive summary

Why we did this study

Generative Artificial Intelligence (GenAI) is no longer just a buzzword, it's becoming a valuable business tool, and organizations are feeling its impact. But how are Canadian organizations specifically navigating this transformation? To find out, Bell partnered with Maru Research to explore how businesses across Canada are adopting generative AI, managing risks and achieving results. Conducted from September 3 to September 13, 2024, our survey gathered insights from over 600 business leaders and information technology (IT)/security professionals across various industries. This report sheds light on how Canadian organizations are navigating the generative AI journey and why their experience differs from the global context.

Key insights

Generative AI is taking root across Canadian organizations, with 55% of early adopters reporting high growth (20% increase or more to revenue since 2022). They're seeing benefits from GenAI including improved customer satisfaction and streamlined operations. While AI's potential is undeniable, businesses are also facing new risks – mainly data privacy concerns and new security vulnerabilities – which have slowed down adoption for many. In our survey, only 6% of businesses felt risk had no impact on adoption, with 20% finding risk had only "somewhat" slowed adoption.

IT professionals and Millennials are leading the charge, using AI tools both at work and in their personal lives more frequently than other groups. Meanwhile, business leaders and older generations are more hesitant, reflecting the slower adoption rates seen in more traditional industries.

AI integration

Bell finds that organizations are embedding generative AI in several core processes, from automating routine tasks to enhancing customer service interactions. IT, finance and security departments are seeing the biggest gains. However, the extent of AI's integration varies by industry and function. For example, banking and insurance companies are more advanced with deployments, using GenAI for fraud detection either in production (42%) or in pilot project (39%). In manufacturing, GenAI is deployed for inventory management, and in retail 47% say they are using it for customer behaviour forecasting. AI integration into company applications and data has benefits but is significantly expanding risk exposure as well.

Business benefits

Early adopters are reaping rewards in key areas, such as 54% reporting improved product quality and 52% reporting faster time to market. Generative AI is already delivering positive outcomes for customer satisfaction and process optimization. Companies that are further along in their AI journey are seeing returns on investment, but most organizations have room to grow.

Risks and how companies are addressing them

With generative AI comes a host of new risks. Canadian companies are particularly concerned about data exposure, unauthorized access and intellectual property issues. While some organizations are taking proactive measures, with 58% putting in place access controls among other measures, others are vulnerable with fewer safeguards in place. Mitigating these risks is critical to unlocking AI's full potential.

Looking ahead

Despite the challenges, many Canadian organizations are optimistic about the future of AI – almost half expect we are just at the beginning of AI model progress and much more progress will take place over the next five years. Looking at how AI adoption influences this sentiment, early adopters expect moderate advancements over the next five years, while mainstream businesses anticipate even more significant breakthroughs. The consensus is clear: generative AI is just getting started and its role in reshaping Canadian industries will continue to expand.

1.2. Introduction

GenAI went from emerging technology to critical business capability at warp speed; many organizations may still be reeling from the shock. Research from McKinsey and KPMG shows that GenAI tools are being rapidly adopted and business leaders expect it will remain a top emerging technology for the next several years. Even as they come to grips with the new landscape, the hype of the initial AI push is fading away and its impact on reshaping business is becoming realized.

This report explores how Canadian organizations are embarking upon the beginning of their GenAI journey and how they plan to navigate the perils ahead as they seek the benefits promised. It investigates the differences between industries, functional roles, levels of adoption and more to reveal critical insights about the opportunities and challenges ahead.

With the help of Maru Research (Harris), Bell conducted this survey to peel back the global hype around AI and ask how Canadian organizations are using GenAI in their day-to-day operations, where their concerns lie and how they are mitigating the risks. As Canada lags the U.S. in productivity, could GenAI be part of the answer in closing the gap? Or will Canadian organizations be too risk-averse to deploy the emerging technology at scale?

Our comprehensive primary research study of business leaders and IT/security professionals was conducted between September 3 to September 13, 2024, collecting 600 survey responses nationally. It brings to light the real picture of AI adoption in Canada and provides an early look at where GenAI delivers positive outcomes. But we'll also see that AI adoption doesn't come without investing in a trusted environment that layers proper governance, security and privacy controls to mitigate the new risks at hand.

This study explores critical differences in AI adoption across industries, between job roles and even between age groups. It brings to light where organizations see measurable outcomes delivered by GenAI and where they face obstacles. Unlike previous studies in the market, this report highlights the Canadian experience with GenAI and provides actionable insights for decision-makers grappling with how to harness the potential of AI while mitigating its risks.

Why you should read this report

- Tailored to the Canadian context: See how Canadian organizations compare to the U.S. and other regions in terms of adoption and learn what challenges must be overcome.
- Adoption trends by role, industry and age group: We'll see how Canadian IT pros and Millennials have embraced GenAI in their personal lives and are using it at work.
- Scope of AI integration: GenAI is past the stage of emerging and is part of many different processes across organizations. Agentic AI is here, as organizations allow systems to autonomously access data and take actions.
- Business benefits and challenges: Early adopters are seeing a return on investment from AI, with some reporting large improvements across various use cases. GenAI is already positively affecting domains such as customer satisfaction and quality control.
- Risks and countermeasures: Organizations seem willing to take on a surprising amount of risk, though concerns have slowed AI adoption to an extent. Some of the early adopters are putting the correct security and privacy measures in place, while others are exposed to new risks.

Join us on the GenAI journey

This report is a data-driven analysis of how Canadians are navigating the GenAI journey, from early adopters who are seeing significant benefits to latecomers who hesitate because of the risks involved. Whether you're a decision-maker wondering if GenAI is more than hype or an IT/security professional who wants to protect their organization from the new threats posed, we have insights to guide you. Together, we'll uncover the future outlook of GenAI and become better equipped to make informed decisions in a rapidly evolving business landscape.

Don't progress in your AI journey alone. Do it with the insights from organizations across Canada as we navigate what's ahead together – with data-driven insights and actionable analysis.

1.3. Key highlights

In this study, we separate organizations into three different groups based on GenAI adoption. The "early adopters" are out front in terms of the most GenAI adoption, the "mainstream" is the bulk of organizations around the average of adoption and the "laggards" are those with the least adoption.

When it comes to the differences in concern for the risks of AI, we see that mainstream adopters have the most concern, early adopters have a moderate level of concern and laggards have a lower level of concern overall.

The graph and table below highlight a few of the different qualities of the groups featured throughout the report, review our graph and table below.

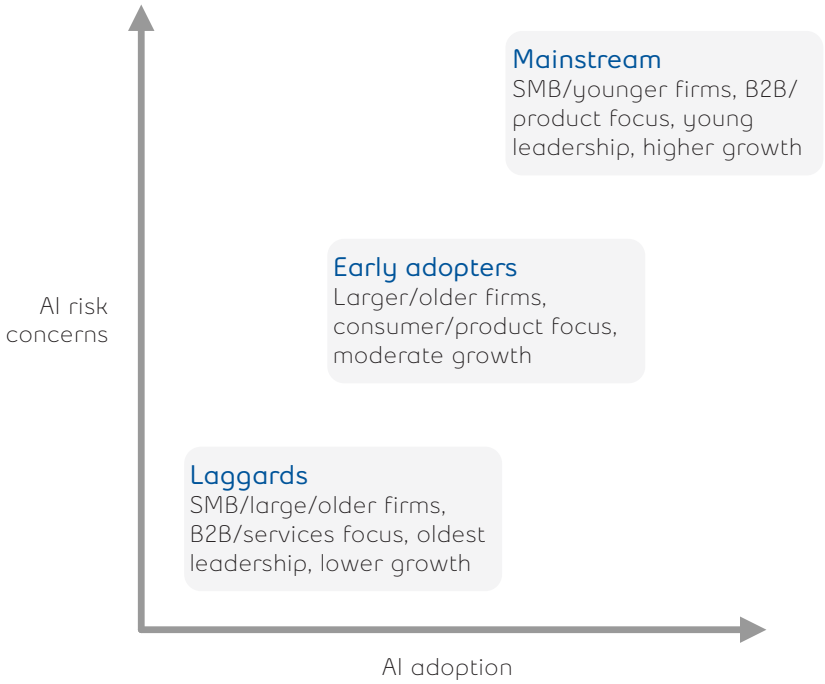

Figure 1: Defining qualities of the three GenAI adopter groups

	Early adopters	Mainstream	Laggards
Adoption of GenAI	**High adoption**: Frequent daily use of GenAI in both work and personal use, using paid external AI services	**Moderate adoption**: Often piloting AI projects or using AI in specific areas, using AI embedded within applications	**Low adoption**: Slower or resistant to adopting GenAI across use cases, using free AI services
Concern of AI risk	**Divided**: Confident but also highly worried about risks (data exposure, copyright violations, bias, model poisoning)	**Moderate concern**: Worry about data exposure, unauthorized access and mistakes	**Low concern**: Most worried about AI being weaponized against them (AI-powered phishing, deepfakes)
Mitigation tactics	**Advanced**: Proactive risk management and focus on data governance, audits and securing input/output	**Developing**: Basic safeguards, planning to implement more security controls	**Limited**: Few formal mitigations in place and reactive to AI security risks
Business benefits realized	**Significant**: See gains in customer satisfaction, cost reduction, quality improvement, time savings and process optimization	**Moderate**: Achieving benefits in efficiency, scalability and cost reduction	**Limited**: Few measurable business outcomes from AI and slower realization of value
Future outlook	**Cautiously optimistic**, expecting moderate progress in the next five years	**Most optimistic** about future advancements in AI, expecting major breakthroughs	**More optimistic** than early adopters but expecting major progress, driven by outside forces

2. Adoption and use cases

2.1. GenAI gets used at work every day

We're past the hype with GenAI and ready to get to work with it – even while new concerns about risks remain at the forefront. While recent headlines ponder whether businesses can find the ROI from GenAI projects, Canadian organizations in our study show that not only is GenAI quickly becoming a habit for most professionals, it's also deployed across various use cases in the enterprise.

For ChatGPT specifically, employees say they are more likely to use it for work than for personal use. IT professionals are more likely to use it more often than business leaders. At work, 50% of IT professionals say they use ChatGPT daily, compared to 41% of business leaders. For personal use, 37% of IT pros are logging into OpenAI's tool on a frequent basis (potentially as often as daily) while business leaders do so a little less often.

Age also factors into GenAI usage habits. As you might expect, Gen Z and Millennials are the most likely to use it every single day – both at work and at home. By contrast, Boomers use it less frequently.

As we'll see later in this section, the known risks of GenAI are affecting adoption, but even still the emergent technology is finding footholds across the enterprise.

71%
of professionals are using GenAI to some degree

41%
of professionals are using GenAI on a semi-regular basis

Data shows that just about everyone is using GenAI at work, and many are using it often.

2.2. What are organizations using GenAI for?

Of course, GenAI goes so much beyond ChatGPT. Over the past couple of years large language models (LLMs) and diffusion models, the engines behind GenAI's creative output, have become integrated into the products of almost every major enterprise software vendor – from Microsoft's Copilot to Salesforce's Einstein to CrowdStrike's Charlotte AI. But which of these new AI-powered features are being deployed or at least piloted in the enterprise?

Professionals deploy GenAI into production across many different areas of the enterprise. The top five most popular areas focus on either enhancing data-driven decision-making (in the form of visualizing data or personalizing responses provided to customers) or optimizing operations by automating tasks, detecting fraud and managing inventory.

Many GenAI pilots are in flight. At this earlier stage, we see GenAI's capability to produce text, summarize unstructured data and process information in a semantically relevant way.

Looking at the aggregate of use cases where GenAI is either already deployed or in a piloting stage, about three-quarters of employees would like to automate tasks; another three-quarters say they are looking to draft and edit documents. Additionally, respondents most often say GenAI is being used to summarize information or documents, detect fraud and search company data.

Top 5 uses for AI at work:[1]

I. Reducing/automating tasks
II. Drafting/editing documents
III. Summarizing information/documents
IV. Fraud detection
V. Searching company data

Which areas are organizations expecting GenAI to drive the most benefits?

With GenAI already factoring into so many professionals' daily lives and seeing broad deployment across multiple different use cases, it's clear that GenAI is quickly transitioning into a mainstay of organizational toolkits. But professionals are split on what area of the business could see the biggest benefit.

On average, IT professionals rank their own department as the greatest potential benefactor. Security is second and is most often a responsibility under the IT department. Customer service and support ranks third, with IT pros possibly looking to chatbots to field support tickets.

Top 3 areas expected to see the greatest benefit from GenAI

Average of functional areas ranked in the top 3

For IT professionals:

I. IT (19.67%)
II. Security (13.33%)
III. Customer Service/Support (10.00%)

For business leaders:

I. Finance (12.67%)
II. Marketing (10.67%)
III. HR (10.67%)

Industry focus: How different industries are deploying GenAI

Some industries are keen to deploy GenAI while others are still in the exploratory phase to determine the best use cases. Where industries have deployed it to broad production use, they often focus on strategic areas to differentiate themselves from the competition.

I. Retail
II. Manufacturing
III. Infrastructure/Media
IV. Public Sector/Education
V. Business Services
VI. Other Industries
VII. Banking/Insurance

IT vs. business: Technical focus vs. Creative exploration

IT/security professionals and business leaders differ on where they see GenAI adopted in the organization. IT is more aware of deployment to technical areas such as data analysis and fraud detection. In general, IT professionals reported more broad production use. Business leaders are experimenting with AI in activities like brainstorming and summarizing documents, driving toward strategic activities.

Business leaders see things differently. They expect GenAI to drive benefits in the finance department, the top-ranking department on average, followed by the marketing department and HR. It seems that IT/security professionals are more likely to associate GenAI with benefits in the areas they are most familiar with and therefore can imagine the unique capabilities of GenAI automating or augmenting a process.

[1] Broad production use and piloting

- Retail is an eager GenAI adopter, with more production use deployment across a variety of use cases including customer-facing applications and inventory management.
- Public sector and education see more adoption in summarizing information but are piloting image generation and customer-facing applications.
- Manufacturing prefers to deploy GenAI for data analysis and inventory management, aligning with operationally intensive processes.
- Banking and insurance are piloting in many areas but lag behind others in terms of broad production use.

2.3. Concerns about GenAI risks disrupt adoption

Despite the embrace of GenAI, there are concerns among organizations across industries as they seek to harness the emerging technology. GenAI brings new security, legal, reputational and other risks. Moreover, there are new threat vectors to defend against where cyber-criminals are weaponizing GenAI. There are privacy concerns as well, with ethical questions lingering on how AI vendors gathered training data for their LLMs and whether sensitive information fed into training of models could be inappropriately exposed.

These concerns are slowing adoption of GenAI in the workplace. About three in four professionals say the risks are significantly slowing adoption. But there are organizations in our study moving ahead with much higher levels of AI adoption; we can learn from both in terms of what to do and what not to do.

Even though they are more often frequent users of AI, Millennials are also the most likely to say concerns are slowing adoption. Surprisingly, business leaders are more likely than their IT peers to say concerns are having an impact on adoption. But the overall trend across all groups is clear: most professionals agree that concerns are causing a significant or major reduction to GenAI adoption. In our study we found significant use of AI, with room to grow if organizations could ameliorate concerns.

With so many workers embracing GenAI and expecting it to drive benefits in many different areas of the business, how can they mitigate the risks involved to encourage faster and wider adoption? To explore this, in the next section we'll look at how AI early adopters are expressing concerns about AI and what exactly they plan to do about it.

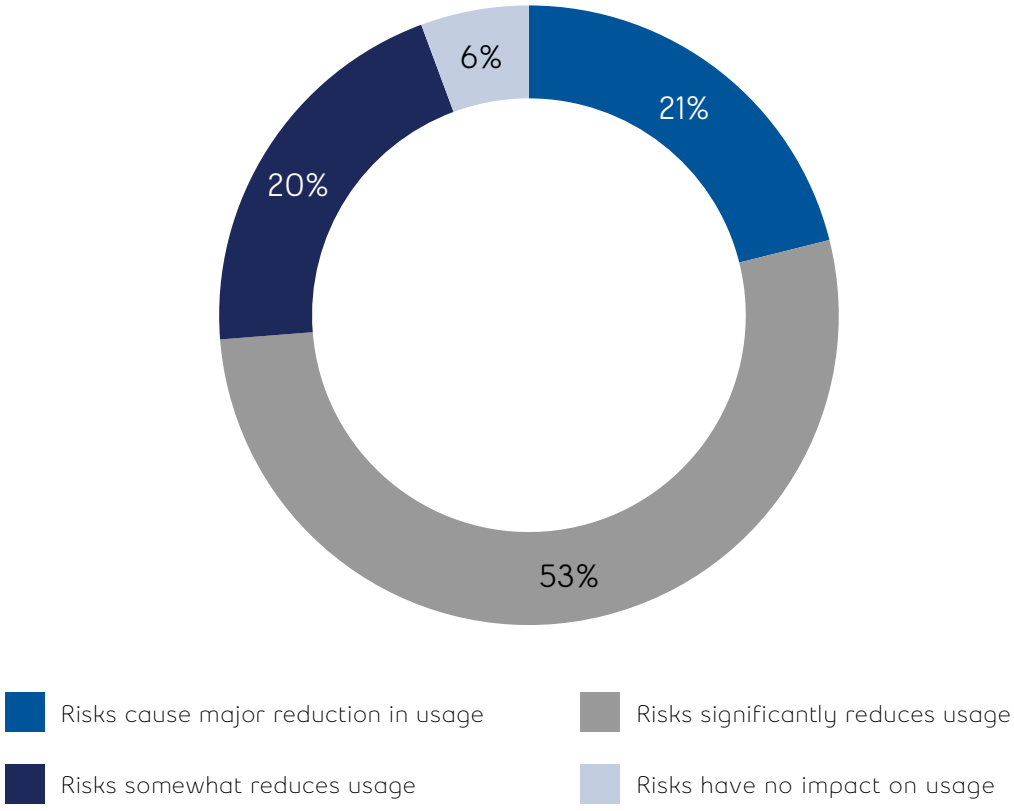

Figure 2: Respondents are most likely to say that GenAI risks are significantly reducing usage in their organizations.

3. GenAI risks and concerns

3.1. What keeps GenAI users up at night?

To explore the concerns of GenAI adopters, we segmented out those who are harnessing AI the most. We'll compare this group of early adopters to the mainstream (the majority of users) and the laggards (the lowest level of adoption).

Of the myriad risks organizations face, they are most concerned about exposing proprietary data through the use of GenAI. Six in 10 organizations consider this likely or almost certain to occur. It's not surprising, as a common reaction among organizations following the release of ChatGPT in November 2022 was to restrict employees from sharing sensitive data with the tool, recognizing that data entered in as prompts could train publicly-available LLMs.

This has translated into concern that vendors of GenAI will harness user data to train their models. Add to that a litany of court cases alleging intellectual property infringements involving celebrity creators and headline news organizations, and users may even be wondering if outputs from AI tools could violate copyright (e.g., using an image created by Google's Gemini in a marketing campaign having legal repercussions).

Very close behind: concern for mistakes made by GenAI and concern for new security vulnerabilities introduced. By now, most GenAI users will be familiar with the concept of "hallucinations" (errors made by GenAI). Similarly, users are hearing about bad actors exploiting the vulnerabilities of these tools to exfiltrate data (e.g., jailbreaking) and attacking unpatched software flaws in and around the LLM, connected application or other related IT infrastructure.

44%
of organizations acknowledge employees are entering sensitive data into services such as ChatGPT

Many professionals think their peers are clandestinely whispering secrets into the ear of OpenAI.

Again, we see that while all organizations demonstrate concerns, the level of AI adoption shapes where those concerns are most acute.

Early adopters of GenAI worry about their own use going awry. Almost six in 10 organizations worry about copyright violations, aware that the models they are using have been trained on copyrighted data. They are also concerned about bias against customers or employees, perhaps due to a non-representative training dataset. Or they worry the models they use will be maliciously influenced by bad actors that seek to poison the training data or otherwise tamper with the intended behavior of AI.

How adoption influences concerns related to GenAI

The mainstream group shows concern that mistakes they make with AI could hurt them. In addition to exposing proprietary data, they also worry about unauthorized access or embarrassment with customers over mistakes.

Laggards are most concerned about bad actors weaponizing AI to harm them. They worry about defending against AI-powered phishing messages and convincing deepfakes. Or they stress about data theft from models.

	Early adopters	Mainstream	Laggards
Top concern	AI will lead us to bad outcomes	We'll use AI the wrong way	AI will be used to harm us
Specific concerns	Copyright violation, bias, model poisoning/tampering	Data exposure, unauthorized access, reputation loss, customer loyalty	AI-powered phishing, deepfakes, stealing data from models
Shared concerns	Exposing proprietary data, incorrect actions, security vulnerabilities		

Figure 3: How adoption influences concerns related to GenAI

Industry focus: Concern lies in strategically important areas

Industries align their concerns about GenAI with strategically important areas for earning revenue or for maintaining compliance. Regulated industries demonstrate the highest degree of consistent concern across different risks.

Tier 1 - Regulated industries (Banking/Insurance, infrastructure/media)	Regulatory requirements drive concerns about data privacy, security and decision-making errors.
Tier 2 - Customer-first industries (Business services, retail, public sector or education)	The need to maintain customer or citizen trust in these sectors leads to concern about bias, poor decision-making and reputation loss.
Tier 3 - Operational industries (Manufacturing, other)	This group is behind other industries in terms of concern about AI, more focused on risks related to fraud or operational errors rather than reputation-oriented factors or regulatory infractions.

IT vs. business: Security and strategy

Profession significantly shapes concerns about GenAI, with business leaders showing consistently more concern across the board than IT leaders. Business leaders focus more prominently on strategic risks including poor decision-making, reputation loss and bias. This reflects the business leader's accountability for the brand and the bottom line.

IT/security professionals tend to be less concerned. Perhaps that's because they are more involved in putting mitigations in place and are more comfortable after using the technology more often (as we saw in the first section on adoption). IT professionals focus more on the risks of bad actors weaponizing GenAI, including security vulnerabilities, fraud and unauthorized data access.

In more detail, here is how organizations rank the likelihood of various risks related to their use of AI (as opposed to other sorts of risks, such as those posed by threat actors). Also included in this chart is the willingness of organizations to accept these risks associated with AI.

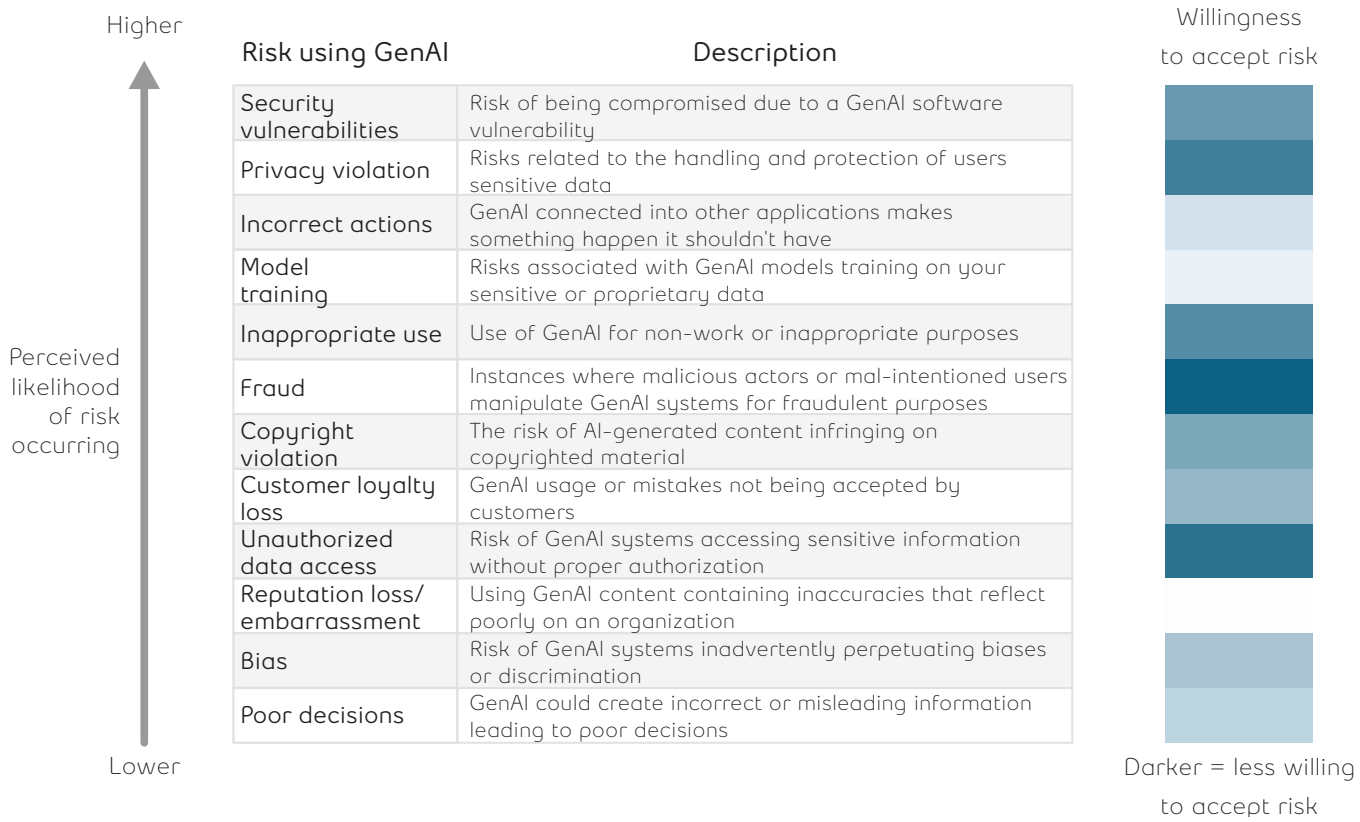

Risk using GenAI	Description
Security vulnerabilities	Risk of being compromised due to a GenAI software vulnerability
Privacy violation	Risks related to the handling and protection of users sensitive data
Incorrect actions	GenAI connected into other applications makes something happen it shouldn't have
Model training	Risks associated with GenAI models training on your sensitive or proprietary data
Inappropriate use	Use of GenAI for non-work or inappropriate purposes
Fraud	Instances where malicious actors or mal-intentioned users manipulate GenAI systems for fraudulent purposes
Copyright violation	The risk of AI-generated content infringing on copyrighted material
Customer loyalty loss	GenAI usage or mistakes not being accepted by customers
Unauthorized data access	Risk of GenAI systems accessing sensitive information without proper authorization
Reputation loss/ embarrassment	Using GenAI content containing inaccuracies that reflect poorly on an organization
Bias	Risk of GenAI systems inadvertently perpetuating biases or discrimination
Poor decisions	GenAI could create incorrect or misleading information leading to poor decisions

Figure 4: GenAI risks and how willing organizations are to accept them

3.2. How concerned are we that threat actors will exploit us?

The differences between early adopters of GenAI and laggards extend to concerns for the risks specific to bad actors. Aside from the risks that may come from using AI services – either internally or unexpected customer activities with AI – organizations must also consider that hackers will seek to exploit new vulnerabilities created by LLMs and other forms of AI. Here we'll highlight some broad differences between early adopters and laggards so that those considering increased adoption can learn from the early adopters.

Early adopters show concern about more advanced hacking techniques that target the AI systems they're using:

- Close to **25%** of early adopters worry about bad actors stealing sensitive data from models.
- **33%** of early adopters worry about model poisoning or tampering.
- About **10%** of organizations worry about prompt injection or the manipulation of input to bypass filters.

Laggards are most concerned about AI amplifying existing attack vectors:

- AI-powered phishing and deepfakes rank highly for laggards, with **45%** saying they are most concerned with AI-powered phishing attacks and **36%** saying they are most concerned with deepfakes.
- Laggards are less aware overall of more sophisticated risks pertaining to AI, with only **14%** stating concern for model jailbreaking.

3.3. Are organizations exposed to AI risk?

Given the concerns relating to GenAI are often about the wrong person (internal or external) accessing sensitive data, we wanted to know how well organizations have classified their data. Classifying data into how it can be accessed or shared versus labelled confidential – or destroyed outright – will be an important indicator of an organization's risk posture when adopting GenAI. Classifying data in this way is no small challenge considering the distributed nature of information across the cloud, devices and organization servers.

Overall, 60% of organizations admit they know where less than half their sensitive data is. In other words, the majority of the data hasn't been classified.

Early adopters show more focus on data governance and classify more data overall. Laggards stand out as being far more likely to simply admit they do not know how much of their data is classified, with more than one in five organizations saying so.

Agentic AI on the rise

Next, we were curious which areas an organization would permit GenAI to autonomously access applications and complete tasks. The more integration with other systems, the more potential risk. Organizations have an unexpectedly high tolerance for AI autonomy, allowing AI services to take actions within business processes and connect into applications and data.

Organizations are pursuing agentic access, or automated integration and action-taking for agents in several areas. They have already put it into production use in IT, customer support and security. They are piloting it more often in procurement, marketing and supplier/partner management.

Overall, organizations are integrating GenAI in areas dealing with supply chain (procurement, supplier/partner management) and revenue-generating operations (marketing, sales). Here, GenAI's strengths in summarizing large text as well as personalizing content can really shine.

Organizations are generally more hesitant to let AI loose in backend environments where key controls or information might reside (IT, finance, security). This may reflect the concerns that AI risks exposing confidential data or it is prone to make mistakes.

Early adopters are willing to connect AI to systems where they expect to see benefits. While restrained in critical areas like IT and finance, they're much more likely to integrate it than the other groups. Even the mainstream group is well ahead of the laggards in deploying AI to many different areas, including backend functions like HR.

What allows early adopters to integrate GenAI into critical business functions despite the concerns? Is it that the benefits are so great that caution is thrown to the wind? Or are the latent risks of legacy environments deftly managed to avoid bad outcomes?

In our next section we'll explore what security and privacy controls are in place to reduce AI-related risks.

What about AI hallucinations?

Early on, many found the responses from GenAI to be fraught with errors. In fact, the term "hallucinations" became associated with the odd and sometimes flat-out wrong output from AI text and image-creation models. However, with rapid advances in model capabilities, the quality of responses has improved significantly – and the models themselves are reviewing their responses for accuracy as an additional form of quality check in some cases. Our study reflects the improvement in quality, with half of organizations reporting high-quality results from GenAI and less than 5% reporting low quality, with the remainder reporting acceptable quality.

4. Mitigating GenAI risk

4.1. Where the mitigations meet the models

LLMs introduce a new type of architecture to secure. Moreover, they aren't quite software and they are more than just data. As foundational AI research papers such as "Attention Is All You Need" have taught us, LLMs are a combination of code and data, vectorized by deep training with a neural network to produce knowledge. This new architecture holds new implications for security: risks include what data is input to the model, output from the model and the model itself. In the same way, we must examine integrations between LLMs and applications for vulnerabilities. Finally, security practitioners will need to step back and consider the new third-party risks introduced from LLM suppliers, hackers and regulatory bodies.

It's a lot to demand from security departments that are already stretched thin. Yet success is crucial for organizations to adopt AI and reap its benefits. Let's examine how our different adoption groups are pursuing risk mitigations.

4.2. Data security

When organizations are taking control of the data flow before it's provided to an LLM through a prompt, the most common approach is to classify data after the basic control of data encryption. To round out the top three controls, organizations are taking the necessary step of monitoring and auditing LLM input. This is likely at a simple level of Layer 3 / firewall and Layer 7 secure web gateway controls. In this case, however, any logs are good logs.

We have listed the controls your organization should consider to promote data security. They are sorted from top to bottom by the amount of deployment we saw overall in our study. Take note that controls such as risk assessments and documenting data accountability were not executed on often enough, particularly with proposed Canadian legislation on AI and privacy. These controls underlie a strong privacy program as well. However, we further break out privacy controls later in this study.

Data security measures taken for LLMs

- Implement data encryption
- Classify data by sensitivity level
- Monitor, log and audit activity/usage
- Increase access controls, role-based permissions and least privilege for data access
- Understand how GenAI uses input data (e.g., whether it's used as training data and whether it's accessible to other third parties)
- Use data loss prevention (DLP) tools
- Train employees
- Examine AI service/vendor contacts for security and privacy concerns
- Use an LLM firewall to manage data inputs/output
- Conduct risk assessments on GenAI security issues
- Limit use/number of queries per user
- Conduct data audits to determine who is using sensitive data
- Document data accountability/record of processing activity (ROPA)
- Apply data-masking techniques such as anonymization and pseudonymization

While organizations are practising some of these measures less often, many have plans to start within the next 12 months. Organizations are looking to close the gap in the bottom three measures, for example, with almost half planning to conduct data audits. And almost the same number of organizations also plan to start documenting data accountability and applying data-masking techniques.

4.3. Application security

The application layer must also be considered to secure GenAI. Like the data layer, application layer controls help avoid prompt injection and myriad other attacks. Moreover, security controls for enterprise applications that are integrating into GenAI are needed.

Organizations are recognizing the importance of both traditional security controls and the potential need for other approaches to securing interfaces to GenAI. This is leading some to consider the evolving LLM firewall technologies, for example.

Application security measures taken

- Monitoring, logging and auditing of activity/usage
- Access controls, role-based permissions and least privilege
- Input validation security check to prevent injection attacks, etc.
- Vulnerability scanning
- API security
- LLM firewall for application layer controls
- Software supply chain security
- Threat modelling
- Penetration test/red-teaming
- Isolate/sandbox model deployment
- Fuzzing
- LLM Firewall to Manage Data Inputs/Outputs

From our survey, we learned almost half of organizations will be monitoring, logging and auditing activity/usage of LLMs. Another 36% plan to adopt this tactic in the next 12 months. These organizations will be better able to maintain accountability for GenAI interactions in the future, as monitoring is a key component of any security program and many regulatory compliance regimes for AI.

4.4. Privacy

With headlines over the past year highlighting gaffes like dumping sensitive intellectual property into publicly available models or a negative backlash from consumers over a brand's use of AI, organizations are taking steps to identify the privacy risks they face – and educate employees about how to avoid ending up in headlines themselves.

In the list of privacy controls below, our research shows that organizations need to increase the focus on privacy assessments related to AI usage as Canadian and international legislation come into force. But the good news is that the call from the privacy industry and regulators for organizations to obtain customer consent for marketing and other efforts has clearly resonated with Canadian organizations, given it is their top focus today for AI privacy controls.

Privacy measures taken

- Put in place policies for customer consent for data collection, use and disclosure
- Communicate how AI may be used to interact with customers, their data, etc.
- Provide staff with GenAI training around privacy/bias
- Establish procedures for handling customer complaints and requests
- Ensure ethical use of GenAI to prevent bias or unfair treatment of customers
- Use security safeguards/technologies
- Conduct a regular (e.g., annual) privacy assessment
- Implement data minimization and retention policies

Implementing these controls to curtail potential privacy infringement related to GenAI should help prepare organizations for compliance requirements that could emerge from Canada's plans to pass the Artificial Intelligence and Data Act (AIDA).

4.5. Prepared to comply

First introduced as Bill C-27 in 2022, AIDA is currently before Parliament and could be passed into law some time in 2025. While potential disruptions to its passage remain, either through voting in the House of Commons or the dissolution of Parliament to prepare for a new election, it's worth understanding the direction that lawmakers

are going. Canada's AIDA is very similar to Europe's AI Act, already passed into law. It also holds similarities with how AI will be regulated by the Securities Enforcement Commission in the U.S.

Overall, organizations are feeling positive about their preparations for the new law. More than eight in 10 organizations believe they've made at least some progress toward being prepared for AIDA. But only 43% of AI adopters are willing to go so far as to say they are fully prepared for the Act. Another one in 10 say they haven't done anything yet, but could be prepared if required.

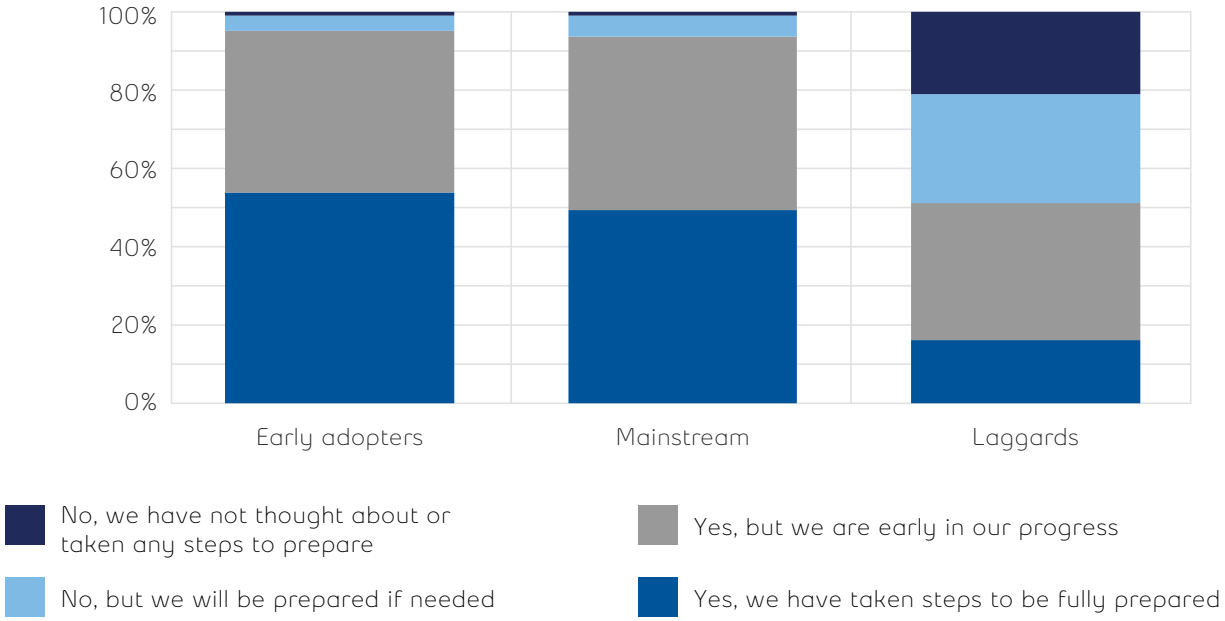

Figure 5: How organizations are preparing for Canada's AI legislation

Laggards are the least prepared group for AIDA. While early adopters are only slightly ahead of the mainstream in preparation, laggards are well behind in this area, too. Likely, they are not pursuing compliance because they haven't adopted AI tools yet and therefore don't anticipate being the target of future enforcement measures. But that lack of preparation could become a hindrance to adopting AI tools and reaping their benefits. Moreover, they may very well use AI within other applications that end up having legal implications under AIDA.

Industry focus: Who's ready for AI regulations?

It's not surprising to see that the sectors with experience with compliance are thinking about how to prepare for AIDA. Banking/insurance is the most proactive, with the highest proportion of respondents reporting they are fully prepared. They are closely followed by the retail sector, which has experience complying with legislation stemming from new technology capabilities, such as the Canada Anti-Spam Legislation.

Public sector/education and business services organizations still have their work cut out for them, with many saying they've made some progress and others saying they can start when required. Only a small number of organizations admit that they've done nothing at all.

IT vs. business: Aligned and ready

Preparing to comply with AIDA is an area where both IT and business agree. They share the perception that their own organizations are at least making progress toward meeting the requirements of the Act, if not already fully prepared.

It's positive that so many firms feel prepared for the new legislation – and many of the mitigations we covered in the previous section will help them toward that. Ensuring data and application layer controls are in place will help satisfy the Act's requirements for data governance, proper risk assessments and mitigations. But what are organizations planning to do to improve overall transparency and keep a human "in the loop"?

Validating the quality of GenAI responses to avoid bias and inaccuracies will be an important component of compliance with AIDA, particularly for high-risk systems. About half of organizations say they are fully validating the outputs from GenAI. Another 42% are only somewhat checking the outputs. These checks are likely shouldered by the average employee rather than rigorous validation.

So, organizations agree on doing at least some validation, but they differ in their approaches. Three in 10 say they have an automated approach or another GenAI tool reviewing the responses. However, 38% rely solely on each employee to ensure the use of GenAI is correct and meets compliance. It's interesting to see that early adopters are actually less likely to involve a human reviewer, with fewer than one in four providing human oversight.

We inquired separately if organizations have a steering committee in place to oversee decisions and policies governing their use of AI. Steering committees are important – particularly early on in adoption – yet fewer than two-fifths of organizations have that oversight body in place to draw from expertise across the organization. Typically, IT/security has the mandate to oversee usage and is leading AI policy and decisions.

AI oversight	% of AI users
IT/security	61%
Executive team/business managers	53%
Steering committee of multiple stakeholders	39%
Individual employees are responsible for their own usage	36%

Having the human in the loop is an important element of AI regulations not only here in Canada, but in other jurisdictions around the world. Organizations will have to do more to demonstrate that it is people exercising judgment about decisions made in high-risk systems. But people alone can't review at the speed of AI and will require automated checks as well.

Requirements set the foundation for success with GenAI, but they are only worth pursuing if there are rewards for making the effort. Are investments made to deploy appropriately governed GenAI going to pay off? Are AI vendors just pushing hype to sell a new wave of so-called solutions? Or will these new tools deliver the business value that will make them worth the investment? We'll explore these questions in the next section.

5. GenAI impact and future outlook

5.1. Where GenAI delivers value

Almost all surveyed organizations report achieving measurable outcomes with GenAI in the areas where they use it. A variety of different benefits are being enjoyed by AI users, with the most common being improved quality of products or services, improved speed at going to market and optimization of business processes. At least four in 10 organizations have experienced these outcomes.

In the next tier of business benefits, organizations are seeing better customer satisfaction, reduced time spent by staff on tasks and better innovation.

Close behind these benefits are reduced cost of operations and improved scale in terms of serving more customers or entering new areas of business.

Business outcomes through use of GenAI

Tier 1 - Most adopted	• **Quality**: Better products or services • **Speed**: Faster time to market • **Business process**: Refinement and optimization of processes
Tier 2 - Close behind	• **Customer satisfaction**: Improved interactions with customers • **Time**: Staff can complete tasks faster than before • **Cost**: Operations cost reductions
Tier 3 - Room to grow	• **Innovation**: Creation of new offerings • **Scale**: Can serve more customers, enter new areas of business.

While both early adopters and mainstream organizations are benefitting from GenAI, early adopters reap significantly more benefits in certain areas. They are furthest ahead in seeing improved customer satisfaction and reduced costs compared to mainstream. This suggests that realizing the most return on investment may mean laggards and mainstream adopters will want to prioritize deploying GenAI in these areas.

Almost all organizations report positive outcomes from using GenAI. A good portion of them also say those positive outcomes are large improvements.

Industry focus: Improving with purpose

Business services emerges as a leader in terms of driving improvements across multiple areas with GenAI. But every industry demonstrates some focus in finding improvements in key areas.

Area	Improvement leaders	Improvement followers
Quality: Better products or services	Business services, retail	Banking/ insurance
Speed: Faster time to market	Manufacturing, banking/ insurance	Infrastructure/ media
Business process: Refinement and optimization	Retail, manufacturing	Business services
Customer satisfaction: Improved interactions	Business services, retail	Public sector/ education
Time: staff complete tasks faster	Manufacturing, public sector/ education	Retail
Cost: Operations cost reductions	Business services	Public sector/ education
Scale: Serve more customers, enter new areas	Manufacturing, business services	Public sector/ education, retail

Overall, organizations most commonly report a small improvement across all areas of AI investment. This indicates many may see incremental improvements with AI rather than disruptive effects. But with some organizations reporting large improvements, it's likely that those seeing underwhelming results could get more value out of their GenAI deployments.

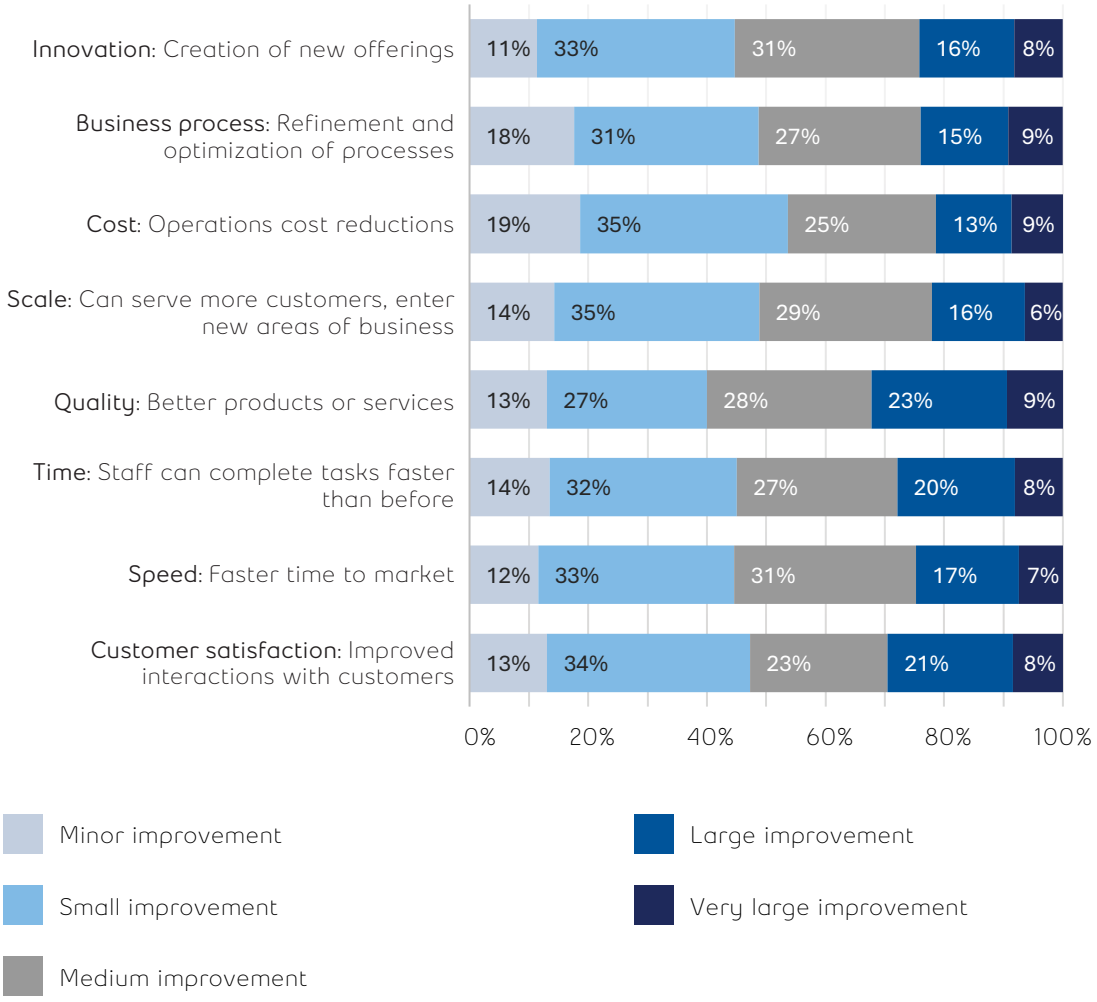

Figure 6: Where professionals see the most improvement from GenAI

The largest number of businesses that we surveyed reported a large to very large gain in improving quality and customer satisfaction. Saving staff time on completing tasks trails closely behind. (Respondents rated their improvement on a scale of 0-100% compared to before they were using GenAI.)

So will GenAI take my job?

With many organizations reporting positive outcomes across several areas and significant improvements in some, it's natural to wonder if GenAI could fully automate enough tasks to either render certain staff positions totally obsolete or at least contribute enough to reduce staff size.

Most organizations anticipate staff reductions due to GenAI, although the majority expect less than a 10% reduction. However, some organizations expect an increase in staff will be required due to GenAI.

At more than 25%, early adopters are most likely to expect a staff increase, whereas 65% of the mainstream anticipate a reduction and almost half of laggards expect no change.

While the optimization and automation capabilities reported from GenAI clearly indicate that some staffing adjustments in certain areas could be needed as GenAI takes on a bigger portion of the workload, organizations may not be considering the big picture. Early adopters may be expecting to increase staff because they've identified higher value roles and new positions that will drive new opportunities. The mainstream should consider where to shift their staff when their time spent on lower value tasks is freed up.

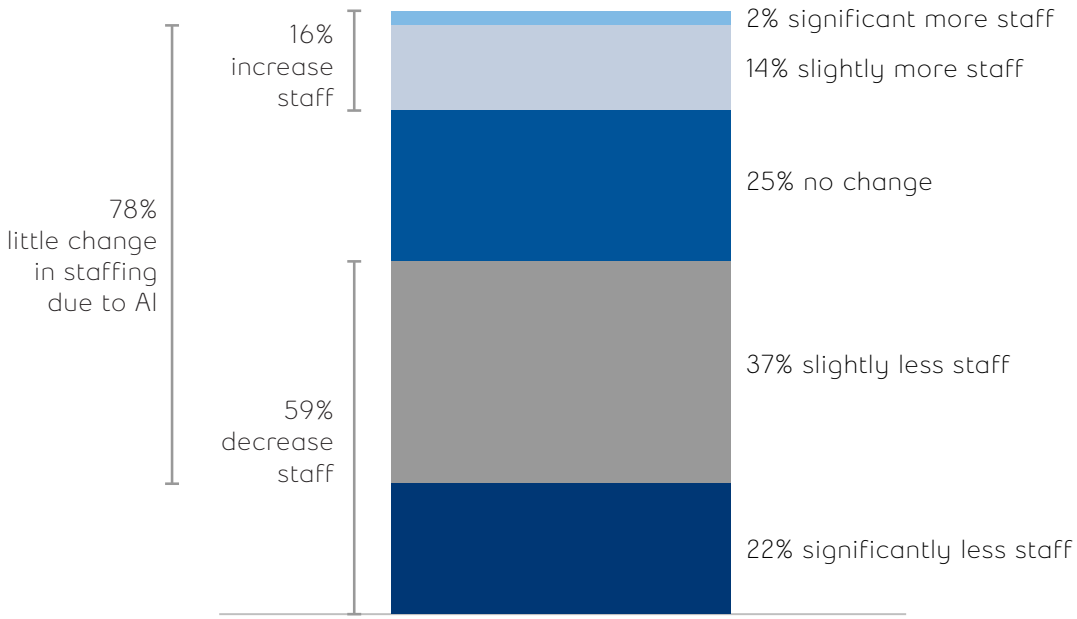

Figure 7: The expected impact of GenAI on staff size

5.2. GenAI in the future

For all the work organizations have already done to deploy AI – and for as much concern as they have about the risks – most see this as the beginning of the journey.

Almost half of organizations say we're just at the beginning of AI model progress and major advancements will occur in the next five years. Another 44% expect the advancements to be moderate and only 8% think that we'll see little-to-no progress.

It's understandable that most view GenAI as a field that will continue to expand and deliver more breakthroughs. Since fall 2022, it feels like an AI firm is announcing some incredible new capability that catches the world's attention on a frequent basis.

Interestingly, it's our mainstream respondents who are most optimistic about future advancements. Even laggards are more likely to think that major progress will be made compared to early adopters. Early adopters are more likely to think that we'll see only moderate progress in the next five years.

Perhaps that's because they know the risks involved with AI and the work it will take to mitigate them.

What's your p-doom?

In a study about AI risks and responsible security controls, we shouldn't dance around the big question: Could AI advance to the point that it's smarter than humans and have autonomy to the point that it could lead to an apocalyptic event?

It may seem incredible to even consider it, but with leading AI experts extolling warnings about the existential risk of such an "AI singularity" event – and many governments acting to regulate AI as a result – it's worth a moment to take our collective temperature on the question.

In Silicon Valley, the casual way to ask someone how likely they think the AI apocalypse will occur is to say, "what's your p-doom?" But we put the following question to Canadian professionals: "Some AI experts worry that AI will continue to advance and eventually be smarter than humans and have its own autonomy, leading to a potentially apocalyptic event. On a scale of 0 to 100 where 0 is no chance and 100 is certainty, what chance do you think there is of AI causing such a catastrophic outcome for society?"

Most organizations are either taking a "not really worried" approach (scoring it 25 or below) or are skittish (scoring it 25–50), with seven in 10 falling into one of these groups. One in five organizations are worried (rating it 51–75) and about one in 10 think doomsday is coming (higher than 75 rating).

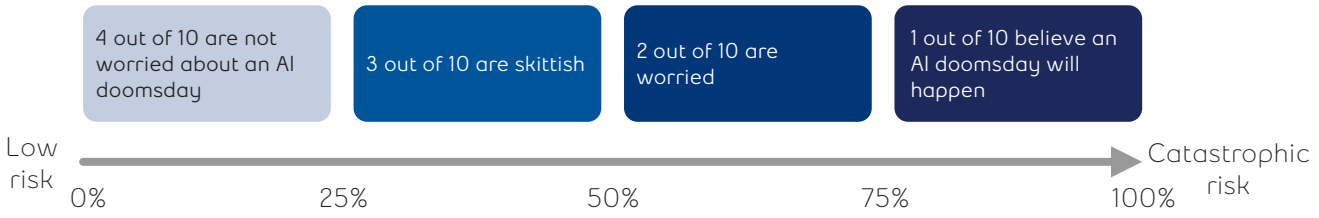

Figure 8: Probability of AI existential risk / p-doom score

The early adopters group exhibits an interesting split in how worried they are about AI competing with humanity. While 46% are in the "no worries" camp, nearly one in four are significantly worried. This is slightly ahead of the majority group and well ahead of the laggards group exhibiting significant worry. It makes you wonder – what do early adopters know that we don't?

5.3. Key takeaways

Deploying AI requires a methodical rather than an ad hoc approach to mitigate the risks discussed in this study. However, more than six out of 10 Canadian organizations that have adopted AI have no AI strategy in place to guide deployment, risks and expected value. If more could be done to mitigate risks and realize the benefits of deployment, Canada could close some of its productivity gap with the U.S. So we offer the following advice to help derisk adoption:

- Start with good governance. There's a tendency to jump to technical controls when curtailing the risks of using AI. Proper governance means security and business working together to understand compliance and other requirements to define responsibilities, policies, access rights and so on. Consider a steering committee of peers from across business functions to make decisions about GenAI usage and risks (at least early on).
- Take a multi-dimensional approach to security risk. As onerous as it sounds, examine exposure and put controls in place both across the layers of the IT stack (e.g., data, applications, access, etc.) and the AI service/model (e.g., input into the model, output from the model, training data if applicable). Moreover, review AI deployments from a privacy, compliance and legal perspective (see the list of controls earlier in this study).
- Data security takes centre stage. There remains a very large data discovery and classification exercise that needs to happen across file shares, user endpoints, email, enterprise applications, cloud services and so on. With the rise of Agentic AI, the creation and flow of data will accelerate. Explore new startups in the AI data security space in addition to the traditional players.
- Don't shy away from AI. There are clear benefits that Canadian organization are achieving as model output quality improves dramatically. Moreover, when proper user settings are established, services such as ChatGPT will not train on company data. As security pros, we need to enable the business to excel with AI using effective guardrails.
- Fight fire with fire. The top two security domains where security professionals are turning to AI to improve their scale and effectiveness are threat detection and data security. These are two key areas to apply multiple forms of AI, including generative, discriminative and reinforcement learning, to defend against a much higher volume of external attacks and inside threat accidents. Our study showed that monitoring is high on the priority list of controls organizations are putting in place to defend against AI risks. AI can also be used for greater detection and response capabilities to reduce the load on security operations.

6. Appendix

6.1. Methodology

Bell conducted this AI security survey with the help of Maru Research, collecting 600 responses from qualified professionals working at companies of at least 100 employees in size. The survey was fielded between September 3 and September 13, 2024, in both English and French. Respondents completed the survey on desktop or mobile devices. The survey is a diverse representation of organizations across the country in various industries. Respondents were evenly split between IT/security professionals and business leaders.

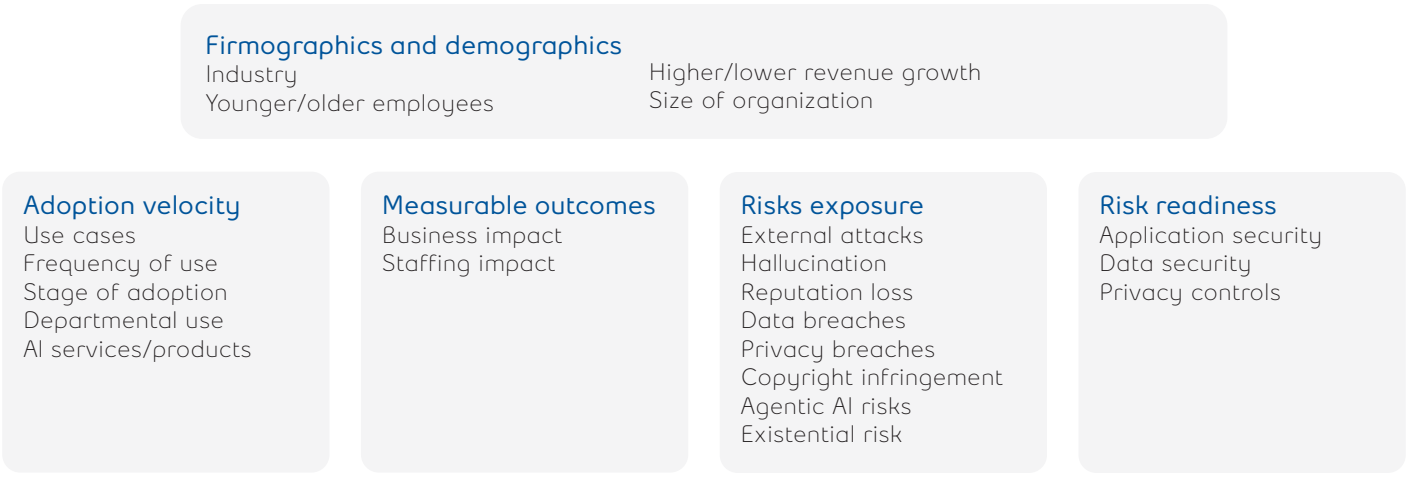

Figure 9: Who we surveyed and what we asked them about

6.2. Firmographics

In what department do you primarily work?

Unweighted base	628
IT/security professional (net)	**51%**
IT (including network ops, software development, etc.)	21%
IT Security (including cybersecurity, data privacy, etc.)	30%
Business leader (net)	**49%**
No department - I am an overall leader	7%
Marketing	3%
Sales	2%
Operations	7%
Finance	23%
Human Resources	4%
Product/Service	3%
Other	0%

Do any of the following match or approximate your current job title?

Unweighted base	310
President/CEO	18%
C-suite (e.g., CTO, CIO, CFO, CMO, etc.)	30%
EVP or other executive	5%
SVP	5%
VP	21%
Director	22%

How many employees work for your organization in total, across all locations?

Unweighted base	628
100 to 499 employees	21%
500 to 999 employees	34%
1,000 to 2,499 employees	30%
2,500 employees or more	16%

In which province(s) does your business operate?

Unweighted base	628
Ontario	58%
Quebec	18%
British Columbia	14%
Alberta	13%
Manitoba	9%
Saskatchewan	7%
Nova Scotia	6%
New Brunswick	5%
Newfoundland & Labrador	3%
Prince Edward Island	3%
Nunavut	2%
Northwest Territories	2%
Yukon	1%

Which of the following categories best describes the industry in which your business operates? If none are a perfect description, please choose the closest match.

Unweighted base	628
Public sector/education (net)	**14%**
Health care and child card (including medical assistants/technicians, home health care and after-school care)	2%
Utilities	6%
Government (federal)	1%
Government (provincial)	1%
Education	3%
Infrastructure/media (net)	**14%**
Infrastructure, construction and land development (including construction managers, workers and contractors)	4%
Media and communications	7%
Transportation	4%
Retail (net)	**9%**
Retail trade (including retail stores, non-store and online retailers, gas stations, and motor vehicle and parts dealers)	7%
Real estate and rental services (including automobile and equipment rentals)	2%
Manufacturing	21%
Banking/insurance (net)	**22%**
Finance or banking	17%
Insurance	5%
Business services (net)	**16%**
Professional services (including lawyers, accountants, engineers, architects, and advertising and public relations specialists)	11%
Business services (including arts and entertainment, social assistances, warehousing and waste management)	5%
OTHERS (net)	3%
Agriculture, forestry, fishing, hunting, mining, quarrying, and oil and gas extraction	1%
Other	2%