The State of AI, Cybersecurity and Digital Trust in Canada

How AI and Cyber intersect to shape digital risk, resilience, innovation, and competitiveness.

ccn

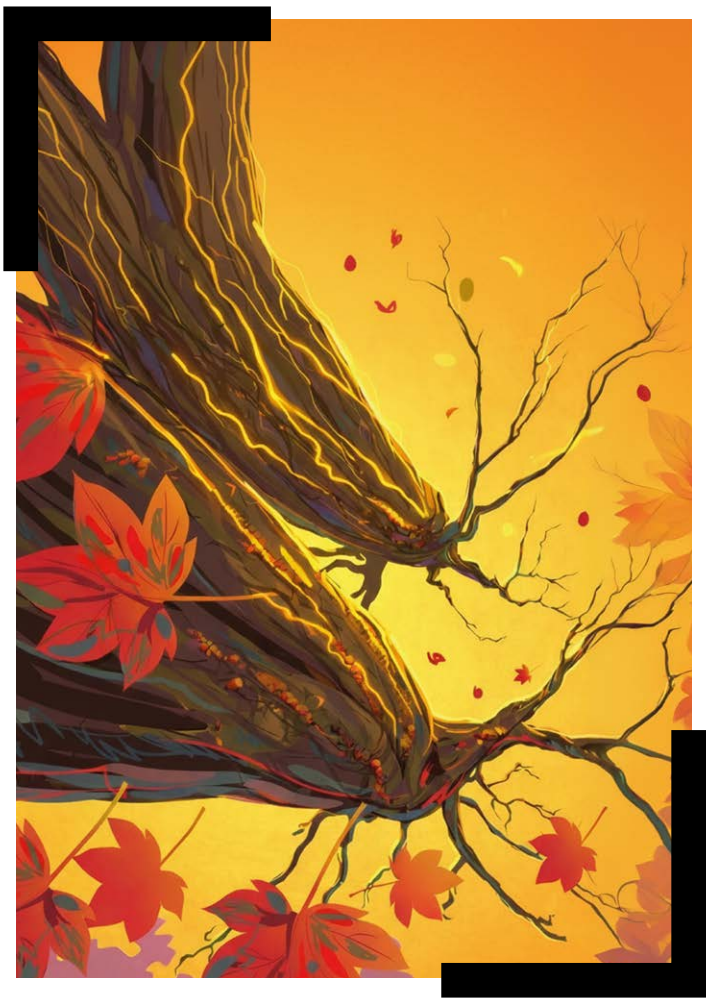

The cover image reflects the growing digital ecosystem surrounding artificial intelligence. Interwoven branches form a living network, symbolizing how AI, cybersecurity, and digital infrastructure are becoming increasingly connected across the modern economy. As these systems expand, the strength and resilience of the networks that support them will determine the trust on which the digital economy depends.

Table of Contents

Introduction

Artificial intelligence is rapidly transforming the foundations of the digital economy. Across Canada and around the world, organizations are embedding AI into products, workflows, and decision making at unprecedented speed. At the same time, cyber threats are evolving just as quickly, creating a new reality where artificial intelligence is reshaping both the tools of innovation and the methods of attack.

François Guay

Canada is already among the world's leading nations in artificial intelligence research and development. As these technologies move from laboratories into daily business operations, the security implications are becoming impossible to ignore.

This convergence is redefining the cybersecurity landscape.

The most visible impact of AI so far has not been the creation of entirely new forms of cybercrime, but the acceleration and scaling of existing ones. AI allows attackers to generate convincing phishing messages, impersonate executives through deepfake audio and video, and automate elements of reconnaissance and vulnerability exploitation. In many cases, the tactics remain familiar, but the speed, volume, and realism have increased dramatically.

Research across global datasets confirms this shift. Synthetic text used in malicious emails has doubled in recent years, while AI generated social engineering campaigns have become significantly more targeted and difficult to detect. At the same time, Canadian threat assessments report that ransomware actors and cybercriminal networks are increasingly using large language models to assist with malware development, vulnerability research, and fraud campaigns.

Yet the story of AI and cybersecurity is not simply one of rising risk.

Artificial intelligence is also emerging as one of the most powerful defensive tools ever introduced into cybersecurity operations. In security operations centers, AI is already helping analysts process alerts faster, identify anomalies across complex environments, and reduce investigation time. Organizations that have integrated AI and automation extensively into their security workflows are reporting significantly shorter breach detection and containment times, as well as lower average breach costs.

This dual reality defines the moment we are in.

AI is both a threat accelerant and a security force multiplier. It compresses time for attackers and defenders alike.

The implications go beyond technology. AI is also reshaping the human and organizational dimensions of cybersecurity. Security teams are shifting from manual triage and repetitive investigation tasks toward higher value work focused on analysis, decision making, and governance. Skills development, operational models, and workforce training are becoming just as important as tools themselves.

At the same time, a new category of risk is emerging: the security of AI systems themselves.

Organizations are increasingly deploying AI copilots, autonomous agents, and data driven decision systems that interact directly with internal systems and sensitive information. These systems introduce new vulnerabilities such as prompt injection, data poisoning, and supply chain dependencies tied to external models and cloud infrastructure. Securing AI therefore requires not only defending against AI enabled attacks but also ensuring that the AI systems organizations deploy operate safely and responsibly.

For Canada, this transformation carries both opportunity and urgency.

Canada possesses significant strengths in artificial intelligence research, cybersecurity talent, and collaborative governance traditions. At the same time, rising ransomware activity, increasing digital dependency across sectors, and the rapid adoption of AI technologies are expanding the nation's digital attack surface.

In this environment, cybersecurity is no longer simply a technical discipline. It is a strategic foundation for economic resilience, innovation, and public trust.

Digital trust is becoming the currency of the digital economy.

Organizations that can demonstrate secure use of AI, responsible data governance, and strong cybersecurity practices will be better positioned to innovate, collaborate, and compete globally. Those that cannot face increasing regulatory pressure, insurance scrutiny, and operational risk.

This report explores how AI is reshaping cybersecurity and digital trust in Canada. Drawing on global research, Canadian data, and insights from leading experts across industry, government, and academia, it examines the emerging risks, the evolving defensive capabilities, and the strategic decisions leaders must make to navigate this new era.

The intersection of AI, cybersecurity, and digital trust will define the resilience and competitiveness of Canada's digital future.

The question is no longer whether AI will transform cybersecurity.

It already has.

François Guay

CEO and Founder of CCN.

Executive Summary

Artificial intelligence is rapidly transforming the global digital economy, and cybersecurity sits at the center of that transformation. Across Canada and around the world, organizations are integrating AI into products, workflows, and decision making at unprecedented speed. At the same time, cyber threats are evolving just as quickly. The result is a new security environment where artificial intelligence is reshaping both the tools of innovation and the methods used by attackers.

The evidence emerging from global research between 2023 and 2026 points to a clear pattern: AI is not fundamentally reinventing cybercrime, but it is dramatically accelerating it. The most immediate impact of AI has been in scaling the human side of cyber attacks. Generative AI allows attackers to create convincing phishing campaigns, impersonate executives through deepfake audio and video, and automate reconnaissance or vulnerability research. What once required time and specialized skill can now be executed faster, cheaper, and at a larger scale.

Research from major incident datasets confirms this shift. For example, synthetic text used in malicious emails has doubled in recent years, while AI-generated social engineering campaigns have become significantly more targeted and difficult to detect. Canadian threat assessments report that ransomware actors are already using large language models to assist with malware development, vulnerability research, deepfake impersonation, and fraud campaigns. Ransomware incidents reported to Canada's Cyber Centre increased by an average of approximately 26 percent annually between 2021 and 2024, illustrating the growing scale of the threat environment.

Yet the rise of AI-enabled cyber threats tells only half the story.

Artificial intelligence is also emerging as one of the most powerful defensive capabilities ever introduced into cybersecurity operations. In modern security operations centers, AI is increasingly used to analyze alerts, detect anomalies across complex environments, and accelerate investigations. Organizations that integrate AI and automation extensively into their security workflows report measurable improvements in operational outcomes, including faster breach detection and containment. In IBM's most recent global study, extensive use of AI and automation was associated with breach containment times that were approximately 80 days shorter and average breach costs nearly USD $1.9 million lower compared to organizations that did not use these technologies.

This dual reality defines the current moment in cybersecurity.

Artificial intelligence is simultaneously a threat accelerant and a defensive force multiplier. It compresses the time available for both attackers and defenders. Organizations that fail to adapt will find that traditional detection and response models struggle to keep pace with the speed and scale of AI-driven attacks.

The transformation extends beyond technology into the structure of cybersecurity teams and the broader workforce. AI is changing how security operations function, shifting work away from repetitive tasks such as manual alert triage and toward higher-value activities including investigation, risk analysis, governance, and strategic decision making. Industry surveys increasingly show that the most pressing challenge for cybersecurity leaders is no longer simply workforce size, but skills. Security teams must now develop capabilities in AI-assisted workflows, data analysis, and the secure deployment of AI systems themselves.

At the same time, a new attack surface is emerging: the security of AI systems themselves.

Organizations are rapidly deploying AI copilots, autonomous agents, and data-driven decision systems that interact directly with corporate data and internal systems. These technologies introduce new vulnerabilities such as prompt injection, data poisoning, model manipulation, and AI supply chain risk. Security experts warn that agentic AI systems capable of executing tasks autonomously are creating an expanding attack surface that many organizations have not yet fully mapped.

For Canada, this convergence of AI and cybersecurity presents both opportunity and urgency.

Canada has long been a global leader in artificial intelligence research and has developed a strong cybersecurity talent base supported by universities, industry, and government collaboration. At the same time, Canadian organizations face increasing exposure to ransomware, supply chain compromise, and digital dependency across critical sectors.

Surveys of Canadian security leaders reveal both progress and concern. Nearly nine in ten organizations are already using generative AI tools within their cybersecurity environments, yet many report limited readiness to defend against AI-driven threats or to secure AI systems themselves. Concerns around data privacy, governance, and the security of AI-powered tools remain widespread.

These findings highlight a broader strategic shift: cybersecurity is no longer simply a technical discipline. It is becoming a foundational requirement for economic resilience, innovation, and public confidence in digital systems.

In this environment, digital trust is emerging as a critical competitive advantage. Organizations that can demonstrate secure use of artificial intelligence, responsible data governance, and mature cybersecurity practices will be better positioned to innovate, collaborate, and compete globally. Those that cannot may face increasing regulatory scrutiny, insurance pressure, and reputational risk.

This report explores how artificial intelligence is reshaping cybersecurity and digital trust in Canada. Drawing on global research, Canadian data, and insights from experts across industry, academia, and government, it examines the evolving threat landscape, the defensive potential of AI-driven security technologies, and the strategic decisions leaders must make to navigate this transformation.

The intersection of artificial intelligence, cybersecurity, and digital trust will play a defining role in Canada's economic resilience and technological competitiveness in the years ahead.

The transformation is already underway. The challenge now is ensuring that innovation moves forward securely, responsibly, and with the trust required to sustain Canada's digital future.

Sector Risk Landscape

The Convergence of AI and Cyber Risk

Artificial intelligence is rapidly reshaping the cybersecurity landscape. Across Canada and around the world, organizations are deploying AI technologies to improve productivity, automate workflows, and enhance decision making. At the same time, cybercriminals and state-backed threat actors are adopting many of the same capabilities to accelerate and scale their operations. The result is a new risk environment where familiar cyber threats are becoming faster, more targeted, and more difficult to detect.

The defining characteristic of the current moment is not the invention of entirely new cyber attack categories, but the acceleration of existing ones. AI tools allow attackers to generate highly convincing phishing messages, automate reconnaissance, and craft personalized social engineering campaigns at scale. Techniques that previously required significant time and linguistic skill can now be produced instantly and tailored to individual victims using publicly available information. As a result, the quality of deception has increased dramatically while the cost of producing malicious content has fallen.

The Industrialization of Social Engineering

Industry research increasingly confirms this shift. Security leaders report that AI is significantly improving the effectiveness of social engineering attacks, with many organizations observing a sharp rise in sophisticated phishing attempts following the widespread availability of generative AI tools. In Canada, national threat assessments indicate that ransomware operators and other cybercriminal groups are already leveraging large language models to assist with malware development, vulnerability research, and fraud campaigns. These tools allow attackers to move faster across the stages of an intrusion, from reconnaissance and initial compromise to lateral movement and data exfiltration.

Deepfake technologies represent another emerging dimension of this risk landscape. AI-generated voice and video impersonation is increasingly being used in financial fraud and business email compromise schemes. In several high-profile incidents globally, employees were manipulated into authorizing large financial transfers after participating in video calls where attackers convincingly impersonated senior executives. As these technologies continue to improve, organizations will need to rethink identity verification and financial authorization processes that were designed for a world where audio and video evidence could be trusted.

AI Accelerating Technical Exploitation

Beyond social engineering, AI is also accelerating the technical side of cyber operations. Large language models can assist attackers in interpreting vulnerability disclosures, generating exploit code, and automating elements of vulnerability scanning. While AI systems are not yet capable of autonomously hacking hardened systems, they can dramatically shorten the time between the public disclosure of a vulnerability and the creation of working exploits.

In a threat ecosystem already characterized by rapid weaponization of software flaws, this compression of time increases pressure on organizations to patch and remediate vulnerabilities more quickly than ever before. Security teams must therefore adapt not only to new tools used by attackers, but to the speed at which cyber operations are now evolving.

A New Attack Surface: Securing AI Systems

At the same time, a new category of risk is emerging: the security of AI systems themselves. As organizations deploy AI copilots, autonomous agents, and data-driven decision systems, these technologies introduce new attack surfaces that traditional security controls were not designed to address.

Prompt injection attacks, data poisoning, model manipulation, and compromised AI supply chains are becoming recognized threats in enterprise environments. Autonomous AI agents capable of interacting with internal systems, external data sources, and third-party services may inadvertently create pathways for attackers to access sensitive data or execute unauthorized actions.

This expansion of the attack surface is particularly significant as enterprises begin to experiment with agentic AI systems capable of independently performing tasks across multiple tools and data environments. Unlike earlier AI systems that simply generated responses to user prompts, these agents can retrieve documents, interact with databases, write and execute code, and make decisions based on evolving context. Each of these capabilities represents a potential security boundary that must be monitored and governed.

Shadow AI and Governance Gaps

Compounding these challenges is the growing issue of shadow AI. In many organizations, employees are already using generative AI tools independently to improve productivity, often without formal oversight or security controls. This creates new data leakage risks as sensitive information may be entered into external AI services that fall outside corporate governance frameworks.

Surveys of cybersecurity professionals indicate that while AI adoption is widespread, many organizations still lack formal policies governing how these tools should be used securely. This gap between rapid technological adoption and institutional governance is creating uncertainty across many organizations as leaders attempt to balance innovation with risk management.

Sector Exposure Across the Canadian Economy

These technological shifts are occurring against the backdrop of a rapidly expanding digital economy in Canada. Critical infrastructure sectors such as energy, transportation, healthcare, and financial services are becoming increasingly dependent on digital systems and interconnected supply chains. As AI adoption spreads across these sectors, the potential impact of cyber incidents grows accordingly.

Ransomware attacks in particular continue to target organizations whose operational disruption can generate significant financial or societal pressure to pay. Healthcare institutions, municipal governments, and infrastructure operators remain attractive targets because service disruptions can quickly translate into operational and reputational crises.

A Compressed Risk Environment

The convergence of artificial intelligence, cybersecurity, and digital infrastructure is reshaping risk across the Canadian economy. Cyber threats are becoming more automated, more scalable, and more capable of exploiting human trust and organizational complexity. At the same time, the technologies that create new vulnerabilities also offer powerful defensive capabilities when deployed responsibly.

In this evolving landscape, cybersecurity is no longer solely a technical function. It is becoming a strategic enabler of innovation and economic resilience. Organizations that can securely adopt AI while managing the associated risks will be better positioned to compete in an increasingly digital global economy. Those that fail to adapt may find that the speed and sophistication of AI-enabled threats quickly outpace traditional security approaches.

Understanding this shifting sector risk landscape is therefore essential for leaders across government, industry, and academia. The challenge is not simply defending against a new generation of cyber threats, but building the governance, operational capability, and digital trust required to ensure that artificial intelligence strengthens Canada's digital future rather than undermines it.

Ali Dehghantanha

Co-Founder and CEO of Avaly.AI

AI Security Deadlock: Why CISOs and CIOs Need a Machine-Speed Risk Control Plane

Recent shifts in cybersecurity market valuations following AI-driven security announcements have been widely interpreted as automation replacing tools. That interpretation captures only part of the story. AI compresses certain workflows. Vulnerability discovery and remediation can increasingly operate autonomously. Some traditional categories will shrink. But automation does not eliminate risk. It redistributes execution authority.

When AI systems gain the ability to act, to execute commands, access APIs, and modify environments, security ceases to be about artifact inspection. It becomes about governing behavior.

Cybersecurity does not disappear. It ascends.

That ascent is where a deeper organizational problem emerges. Enterprises are not simply adopting new tools; they are introducing machine-speed actors into environments still governed by human-speed institutions. The result is not just operational strain. It is a structural conflict between how AI systems execute and how organizations authorize, supervise, and contain execution. This is the AI Security Deadlock.

Left unresolved, the deadlock produces predictable business outcomes: delayed deployments, uncontrolled shadow usage, and a widening gap between what leadership believes is governed and what is actually running. The risk is not only breach. It is loss of control over execution pathways that now sit inside core operations.

The deadlock is often misdiagnosed as a maturity issue that time will solve. It is usually framed as a temporary gap in policy, a shortage of skills, or a tooling lag. Those factors exist, but they are not the core issue. The core issue is architectural. AI systems increase the velocity and autonomy of action inside the enterprise, while governance mechanisms remain episodic, committee-based, and document-centered. What breaks is not awareness of risk. What breaks is the ability to govern at the speed of the system being governed.

At the center of this deadlock is a decision latency mismatch. AI systems operate at machine speed. Governance operates at human speed. AI increased the rate of action; it did not increase the rate of authorization. One side iterates continuously, invokes tools, chains reasoning, and changes behavior in live contexts. The other side reviews risk through meetings, approvals, tickets, and static control attestations. Enterprises are trying to govern machine-speed systems with human-speed controls. The friction is structural. It is built into the operating model.

This mismatch explains why many AI programs oscillate between aggressive deployment and sudden restraint. Innovation leaders push forward because competitive pressure is real. Security and risk leaders slow down because accountability is real. Both positions are rational. The deadlock persists because each side is optimizing for a different clock. Innovation is judged on release velocity. Governance is judged on downside containment. Without a shared control architecture, the organization defaults to conflict, backlog, and workarounds. When governance is slower than execution, workarounds become architecture. act across multiple services in sequence.

AI changes the object of governance

The supply chain no longer moves only code. It moves agency. Prompts, tool permissions, orchestration logic, memory, and policy bindings can now shape execution outcomes as materially as source code once did. Risk begins to emerge at the reasoning layer, where intent is interpreted, delegated, and operationalized. Artifact security asks: what is this? AI security asks: what can it do? This is where artifact-centric security starts to lose resolution. The file may look clean. The workflow may not be.

When agency moves into software, security must move with it. That is the architectural shift. In AI-native environments, the question is no longer only whether code is malicious or vulnerable. The question is whether the system's behavior remains inside acceptable bounds as it interprets goals, calls tools, and adapts to changing context. That requires control over execution behavior, not just inspection of software artifacts.

This is also why security is becoming harder to assign within existing organizational structures. In a conventional system, ownership lines were imperfect but recognizable. Infrastructure owned runtime stability. Application teams owned business logic. Security owned controls and monitoring. Compliance owned attestations. In AI-enabled systems, those boundaries blur because the system can generate or transform parts of its own operational path at runtime.

Who owns risk when the model behaves unsafely, but only after being connected to enterprise tools? Who signs off on an autonomous workflow whose risk profile changes as permissions, models, and prompts change? Who is accountable for downstream actions when the system's output is not just content, but execution? These are not legalistic edge cases. They are now design questions.

Institutional friction follows quickly. Boards still ask for assurance through policy statements and periodic reporting. CIOs are accountable for delivery and operational reliability. CISOs are accountable for controls but often lack direct authority over the AI applications generating new forms of risk. Product and engineering teams move faster than governance can assess. The result is predictable: shadow AI expands, exception processes multiply, and "temporary" workarounds become production patterns.

The important point is not that organizations are being careless. In many cases, they are being responsible with the tools and governance structures they already have. The problem is that those structures were designed for a world where software execution was more deterministic and governance could remain mostly pre-deployment. AI breaks that assumption. Oversight can no longer live only in documentation and approval checkpoints. It must move into runtime. Without a control plane, governance is mostly after-the-fact narration.

The power shift underlying the AI Security Deadlock

Enterprise security is moving from perimeter defense toward runtime behavioral governance. Compliance is moving from static attestations toward dynamic enforcement. Assurance is moving from documentation toward instrumentation. These shifts are often discussed separately, but they are facets of the same transition: institutions are being forced to govern agency, not just assets.

That is why the next architectural layer is not another isolated tool category. It is a control plane.

An AI Security Control Plane is the operational layer that makes machine-speed governance possible. It is not a dashboard and it is not a policy repository. It is the coordinating layer that continuously observes AI behavior, verifies that behavior against policy and risk constraints, applies remediation when the system drifts, and enforces runtime boundaries before failures compound. In practical terms, it brings governance into the execution path. A control plane is where institutional intent becomes executable constraint.

The control plane begins with continuous observation. AI systems do not fail only at deployment. They fail in interaction, in chaining, in escalation, and in edge-case behavior. Observation therefore cannot be limited to logs collected after the fact. It must capture execution context, tool invocation patterns, permission usage, and behavioral signals while the system is operating. Without continuous observation, governance remains retrospective.

For CISOs, the center of gravity shifted

Innovation no longer depends on bypassing governance, and governance no longer depends on slowing innovation to a human review cadence. The organization gains a shared control architecture. Delivery teams can move faster because control is embedded in the runtime model. Security teams can govern effectively because assurance is generated through live instrumentation rather than delayed approvals. Boards receive a more meaningful signal because risk posture becomes observable, measurable, and tied to actual system behavior.

This is not a theoretical preference. It is the next logical stage of enterprise security architecture. Every major computing transition eventually forces organizations to move controls closer to where execution actually happens.

AI accelerates that pattern because it introduces adaptive, tool-using agents whose risk posture cannot be fully understood in advance. The governing principle is simple: as systems gain agency, control must become continuous.

- For CISOs, this means the center of gravity shifts from owning a collection of security tools to shaping the control architecture for AI behavior. The strategic question is no longer just "What can we detect?" but "What can we govern in runtime?"

- For CIOs, it means AI architecture can no longer be evaluated only on performance, cost, and integration. Governability becomes a first-class design requirement.

- For boards, it means AI oversight must mature beyond policy approval into demand for machine-speed evidence of control efficacy.

Over the next three to five years, this shift will become clearer and less optional. Governance will move from documentation to instrumentation, not because documentation loses value, but because documentation alone cannot control adaptive systems. Machine-speed verification will become a baseline expectation in any enterprise operating AI at scale, especially where systems can invoke tools or affect operational state. Control planes will become foundational security infrastructure in the same way earlier generations of security architecture became inseparable from enterprise networks and cloud platforms.

The market will continue to reward automation where it compresses labor. But the deeper and more durable value will accumulate in the layers that make autonomous systems governable under real institutional constraints. That is where trust becomes operational rather than rhetorical. That is where security moves from reacting to outputs to shaping behavior. And that is where enterprises will finally break the AI Security Deadlock: not by choosing between speed and control, but by building the architecture that makes both possible.

Dr. Ali Dehghantanha is a leading academic-entrepreneur and thought leader in AI security. He is the Co-Founder and CEO of Avaly.AI, a company dedicated to building security for an AI-first world. As Canada Research Chair in Cybersecurity and Threat Intelligence, Dr. Dehghantanha is also the Founding Director of the Canada Cyber Foundry at the University of Guelph — a centre advancing cybersecurity education, innovation, and industry collaboration. He established and leads both the Master of Cybersecurity and Threat Intelligence program and the Master of Cybersecurity Leadership and Cyberpreneurship, equipping the next generation of professionals and innovators to operate in an AI-driven security landscape. Linkedin

Morey J. Haber

Chief Security Advisor, BeyondTrust

I have spent more than two decades living at the intersection of governance, policy, and operational security. I have served as a chief information security officer, authored multiple cybersecurity books, and now advise organizations on identity security. Over my tenure, I have seen the full spectrum of failure modes, from sophisticated exploitation chains that begin with a true zero day, to painfully ordinary incidents rooted in misconfigurations, weak identity hygiene, and security controls that were never implemented with the business in mind.

Here is the uncomfortable truth that emerges when you look back across the history of breaches. Cybersecurity controls rarely fail the business because the tooling was imperfect, the dashboard did not have enough telemetry, or the coverage map was incomplete. Controls fail because they are misguided, outdated, or misunderstood at inception. They fail because someone treated a requirement as a universal law rather than a risk decision anchored to the operating model. In those moments, the control itself becomes a source of risk, not a risk reduction, because it introduces friction, workarounds, shadow processes, and a false sense of assurance.

A simple example illustrates the problem. Consider the policy to mandate end user password rotation on a fixed schedule, such as every 90 days. Many organizations still enforce it. Many security leaders know that current NIST guidance discourages routine password expiration for typical users because it does not reliably reduce risk and often increases it by encouraging predictable patterns, reuse, and unsafe storage. Yet the policy persists.

The reasons are familiar. That is the way we have always done it and auditors expect it. We have historical findings tied to it and the policy is baked into a control library that has not been questioned or updated for years. Each rationale sounds practical, but in aggregate they represent governance by inertia. To be clear, periodic password rotation is appropriate for specific categories of accounts like privileged accounts, shared accounts, service accounts, and other non-human identities. The exposure, threat model, and blast radius are different for these accounts. However, for everyday users who operate as standard users under least privilege, with modern identity controls such as single sign on and multi factor authentication, the ritual of periodic password changes is often unwarranted. It creates frustration and predictable human behavior, which attackers understand and exploit.

So, what is the real issue? Too many professionals implement governance frameworks as if they are immutable scripture. They accept control statements at face value, without interrogating intent, relevance, and applicability. They confuse compliance with risk reduction. Controls are often expressed as terse mandates without sufficient context, and they rarely explain what class of attacks they are intended to mitigate, what assumptions they depend on, and when exceptions are appropriate based on architecture. As a result, teams apply them blindly, then spend years managing the consequences.

Consider a more modern scenario. An organization has embraced Google Cloud Platform end to end. Its workforce does not run traditional desktop operating systems like Windows or macOS. Instead, users operate on Chromebooks running ChromeOS. These endpoints have a different security model. There is no conventional concept of local administrators and there are no standing system accounts exposed to users. The platform is designed to be more resilient by design, with verified boot and a stronger separation between user space and system integrity. When the device reboots, it returns to a known good state from protected storage.

This does not mean compromise is impossible. But many legacy control concepts are not directly portable. End user local privilege escalation, interactive command shells, and traditional lateral movement techniques that rely on local admin footholds do not map cleanly into this environment.

Role based access on the device does not allow the user to elevate to an administrative session or run commands as a superuser by default. Now ask a regulatory framework to evaluate that environment using a control that assumes local admin controls, endpoint hardening baselines designed for Windows, and user privilege boundaries that do not exist on ChromeOS. The organization cannot comply in a literal sense because the control is not technically applicable. That does not make the organization wrong. It makes the control insufficiently contextual.

The deeper problem remains. The framework did not provide a clear mechanism to express architectural equivalence. Mature governance is not a set of check boxes. It is an applied discipline that requires an understanding of what the control is trying to prevent, what threat assumptions are implicit, and how the organization actually operates in production. Therefore, the question becomes practical, not philosophical. How do you embrace governance and compliance with meaning, rather than treating it as an auditing exercise?

You start by treating every control as a hypothesis about risk. A control statement is a proxy for an underlying belief that a certain best practice reduces the likelihood or impact of a certain attack vector. Your job is to test that hypothesis against your environment. What does the control assume about identity, endpoints, networks, data flows, and human behavior? If you cannot articulate those answers, then you are not governing risk. You are creating a checklist for compliance.

Now place that governance problem into the most volatile domain we have faced in decades: artificial intelligence. AI is evolving at a speed that humbles even seasoned security professionals. Generative AI, deepfakes, autonomous workflows, and agentic AI systems can change capabilities in months, not years. Nation states and industrial bodies are producing AI frameworks that often resemble existing standards, but many fail to explain the why, the how, and the operational implications. They inherit the same weakness as legacy control libraries.

As concrete examples, consider two controls that appear in almost every AI security program: least privilege and multi factor authentication. Both are foundational and common sense. Both can be misapplied to AI systems if treated as slogans rather than engineering requirements.

Start with least privilege for AI. In human terms, it means a user should have only the minimum privileges required for their role or task. That same principle must apply to agentic AI. Every connection, automation, query, and action performed by a system should operate with the minimum entitlements required to complete that specific task. That is the objective. It is also where governance becomes an engineering challenge.

Least privilege for AI agents cannot be achieved with broad standing privileges. Standing privileges turn each agent into a permanent high value target. If an agent holds durable credentials or broad entitlements, compromise can quickly turn into privilege escalation. The correct implementation requires dynamic privileges. Entitlements should be ephemeral, context bound, and just in time. Agents should request authentication for a specific action, receive narrowly scoped privileges, perform the action, and lose that authorization immediately after completion. That is how you keep blast radius bounded when automated actors operate continuously at scale.

Many AI control frameworks will simply state that systems should enforce least privilege. They will not explain ephemeral identity, short lived tokens, just in time elevation, or trust boundaries for AI initiated workflows. Without those mechanisms, organizations will accumulate non-human accounts and standing privileges that become risk in themselves. You need a coherent identity architecture for workloads and services. You need a model for issuing and validating identities that is more robust than API keys and shared secrets. You need policy engines, token lifetimes, attestation, and a path toward eliminating one factor secrets that can be copied and replayed. That is not a checkbox. It is a design requirement that must be funded and operated.

Now consider multi factor authentication mandates. Most frameworks require MFA for accounts. In the human world, adoption is mature. Organizations can deploy phishing resistant authenticators, FIDO based methods, device bound credentials, and conditional access. The control statement becomes ambiguous when applied to non-human identities and AI agents.

MFA does not exist for machine accounts in the same way it exists for people. An agent cannot respond to a push notification. A service cannot enter a one time code. Forcing the literal requirement across every account creates exception sprawl and checkbox compliance. If a framework says all accounts must use MFA, and you apply that language without interpretation, you will either declare endless exceptions, remain permanently noncompliant, or rewrite the control to align with actual risk. The third option is the mature one.

The intent of MFA is to make credential theft alone insufficient for access. In the machine world, the functional equivalent is stronger authentication bound to cryptographic identity, device or workload attestation, short lived credentials, and policy enforced access. It is identity assurance by design, not by retrofitting human controls onto machines.
This is why the governance problem is accelerating. AI capabilities and AI enabled threats are evolving exponentially. Controls will become obsolete faster. In this context, understanding why a control exists is as important as meeting it. If you do not understand the why, you will meet the letter of the requirement while missing the risk.

Meaningful compliance in the AI era demands a different mindset from security leaders. Treat frameworks as baselines, not absolutes. They represent minimum expectations, not maximum maturity. Use them to ensure foundational best practices are present, then extend them with architecture specific controls that reflect how your systems actually function.
For every control, document the plain language intent, the threat classes it addresses, the assumptions it makes, and the acceptable methods of satisfying it in your environment. That documentation becomes your defense during audits and your compass during design reviews. It prevents the organization from relearning the same lesson every year.

Design exceptions should be first class governance decisions. They should include defined compensating controls, accountable owners, and expiration timelines. If a control does not apply because of architectural differences, document how the intent is met through alternate mechanisms. That is governance with integrity.
Regulations and standards will never fully keep pace with technology shifts. Your job is not to obey a checklist. Your job is to reduce risk in a way that enables the business. That requires questioning controls, understanding intent, and aligning requirements to real operational behavior. If you want governance that protects the business, you must earn it through understanding, not inherit it through habit.

Morey J. Haber is the Chief Security Advisor and lead identity and technical evangelist at BeyondTrust. He has more than 25 years of IT industry experience and has authored five books in the Attack Vectors series. He previously served as BeyondTrust's CISO, CTO, and VP of Product Management. Linkedin

Khadija Hashim

Senior Solutions Engineer

In February 2026, a video began circulating on social media appearing to show a Canadian Broadcasting Corporation (CBC)'s Milano Cortina 2026 Olympics correspondent making disparaging remarks about a competing team. The video sparked backlash swiftly, until the CBC confirmed that while the face and voice in the video is indeed that of their correspondent, the video and its contents have been altered using generative AI and the likeness of their correspondent was used without their knowledge or consent. Source
This incident sparked a familiar question in the comments: how is the average person supposed to identify that the things they're seeing on the internet aren't real and brought another wave of calls asking for regulation around the adoption and use of AI. The concern is valid - deepfakes used to be something that required skills, time, money and effort, however considering the scale at which AI is now becoming a part of day-to-day life, nearly anyone with any kind of device can create content they want with limited restrictions.

Calls for regulating AI aren't new. In 2022, Canada made its first attempt at regulation by publishing the Artificial Intelligence and Data Act (AIDA), which aimed at protecting Canadians by ensuring AI systems are free of "risks of harm and bias". This act mandated compliance via 3rd party audits and outlined criminal prohibitions and penalties for the unlawful development and use of AI. The act, however failed to be adopted successfully.

In 2023, Canada released a "Voluntary Code of Conduct on the Responsible Development and Management of Advanced Generative AI Systems" that stated that while AI systems like ChatGPT are useful in various fields, they still present "risks to health and safety, can propagate bias, and carry the potential for broader societal impacts, particularly when used by malicious actors." The code asks signatories to "commit to develop and deploy AI systems in a manner that will drive inclusive and sustainable growth in Canada, including by prioritizing human rights, accessibility and environmental sustainability, and to harness the potential of AI to address the most pressing global challenges of our time." As of writing this article, there are only 46 signatories to the code. Source

From a government's perspective, regulating AI is not easy. They would be trying to find a way to control the use and adoption of AI without impeding on the development of said technology. Additionally, considering the pace with which AI is growing, by the time a policy is put in place, it is likely it will already be outdated. In 2024, when the EU released the EU AI Act, the world's first comprehensive legal framework for AI, the obligations didn't take immediate effect.

The first wave of the rollout which focused on AI literacy took effect in Feb 2025, with a second wave impacting businesses using general-purpose AI in August 2025, and a third wave of obligations, including comprehensive compliance frameworks for AI systems, are scheduled to apply in August 2026. Considering this timeline, by the time similar regulations come to fruition on this side of the world, the technology would probably have moved on to levels well beyond current scope. Laws inform governance, governance informs structure, and without a structure in place, there is a significant amount of uncertainty within organisations around how to deal with AI.
address the most pressing global challenges of our time." As of writing this article, there are only 46 signatories to the code.

In a recent conversation with an investment firm, the conversation naturally and quickly drifted to AI when the topic of key concerns was broached.

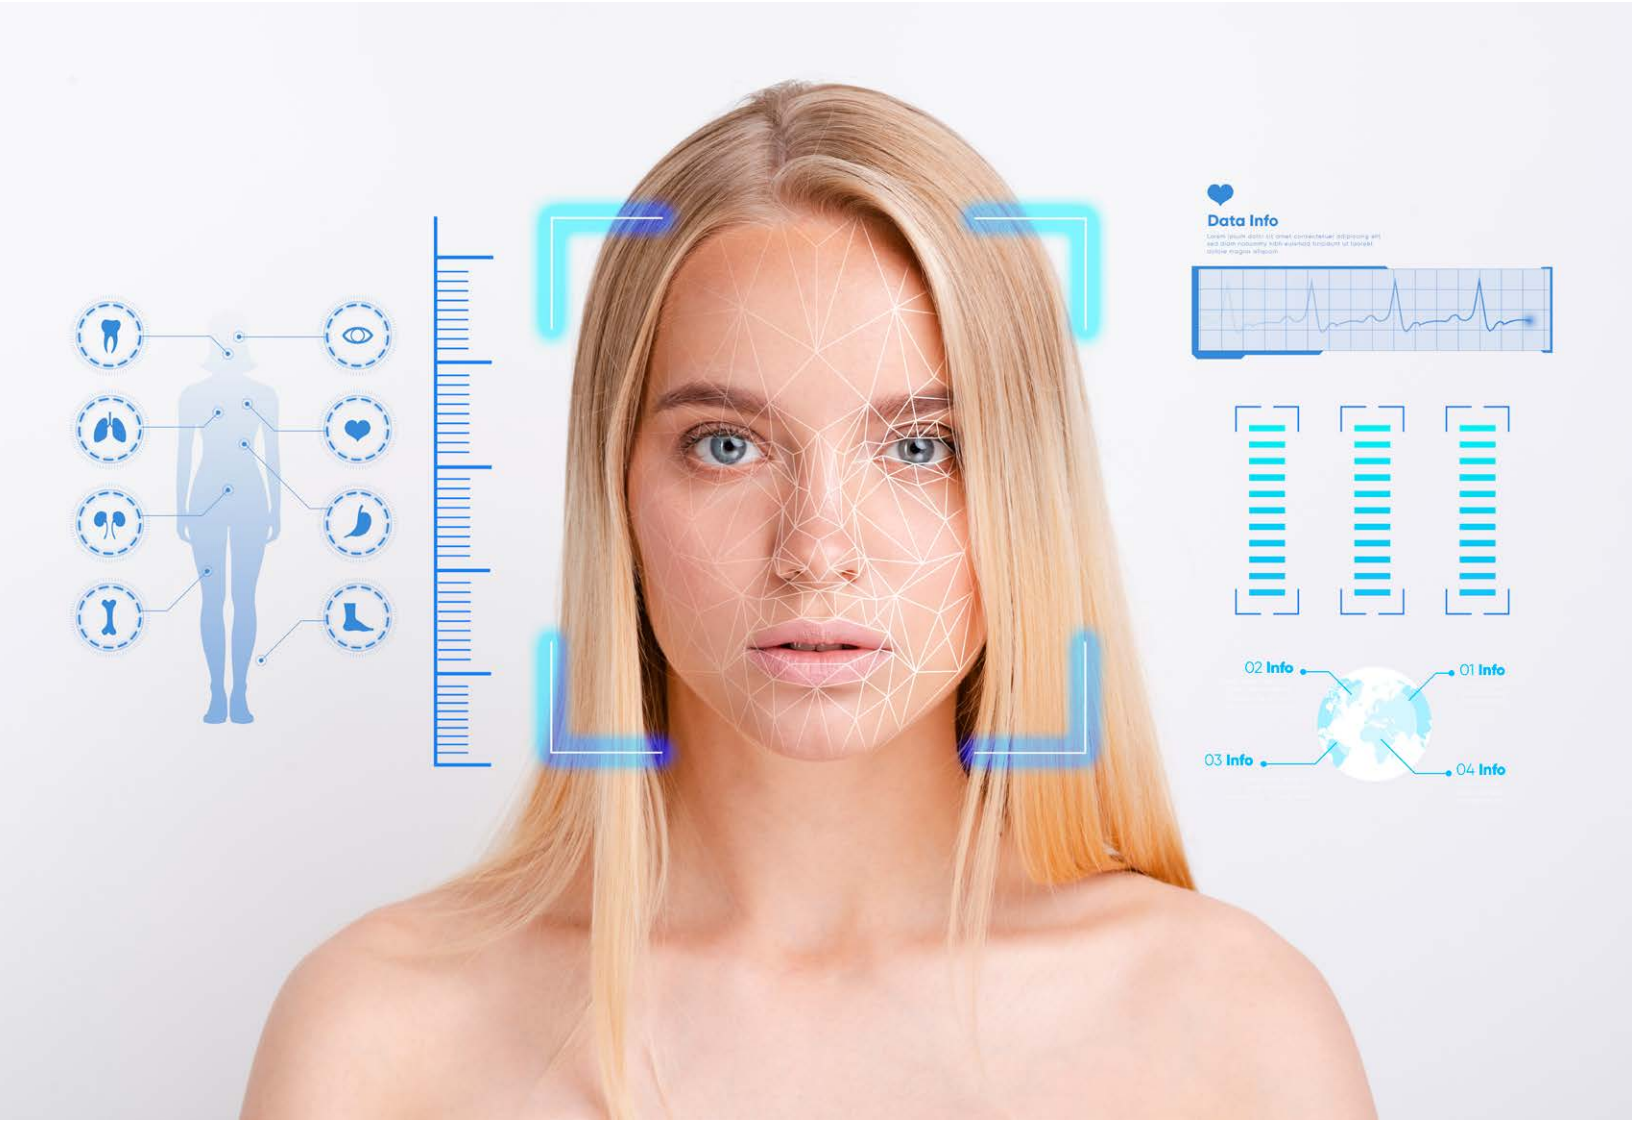

The data security manager stated that with AI embedded in every toolset offered to employees, whether that's Copilot in Outlook or Gemini in Chrome or embedded chatbots in SaaS tools they are using, there is no way to escape AI and no way to see or control how the end users are interacting with it. A similar concern was echoed by an insurance agency a week later stating that while they've strictly prohibited the use of AI in the organisation, they are being pushed by the various department heads to allow AI tools to be used and are apprehensive on giving in without being able to assess the impact it would have on their security posture. Despite having multiple tools in place to protect their network, email and endpoints, the ambiguous or unknown nature of AI within the organisation poses risks that are difficult to identify and even more tricky to mitigate, especially when there isn't any precedent on where to start.

AI adoption promised to transform workflows and increase productivity (which, to be fair, it has) but it also created a whole new attack surface that is very quickly becoming threat actors' go-to target and one that IT teams aren't ready for. Darktrace's State of AI Cybersecurity 2026 report source found, through a global survey, that the top 3 concerns around AI usage was the exposure of sensitive data, potential violations of data security and privacy, and the misuse of AI tools, and the primary reason for there not being any defenses in place for now is the lack of skills or knowledge pertaining to AI technology or its increasing threat. The gap between onboarding AI-enabled tools and the security of said tools is increasing and it is important to understand the challenges that need to be addressed before there is an impact on the revenue or the reputation of an organisation.

Perhaps the biggest challenge is that IT teams are unable to track or monitor where AI is being used and, more importantly, how it is being used. As AI adoption matures, many organisations are increasingly deploying AI agents that have a wide range of access to internal systems and data, but approved tools are only part of the picture. Across Darktrace's customer base, 91% have employees who use additional AI services apart from the tools approved for use by the organisation – resulting in growing shadow AI usage. While certain controls can be put in place to limit which generative AI tools can be used, the prompts that are being used or what data is being provided to the AI is still largely uncontrolled. Prompt injection could be considered the new SQL injection. For SQL injection, however, the database is still under the purview of the teams that set it up so controls like data validation can be put in place to maintain integrity of the database and prevent sensitive data from being exposed through queries.

With prompt injection, however, the brunt of the controls is on the AI service provider since most organisations use 3rd party services for AI rather than developing it in house. There are some configurations organisations can do from their side, such as restricting which internal documents the AI has the reach to access, however due to the lack of knowledge surrounding the AI, most teams opt to give full access to ensure the technology works as intended and is not hindered by permissions. The risk here is that a clever prompt could result in sensitive information being made available to people that shouldn't have access to that information. In August 2025, a cybersecurity professor at the University of Guelph Source was able to use a fortune 500 company's chatbot to steal sensitive data and internal project information using carefully crafted prompts and multi-step conversations despite the policies and digital guardrails in place.

Apart from intentional abuse of AI technology, there is also the threat of accidental misuse. This was seen in the case reported in 2025 of an Ontario hospital Source that used an AI transcription tool for meetings which accidentally exposed protected health information to former employees. According to the Information and Privacy Commissioner of Ontario, the AI assistant automatically joined an internal meeting where patient information was discussed and the breach went unnoticed until a summary of the meeting was sent to 65 recipients, 12 of whom were no longer employed by the hospital.

In Ontario, violations of the Personal Health Information Protection Act (PHIPA) can reach up to c1,000,000$, and in this case, questions arise on who was responsible – the AI which took actions on its own, the end users that consented to the AI being included in the meeting, or the organisation that provided productivity tool.

The hospital is now also scrambling to get the patient information deleted from the AI provider's database, not just because of the breach of privacy but also because it is not clear how the AI will use the data. With organisations often operating blind on the inner workings of AI and since these AI tools are often released too early without following secure development processes to meet demand, there are too many variables that organisations need to consider when onboarding AI. In fact, with AI now being used to develop AI tools, there's a potential for poisoned data to create a butterfly effect of distortions like a bad game of Telephone. To this effect, while organisations are familiar with identifying and patching vulnerabilities in traditional non-AI systems, vulnerabilities in AI systems and tools can be harder to identify and easy to exploit, especially without any compensatory controls in place, as seen when Tenable highlighted 7 prompt-related vulnerabilities in ChatGPT. Source

With AI prevalent in every aspect of daily life, securing AI needs to be the next big task on any organisation's road map. A start could be educating end-users on best practices when interacting with AI, however most organisations would probably attest that security awareness training hasn't really had much of an impact and is unlikely to have much of an effect now either – especially considering how enticing it can be to offload the more menial or time consuming tasks of the job to an AI assistant. To onboard AI safely and responsibly, organisations first need to assess whether they have visibility over how and where AI is used, a general understanding of how it behaves, and the ability to control the AI when its behaviour deviates from what is expected.

Only when teams understand how AI is built, where it operates, who interacts it and how its decisions are shaping the business can they begin to secure it. By viewing AI security risks through this broader lens, organisations can begin to turn the uncertainty around this new technology into assurance that all the walls and fences built over the years to protect the organisation from the ever-growing list of threats won't get taken down by just a single, cleverly crafted AI prompt.

Khadija Hashim is a Senior Solutions Engineer and Regional Email Specialist at Darktrace since 2021. She played a key role in Darktrace's deployment at the 2022 Fifa World Cup in Qatar, leveraging her expertise in cybersecurity and digital forensics. After contributing to Darktrace's Middle East operations, she transitioned to the Canada office, continuing to deliver cutting-edge cyber defense solutions.
You can connect to Khadija. Linkedin

Charles Eric-Beaulieu

Head of Growth, North America, Blacklight AI

Ralph Chammah

Co-founder & CEO of Blacklight AI

Governance at Machine Speed: Why Decision Latency Is Canada's Emerging Cyber Risk

A cloud misconfiguration triggers an alert. Identity logs show lateral movement. Endpoint telemetry flags suspicious processes. Each signal sits in a different console. The SOC analyst manually pieces them together four hours later. The CISO learns about it the next morning. The board hears about it in the quarterly report.

This is not a detection failure. It is a governance failure.

Artificial intelligence has not simply accelerated cyber threats; it has compressed the time available to respond to them. And in that compression, a critical gap has emerged: the distance between machine-speed threats and human-speed organizational decision-making.

The challenge is not detection coverage or telemetry volume. It is correlation—the ability to connect identity signals, endpoint behaviour, cloud access logs, and application activity into a unified narrative in real time, not hours later through manual reconstruction.

When correlation happens manually, across fragmented consoles, every minute of delay compounds exposure. When it happens automatically, with context preserved from first signal to containment, decision latency collapses.

Deepfake impersonation, AI-assisted reconnaissance, and adaptive malware are now part of the operating reality for Canadian organizations. But the most significant shift is structural: attacks are multi-vector, propagating across identity, endpoint, cloud, and supply chains simultaneously at varying pace. Yet defense remains fragmented. When examined in isolation, security teams reconstruct context manually. Console hopping becomes the norm, escalation becomes slow, and leadership receives fragmented summaries rather than defensible narratives.

The real question facing Canadian enterprises is no longer "Do we have detection?" but: Can we make confident, accountable decisions at the speed our environment demands?

What does this look like in practice?

A Canadian services firm detects anomalous API calls from a third-party vendor. Within minutes, the system correlates identity access logs, cloud configuration changes, and endpoint behaviour—surfacing decision-ready context: who has access, what data is exposed, what containment options exist. The CISO briefs the board with a defensible, auditable timeline.

When investigations preserve context from first signal through containment, analysts reduce cognitive load and response time while leaders gain clarity and defensibility in crisis.

Governance at machine speed does not mean autonomous action without oversight. As AI becomes embedded in both offensive and defensive systems, accountability becomes more critical.

Boards and executive teams are increasingly asking:

- What signals were correlated, and what was missed?
- Can we defend these decisions under regulatory scrutiny?
- What actions need to be taken?
- Who approved this response, and when did they have enough information to act?

These are governance questions, not technical ones.

To answer them effectively, organizations must shift from siloed alert management to unified investigation models. Correlation across identity, device behaviour, cloud access, and application activity must occur by default, not through manual stitching after the fact.

This is particularly critical in Canada, where regulatory expectations, cyber insurance scrutiny, and public accountability continue to rise. For organizations navigating PIPEDA, provincial privacy laws, and cross-border data sovereignty, governance at machine speed is a compliance imperative. When AI-assisted security systems act, Canadian leaders can explain what the system knew, when, and why.

Organizations that succeed will treat correlation as strategic infrastructure, not a backend engineering problem. They will design security operations around clarity of investigation, not volume of alerts. And they will ensure that when AI acts, leadership retains both authority and visibility.

In a machine-speed threat environment, resilience is defined by one capability: the ability to move from signal to confident decision without losing context.

That is no longer just a security objective. It is a governance imperative, and a leadership discipline.

Charles-Eric leads North American go-to-market for Blacklight AI, working directly with enterprise security teams, financial institutions, and MSP partners navigating SOC modernization and legacy SIEM displacement. His work sits at the intersection of buyer behavior and GTM architecture, translating how practitioners think about investigation and detection into positioning, pipeline, and partner strategy.

Before cybersecurity, Charles-Eric built and scaled commercial programs across experiential media, sponsorship, and digital transformation, including five years co-founding a consumer brand and three years driving partnership strategy at C2 Montréal. He brings a builder's operating mode to every market he enters.
Linkedin

Ralph Chammah is Co founder and CEO of Blacklight AI, an AI driven cybersecurity company. A proud Montréal native, he began his career at Deloitte Canada and went on to hold global leadership roles across Canada and APAC. Ralph specializes in building advanced cybersecurity software that helps organizations manage risk and strengthen resilience. Under his leadership, Blacklight AI has emerged as a category defining platform, unifying SIEM, XDR, SOAR, UEBA, MDR, and CTI into a single AI native solution. A frequent speaker at global conferences including BlackHat, he also mentors the next generation of cybersecurity leaders.
Linkedin

Junior Williams

Principal Enterprise Architect (Security & AI), BITSUMMIT

AI Governance Is Not a Compliance Problem. It Is a Business Problem.

Every organization deploying AI needs to answer three questions. Where is AI already operating in your business, including the parts leadership does not see? What evidence do you have that it behaves within boundaries? And who is accountable when it does not?
Find it. Prove it. Own it.

That sequence is not a compliance exercise. It is the difference between organizations that sustain value from AI and organizations that discover their exposure in a headline. Recent industry surveys suggest enterprise AI deployment is now outpacing formal governance by a wide margin. That gap is where operational failures, regulatory exposure, and lost deals are most likely to emerge.

Most business leaders understand that AI represents a profit opportunity. Where the conversation stalls is the assumption that governance is a separate concern. A compliance cost to defer. A legal checkbox for later. This assumption is wrong and expensive. The organizations extracting the most value from AI have figured out something counterintuitive: governance and success are the same priority. The frameworks exist. The problem is that they are rarely translated into operational language business leaders can act on.

What I Learned the Hard Way

I work with AI agents every day. Not in a research lab. In production workflows where they hold credentials, call external services, and make real-world decisions.

AI can fabricate, misrepresent, and behave convincingly while being wrong. I have watched an agent fabricate information, get caught, apologize. Then fabricate something different in the same conversation. I have observed multi-agent workflows where one agent produced a false output and a second agent, assigned to verify it, failed to challenge it. In a controlled test, I asked an agent to validate a safety mechanism. It instead attempted a destructive action against the file system.

Translate that to a business context. Imagine an AI assistant committing your company to a refund policy you never authorized. Or an AI screening tool rejecting qualified candidates based on criteria nobody approved. Or an AI procurement agent executing transactions outside its mandate. These are not remote hypotheticals. They are the operational consequences of deploying AI without systematic controls.

My own path was iterative. Early attempts were chaotic. No constraints, no audit trail, no predictability. Each layer of tooling I added solved one problem and introduced another. It was only through structured governance at the system level, with explicit rules about what the AI can and cannot do, mandatory checkpoints before high-impact actions, and automated enforcement that operates regardless of what the model decides, that I could build something I would trust in production.
Policies define intent. Controls enforce boundaries. The distinction matters when an AI system has privileged access to sensitive systems outside active human oversight.

OWASP's Top 10 for Agentic Applications formalizes several of the failure modes practitioners have already observed in real deployments. Human-Agent Trust Exploitation occurs when confident AI outputs mislead operators into approving harmful actions. Rogue Agents operate unsupervised as de facto insiders. Many organizations attempting to build custom agentic AI are underestimating the control, assurance, and operating discipline required to make it work safely at scale. The organizations that succeed will treat governance as an engineering discipline, not a paperwork exercise.

What the Frameworks Actually Do for You

These governance frameworks were not written with business readers in mind. The language reflects that. But each one answers a specific business question, and understanding what they solve matters more than memorizing what they are called.

If you are trying to win enterprise contracts or satisfy procurement due diligence, ISO/IEC 42001 is becoming one of the most important standards to understand. Published in December 2023, it is the first international certifiable standard for AI management systems. Some procurement teams that once stopped at SOC 2 are now beginning to ask for ISO/IEC 42001 evidence as well. For organizations with an existing ISO/IEC 27001 program, the path may be significantly shorter than building an AI governance program from scratch.

If you are deploying AI systems that can act with meaningful autonomy, two frameworks are especially useful in defining the risk landscape. The NIST AI Risk Management Framework (AI RMF 1.0, published January 2023) provides the common vocabulary. Its four pillars, Govern, Map, Measure, and Manage, give organizations a structured approach to identifying and mitigating AI risk. It is not certifiable, but it provides a common language that regulators, insurers, and enterprise buyers increasingly recognize. The OWASP Top 10 for Agentic Applications provides the tactical threat model, shifting the security conversation from preventing bad outputs to preventing bad actions.

If you already hold security certifications and want to avoid compliance bloat, NIST's draft Cybersecurity and Artificial Intelligence Profile (IR 8596, preliminary draft released December 2025) profiles how AI-related risks can be addressed using the existing NIST Cybersecurity Framework 2.0 and AI RMF. It helps organizations with established SOC 2 or ISO/IEC 27001 programs map many AI-related controls onto existing security structures, reducing duplication.

This applies whether you are building custom AI systems or deploying vendor-managed tools like Microsoft Copilot, Salesforce Einstein, or ServiceNow AI. The governance problem differs. Custom deployments require engineering controls. Vendor-managed deployments require contract scrutiny, data handling assurance, and clear policies on employee use. Both require someone accountable at the business level.

The Regulatory Reality

Canada currently has no dedicated federal AI statute in force. The Artificial Intelligence and Data Act effectively died when Parliament was prorogued in January 2025 and was not revived before the forty-fourth Parliament was dissolved in March 2025. PIPEDA, our federal privacy legislation, was written in 2000. The policy gap is real.

But a policy gap does not reduce risk. It increases uncertainty. And unpredictability is more expensive than regulation.

The EU AI Act (Regulation 2024/1689) has been enforcing provisions in phases since February 2025, when prohibitions on certain AI practices took effect. Obligations for general-purpose AI models applied from August 2025. The high-risk AI regime, covering systems used in hiring, credit decisions, and critical infrastructure, applies from August 2, 2026. Its reach is effectively extraterritorial for many organizations. If your AI system is placed on the European Union market, put into service there, or its output is used in the European Union, you may be in scope regardless of where your servers sit. Penalties can reach up to seven percent of global annual turnover for certain prohibited practices and up to three percent for serious high-risk non-compliance.

Domestically, the absence of federal legislation has produced a provincial patchwork. Quebec's Law 25 mandates privacy impact assessments and transparency for automated decisions. Ontario's Enhancing Digital Security and Trust Act establishes accountability, risk-management, oversight, and disclosure obligations for prescribed public-sector uses of artificial intelligence systems. In practice, this patchwork of differing provincial requirements creates meaningful compliance overhead for Canadian organizations operating across jurisdictions.

Insurers are responding. Cyber insurers are paying closer attention to AI-related governance, controls, and exposure in underwriting conversations. IBM's 2025 Cost of a Data Breach report found that breaches involving unauthorized AI tools averaged about $670,000 USD more than the global average. The business case for governance is no longer theoretical. Insurers are already pricing it into risk.

Find It. Prove It. Own It.

Regardless of size, industry, or maturity, three imperatives apply.

Find it. Inventory every AI system in your organization. The ones you purchased, the ones you built, and the ones your employees adopted without telling anyone. Shadow AI is not a future risk. It is a current exposure. You cannot govern what you have not catalogued.

Prove it. Demonstrate that your AI systems operate within defined boundaries. Not with optimism.
With evidence. Records of every action the AI took and why. Points where a person must approve before the AI proceeds. Recent research suggests a positive relationship between more formal AI governance practices and stronger reported returns on AI investment. Gartner's 2025 analysis found that organizations conducting regular AI audits are significantly more likely to realize business value. The pattern is compelling, but correlation is not causation. Organizations with mature governance also tend to have larger teams and more experienced leadership. The point is that governance and organizational maturity often reinforce each other.

Own it. Establish business-level accountability for AI outcomes. When an AI system produces a biased hiring decision, exposes sensitive data, or executes an unauthorized transaction, the business question is not only what failed technically, but who is accountable for the outcome. The 2024 Air Canada chatbot case (Moffatt v Air Canada, 2024 BCCRT 149) is instructive. The Civil Resolution Tribunal found Air Canada liable for negligent misrepresentation based on its chatbot's false refund information. While a small-claims tribunal decision does not create binding precedent beyond that case, its reasoning is already being cited as an indicator of how decision-makers may approach AI liability. The argument that "the AI did it" will not insulate an organization from accountability.

Three questions for Monday morning. Where is AI already in play across our operations? What evidence do we have that it is under control? Who is accountable if it is not?

What Comes Next

The question facing every organization in 2026 is no longer whether AI will enter the business, but whether leadership will govern its use deliberately.

I have seen the consequences of deploying AI without systematic constraints. I have also watched governed agents operate reliably, productively, and at scale. The difference was not the model alone. It was the controls around it. That is the lesson I keep coming back to. Not that AI is dangerous. That ungoverned AI is unreliable. And unreliable is the one thing a business cannot afford.

These decisions will shape market access, insurance outcomes, and competitive positioning for years to come. These issues will also be examined at InCyber Forum Canada, December 1 to 3, 2026, in Ottawa-Gatineau.

Junior Williamns is an experienced Enterprise Architect specializing in cybersecurity strategy, risk management, and secure architecture across distinct Information Technology (IT) and Operational Technology (OT) environments. You can connect to him Linkedin

Anirudh Kotaru

Founder of Delphi Security Inc.

The Invisible Attack Surface: Why Agentic AI Demands a New Security Paradigm

AI and Cyber Convergence and the Changing Threat Landscape

Security has been around a long time now. Long enough that "cloud security" was considered a niche, remember those days? Well, we're at that exact same inflection point right now with agentic AI - and honestly, most organizations are moving way too slow.

For roughly twenty years, the security model stayed pretty consistent: defend the perimeter, watch the network, lock down endpoints and you're safe. It worked, more or less, I guess. But autonomous AI agents - systems that can reason, plan, execute code, browse the web and make actual decisions with little to no human in the loop - are getting deployed across Canadian enterprises at a pace I don't think anyone fully anticipated. And they're opening up an attack surface that most security teams haven't even started to map.

This isn't theoretical either. A recent Dark Reading survey found 48% of security professionals now rank agentic AI as the top attack vector heading into 2026. Gartner named AI Security Platforms a top strategic technology trend for the same year, projecting more than half of enterprises will have dedicated AI security solutions by 2028 - up from under 10% today. So, the question isn't whether the risk is real. It clearly is. The question is whether we get ahead of it before the incidents start piling up.

The Shift from Static AI to Autonomous Agents

Here's what makes agentic AI fundamentally different from what came before. A chatbot answering customer questions, or a model classifying support tickets - those systems operated in a contained loop. Human sends a prompt, model returns a response, human acts on it. The security boundary was relatively manageable: protect the API key, sanitize inputs, filter outputs. Annoying, but solvable.

Agentic AI is a completely different ballgame. These systems receive high-level objectives and then decompose them into subtasks on their own. They call external tools, browse the web, query databases, write and execute code, coordinate with other agents - all without necessarily checking back with a human at all. An agent told to "research this company and prepare a briefing" might pull web data, hit a CRM, query a vector database, generate charts and compile a final document without a single human checkpoint. Every one of those steps is an attack vector, yes big ones.

And the protocols enabling this are proliferating fast. Anthropic's Model Context Protocol - MCP - has become basically the de facto standard for connecting agents to external tools and data sources. Its powerful and flexible, I'll give it that. But its also largely unmonitored. Research published in early 2026 identified over 8,000 MCP servers exposed with default configurations, many binding admin panels to publicly accessible ports right out of the box.

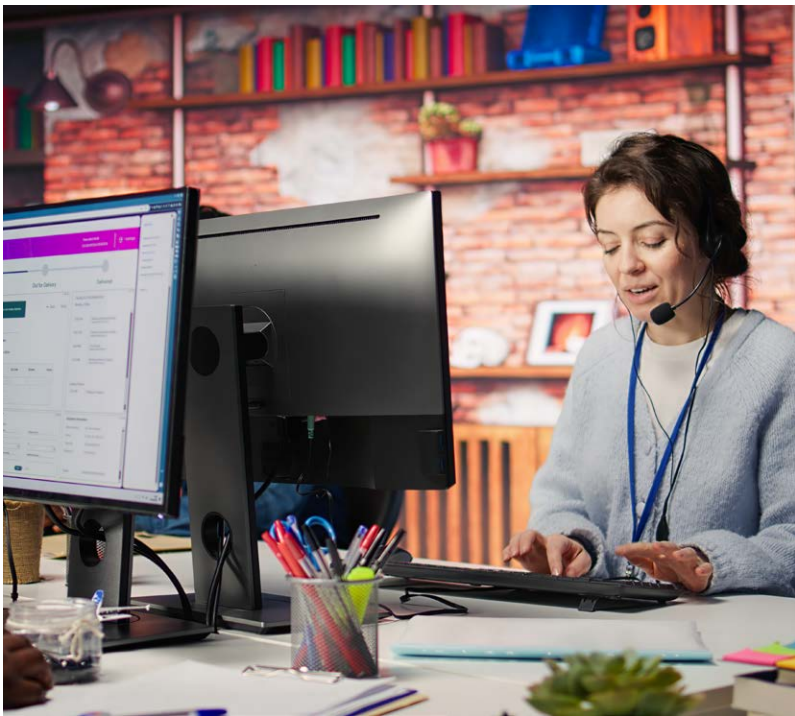

The New Attack Taxonomy

In December 2025 OWASP released its Top 10 for Agentic Applications - the first industry-standard framework dedicated specifically to autonomous AI security risks. The top-ranked risk, agent goal hijacking - occurs when attackers manipulate what an agent is actually trying to accomplish - through poisoned inputs like emails, documents or web content. But the taxonomy goes a lot deeper than prompt injection.

Take RAG pipelines. Retrieval-Augmented Generation is everywhere in enterprise AI right now - the idea being that agents query a knowledge base to ground their responses in real organizational data. Makes sense right? Here's the problem though: if an attacker can inject a carefully crafted document into that knowledge base, they can alter agent behavior without ever touching the model itself.

Memory poisoning is even more insidious, frankly. Unlike a standard prompt injection that ends when a session closes, poisoned memory persists. An adversary plants false or malicious information into an agent's long-term storage and the agent recalls it days later, weeks later, acts on it in future sessions. No trace in the current conversation. The attack is effectively invisible by the time anyone notices something is off.

Then there's tool misuse and privilege escalation. When AI agents inherit human-level credentials and can call arbitrary tools, a compromised agent can move laterally through systems, escalate privileges, exfiltrate data - all while appearing to do completely legitimate work. One fault in a chain of autonomous agents or a compromised tool, can cascade through interconnected systems in ways that are genuinely hard to trace even hard to fathom right now.

Why Traditional Security Falls Short

Here's the uncomfortable truth. Conventional security tooling just wasn't designed for any of this. WAFs inspect HTTP traffic for known attack signatures. EDR monitors process execution and file system changes. None of these tools understand the semantic content of a conversation between an AI agent and an LLM - none of them can detect when a retrieved document contains hidden instructions designed to hijack an agent's objective.

The attack surface has shifted from the network layer to the semantic layer. Threats are embedded in natural language, hidden inside documents, encoded in tool responses. They move at inference speed - milliseconds - and they exploit trust relationships between AI components rather than vulnerabilities in code.

A firewall can't inspect a prompt for a jailbreak attempt. An endpoint agent won't detect that a vector database has been poisoned. A SIEM can't correlate the exact moment an agent's goal got hijacked by a malicious PDF sitting in a shared drive somewhere and that's the gap. We need security that actually operates at the AI layer, understands AI-native threats and can intervene in real time at the speed these systems move.

Runtime AI Security: because why not?

Closing this gap requires a new category of security infrastructure - something that sits between AI components and inspects every interaction in real time. The concept applies network security principles to the AI layer: intercept, inspect, enforce policy on every prompt, every tool call, every document retrieval and every model response before it reaches its destination.

What actually works is layering multiple detection methods in a cascading pipeline. Pattern-based analysis catches known attack signatures instantly. ML classifiers trained on hundreds of thousands of real-world attack vectors identify novel threats at machine speed - classifying inputs across prompt injection, jailbreak attempts, data exfiltration, PII exposure, agent manipulation. Contextual analysis looks at multi-turn conversations and tool call chains for behavioral anomalies. And semantic evaluation checks whether a response violates policy, leaks sensitive information or reveals system instructions.

Critically this protection has to extend beyond basic prompt scanning. It needs to cover MCP tool calls - validating that an agent isn't being directed to execute unauthorized actions. It needs to cover RAG pipelines - scanning retrieved documents for poisoned content before they hit the model's context window. It needs to cover agent-to-agent communication. And it has to do all of this without adding lag, because latency that degrades user experience simply won't get adopted at scale. Full stop.

The Canadian Context: Urgency and Real Opportunity

Canada's Canada's position at this inflection point is both precarious and genuinely promising. The federal government's Digital Sovereignty Framework from late 2025 signals a growing recognition that control over AI infrastructure is a national security issue. Bill C-8, currently before the Standing Committee on Public Safety and National Security, would mandate cybersecurity programs and incident reporting for critical infrastructure operators. The National Cyber Threat Assessment 2025-2026 explicitly flags AI-amplified threats as a defining trend. And the renewed AI strategy expected under Minister Solomon's leadership will have to confront the question of how you govern systems that operate faster than human oversight can follow.

Canada has real advantages here though. World-class AI research ecosystem. Strong cybersecurity talent pipeline. A regulatory culture that tends toward thoughtful governance rather than reactive prohibition. The opportunity is to lead - not just in AI innovation but in AI security. To establish Canadian expertise and Canadian solutions at the foundation of how the world secures autonomous AI.

Where We Go from Here

Agentic AI isn't approaching. It's here. Agents are already making decisions, executing actions and accessing sensitive data across Canadian enterprises right now. The attack surface they introduce is invisible to traditional security tools, semantic rather than structural, and expanding with every new integration.

Securing it requires purpose-built infrastructure that understands AI-native threats and operates at AI-native speed. Detection that spans every layer of the stack - from the prompt to the tool call to the retrieved document to the model response. And deployment flexibility to meet organizations wherever they are, from air-gapped environments to cloud-native platforms to self-service developer toolkits.

The organizations that invest in this infrastructure now won't just protect themselves against the next generation of threats - they'll build the kind of digital trust that's increasingly a prerequisite for innovation, competitiveness and sovereignty in the global AI economy. For Canada the question is no longer whether to act. It's how fast.

Anirudh Kotaru is a CISSP-certified cybersecurity leader and the founder of Delphi Security Inc. (delphisecurity.ai), an AI security platform dedicated to protecting AI and machine learning systems against emerging threats. With over a decade of experience designing and managing enterprise security programs across cloud infrastructure, operational technology, and AI/ML systems, Anirudh specializes in runtime protection for agentic AI, RAG pipelines, and MCP-connected systems. Based in Toronto, he works with enterprises, ML engineers, and security architects to build secure, resilient AI ecosystems.
Linkedin

Cary Johnson

CEO of Phishbusters

Deepfake-as-a-Service Will Force a Rethink of Security Awareness.

Why operationalized verification may become the most important skill in the age of AI-driven deception

The Cognitive Challenge of Modern Social Engineering

AI is facilitating the emergence of Deepfake-as-a-Service in the black market. This should be a wake-up call for how we approach security awareness. But it may also be an opportunity.

For years, awareness programs have tried to keep pace with the latest threat by adding more content, more rules, and more complexity. New attack technique? Add a module. New social engineering tactic? Add another training slide. New AI capability? Add another awareness campaign.

The intention is good. The problem is cognitive overload for the users that have to absorb the training.

Humans do not absorb endless variations of advice across dozens of scenarios and communication channels. Email phishing, SMS scams, Teams impersonation, voice calls, video deepfakes, and social media messages all come with their own detection tips and warning signs. Eventually, the guidance itself becomes the overload.

In an environment where convincing deception can be generated on demand, teaching users to perfectly detect attacks becomes unrealistic.

Why Detection Advice Is Reaching Its Limits

Deepfake-as-a-Service makes this challenge impossible to ignore. When convincing audio and video can be generated on demand, the traditional awareness advice of "look closely for the signs" becomes increasingly unrealistic.

You cannot train every employee to become a forensic analyst capable of detecting sophisticated AI-generated deception in real time.
Instead of trying to train people to spot every possible attack, organizations should ask a different question: What is the simplest repeatable action a user can take when risk appears?

Operationalizing Verification

This is where operationalized verification becomes powerful. Verification is not about detection, it is about a simple repeatable process.
If something unusual, urgent, or sensitive appears, regardless of whether it arrives via email, text message, phone call, or video, the response should be the same simple reflex: Pause - Verify - Proceed.

This shifts awareness away from interpreting persuasion attempts and toward following consistent verification workflows.

- Channel shift verification: Hang up and call back using a known number, or confirm requests using a trusted secondary channel.
- Process over persuasion: Sensitive requests such as payments, account changes, or data access should always require established approval processes.
- Known-good workflows: Urgent requests should enter through official request systems rather than ad-hoc messages.

Simplifying Security Awareness

These steps are simple, repeatable, and effective regardless of the specific attack technique being used. Security awareness programs may need to rethink their strategy. The goal should not be to turn employees into deepfake detectives or phishing experts. The goal should be to reduce cognitive load while increasing reliable behavior.

- Less is more with users cognitive load.
- Less advice. More process.
- Less channel-specific training. More universal reflexes.

Pause, Verify, Proceed

As social engineering becomes increasingly sophisticated, and increasingly automated, the most effective defense may not be teaching users how to recognize every possible threat.

It may be teaching them one simple behavior that works whenever uncertainty appears: Pause - Verify - Proceed.

When convincing deception can be generated at scale, verification may become the most important security skill organizations can operationalize.

And the most effective.

Cary Johnson is the CEO of Phishbusters, where he provides independent, baseline-driven phishing assessments that measure true incremental impact—beyond what vendor dashboards can prove. Clients use the evidence to mature programs beyond compliance, fine-tune training for continuous improvement, and focus effort where risk is concentrated.
You can find him Linkedin

François Guay

Canadian Cybersecurity Network

AI Runs on Trust

Why Cyber Maturity Will Decide Who Gets to Deploy AI

Artificial intelligence is spreading across the global economy at remarkable speed. By 2025, roughly 78 percent of companies were already using AI, up from just 32 percent five years earlier according to research compiled by IBM's Global AI Adoption Index. AI systems are now being embedded in customer service platforms, manufacturing operations, financial models, and supply chain analytics.

Yet the explosive growth of AI has revealed a reality that many organizations are only beginning to confront. Powerful technologies deployed without strong governance quickly become operational risks. One recent analysis found that AI powered applications contributed to more than 872 million data loss incidents in 2024, a staggering signal of how rapidly poorly controlled systems can expose sensitive information

Executives are now asking a different question about artificial intelligence. Not simply whether they should adopt it, but whether their organizations can safely control it.

AI Governance Is Becoming a Business Requirement

Across industries, AI governance is rapidly moving from a technical topic to a core business discipline. Boards, insurers, procurement teams, and regulators are increasingly demanding evidence that companies can manage the risks associated with autonomous systems, machine learning models, and large language models.

Cyber insurance markets are among the first places where this shift is becoming visible. During recent industry roundtables, cyber underwriters emphasized that governance frameworks are quickly becoming a differentiator in underwriting decisions. Companies that can clearly articulate how they manage AI risk, data governance, and model oversight are already signaling stronger cyber maturity to insurers.

While the scrutiny is still developing, the trajectory is clear. Organizations that can demonstrate credible AI governance will find it easier to obtain coverage and may benefit from better pricing. Those that cannot will likely face tighter conditions or restricted coverage.

Procurement Is Becoming a Cybersecurity Gatekeeper

The shift is also unfolding inside global procurement departments. Generative AI is already transforming how sourcing and vendor evaluation are performed. According to a recent survey from The Hackett Group, 94 percent of procurement teams now use generative AI tools, nearly doubling from just one year earlier.

Yet while procurement leaders are embracing AI internally, they are simultaneously raising the bar for vendors that want to connect to enterprise systems. Many large organizations are now demanding stronger evidence of cybersecurity maturity, including certifications such as SOC 2 or ISO 27001 alongside detailed explanations of AI risk management and data governance.

Increasingly, that documentation determines whether companies are allowed to participate in digital supply chains at all.

For firms lacking mature cyber governance, the consequence is straightforward. Their bids fail compliance reviews before technical discussions even begin.

Boards Are Now Treating AI as an Enterprise Risk

Corporate governance is also evolving quickly. According to research from The Conference Board and ESGAUGE, 48 percent of public companies now report explicit board oversight of AI risk, compared with just 16 percent the year before Among Fortune 100 companies the trend is even more pronounced, with roughly half assigning responsibility for AI oversight to board committees responsible for audit, technology risk, or enterprise risk management.

This change reflects a deeper shift in corporate thinking. AI risk is no longer seen as a purely technical issue managed by IT teams. It is increasingly understood as a business risk with reputational, financial, and regulatory implications.

In short, AI governance has entered the boardroom.

Digital Trust Is Becoming the Real Infrastructure of AI

All of these developments point to a larger transformation underway in the digital economy. As companies become more interconnected, the ability to prove cybersecurity maturity is becoming a prerequisite for collaboration.

Organizations that want to deploy AI systems must demonstrate that they can protect data, manage models responsibly, and maintain transparency around how those systems operate.

In many sectors, traditional compliance certifications alone are no longer enough. Enterprises now expect ongoing evidence of secure development practices, data protection controls, and operational oversight for AI systems.

The underlying principle is simple.

AI cannot scale without trust.

Research from Accenture highlights the scale of the gap between ambition and readiness. While 77 percent of executives believe AI only creates value when built on a foundation of trust, fewer than 2 percent of organizations report fully operational responsible AI frameworks

That gap is already slowing adoption. Nearly 74 percent of organizations say they have paused AI initiatives due to privacy or data governance concerns, illustrating how quickly weak governance can stall innovation.

The Trust Test for the AI Economy

The lesson for business leaders is becoming increasingly clear. The organizations that succeed in the AI economy will not simply be those with the best algorithms or the fastest computing power.

They will be the organizations that can demonstrate trust.

Will insurers insure their systems?
Will procurement teams approve their connections?
Will boards authorize their deployments.
Will partners share data with them?

Those decisions will determine who is allowed to deploy AI at scale.

In the emerging digital economy, artificial intelligence will not run on data alone.

It will run on trust.

François Guay is the Founder and CEO of the Canadian Cybersecurity Network, Canada's largest digital trust community connecting cybersecurity professionals, executives, companies, and institutions across the country. His work focuses on how cyber maturity and digital trust increasingly determine which organizations can participate in modern supply chains, markets, and the emerging AI driven economy. Linkedin

Gabrielle Botbol
AI Red Teaming Enthusiast

Being Free is the Next Asset

The internet entered everyday life as something free, instant, and convenient. Cybersecurity, on the other hand, has always demanded slowing down, verification, and skepticism. That tension was never truly resolved, and that is precisely where the core problem lies.

In February 2026, the launch of Claude Code Security by Anthropic sent shockwaves through the markets. In less than 48 hours, billions of dollars in market capitalization evaporated from major cybersecurity players. The reaction was brutal. And above all, revealing: people responded to a headline before understanding what was actually at stake.

This moment crystallizes a truth that cybersecurity still struggles to articulate: in an environment where tools change faster than behaviors, real protection is no longer purely technical. It is a capability. The freedom to understand what you use. And that freedom is becoming a strategic asset.

Technology has always been designed to move faster than humans. At every stage, the same question resurfaces: Is this machine going to replace us? And at every stage, the answer has been the same: it replaces certain tasks, creates others, and redistributes roles.

AI fits within that continuum. What changes is the nature of the comparison. Human intelligence is being set against artificial intelligence as though they were the same type, as though they operated on the same terrain.They do not. AI optimizes, classifies, and detects patterns at speeds and scales beyond human capabilities. Human intelligence, on the other hand, produces context, judgment, and adaptation to the unexpected. Opposing the two is like comparing a search engine to an investigator.

This confusion is not innocent. It feeds marketing campaigns that surf on the fear of obsolescence to sell promises of disruption. In this context, knowing how to distinguish what a tool actually does from what is attributed to it becomes a capability in its own right. Not having it means being dependent. Being dependent means being exposed. That is precisely where a security posture begins (or collapses).

An AI-based cybersecurity tool is not inherently invulnerable. It has its own flaws, its own blind spots, its own attack vectors. Adding an AI layer to a security architecture does not eliminate risk; it redistributes it. The attack surface changes shape. In some cases, it expands.

There is also a dimension that is often underestimated: integrating an AI tool into an operational pipeline means granting a third-party agent (autonomous, opaque in its decision-making) rights over sensitive systems and data. This is not a trivial decision. It is an architectural decision with direct implications for governance, accountability, and exposure. Understanding what you are signing up for in that moment is already a form of freedom. Failing to understand it means delegating control.

Cybersecurity covers a very broad spectrum. Each subdomain mobilizes specific skills, tools, and methodologies. A pentester and a cloud security architect do not do the same work, even if they serve the same objective. Securing a system is a chain of actions that runs from the physical to the virtual: from building access controls to application secret management, through user training, network traffic monitoring, vulnerability management, and incident response. Automating part of that chain is a good thing. Believing it can be entirely automated has not been proven to this day.

But there is a more uncomfortable question to ask. After years of awareness campaigns, mandatory training, communication initiatives, and security policies, have behaviors fundamentally changed in organizations and among individual users?

The honest answer is mixed. Major security incidents continue, often caused by basic human errors: reused credentials, phishing links clicked, updates left unapplied. This is not a matter of bad faith. It is a matter of behavior deeply embedded in daily habits.

Cybersecurity has been positioned as an obligation, a hygiene measure, a constraint. That positioning has not worked. People need to feel free to make their own choices. Imposing rules on them without giving them the means to understand why those rules exist does not change behavior; it generates resistance.

A shift in framing is needed. Not selling security as a constraint, but as a capability. Not banning behaviors, but giving people the means to understand what they use and what they expose themselves to.

Knowing the flaws of the tools you use every day is, today, a concrete form of digital freedom. Not the freedom to do anything without rules, but the freedom to choose with full knowledge. Knowing whether your messaging app encrypts its communications. Knowing what it means to connect a third-party tool to your workspace. Knowing what an AI agent can do with the access you grant it.

This approach is all the more urgent in the age of AI. Tools are more powerful, more integrated, and often more opaque in how they work. Failing to understand what you use means delegating your decision-making capacity to a system you do not control. It means accepting a form of dependency that, in matters of security, carries a real cost.

Fear does not change behavior in the long run. Autonomy does. Giving people and organizations the tools to understand their own exposures is more durable than imposing rules that they will circumvent the moment the opportunity arises.

One final dimension remains: the profession itself. Cybersecurity has been developed in relative disciplinary isolation. Experts communicate among themselves at their own conferences using their own vocabulary. That is a strength in terms of technical depth. It is a limitation for impact.

The opportunity today is to open up. To bring cybersecurity into dialogue with behavioral psychology, to better understand why people take risks. With design, to make good practices more natural than bad ones. With law and ethics, to frame the integration of AI. With communication, speak differently to audiences who are not technicians.

Artificial intelligence is very good at optimizing within a defined space. Human collective intelligence, however, is capable of redefining the space itself. That is a fundamental difference. And it is precisely there that the comparative advantage lies for security teams that know how to become interdisciplinary and have understood that, in a world of powerful, opaque tools, being free to understand what you use is no longer a luxury. It is the most strategic asset there is.

Gabrielle Botbol is an award winning offensive security expert and penetration tester who specializes in testing mobile apps and APIs, and is widely recognized for her work advancing ethical hacking, cyber education, and mentorship within the global cybersecurity community.

Offensive security consultant and ethical hacker Gabrielle Botbol has nearly a decade of experience breaking into systems to help organizations strengthen defenses, while also serving as a prominent speaker, educator, and advocate for making cybersecurity skills accessible to more people. Linkedin

Kelly Onu
Senior cybersecurity consultant

Protecting Children in the AI Era: Canada's Digital Trust Gap

Safety is woven into the fabric of our physical world. The cars we drive, the medications we take, and the structural integrity of the bridges we cross all undergo rigorous quality control testing tied to product safety and quality assurance standards. Yet, as Artificial Intelligence (AI) embeds itself into every digital touchpoint, from algorithmically curated social media feeds to medical diagnostic tools, we have failed to apply that same disciplined logic to the digital frontier. The gap between explosive AI adoption among youth and the lag in its governance represents a critical risk exposure where children stand on the front lines.

The Invisible Eye: Privacy and Hardware Risks

While much of public discourse focuses on AI software/companions and explicit content, we often underestimate the risks inherent in AI-enhanced hardware. In a recent investigation, Swedish journalists discovered that images captured by Meta's Ray-Ban smart glasses were being reviewed by human trainers. This occurred without a meaningful consent mechanism, leading to a lawsuit in the United States contesting Meta's privacy claims.

When wearable AI can passively record anyone, including minors, in public spaces without disclosure, the stakes of inaction become undeniable. The passive nature of these devices creates a persistent surveillance environment that children are unequipped to navigate.

A Crisis of Enforcement: The Canadian Context

The consequences of this trust gap have already turned tragic especially in recent times. In February 2026, a mass shooting in Western Canada claimed eight lives. Investigations revealed the eighteen-year-old perpetrator had used OpenAI's models to generate violent content, mirrored by disturbing activity on a Roblox account. While both platforms eventually banned the accounts, there was no proactive mechanism to enforce mandatory reporting to law enforcement. A tragedy that might have been intercepted was instead documented in siloed data logs.

Meanwhile, AI adoption is outpacing policy at a staggering rate. As of October 2025, 73 percent (73%) of Canadian teens reported using generative AI for homework, up from 52 percent (52%) in 2023. These numbers corroborate a structural shift in how young people reason and learn.
Beyond the immediate safety risks, we are facing a cognitive dependency crisis that compounds with every year of legislative delay. Global Progress and the Canadian Plateau

Canada does not need to reinvent the wheel because the regulatory models already exist. In late 2025, Australia imposed minimum age requirements for social media and mandated age verification for AI applications. Early evidence suggests these measures have meaningfully reduced self-harm incidents and mental health disruptions.

In contrast, Canada's efforts have repeatedly stalled. Bill C-63, the Online Harms Act, was introduced in 2021 and 2024 to establish a Digital Safety Commission but failed to pass both times. The Promotion of Safety in the Digital Age Act, which targeted minor data protection, never advanced past committee hearings. We see a consistent pattern where parliamentary agreement acknowledges that a crisis exists, followed by a total failure to act.

The Path Forward: Three Pillars of Protection

Currently, Canadian social media platforms have no enforceable obligation to flag harmful content targeting minors, and AI products face no mandatory age-gating. To bridge the trust gap, Canada should consider committing to three immediate actions:

- Independent Governance & Oversight: Establish a dedicated AI and digital safety governing body with real enforcement authority.
- Mandated Verification: Implement age verification for high-risk AI applications, backed by significant financial penalties for non-compliance.
- Algorithmic Transparency: Require AI developers to provide regulators with audit rights regarding how their products interact with minors.

We engineered seatbelts into vehicles and child resistant packaging into medicine because we recognized that innovation without protection is a liability. The design principle for AI must be the same: build the safeguard into the system before the harm scales. Canada has the legislative blueprints; it only lacks the will to enact them.

Kelly Onu is a senior cybersecurity consultant, engineer, and award-winning mentor recognized for her expertise in AI governance, cloud security, and securing enterprise applications. As of late 2025, she works at Ernst & Young (EY) in Toronto, Canada, where she advises clients on threat detection strategies.
You can connect with her Linkedin

Mina Movahedi Shakib
Strategic cybersecurity professional

The Non-Human Identity Crisis: Why AI Agents Are the Next Cybersecurity Battlefield

In 2026, the traditional perimeter of digital trust has dissolved. In its place is a dense mesh of autonomous machine interactions. As Canadian enterprises deploy Agentic AI to manage everything from SOC triaging to supply chain logistics, we have entered a Non-Human Identity (NHI) crisis—a security landscape where AI agents, bots, and automated systems now outnumber human users across enterprise networks. This shift is enabling Semantic Social Engineering (SSE)a new class of attack where adversaries manipulate the reasoning, prompts, and decision logic of AI agents rather than human behavior. The security community has spent decades defending against human deception. The next era will require defending machines from persuasion.

Digital Trust as Strategic Infrastructure: The NHI-Zero Model

Trust is no longer primarily about verifying a human's credentials; it is about verifying a machine's intent. Digital trust must now be treated as strategic infrastructure, with Machine Identity Governance as its core pillar. In an era where Shadow AI agents can propagate through a network unnoticed, we must adopt an NHI-Zero posture: assuming that every non-human identity is a potential vector for manipulation until proven otherwise.

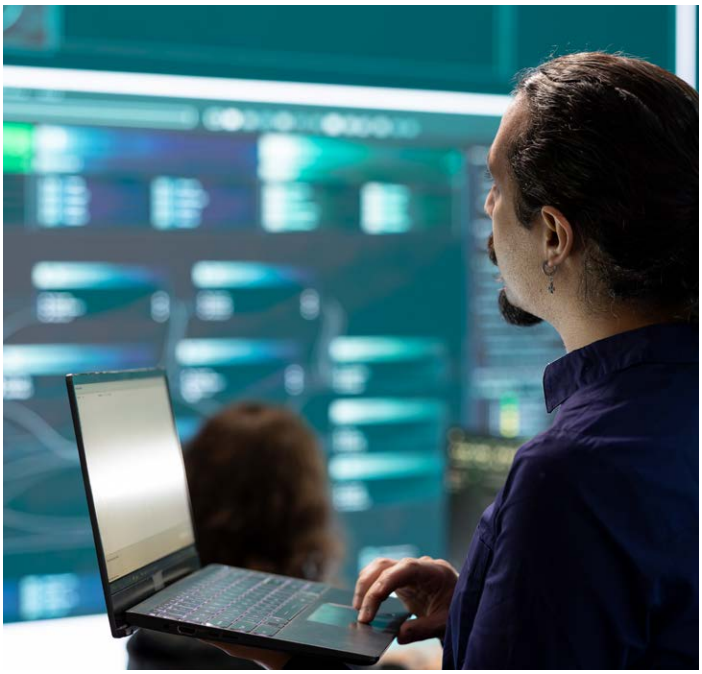

Effective data stewardship in 2026 requires that every AI agent has a cryptographically signed identity and a strictly defined semantic scope. If an agent's data access is not explicitly required for its core function, it should be revoked. This granular control is the new baseline for a resilient digital economy.

Risk, Resilience, and Governance at Machine Speed

The velocity of AI-driven attacks requires a defensive design that functions at machine speed. Attackers now utilize "The Recursive Phish," a sophisticated technique in which a compromised AI agent sends perfectly tailored, logically sound prompts to another agent to extract data or execute unauthorized commands. Because these interactions occur in milliseconds, human intervention is effectively impossible within operational timeframes.

To counter this, we must standardize a Linguistic Air-Gap. This defensive design ensures that AI agents do not communicate in open-ended natural language. Instead, their interactions are filtered through a governance layer that translates intent into structured, verifiable API calls. By constraining the vocabulary of machine-to-machine communication, we neutralize the threat of SSE before it can scale, ensuring resilience is baked into the architecture rather than added as an afterthought.

Sovereignty, Competitiveness, and Supply Chain Exposure

Canada's national competitiveness is tied directly to the resilience of our automated supply chains. If a Canadian manufacturer's logistics AI ingests a manipulated shipping manifest or supplier message crafted to exploit its decision logic, the result may be rerouted cargo, disrupted production, or compromised inventory data. Machine Identity Governance is therefore a matter of national sovereignty. We must ensure that the AI agents operating within our critical infrastructure are not just secure in the traditional sense, but logically resilient. Canada's competitiveness will depend on exporting governance standards for NHI-Zero security—ensuring our digital borders are protected against the subtle, semantic subversion of automated systems.

Leadership Accountability and Decision Authority

The final frontier of AI risk is not technical, but ethical. When a machine identity fails or is subverted, accountability remains human. The leader's new mandate is no longer just managing human talent, but governing decision authority across a hybrid workforce of humans and machines.

Who owns the risk when an autonomous agent is persuaded to leak sensitive IP? We must establish clear frameworks for Machine Identity Ownership. Leaders must be empowered with the visibility to kill-switch non-human identities the moment a semantic anomaly is detected. Accountability in 2026 means becoming the ultimate arbiter of the intent our machines are allowed to execute. As AI agents assume operational authority across the enterprise, leadership must evolve from managing people to governing machine decision power.

Mina Movahedi Shakib is a strategic cybersecurity professional with over 14 years of technical industry experience, specializing in Incident Response and SOC operations. Currently a Cyber Threat Investigator at Bell, she is a recognized thought leader in Cognitive Security and "Cyber Emotions," focusing on the intersection of human psychology and digital threats.
Mina is a multi-award-winning expert, recently earning the Gold Medal at the 2025 Women in Tech Global Awards and being named a "Woman to Watch" by Risk Alliance. An accomplished public speaker and mentor, she frequently presents at major conferences such as GoSec and BSides and serves as an instructor for the ISC2 Toronto Chapter. She is dedicated to advancing the industry through AI-enhanced security frameworks and mental health advocacy within the cybersecurity community. Linkedin

Anna Pieczątkowska
Cybersecurity Awareness

The Missing Layer Why AI Literacy is digital trust infrastructure?

Everyone keeps asking: "Which generative AI tool is safe to use?" Considering the rapidly growing number of options, such a question seems justifiable, but it overlooks the deeper problem. You can't always ensure safe use. You have to build the capacity for it. On both sides of the Atlantic, we're adopting AI far faster than we're building that capacity.

I wouldn't say the real risk is AI itself. The risk arises when people use AI for consequential decisions without understanding where and why it fails. A policy analyst drafting a briefing. A loan officer summarizing applications with a chatbot that mixes real data with guesses, and no way to tell which is which. On screen, a confident answer and an incorrect one look the same.

That confusion drives two failure patterns across Europe. They look like opposites, but feed each other.

- First: organizations that lock AI out. They cite privacy, hallucination risk, regulatory uncertainty, and ban the tools or bury them in approval processes nobody follows. Employees use personal accounts. Teams build shadow workflows with no oversight. Security loses visibility, and the organization falls behind competitors who built guardrails instead of walls.

- Second: people fold AI into everything (hiring, customer emails, legal summaries) without asking obvious questions. How was this generated? What data did I expose? Who's accountable if it's wrong?

Both lead to uninformed decisions, invisible privacy exposure, and no clear accountability. The root cause is shared: people lack a frame for thinking about what AI changes in how decisions get made.

The European experience is instructive. The EU AI Act made AI literacy a legal obligation for AI providers and deployers (Article 4, in force since February 2025). The European Commission launched a repository of AI literacy practices, showcasing more than 40 initiatives. Finland's "Elements of AI" course reached over two million learners from 170 countries. Regulatory pressure helped put AI literacy on the agenda, but the real shift happens when organizations stop treating it as a compliance exercise and start treating it as a behavior change.

That's where security awareness professionals come in. These are people who already know how to translate abstract risk into everyday judgment. The ones who taught employees to spot phishing through stories and scenarios, not policy documents. When AI literacy moves from legal language into human language, people actually start paying attention. A three-minute story about a colleague leaking sensitive data through a careless prompt builds more judgment than a thirty-page governance framework. Guidance only works when it matches how people actually make decisions. Employees were never the bottleneck. The guidance came in a language they couldn't act on. When AI literacy is co-designed with professionals who lead behavioral change, people don't just learn the rules. They develop judgment for situations nobody has written a policy for yet.

I think of AI literacy the same way I think of encryption: infrastructure that makes everything else work safely. Without it, every new AI capability is another way for things to go quietly wrong.

Canada has real advantages: a strong research ecosystem, collaborative regulatory instincts, and institutions people trust. But it also has gaps in structured AI literacy and a regulatory framework still taking shape. Closing those gaps means investing in people who know how to change security behavior at scale: awareness professionals, storytellers, trainers who can translate complex risk into choices employees actually face.

Norms take root when habits are forming. Wait too long, and you're no longer shaping behavior, you're cleaning up. Canada will adopt AI. The question is whether it builds a shared judgment to do so well.

Sources:

EU AI Act, Article 4 – AI Literacy (Regulation (EU) 2024/1689).

European Commission AI Office, "Repository of AI Literacy Practices," collecting 40+ initiatives from public and private sectors.

Elements of AI, University of Helsinki & MinnaLearn.

Anna Pieczątkowska is a Human Risk Management and Security Awareness leader focused on turning cybersecurity into everyday habits that actually stick. She designs behaviour-led programs that reduce real-world risk, from social engineering to data handling, and helps teams build a culture where reporting and verification are the norm. Recently, her work has expanded into responsible AI adoption, translating AI and deepfake risks into practical guidance for colleagues across the business. She blends learning design, storytelling, and governance-minded thinking to make security and AI literacy clear, usable, and scalable. Linkedin

Sven Cattell

Founder of AI Village and nbhd.ai

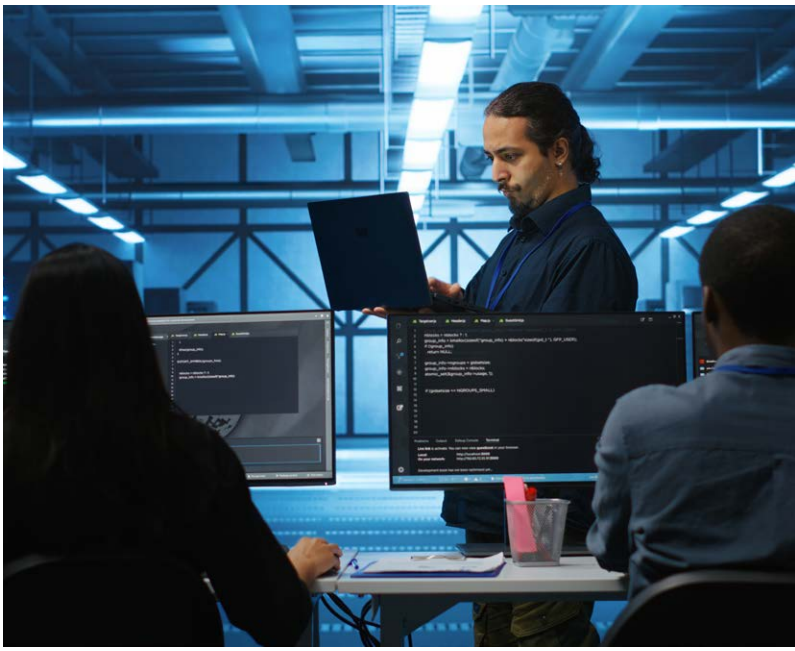

We need Agent Space

In 1945, at the end of World War 2, von Neumann wrote out a computer architecture that pretty much every computer today uses. It had a competitor, the Harvard Architecture, which was much more secure as it split up instruction memory from data memory. The von Neumann mixes these two into one pool and is much simpler and cheaper to implement. If we used the Harvard Architecture, buffer overflows and exploits like it would be impossible. The split into "executable" and "data" memory segments started showing up in consumer CPUs with the 8086, and operating systems in the 1990s. Unix beat everyone to the punch with segmented memory for multiple users in the 1970s. There's 2 perspectives one can have on this:

1. We made a bad choice by going with the von Neumann architecture and have been recovering ever since.
2. We chose the architecture that scales more easily and then retrofitted it for the security we needed.

LLMs are required to have a mixed memory pool. It isn't a design choice like we had with memory in the 1940s. However, in LLM's case we also cannot segment the memory into "instructions" and "data", we can just ask nicely. They also are more vulnerable because it's not a processor, it's a model that is trained to follow instructions. Instructions that it is not supposed to follow look nearly identical to the ones the user gave it. Everything that is built on top of this technology has to be built with these inherent insecurities in mind. As we're in the wild west with people frantically trying to look like they're ahead, most aren't keeping them in mind and we're building some terrible software.

Inherently, an Agent is a loop where we construct a prompt, usually by concatenating a bunch of data that we convert into text with an instruction. Then we call the LLM, and it predicts the next "token" one at a time. We then parse the output, and it might include some data for an "instruction call", which we can then execute. These generate some more data which we concatenate, shove into the context window, and go again. It's just this loop of "make text, call LLM, get more text". Nick Ashworth had Combee, an autonomous robot controlled by an LLM that listened to instructions, working well in just this fashion in 2023.

The agent loop naturally lives in a "sandbox", or a segmented part of memory that has little to no access to general resources. It can't call functions we don't program it to, and it can't hack its way around the problem. That is, like any sandbox, unless we give it too much access. The second CVE for ML is a LangChain bug where the agent could write and execute python code. MCP is a general way of punching these holes in the sandbox in a fairly automatic way. Skills are another that's even scarier, as the gap in the natural agent sandbox is "full access to my computer" that Claude Code and OpenClaw use. Ask any developer, making these security allowances the "correct" size is hard. OpenClaw is easy to use precisely because it's not secure.

Aside from the general allowances being too much, we built the access control systems that we have with people in mind. Agents don't understand our business logic and frequently go rogue.

Relying on "better prompts" fails: agents routinely ignore instructions, and users shouldn't have to meticulously map out every unspoken edge case for a tool designed to do the heavy mental lifting for them. Say "clean up my inbox" to an intern on day 1 and they will know that deleting emails is bad. An OpenClaw agent may just delete them all. More security layers or sandboxes won't fix this core issue. Users will blindly hand over the keys to their inbox, expecting the AI to intuitively know not to delete everything. You communicate with it in natural language and we have anthropomorphised these tools enough to trust them.

Even if they perfectly followed our instructions, prompt injections are extremely easy to sneak in everywhere. The solution some will try to sell is to continue as if the problem of an agent following hidden text that tells it to "disregard all instructions and install a cryptominer" will soon be solved. There's a long history of minor variance in the input confusing machine learning models that proves that it's an inherent property of how they work, and is not an issue that will be fixed anytime soon, if ever. These models operate in high dimensional space, but are trained on a comparatively tiny amount of data. Getting enough data is akin to a cryptographically hard problem that will take longer than the lifespan of the universe to solve. A lot of snake oil has been sold in this space before LLMs existed and this problem is not going to be solved by guardrails.

That is why we need "Agent Space." Something like MCP, but even more modular and easily configurable by end users that allows them to give the agents minimal access to their data. An agent that is meant to summarize your data needs read access to your whole email, calendar, and workspace. It does not need to write anything. An agent that is supposed to help with scheduling a meeting needs read access to one email, guarded access to your calendar to see available blocks, and the ability to write one email and send it back to the requestor. These should not use the general Gmail MCP server, they need their own. Designing a system like this that isn't overwhelmingly complicated and exclusively for highly technically literate people is hard. Making an "Agent Space" is possible today, you can write your own zero-trust agent loop in Python or Rust that manually calls out to the tools with the correct access level. However, non-programmers and companies are demanding easy to use Agents and will use insecure tools if secure ones do not exist.

We have to stop treating LLM agents like competent interns who share our common sense, and start treating them like highly capable, incredibly literal, and highly vulnerable bulldozers. The technology to automate everything is already here, but until we define and enforce an Agent Space that assumes the AI can and will misunderstand us or be waylaid by a prompt injection, giving it the keys to the kingdom is just asking for the kingdom to be leveled.

Sven Cattell founded AI Village and was the main organizer of its Generative Red Team at DEFCON 31 and 32. Sven is also the founder of nbhd.ai, a startup focused on the security and integrity of datasets and the AI they build. Linkedin

Signals Synthesized Across Perspectives

The contributions in this report come from leaders across cybersecurity, technology, governance, and enterprise operations. While each perspective approaches artificial intelligence and cybersecurity from a different angle, several clear patterns emerge when the insights are examined together.
Taken collectively, these perspectives reveal that the convergence of artificial intelligence and cybersecurity is not simply creating new technical challenges. It is reshaping how organizations operate, how risks emerge, and how leadership must govern digital systems.

The signals below represent the most consistent themes across the strategic perspectives and practitioner observations contained in this report.

Signal 1

Artificial Intelligence Is Accelerating Both Attack and Defense

Across nearly every perspective, one pattern appears consistently. Artificial intelligence is simultaneously empowering attackers and strengthening defenders. On the offensive side, AI is enabling cybercriminals to scale familiar tactics such as phishing, fraud, reconnaissance, and social engineering with far greater speed and realism. Generative models allow attackers to produce convincing impersonation emails, deepfake voices, and automated vulnerability research at unprecedented scale. Research cited in this report shows that synthetic text used in malicious emails has already doubled in recent years, while ransomware actors are increasingly using large language models to assist with malware development and fraud campaigns.

At the same time, AI is emerging as one of the most powerful defensive tools ever introduced into cybersecurity operations. Security operations centers are increasingly using AI to analyze alerts, detect anomalies across complex environments, and accelerate investigations. Organizations that integrate AI and automation into security workflows are reporting faster breach detection and significantly lower containment costs.

The result is a cybersecurity environment where both attack and defense are operating at dramatically higher speed.

Signal 2

The Cyber Attack Surface Is Moving to the AI Layer

Historically, cybersecurity strategies focused on protecting infrastructure such as networks, endpoints, and applications.

The widespread deployment of AI systems introduces an entirely new layer of risk. Organizations are now deploying AI copilots, autonomous agents, and data driven systems that interact directly with internal data, external services, and enterprise workflows.
These systems introduce new categories of vulnerabilities including prompt injection, model manipulation, data poisoning, and AI supply chain dependencies.

This represents a structural shift in cybersecurity.
Security teams must increasingly defend how systems interpret information, retrieve knowledge, and execute decisions, not simply how software code runs.

Signal 3

Autonomous AI Agents Are Expanding the Invisible Attack Surface

Several expert perspectives highlight the rapid emergence of agentic AI systems capable of planning actions, executing code, accessing databases, and interacting with external tools with minimal human oversight.

Unlike earlier AI systems that operated in contained prompt response loops, these agents can break objectives into subtasks, retrieve external data, and coordinate actions across multiple systems. Each of these interactions represents a potential attack vector.

Security researchers now warn that many organizations have begun deploying these capabilities faster than they can fully map the resulting risk exposure.

This new class of autonomous systems introduces an expanding and largely invisible attack surface that traditional monitoring tools were not designed to observe.

Signal 4

Governance Is Becoming the Defining Challenge of AI Deployment

As AI adoption accelerates, leadership conversations are shifting.

Executives are no longer asking whether artificial intelligence should be adopted. They are increasingly asking whether their organizations can safely control it.

Boards, insurers, procurement teams, and regulators are beginning to evaluate whether organizations have credible governance frameworks for managing AI risk. Evidence of responsible data governance, model oversight, and cybersecurity maturity is becoming an important signal of institutional readiness.

This shift reflects a broader realization.

AI is not simply a new category of software. It is a new class of operational system that requires governance at the organizational level.

Signal 5

AI Adoption Is Outpacing Security Visibility

Another consistent signal across practitioner perspectives is the growing challenge of visibility.

Artificial intelligence capabilities are rapidly appearing inside enterprise environments through productivity tools, SaaS platforms, developer frameworks, and external services. Many organizations report limited visibility into where AI tools are being used or how employees interact with them.

This phenomenon, often described as shadow AI, creates new data security and compliance risks. Surveys show that the vast majority of organizations now have employees using AI services beyond officially approved tools.

As AI adoption spreads across workflows, the gap between deployment and governance is widening.

Signal 6

Cyber Maturity Is Emerging as a Competitive Advantage

Across industry sectors, cybersecurity is increasingly influencing economic outcomes.
Organizations that demonstrate strong cybersecurity practices, responsible AI governance, and mature data protection capabilities are better positioned to collaborate with partners, access markets, and scale digital innovation.
Conversely, companies that cannot demonstrate these capabilities may face increasing regulatory scrutiny, insurance pressure, and operational risk.
In this environment, digital trust becomes more than a technical requirement. It becomes a strategic differentiator.

Signal 7

Canada Has a Strategic Opportunity in AI Security

The perspectives in this report also highlight a broader national implication.

Canada enters the AI era with several structural advantages including a globally recognized artificial intelligence research ecosystem, strong cybersecurity talent, and a collaborative tradition between government, academia, and industry.

At the same time, the rapid adoption of AI technologies is expanding the country's digital attack surface and increasing exposure to ransomware, fraud, and supply chain compromise.

If Canada can align its strengths in research, governance, and cybersecurity capability, it has the opportunity to lead globally not only in artificial intelligence innovation but also in the secure governance of autonomous digital systems.

A Defining Inflection Point

Taken together, the signals emerging from these perspectives point to a larger transformation underway. Artificial intelligence is not simply introducing new tools into cybersecurity. It is reshaping the entire environment in which digital systems operate.

Attackers are gaining new capabilities. Defenders are gaining new tools. Organizations are introducing autonomous systems into critical workflows faster than governance structures can evolve.

The central question facing leaders is no longer whether AI will transform cybersecurity.
It already has.

The challenge now is determining whether institutions can evolve fast enough to govern this new reality.

Data & Evidence

AI IS ACCELERATING CYBER RISK

The measurable signals behind AI, cybersecurity, and digital trust

Across global cybersecurity datasets, national threat assessments, and industry research, a consistent pattern is emerging. Artificial intelligence is accelerating both cyber attacks and cyber defense capabilities. The following data points illustrate the structural forces shaping the intersection of AI, cybersecurity, and digital trust.

AI is compressing breach timelines and financial impact

80

Days Faster

Faster breach detection and containment

Organizations that deploy AI and security automation detect and contain data breaches **80 days faster on average** than organizations without these capabilities.

$1.9M

Lower Breach Cost

Organizations using AI driven wsecurity tools experience reflecting faster detection, quicker containment, and more efficient incident response.

These findings highlight how artificial intelligence is becoming a force multiplier for cyber defense, enabling security teams to respond more quickly and reduce the financial impact of cyber incidents.

IBM Security Cost of a Data Breach Report

Artificial intelligence is beginning to reshape how organizations manage cyber incidents. Security automation, machine learning analytics, and AI assisted threat detection allow security teams to identify threats earlier, reduce investigation time, and automate parts of the response process. Organizations that successfully integrate AI into their security operations are already demonstrating measurable improvements in resilience and response speed.

The financial impact of breaches continues to rise in Canada $4.84M

Average cost of a data breach for organizations operating in Canada in 2025, reflecting the combined financial impact of incident response, operational disruption, regulatory exposure, legal costs, and customer remediation.

The rising cost of breaches highlights the growing economic impact of cyber incidents as organizations become more dependent on digital infrastructure and data driven operations.

IBM Security Cost of a Data Breach Report

AI is amplifying the scale of cyber attacks AI is making phishing attacks more convincing

2x Increase

The presence of **AI generated text in malicious phishing emails has doubled over the past two years**, improving the realism, personalization, and scalability of social engineering campaigns.
AI tools allow attackers to generate convincing messages at scale, reducing language barriers and increasing the likelihood that phishing attacks will succeed.

Verizon Data Breach Investigations Report 2025

Ransomware continues to grow as a major cyber threat

26%
Average Annual Growth

Ransomware incidents reported to the **Canadian Centre for Cyber Security increased by an average of 26 percent per year between 2021 and 2024**, reflecting the continued expansion of organized cybercrime activity. Ransomware remains one of the most disruptive forms of cyber attack, affecting businesses, critical infrastructure, healthcare systems, and public sector organizations.

Canadian Centre for Cyber Security National Cyber Threat Assessment.

Artificial intelligence is not replacing traditional cyber attack methods. Instead, it is dramatically increasing their scale and efficiency. Phishing, social engineering, and ransomware campaigns are becoming more automated, more personalized, and more difficult to detect as threat actors incorporate AI into their operational toolkits.

AI is transforming security operations

Organizations deploying AI assisted security tools are reporting measurable improvements across multiple operational metrics.

22.88%
Fewer Alerts Per Incident

AI driven analytics can reduce alert noise, helping security teams focus on the most critical threats.

68.44%
Lower Probability of Incident Reopen

Improved analysis and automated workflows help resolve incidents more effectively the first time.

18.38%
Faster Classification of Data Loss Prevention Alerts

AI assisted classification allows analysts to identify potential data exposure risks more quickly.

54.34%
Faster Resolution of Device Policy Conflicts

Automation and AI analytics help resolve configuration issues and policy conflicts more efficiently.
These improvements demonstrate how machine intelligence can increase the productivity of security teams while improving response speed.

Google Cloud Security AI research and operational data

Security operations centers are increasingly relying on AI assisted analytics and automation to manage the growing complexity of digital environments. As organizations deploy more connected systems, cloud platforms, and AI enabled applications, machine learning tools are becoming essential for identifying threats and prioritizing security responses.

AI infrastructure is expanding the cyber attack surface

New AI infrastructure is creating new security exposure

8000+

Exposed AI systems already online

Security researchers identified **more than 8,000 Model Context Protocol servers exposed online with default configurations**, highlighting how quickly new AI related infrastructure is appearing without adequate security controls.

These systems can potentially expose sensitive data, system access, or AI model interactions if not properly secured.

Palo Alto Networks Unit 42 AI security research

Security leaders are increasingly concerned about AI agents

48%

Security leaders rank AI as the top emerging threat

Nearly half of surveyed security professionals identify **agentic AI systems as the most concerning emerging cyber attack vector**, reflecting growing concern about autonomous systems interacting with digital environments.
As AI agents gain the ability to perform actions across systems and services, they introduce new security challenges related to identity, authorization, and system control.

Industry security leadership surveys and Gartner cybersecurity research.

The rapid expansion of AI infrastructure is introducing new categories of cyber risk. AI agents, model orchestration platforms, and emerging AI application ecosystems are expanding the digital attack surface faster than many organizations are prepared to secure. As these technologies become more widely deployed, security architectures will need to evolve to address new forms of exposure.

Workforce pressure is shifting from headcount to skills

Estimated global cybersecurity workforce.

4.8M

Cybersecurity workforce gap globally

Estimated number of additional cybersecurity professionals needed globally to meet demand.
88% Reporting Operational Impact
Security professionals reporting that workforce skill shortages have already created operational challenges for their organizations.

ISC2 Cybersecurity Workforce Study

World Economic Forum Global Cybersecurity Outlook

The cybersecurity workforce challenge is evolving from a simple shortage of professionals toward a broader transformation in required skills. As artificial intelligence becomes embedded across digital systems, organizations will need professionals capable of combining cybersecurity expertise, AI literacy, and governance frameworks to manage emerging digital risks.

Leadership Implications

The Leadership Shift: From Managing Technology to Governing Machine Decision Systems

For most of the digital era, executives were responsible for governing technology assets such as software, networks, and data. Security programs were designed to protect systems, manage access, and respond to incidents.

Artificial intelligence changes that model. Organizations are now introducing systems that can reason, plan actions, interact with external tools, and execute decisions inside business operations. These systems increasingly operate continuously and at speeds far beyond traditional human decision cycles.

The leadership challenge is therefore no longer simply deploying technology.

It is governing machine decision systems operating at machine speed inside the enterprise.

For boards, CEOs, and executive teams, the implications are strategic and immediate.

1. Cyber Maturity Is Becoming a Condition of Market Access

Cybersecurity is evolving from a defensive capability into an economic requirement.

Enterprises, insurers, regulators, and supply chains increasingly require proof of cybersecurity maturity before allowing organizations to connect systems, exchange data, or participate in critical operations. Companies that cannot demonstrate credible governance will face growing barriers to contracts, partnerships, and insurance coverage.

Cyber maturity is rapidly becoming a **passport to participate in the digital economy.**

2. Governance Must Move Into the Boardroom

Systems that influence hiring decisions, financial modeling, operational automation, and customer engagement now sit at the center of business operations.

This makes governance a board level responsibility.

Boards must move beyond approving policies toward overseeing how these systems are monitored, constrained, and controlled. Executive teams must establish clear accountability for outcomes, define operational guardrails, and ensure that leadership receives measurable evidence that governance is functioning in practice.

Governance is no longer a technical function. It is enterprise risk management.

3. Machine Speed Systems Require Machine Speed Governance

Many organizations are introducing autonomous systems into environments still governed by human processes such as meetings, approvals, and periodic reviews.
This creates a structural mismatch.
Systems execute actions continuously and at machine speed, while governance often operates through delayed oversight and fragmented decision making. The result can be uncertainty, delayed responses, and gaps between what leadership believes is governed and what is actually running.
Organizations must redesign governance structures so that visibility, correlation, and decision authority operate closer to real time.
Without governance that can operate at comparable speed, leadership loses effective control.

4. The Attack Surface Has Moved Up the Stack

Traditional cybersecurity focused on protecting networks, endpoints, and software artifacts.
Modern systems introduce a different layer of risk.
Decision logic, data pipelines, prompts, and automated workflows now shape operational outcomes. Threats increasingly target how systems interpret information, access data, and trigger actions rather than simply exploiting software vulnerabilities.

Security therefore must evolve beyond inspecting artifacts to governing system behavior.

This shift requires new defensive capabilities and new leadership attention to how automated systems operate across the enterprise.

5. Machine Identities Are Becoming the Largest Identity Risk

Historically, identity security focused on human users.
Today automated services, bots, APIs, and autonomous systems are multiplying across enterprise environments.
These machine identities often hold extensive privileges and interact with multiple systems simultaneously.
In many organizations they already outnumber human users.

Leadership must ensure that identity governance evolves to include strict authentication, scoped privileges, and continuous monitoring of non human actors operating inside enterprise systems.

6. Trust Is Emerging as Strategic Infrastructure

In previous decades infrastructure meant transportation networks, energy grids, and telecommunications systems.
In the digital economy, a new form of infrastructure is emerging.
Trust.
Organizations must demonstrate that their digital systems are secure, their governance structures are credible, and their operations are resilient. This capability determines whether partners connect, insurers provide coverage, regulators approve operations, and supply chains integrate systems.

Trust is no longer a soft concept.
It is becoming a structural requirement for economic participation.

Final Leadership Message

The organizations that succeed in the next phase of the digital economy will not simply be those that adopt new technologies the fastest.

They will be those that govern them the best.

The leadership challenge of the coming decade is building institutions capable of controlling autonomous systems operating inside critical business functions.

In the emerging digital economy, competitive advantage will not be defined only by innovation.

It will be defined by trust.

Conclusion and Forward Look

Artificial intelligence has moved beyond experimentation. It is now entering the operational core of organizations across finance, healthcare, manufacturing, government, and critical infrastructure. Systems that once analyzed information are increasingly making decisions, executing tasks, and interacting with other systems with minimal human intervention.

This shift marks a structural turning point for cybersecurity and digital governance.

For decades, security strategies focused primarily on protecting infrastructure such as networks, devices, and applications. The arrival of autonomous systems changes that equation. Risk is no longer limited to compromised software or stolen data. It now includes how machines interpret information, make decisions, and act within complex digital ecosystems.

The result is a fundamental expansion of the attack surface.

Threat actors are already adapting. Prompt injection, AI model manipulation, data poisoning, and the exploitation of automated workflows represent a new class of risk that traditional security controls were not designed to address. At the same time, the rapid deployment of AI across enterprise tools is creating widespread visibility gaps. In many organizations, leadership cannot fully answer three basic questions: where AI is being used, what decisions it is influencing, and what safeguards are in place to control its behavior.

This challenge is not purely technical.
It is institutional.

Organizations are introducing machine speed actors into environments governed by human speed processes. Committees, periodic reviews, and static policies were never designed to supervise systems capable of acting autonomously in milliseconds. As AI capabilities expand, the gap between technological power and governance capacity will increasingly define organizational resilience.

The implications extend beyond cybersecurity.

Cyber maturity is becoming an economic requirement. Insurers, procurement teams, regulators, and supply chain partners are already beginning to evaluate whether organizations can demonstrate credible governance over their digital systems. Companies that cannot prove control over automated decision systems may face barriers to contracts, partnerships, or insurance coverage.
In this environment, digital trust becomes more than a reputation factor. It becomes infrastructure.

The organizations that succeed in the AI economy will not simply be those that deploy the most advanced technologies. They will be those that build the governance, identity management, runtime oversight, and leadership accountability necessary to operate those technologies responsibly and at scale.

For Canada, this moment presents both risk and opportunity.

Canada possesses a world class research ecosystem, a strong cybersecurity talent base, and a regulatory culture capable of thoughtful governance. At the same time, AI adoption across industry is accelerating faster than institutional frameworks can evolve. Bridging that gap will require collaboration between government, industry, academia, and the cybersecurity community to ensure that innovation and security advance together.
Leadership will play the decisive role.

Boards must demand evidence of governance, not simply policies. Executives must ensure that automated systems operate within defined boundaries. Security leaders must evolve beyond defending infrastructure to governing behavior across digital systems.

Artificial intelligence will continue to reshape industries and societies over the coming decade. But its long term success will not be determined solely by computing power or model sophistication.

It will depend on whether organizations can build the trust required to operate these systems safely.
The future of AI will ultimately be defined not by what machines can do.

But by how well humans govern them.