AI and the Future of Canadian Policing: Implications and Strategic Considerations

Executive Decision Report for Sponsoring Chiefs of Police

A companion to the Comprehensive Report (October 2026)

November 2026

Community Safety Knowledge Alliance

cskacanada.ca

This page left intentionally blank

Table of Contents

GLOSSARY 4

1. INTRODUCTION 5

2. WHY AI GOVERNANCE IS A LEADERSHIP ISSUE NOW 6

3. KEY FINDINGS AND IMPLICATIONS 7

4. LEGAL, RIGHTS AND ACCOUNTABILITY BOUNDARIES 15

5. RECOMMENDATIONS AND PROPOSED GOVERNANCE APPROACH 17

6. IMPLEMENTATION, CAPACITY AND PHASING 22

7. DECISION GUIDELINES FOR AI ADOPTION 24

8. FURTHER STUDY 25

9. SUGGESTED IMMEDIATE NEXT STEPS 25

REFERENCES 27

ACKNOWLEDGEMENTS 32

REPORT AUTHORS 33

ABOUT CSKA 34

AI DISCLOSURE 34

Glossary

Agentic AI: AI systems that can independently plan and take actions to achieve a goal, rather than simply responding to a user's request. They can perform multi-step tasks, use tools and make decisions with limited human intervention.

AI governance: The set of rules, processes, oversight structures and accountability mechanisms that guide how AI systems are designed, acquired, used, monitored and reviewed, so that they are deployed lawfully, ethically, transparently and securely, not just effectively.

Algorithmic bias: When a computer system produces unfair or skewed results, often because the data it was trained on reflects existing inequalities, gaps or patterns of discrimination.

Automation bias: The tendency to over-trust automated outputs, accepting a system's recommendation, match or summary without sufficient independent scrutiny, particularly under time pressure.

Data sovereignty: The principle that individuals, communities, organizations or nations have the right to govern how data about them is collected, used, stored, shared and protected.

Function creep: The gradual expansion of an AI tool beyond its originally approved purpose, scope or data sources, often without renewed review or approval. Also known as scope creep.

Generative AI: A type of AI that creates new content, such as text, images, audio, video or code, in response to a prompt, based on patterns learned from large amounts of training data.

Shadow AI: The use by personnel of AI tools that have not been approved by the organization, such as personal generative AI accounts. A growing subset of "shadow IT."

Zero-option: The retained authority, supported by predefined conditions, to decline, restrict or withdraw an AI tool or system at any stage, regardless of prior investment or vendor commitments.

1. Introduction

Purpose

This report supports decisions by chiefs of police, police boards and policymakers on the adoption and governance of artificial intelligence (AI) in Canadian policing. It condenses the Comprehensive Report, AI and the Future of Canadian Policing: Implications and Strategic Considerations (October 2026), which remains the full record of evidence, methods, scenarios and analysis.

The findings draw on a narrative synthesis of scholarly and grey literature, confidential semi-structured interviews with seven subject matter experts conducted between February and May 2026, and review of the draft by a multidisciplinary expert panel in September 2026. The analysis reflects conditions as of June 2026 and was revised in October 2026 to incorporate the panel's feedback. Given the pace of change, the guidance should be revisited at least annually.

Strengths and limits of the evidence

The review is a narrative synthesis of peer-reviewed research and grey literature, including investigative journalism, supplemented by a small purposive sample of expert interviews and an expert panel review. Independent validation of many police AI tools is limited, and several illustrative cases come from jurisdictions whose legal and social conditions differ from Canada's. The findings identify risks and the governance conditions that address them. They do not establish that a particular tool will produce a particular outcome in a Canadian service. Because law and technology are changing quickly, the guidance is principles-based so that it remains useful as specific tools, vendors and laws change.

Bottom line

Some AI adoption in Canadian policing is unavoidable. Organized crime and other malicious actors are already using it, and police cannot afford to disengage. But AI adoption is not a neutral modernization step. AI systems reshape what is visible, prioritized and acted upon. Adopted without governance, they can erode the public trust, legal legitimacy and professional judgement on which democratic policing depends.

The report recommends treating AI as a governed capability: embedded within professional judgement, constrained by law and accountable to the communities the police serve. It sets out 21 recommendations. Five form a **Year-1 Core** of foundational controls that every service should put in place first, regardless of size. The remainder can be phased in as capacity matures, with priority given to those that apply to medium- and high-risk tools.

Decisions for consideration

Chiefs and boards should consider the following measures:

1. **Commit to the five Year-1 Core measures:** a named executive accountable for AI governance, a maintained inventory of AI tools and quarterly reporting (Recommendation

#1); a tiered risk-classification scheme, with board approval before any high-risk tool proceeds to procurement (#5); a standard contract template protecting data ownership, evidentiary disclosure and exit (#10); an acceptable-use policy, updated within 90 days, that prohibits unapproved "shadow AI" (#11); and mandatory AI-literacy training within 12 months, backed by safeguards against automation bias (#14).

2. **Define the "zero-option" in advance:** the conditions under which a tool must be rejected, restricted or withdrawn, with a published register of tools in use and tools withdrawn (#8).
3. **Commit to structured community consultation** for medium- and high-risk tools, supported by a dedicated budget line and executive lead (#18, #19).
4. **Require documented joint legal, privacy and IT security sign-off** before any operational deployment (#6), extended through formal agreements where IT or procurement functions are shared with the municipality (#13).
5. Advocate for a **shared regional or provincial AI evaluation body** through the relevant chiefs' association (#20).

Section 5 sets out all 21 recommendations with accountable parties and timing.

2. Why AI Governance is a Leadership Issue Now

AI already influences Canadian policing through investigative analytics, automation, digital evidence management and emerging decision-support tools. It is also a threat multiplier: criminal networks, violent ideological actors and foreign adversaries are using it to scale fraud, disinformation, recruitment and cyber operations. A key strategic question is whether adoption will be governed in ways that strengthen effectiveness while preserving public trust, legal legitimacy and institutional resilience, and whether malicious use will be understood and countered effectively.

Four conditions underscore the importance of this question:

- **Pressure to adopt.** Labour shortages, recruitment/retention challenges, fiscal restraint and rising public expectations push agencies toward tools that promise cost savings and faster operations. Vendors build their value propositions on such claims, and police culture tends to be quite receptive to compelling new tools. Adoption should therefore begin with a clearly defined operational or community safety problem, and the question "safety for whom?", rather than with a vendor's proposal.
- **Public skepticism.** Recent polling suggests that 68% of Canadians believe government needs to heavily regulate AI even if it slows development, and 74% doubt any government can keep pace with the technology.[1] The public is wary not only of AI tools but of institutions' capacity to oversee them.

[1] Angus Reid Institute (2026, June 1)

- **Regulatory lag.** Legislation, procurement rules and vendor auditing are struggling to keep pace. The February 2026 Tumbler Ridge shooting shows the stakes at hand. According to media reports, OpenAI staff had internally flagged the alleged shooter's ChatGPT conversations about gun violence months earlier, but neither government officials nor local RCMP were notified. Canada currently has no mandatory reporting requirement for platforms, and the threshold for such reporting remains undefined.[2]
- **Broader disruption.** AI-related job displacement, arriving amid a wider "polycrisis" of climate, economic, geopolitical and information pressures, may produce localized insecurity, opportunistic crime and greater demand for health and social services.[3]

Without proactive governance, agencies risk gradual erosion of trust inside and outside the organization; loss of critical thinking and professional judgement; disproportionate impact on racialized and equity-deserving communities; normalized surveillance practices; evidentiary risk in court; and vendor dependence and data sovereignty failures.

These harms rarely appear as single dramatic failures. They accumulate, and they are difficult to reverse.[4] AI governance is therefore not a communications strategy. It is institutional infrastructure, and it depends as much on people, organizational culture and community relationships as on technology.

3. Key Findings and Implications

Four key findings capture the most significant ways AI is affecting Canadian policing now and in the near term. Each presents both opportunity and significant risk. Each also carries governance implications that leaders and boards should not treat as purely technical or operational matters.

3.1 Operational benefits are real but conditional

AI's most immediate value lies in managing scale, complexity and speed:

- **Efficiency:** interview transcription, report drafting, call triage subject to safeguards, administrative workflows, training and human resources functions, and imaging that supports uncrewed aerial vehicles, collision reconstruction and public-order safety monitoring.
- **Investigative capacity:** analytics, pattern recognition, image matching and language processing across large volumes of financial, video and open-source data. These are especially useful in financial crime, organized crime, cybercrime and child sexual exploitation cases, where automation can also reduce investigators' exposure to traumatic imagery.

---

[2] Globe and Mail Editorial Board (2026)

[3] e.g., World Economic Forum (2025); Mehdi & Morissette (2024); Lawrence, et al. (2024)

[4] cf. Kasirzadeh (2025) on "accumulative" AI risk

- **Deployment and risk awareness:** forecasting and scenario tools for planning, and robotics that extend officers' reach into dangerous spaces during critical incidents. These uses aim to reduce risk, not increase force.

These benefits are contingent. AI adds value where it augments professional judgement and addresses a clearly defined problem. It creates risk where it displaces accountability or arrives as a solution in search of a problem. For instance, in an investigation, it remains for the investigator to determine whether and what evidence supports laying a charge.

Two qualifications matter for decisions:

- First, administrative uses such as transcription or scheduling generally carry lower risk than operational uses, and the two should be assessed against different thresholds. AI used as a decision-support system in operations, intelligence and risk assessment is a distinct and more consequential risk category.
- Second, vendor efficiency claims warrant scrutiny. Participants in a recent University of British Columbia roundtable noted that independent validation of many police AI tools remains limited, that operational results have been inconsistent, and that some tools added work, such as correcting inaccurate report content. They also noted accumulating evidence that facial recognition accuracy is improving, especially under controlled conditions.[5]

In summation, the case for any AI tool rests on independent evidence of what it does in practice, not on what it promises. The challenge then is to look beyond whether a tool may be useful, to whether its benefits can be demonstrated, independently verified and sustained under proper governance.

3.2 Ethical and practical risks are significant

The same features that make AI useful to policing (e.g., speed, scale, managing complexity and the appearance of objectivity) also create its primary risks. Some are ethical: bias, opacity and the gradual erosion of human judgement. Others are practical: unreliable outputs, incremental expansion of use, hidden dependencies and workforce strain. Few, if any, of these risks show up at the point of purchase. They emerge through application and use, which is why they must be governed from the outset rather than addressed after the fact.

Bias and systemic risk

AI trained on historical data can reproduce past patterns of over-policing, unequal protection and systemic discrimination, particularly affecting Indigenous, racialized and marginalized communities. Tools trained or tested on unrepresentative or biased data can produce false-positive matches or invalid risk predictions. Without auditing and contextual safeguards, AI may formalize inequity under the appearance of technical neutrality and run counter to police work on race-based data and human rights-based policing.

[5] Perrin (2026)

The Clearview AI case is instructive. The Office of the Privacy Commissioner of Canada (OPC) found significant shortcomings, ranging from accuracy and data collection to a real risk of harm.[6] Police had also submitted images of minors, potentially creating permanent records with no right to erasure.

The governance lesson is that legitimacy failures arise from transparency gaps, unclear legal authority and unauthorized use, not only from technical error. Community response may hinge less on effectiveness than on whether an agency can articulate its legal basis, demonstrate safeguards against bias and show robust oversight. Fear of surveillance is itself a harm: it can deter people from reporting crime, seeking help or exercising lawful rights, even where no misuse occurs. Some uses may also conflict with the Truth and Reconciliation Commission of Canada's (TRC) Calls to Action or the United Nations Declaration on the Rights of Indigenous Peoples (UNDRIP).[7]

Human oversight

A "human in the loop" does not by itself guarantee meaningful control. As one interviewee noted of facial recognition, the cognitive bias that can accompany a reported match "is pretty high." Confident-looking outputs can anchor, rather than inform, a reviewer's judgement under time pressure. One participant described decision-makers facing heavy case volumes who relied on AI-generated summaries instead of the source files. Reviewers need the authority, training, information, supervision and time to challenge outputs; otherwise, oversight becomes box-checking. Practical controls include training on automation bias, a documented rationale whenever a reviewer accepts or overrides a consequential output, and supervisory audits that confirm primary sources are still being examined.

Output quality

Tools built on other flawed models, code or algorithms are more prone to faulty output and to malicious penetration. Where AI replaces human operators in consequential functions, such as emergency or non-emergency call triage, misunderstood or misclassified calls can have serious consequences.

Transparency and explainability

Many AI systems are opaque. Tools that cannot be explained to courts, boards or the public create legal and reputational exposure, so explainability is an operational requirement, not a technical preference. The challenge is not only whether a vendor will open the "black box," but whether the agency and its oversight bodies have the literacy and authority to understand what is disclosed. Access to information is not the same as explanation.[8] Transparency should be layered and calibrated: public disclosure at the level of the tool (what it is, what it is used for and how outputs can be challenged), and confidential disclosure of method-level detail to an

[6] Office of the Privacy Commissioner of Canada (2021)

[7] Truth and Reconciliation Commission of Canada (2015); Office of the United Nations High Commissioner for Human Rights and Asia Pacific Forum of National Human Rights Institutions (2013)

[8] Wischmeyer (2019)

independent oversight body where public release could compromise investigations. Both layers are necessary, however neither is sufficient on its own.

Scope creep and cumulative effects

Interviewees were less concerned about initial deployment than about the absence of structural limits on incremental expansion: digital evidence analytics repurposed for live surveillance, call screening extended to emergency calls, or image recognition approved for static images extended to live video. Aggregating data across investigations can also erode privacy protections that apply within individual investigations. Every deployment needs documented objectives and scope, and any material change should trigger a new risk assessment and fresh approval.

Vendor reputation and hidden labour

Some platforms also serve military intelligence, and local adoption may alarm segments of the community, although prevalence and interoperability can make these choices difficult. Some tools presented as autonomous rely on offshore human operators or reviewers, which also raises error, security and ethical questions.[9]

Workforce and labour relations

How AI is introduced matters as much as what it does. London's Metropolitan Police Service began piloting AI tools in 2026 to identify officers who may be at risk of misconduct; its police federation criticized the approach as "automated suspicion."[10] Research suggests officers accept AI oversight more readily when they see the process as fair and transparent, while some US departments have dropped AI body-camera analytics in the face of union opposition.[11] Services would benefit from engaging early with employees and labour representatives. Officers also face doxxing and coordinated online abuse, risks that synthetic media such as deepfakes are likely to intensify.[12]

3.3 AI introduces new institutional and security risks

The risks discussed in Section 3.2 above arise mainly in how AI affects individuals and communities. The risks in this section can affect the police institution itself (i.e., its security, its independence and its capacity to stay in control of the tools it adopts). These risks include exposure to cyber-attack and foreign reach, dependence on vendors, unsanctioned use by personnel, and the quieter risk of platforms beginning to shape how the organization thinks and acts.

---

[9] Sims (2026); TechCabal (2026)

[10] MPS Confirms the Use of AI (2026); Booth (2026)

[11] Adams (2025); Watts, et al. (2025); Du Sault (2023)

[12] Huey et al. (2025); Watson (2023)

Because these risks are systemic, they extend well beyond any single police service. They have the potential to weaken public confidence in Canadian policing as a whole.

Cyber exposure

AI concentrates sensitive data and expands the attack surface. Adversaries can spoof biometric systems, poison training data or exploit automated decision thresholds, and model tampering can quietly distort outputs. Breaches at the vendor, cloud or agency level can compromise investigations, prosecutions and public trust. Vulnerabilities in upstream suppliers and subcontractors matter as much as those of the vendor itself.

Shadow AI

Personnel using unapproved AI tools or personal accounts is an under-recognized compliance and security risk, including potential privacy violations.[13] Prohibition works best when paired with approved, secure alternatives that meet the demand driving these workarounds.

Organizational capture

As tools become embedded in workflows, agencies may become dependent on vendor-owned systems they do not control and cannot fully audit. Capture can deepen after a contract is signed. Vendor updates can change model behaviour without notice, a change of ownership can change who has access to data and under which jurisdiction, and escalating licensing and integration costs can make exit impractical. Inadequate knowledge transfer compounds the dependence.

Data sovereignty and foreign reach

Where a widely used application carries cybersecurity or foreign-ownership exposure, participating agencies may become part of a systemic vulnerability across Canadian law enforcement. China's *National Intelligence Law* (2017) requires organizations and citizens to assist and cooperate with state intelligence activities, and Canadian security and intelligence authorities have, in some documented cases, assessed Chinese technology platforms and equipment as posing potential security or supply-chain risks.[14],[15] More recently, concern has grown that many enterprise platforms used in Canada are US-owned, or route data through the United States, even when that data is nominally stored in Canada. This creates potential operational and privacy risks in the event of unilateral US policy or regulatory action.

Sector-wide capacity gap

Many police agencies, particularly small and mid-sized ones, lack the in-house expertise to evaluate AI systems on their own terms. As one interviewee put it, police "inevitably end up going by the vendors' claims on how their tools work." Ontario alone has 45 police services of vastly different sizes and capabilities, and even well-resourced agencies struggle with independent technical evaluation. Without shared evaluation mechanisms, agencies risk

---

[13] Perrin (2026)

[14] Standing Committee of the National People's Congress of the People's Republic of China (2017)

[15] Government of Canada, Ministry of Innovation, Science and Industry, & Ministry of Public Safety (2022); Canadian Centre for Cyber Security (2024)

repeating one another's failures and damaging sector-wide legitimacy. Perrin (2026) proposes building on the existing "multi-layered oversight ecosystem," such as Ontario's Inspectorate of Policing, police service boards and British Columbia's Director of Policing and Law Enforcement Services, rather than starting from scratch.[16] In the interim, the International Association of Chiefs of Police (IACP) Connector program offers peer experience and AI-specific guidance.[17] Where IT and procurement sit within shared municipal services, key decisions may be made outside the police service, while the legal, privacy and evidentiary consequences remain with the chief and board.

Decision shaping and practice drift

Over time, a platform can shape what an organization sees, prioritizes and acts upon, so that problems are framed in terms of what the system can measure or automate. A multi-year study of the LAPD's use of Palantir platforms found some surveillance practices drifting toward approaches resembling military intelligence, and away from information gathering based on probable cause toward predictive processes.[18] Dutch research suggests such risks are better managed through ongoing dialogue with a diversity of stakeholders than as a one-time adoption decision.[19] Section 4 sets out the Canadian legal threshold relevant to this drift.

Misaligned AI agents

Media reports in 2026 described AI agents acting beyond their design parameters, including taking unauthorized actions against external networks without human direction.[20] One example was a recent incident in which a group of Anthropic's agents acting independently of human direction sent a fake homicide tip to the Philadelphia Police Department.[21] Research also suggests self-replicating agents running on everyday connected devices could become new attack vectors.[22] Several international experts have called for a slower pace of AI development until stronger safeguards are in place.[23]

---

[16] Perrin (2026); Perrin, et al. (2026)
[17] IACP (2026)
[18] Brayne (2017), reported in Crawford (2021)
[19] Jacobs, et al. (2024)
[20] Bhuiyan (2026); Tuquero (2026)
[21] Conger, et al. (2026)
[22] Guan, et al. (2026)
[23] Castaldo & Posadzki (2026); Cecco (2026)

3.4 AI-enabled threats are escalating

Criminal, ideological and state actors are already exploiting AI. Police services should anticipate:

- **Fraud and impersonation**, including deepfake extortion and business email compromise, increasingly supplied through "crime-as-a-service" markets that make sophisticated offending accessible to a wider range of offenders.[24]
- **Exploitation and abuse**, including AI-enabled grooming, sextortion and synthetic child sexual abuse material, sometimes created from images of real children, together with online sexual exploitation and abuse against women – including intimate partner violence. Experts warn that these may lower barriers to offending, complicate victim identification and make authentic and synthetic imagery harder to distinguish.[25] The federal government has indicated that laws specific to AI-generated deepfakes are being considered.
- **Disinformation and radicalization** targeting both the public and police personnel, including campaigns that exploit social divisions and AI-driven job displacement.[26]
- **Hybrid threats** where cybercrime, disinformation and national security converge, including the online recruitment or coercion of "disposable" agents to carry out violence or sabotage.
- **Insider threats**, where AI-curated online environments facilitate the covert radicalization of civilian or sworn members.

These harms will converge rather than arrive neatly categorized, and they will outpace traditional responses unless countermeasures adapt at the speed of cyber. An effective response requires integrated intelligence; close coordination with cybersecurity, national security and public safety partners; human judgement in the final review of intelligence products; workforce integrity measures that balance freedom of expression and due process with early intervention; and the positive community relations needed to withstand information operations targeting police–community trust.

[24] UNODC (2025)

[25] United Nations Children's Fund (UNICEF) (2025); Internet Watch Foundation (2026)

[26] Mosolova (2026)

3.5 Principal risks and controls

The following figure (Fig. 1) groups the ten key risks identified above by where they arise, and maps each to its principal controls and the recommendations in Section 5 for implementing the controls.

Figure 1 - Principal Risks



Two patterns in Fig. 1 are noteworthy:

1. Most risks are addressed by more than one recommendation, and several recommendations, notably the zero-option and public register (#8) and standard contract terms (#10), address several risks at once.
2. The controls that matter most largely fall within an agency's authority (e.g., what it buys, on what terms, under whose approval and with whose input).

Section 4 sets out the legal boundaries within which those decisions must be made.

4. Legal, Rights and Accountability Boundaries

AI adoption is never simply a procurement decision. It engages constitutional, privacy, evidentiary and cybersecurity obligations that require structured review before deployment and governance that adapts as tools, vendors and the law change. Leaders should keep three categories distinct: established legal obligations; law that is proposed or evolving; and the governance measures this report recommends, which go beyond legal minimums.

Established obligations

Three bodies of established law have application on police decisions about AI. Under the *Canadian Charter of Rights and Freedoms*, AI that influences surveillance, search, detention, risk assessment or investigative prioritization may engage sections 7, 8, 9 and 15. Courts increasingly examine the processes and technologies that shape police decisions, not only their outcomes, so tools that affect investigative thresholds or evidentiary pathways must withstand scrutiny for reasonableness, proportionality and procedural fairness.

Privacy law applies in parallel. Federal policing bodies are governed by the *Privacy Act*; municipal and provincial services by provincial privacy legislation and police services acts, overseen by provincial commissioners; and private-sector vendors by the *Personal Information Protection and Electronic Documents Act* (PIPEDA). Transparency, necessity, proportionality, data minimization and algorithmic accountability are expected, and vendor assurances do not relieve police of their own statutory obligations.

The most immediate operational exposure lies in criminal procedure and evidence. AI outputs, and the prompts used to generate them, may be subject to disclosure under *R. v. Stinchcombe* and to reliability scrutiny under *R. v. Mohan*.

Courts have made clear that the Crown cannot avoid disclosure because a technology is hard to explain or kept confidential, and claims that a tool must be protected as an investigative technique are likely to be harder to sustain. Where a vendor will not disclose system logic, training data provenance or validation results, or support expert testimony, the risk of a failed prosecution shifts to the police service and the Crown.

Courts are also setting expectations for AI use in their own proceedings, requiring certification of the authenticity of facta and expert reports, or disclosure of AI use in court materials. In *Ko v. Li* (2025), where a factum cited non-existent, AI-generated cases, the court held that counsel must verify the authorities they cite. Professional organizations are expected to use AI in informed and appropriate ways.[27]

Benchmarks and evolving law

Beyond the law that already constrains police, a second set of instruments is shaping expectations. Some are federal standards that do not formally apply to municipal services but are increasingly treated as benchmarks for good practice. Others are bills and proposals that may or may not become law. None imposes new obligations on police services today. However,

[27] Lockwood et al. (2025); Ko v. Li (2025)

considered together, they show where regulation is likely heading. Agencies that govern AI to these standards now will be better positioned when the law catches up.

- The Treasury Board *Directive on Automated Decision-Making* binds federal institutions. Municipal police services are not formally bound, but it offers a benchmark for impact assessment, explainability, transparency and human oversight.
- Bill C-36, the *Protecting Privacy and Consumer Data Act*, was introduced in June 2026. If enacted, it would replace Part 1 of PIPEDA and create a Digital Safety and Data Protection Commission of Canada, replacing the Office of the Privacy Commissioner, with expanded enforcement powers including significant administrative monetary penalties.[28]
- The proposed *Artificial Intelligence and Data Act* (part of Bill C-27) and the *Critical Cyber Systems Protection Act* (Bill C-26) did not become law before the 2025 election. They signal a direction toward risk-based regulation and heightened diligence on infrastructure, supply-chain and cross-border data risks; they do not impose obligations.
- Canada may legislate AI generally or by sector, and the *Criminal Code* may be amended to fill gaps created by AI. Where information sharing crosses national boundaries, different laws and conventions may apply.

Predictive analysis and the Charter (section 8) threshold

Police in Canada do not need probable cause to conduct predictive intelligence. The legal trigger is whether the activity intrudes on a reasonable expectation of privacy under the Charter.

Analyzing lawfully obtained data, open-source information or historical crime records is generally permitted. When predictive work involves a search or seizure, such as accessing private digital data or conducting intrusive surveillance, police must have or obtain judicial authorization, typically based on reasonable grounds to believe an offence has been or will be committed. Because AI makes deeper and more intrusive analysis easy, scalable and routine, leaders should anticipate and guard against drift from lawful information gathering toward more legally questionable practices.

Demonstrable public benefit

From a Charter perspective, efficiency alone is unlikely to justify AI adoption that affects privacy. The UBC roundtable positioned necessity, proportionality and demonstrable public benefit as foundational to governing police AI.[29] Demonstrable public benefit means being able to explain to the community how a technology advances its safety and how privacy and liberty are protected in its use.

Chiefs must also consider whether cumulative use, or a vendor's reputation, alters the balance between state power and individual rights in ways that courts or oversight bodies may later reject, or that may undermine public trust. Throughout, accountability for investigative, evidentiary and operational decisions remains with people, not systems.

---

[28] Parliament of Canada (2026); Osler, Hoskin & Harcourt (2026)

[29] Perrin (2026)

5. Recommendations and Proposed Governance Approach

The recommendations rest on five principles:

1. **Start with the problem, not the product.** Every acquisition begins with a documented problem statement and a “safety for whom?” assessment.
2. **Match governance to risk.** Administrative uses and operational uses that directly affect the public are classified and approved differently.
3. **Approve within a defined scope.** Any material change in use, data sources or functionality triggers re-assessment and renewed approval.
4. **Keep the zero-option real.** The decision not to adopt, or to withdraw, is always on the table, supported by predefined conditions so that it does not depend on a crisis.
5. **Be realistic.** A comprehensive framework that stalls under its own weight provides less protection than a smaller set of controls that is actually implemented.

Controls and mitigations by level of risk

Recommendation #5 calls for a tiered scheme (e.g.,, low, medium and high intrusiveness). Particularly consequential uses involve AI for decision support. The mitigations and controls below attach to that classification. Consultation may result in some tools initially classed as medium risk being designated high risk.

Table 1 - Risk Control and Mitigation

Control/Mitigation	Applies to	Rec.
Documented problem statement and “safety for whom?” assessment	Every AI acquisition	#9
Inclusion in the AI tool inventory and public register	All tools in use or withdrawn	#1, #8
Annual independent performance, bias and social compliance audits	Every AI vendor contract	#4
Joint legal, privacy & IT security sign-off	Any tool before operational deployment	#6
Plain-language explanation and complaint channel	Each deployed tool	#17
Structured community consultation	Medium- and high-risk tools	#18
Impact assessment pilot	Next medium- or higher-risk procurement	#21
Board approval before procurement	High-risk tools	#5
Annual external panel review, reported publicly	High-risk deployments	#2
Sequencing to internal capacity	High-risk tools (lower-risk administrative tools may proceed on a faster track)	#15

Recommendations

The following table summarizes the full comprehensive report's 21 recommendations. Note that:

- Year-1 Core recommendations are marked ●.
- Timing reflects the Comprehensive Report.
- "Not specified" indicates that no timeframe was proposed.

Table 2 - Summary of Recommendations & Suggested Timing

#	Recommendation	Suggested Lead	Timing
	Governance and Oversight		
1 ●	Make AI adoption and use a standing item on the executive governance agenda, with a named executive accountable for AI governance, a maintained inventory of all AI tools in use, and quarterly reporting on deployments, incidents and review findings.	Chief and board	Year 1; quarterly reporting
2	Establish a standing external review panel (a technical expert, a human rights or privacy specialist and a community advocate) to review all high-risk deployments annually and report publicly. Except in exceptional circumstances, the panel should have confidential access to method-level information that cannot be disclosed publicly without compromising investigations.	Board	Annual reviews
3	Review investigation protocols annually against current case law and emerging standards on AI-related evidence integrity and Crown disclosure obligations.	Legal counsel, ideally specialized external counsel	Annually
4	Include a standard contract clause requiring annual independent performance, bias and social compliance[30] audits for every AI vendor contract, with re-validation after any material update and audit reports delivered to the board.	Procurement	From next contract renewal cycle
	Decision Gates		
5 ●	Adopt a tiered risk-classification scheme (e.g., low/medium/high intrusiveness) that distinguishes administrative uses from operational uses that directly affect the public, and require board approval before any high-risk tool proceeds to procurement. Each approval defines the tool's permitted scope; any material change in use, data sources or functionality triggers a new risk assessment and renewed approval.	Chief, with the board	As a priority; Yr 1

[30] A social compliance audit involves an assessment of whether an organization and its suppliers adhere to laws, standards and ethical practices regarding labour conditions, human rights, worker health and safety, etc.

#	Recommendation	Suggested Lead	Timing
6	Complete a documented joint sign-off covering legal authority, privacy impact, cybersecurity risk and ethical considerations before any AI tool is deployed operationally, with no exceptions absent board-level emergency authorization. Where IT or procurement are shared municipal services, extend the sign-off to those partners through formal agreements. The process may consider external review panel findings (#2).	Legal counsel, privacy officer and IT security	Before each operational deployment
7	Identify and join, or help establish, a provincial or national vetted-tools registry, developed in step with the shared evaluation body (#20), so that evaluations can draw on other agencies' due diligence.	Chief directs staff	Within 12 months
8	Define in advance the conditions under which a tool must be rejected, restricted or withdrawn (the "zero-option"), including failed audits, unmitigated bias, unauthorized scope expansion, material vendor non-compliance or a significant loss of public confidence. Retain the option to reject any tool at any stage, regardless of prior investment or vendor commitments, and publish a register of tools in use and tools withdrawn.	Board	In advance; not otherwise specified
	Procurement and Risk Management		
9	Require a documented problem statement, including a "safety for whom?" assessment, before any AI acquisition proceeds. Score cybersecurity, inspectability, explainability, data sovereignty and local control in every AI vendor request for proposals, weighted explicitly alongside cost and performance.	Procurement	From next procurement cycle
10 ●	Develop a standard contract template covering data ownership; minimum cybersecurity standards; disclosure of system logic and training data provenance sufficient for court-ordered and Crown disclosure; vendor expert testimony when AI-derived evidence is challenged; advance notice and re-validation before updates are deployed; protections on change of vendor ownership; cost-escalation controls; and exit and data portability. Apply it to every new and renewed AI agreement, recognizing that large foreign vendors may refuse these terms.	Legal counsel; procurement applies	Year 1; all new and renewed agreements
11 ●	Update the acceptable-use policy to identify unapproved AI tools (e.g., personal ChatGPT, Claude or Gemini accounts) as prohibited shadow AI, identify approved and secure alternatives, and incorporate the policy into onboarding and annual training.	IT security	Within 90 days

#	Recommendation	Suggested Lead	Timing
12	Include contract provisions requiring vendors to deliver comprehensive knowledge transfer, including documented procedures, hands-on training and competency development, to sustain in-house capability and avoid unnecessary dependency.	Technology contracts (lead not specified)	Not specified
13	Where IT, data or procurement functions are delivered through shared municipal services, establish formal service level agreements and a cross-governance framework so that the police service retains control over Charter, privacy and evidentiary compliance.	Chief and board	Not specified
	Workforce and Organizational Preparedness		
14 ●	Deliver mandatory AI-literacy training, including explicit instruction on automation bias, to sworn and civilian employees, refreshed annually or as tools change. Pair it with a documented rationale whenever an employee accepts or overrides a consequential AI output, supervisory audits confirming that primary sources are still reviewed, and updated supervision and performance monitoring. Agencies may pool resources for joint training.	Training and professional development	Within 12 months; annual refresh
15	Sequence high-risk deployments to internal capacity: deploy no new update to a high-risk tool until the prior one has completed its first independent review, and plan for prompt evaluation where several high-risk tools are proposed at once. Lower-risk administrative tools may proceed on a faster track. The board retains explicit authority to pause a deployment if review findings, public concern, independent research or vendor conduct warrant it.	Chief; board holds pause authority	Ongoing
16	Add online radicalization indicators to insider-threat screening, recruit vetting and periodic employee review, in consultation with associations and unions on scope and process.	Professional standards	Not specified
	Transparency and Accountability		
17	Produce a plain-language explanation for each deployed tool (what it does, what data it uses and how a person can challenge an output) for public use and court disclosure, published through the public register (#8) with an accessible complaint channel. Any exigent-circumstances exception must be narrowly defined, approved by the board and subject to a timeframe for eventual disclosure.	Legal counsel and corporate communications	Each deployed tool

	Stakeholder Engagement		
18	Treat structured community consultation as a non-waivable requirement for medium- and high-risk tools, completed before board approval of high-risk tools. Prioritize the communities most likely to be affected, including internal personnel and racialized, marginalized and Indigenous communities; address fear of surveillance as well as demonstrated harms; give input meaningful decision weight; and publish how it shaped the decision. Community concerns can trigger external review (#2) or the zero-option (#8). Exigent exceptions must be narrowly defined, board-approved and followed by consultation after the fact.	Board publishes outcomes (lead not otherwise specified)	Before approval
19	Allocate a specific budget line and executive lead for community engagement within the AI governance framework, and review the adequacy of that resourcing annually.	Board	Annual review
	Fragmented Capacity		
20	Advocate for a shared regional or provincial AI evaluation body at the next relevant chiefs' association meeting. The body would consolidate assessment information, maintain the vetted-tools registry (#7) and give smaller agencies access to shared technical, legal, privacy and academic expertise.	Chiefs	Working group within 6 months; feasibility study support within 12 months
21	Explore and pilot a relevant impact assessment tool on the next medium- or higher-risk AI procurement, such as the Law Commission of Ontario's *Human Rights AI Impact Assessment* or the Treasury Board's *Algorithmic Impact Assessment*,[31] and report results to the board to inform whether it should become standard for all high-risk tools.	Board directs staff	Next medium- or higher-risk procurement

The above recommendations are designed to work as a system versus a mere checklist. Each control reinforces the others, from the first problem statement through to the decision to withdraw a tool. Section 6 sets out how services can put them in place in stages, starting with the Year-1 Core.

[31]Law Commission of Ontario & Ontario Human Rights Commission (2024); Treasury Board of Canada Secretariat (n.d.)

6. Implementation, Capacity and Phasing

This section sets out how the 21 recommendations can be phased so that agencies of any size can begin with the Year-1 Core and build from there as capacity allows. It also proposes who should be accountable for each measure, how smaller agencies and those in shared municipal environments can draw on shared resources, what the approach will require, and how to tell whether it is working.

Sequencing

The five Year-1 Core measures are designed to be achievable for services of any size and to form the foundation for the rest. The source timelines produce the following sequence:

- **Within 90 days:** updated acceptable-use policy prohibiting shadow AI (#11).
- **Within six months:** a working group toward a shared regional or provincial evaluation body (#20).
- **Within 12 months:** the remaining Year-1 Core measures (#1, #5, #10, #14); joining or helping establish a vetted-tools registry (#7); and support for a feasibility study on the shared evaluation body (#20).
- **At the next procurement or renewal cycle:** scored procurement criteria (#9), audit clauses (#4) and the impact assessment pilot (#21).
- **Ongoing and annual:** external panel review (#2), legal review of investigation protocols (#3), sequencing of high-risk deployments (#15), training refresh (#14) and review of community engagement resourcing (#19).

Suggested Accountabilities

Board	Chief and executive	Functional leads
Approves high-risk tools before procurement (#5, with the chief) Defines zero-option conditions and publishes the register (#8) Establishes the external review panel (#2) Retains authority to pause any deployment (#15) Approves emergency and exigent exceptions (#6, #17, #18) Funds community engagement (#19) Directs the impact assessment pilot (#21) Receives quarterly reports and audit results (#1, #4)	Designates the accountable executive and maintains the inventory (#1, with the board) Adopts the risk-classification scheme (#5, with the board) Directs participation in a vetted-tools registry (#7) Sequences high-risk deployments (#15) Establishes municipal service agreements (#13, with the board) Advocates for a shared evaluation body (#20)	Legal counsel: #3, #6, #10, #17 Privacy officer: #6 IT security: #6, #11 Procurement: #4, #9, #10 Training and professional development: #14 Professional standards: #16 Corporate communications: #17

Smaller agencies and shared environments

No single small or mid-sized police service can sustain the technical, legal, privacy and academic expertise that rigorous AI evaluation requires. The shared evaluation body (#20), the vetted-tools registry (#7) and pooled training (#14) are intended to let these services draw on one another's due diligence rather than start from zero. Peer resources such as the IACP Connector program offer a practical interim step.

Where IT, data or procurement sit within municipal shared services, formal service level agreements and a cross-governance framework (#13) are needed to keep the police service in control of its Charter, privacy and evidentiary obligations.

Resources and constraints

The Comprehensive Report does not estimate costs. The material resource requirements are specialized legal counsel, preferably external, for annual review and contract development; an external review panel; independent annual vendor audits; mandatory training with annual refresh; a dedicated budget line and executive lead for community engagement; and staff time for sign-off, registers, explanations and consultation. Recommendation #10 recognizes that large foreign vendors may refuse the proposed contract terms, which may narrow procurement options.

Measuring success

If the recommendations are implemented:

- Every AI tool starts with a clearly defined problem, is matched to its level of risk, is approved within a defined scope and is reviewed again whenever that scope changes.
- Boards know what tools are in use, what they are used for and when to say no.
- Chiefs have contracts that protect evidence in court, data ownership and the ability to exit.
- Frontline staff use AI as an aid to judgement, with simple checks that keep the original evidence in view.
- Affected communities have influence and a clear way to raise concerns, courts receive the disclosure they need, and smaller agencies no longer assess complex technology alone.

Agencies should track indicators of both effectiveness and trust and legitimacy, report quarterly to the board and review the framework annually.

7. Decision Guidelines for AI Adoption

Governance frameworks often emerge only after failures or controversies have damaged trust and narrowed the available options. The following questions, condensed from Appendix A of the Comprehensive Report, support anticipatory governance. They should be applied before procurement and whenever a tool's scope changes.

Area	Key questions
A. Justification and necessity	What operational or community safety problem does this solve, and for whom? Could non-AI alternatives achieve similar results? Is AI necessary or merely convenient, and is the scope proportionate? Could the tool improve efficiency while degrading trust, or shift effort from upstream prevention to downstream enforcement?
B. Evidence and validation	Has the tool been independently evaluated, and how? What are its known error rates and its accuracy across diverse Canadian populations? Who conducted bias testing, across which groups, and can the results be verified? Are foreign benchmarks being relied on inappropriately?
C. Human judgement	Where exactly does human oversight occur, and can the reviewer override the output? Are overrides documented? Have personnel been trained on automation bias and anchoring? Who is accountable if the output is wrong?
D. Legal and rights	Has a formal rights impact assessment been conducted? Does the tool aggregate or link datasets, or produce harms visible only at scale? Is its role in investigations discoverable, and will the vendor disclose system logic and training data and provide expert testimony if evidence is challenged?
E. Governance and transparency	Is this an administrative use or an operational use that directly affects the public, and which approvals and reviews apply? What prevents function creep, and who monitors scope? What will be publicly disclosed, and will the tool appear on the public register? How will the community be meaningfully involved?
F. Procurement and vendor risk	Who owns the vendor and the data? Where is data stored and processed, and where does it travel in transit? Can the vendor use it for training? What is the vendor's cybersecurity, legal and reputational record? Are exit, data portability, change-of-ownership, cost-control and knowledge-transfer protections in place? Can the system be independently audited, and will updates be notified and re-validated? If First Nations data are involved, how are relevant principles applied?
G. Monitoring and exit	Which indicators will track effectiveness, trust and legitimacy? What predefined conditions require a pause, restriction or withdrawal, and is there a mandatory review period? How will historical data be returned on exit? Do affected communities have a way to raise concerns and meaningful influence over decisions?

8. Further Study

Two issues require further study: Indigenous AI-related data sovereignty, and executive-level accountability for AI governance. Suggested next steps for chiefs and boards follow.

Indigenous data sovereignty

AI raises distinct and complex questions about data collected from, or generated about, Indigenous people and communities. The OCAP® principles (Ownership, Control, Access and Possession) for First Nations, and similar frameworks reflecting the concerns and aspirations of Inuit and Métis peoples, offer a strong starting point, but standardizing their application remains unaddressed.[32]

Further research is needed on how data sovereignty principles apply to and differ among First Nations, Métis and Inuit communities; on governance and partnership models for Indigenous people living in urban areas; on technical pathways to embed sovereignty rules in AI data architecture, including cloud storage, training datasets and access controls; and on meaningful co-design with Indigenous organizations and rights-holders.

Executive accountability for AI

Some management scholars suggest anchoring AI accountability in a Chief Compliance Officer who is part of the executive group, reports directly to the board, acts independently of the units using AI tools and consolidates decisions now spread across legal, privacy, IT security, procurement and professional standards.[33] For smaller agencies, this capacity might be provided through the shared evaluation body proposed in Recommendation #20. We note that this idea is largely conceptual and has yet to be assessed in police agencies.

9. Suggested Immediate Next Steps

The following steps are suggested starting points. Each is within the authority of the chief and board, and can begin without new legislation or major investment, and lays the groundwork for the Year-1 Core. Taken together, they allow an agency to move from strategic intent to action within the first months, and to show its community that AI is being governed deliberately rather than taken up piecemeal.

1. Chief and board confirm the five Year-1 Core measures and name the executive accountable for AI governance.
2. IT security begins the 90-day acceptable-use policy update.
3. The chief, with the board, begins defining the risk-classification tiers and zero-option conditions.

---

[32]First Nations Information Governance Centre (2019; 2021); Inuit Tapiriit Kanatami (2018); Métis National Council (2025)

[33]Oxéus, et al. (2026)

4. Legal counsel begins the standard contract template ahead of the next new or renewed AI agreement.
5. The chief places a shared evaluation body on the agenda of the next relevant chiefs' association meeting.
6. The service schedules an annual review of this guidance.

Police leaders know that public trust is built slowly and can be lost very quickly. The potential harms described in this report rarely arrive as a single crisis. Rather, they build-up quietly, through a series of decisions that each seems quite reasonable, until they are hard to undo. Acting now, beginning with the Year-1 Core, allows agencies to gain the operational benefits of AI while protecting the legitimacy and independence of Canadian public policing and the professional judgement of its members.

The AI agenda should be set by police leaders, not by vendors, and not by a crisis that forces their hand.

References

Adams, I.T. (2025). Automation and artificial intelligence in police body-worn cameras: Experimental evidence of impact on perceptions of fairness among officers. *Journal of Criminal Justice*, *97*, Article 102330. https://www.sciencedirect.com/science/article/pii/S0047235225000224

Angus Reid Institute. (2026, June 1). Canadians call for heavy AI regulation, but three-quarters doubt any government can keep up with the technology. https://angusreid.org/ai-regulation-canada/

Bhuiyan, J. (2026, August 18). OpenAI announces slowing pace of development after hack by rogue agent. *The Guardian*. https://www.theguardian.com/technology/2026/aug/18/open-ai-pause-hack

Booth, R. (2026, February 22). Met police using AI tools supplied by Palantir to flag officer misconduct. *The Guardian*. https://www.theguardian.com/uk-news/2026/feb/22/met-police-ai-tools-officer-misconduct-palantir

Brayne, S. (2017). Big data surveillance: The case of policing. *American Sociological Review*, 82(5), 997-1008. https://doi.org/10.1177/0003122417725865

Canada. (1982). Canadian Charter of Rights and Freedoms (Part I of the Constitution Act, 1982, being Schedule B to the Canada Act 1982 (UK), 1982, c. 11). Justice Laws Website. https://laws-lois.justice.gc.ca/eng/const/page-12.html

Canada. (1985). *Privacy Act*, R.S.C. 1985, c. P-21. Justice Laws Website. https://laws-lois.justice.gc.ca/eng/acts/P-21/

Canada. (2000). *Personal Information Protection and Electronic Documents Act*, S.C. 2000, c. 5. Justice Laws Website. https://laws-lois.justice.gc.ca/eng/acts/P-8.6/

Canadian Centre for Cyber Security. (2024, October 30). National cyber threat assessment 2025–2026. Communications Security Establishment. https://www.cyber.gc.ca/en/guidance/national-cyber-threat-assessment-2025-2026

Castaldo, J. & Posadzki, A. (2026, October 13). Scientists push for creation of international body to regulate artificial intelligence. *The Globe and Mail*. https://www.theglobeandmail.com/business/article-scientists-international-body-regulate-artificial-intelligence/

Cecco, L. (2026, October 13). AI godfather Geoffrey Hinton: 'We're creating alien beings. I'm willing to be selfish and choose humans'. *The Guardian*. https://www.theguardian.com/us-news/2026/oct/13/geoffrey-hinton-trump-ai-tech

Crawford, K. (2021). *Atlas of AI*. New Haven, CT: Yale University Press.

Du Sault, L. (2023, July 9). Under union pressure, Vallejo police chief ends body camera analysis. *Open Vallejo*. https://openvallejo.org/2023/07/09/under-union-pressure-vallejo-police-chief-ends-body-camera-analysis/

First Nations Information Governance Centre. (2019). First Nations data sovereignty in Canada. *Statistical Journal of the IAOS, 35*(1), 47–69. DOI 10.3233/SJI-180478

First Nations Information Governance Centre. (2021). The First Nations principles of OCAP®. First Nations Information Governance Centre. https://fnigc.ca/ocap-training/

Globe and Mail Editorial Board. (2026, February 27). Don't leave AI safety to the tech bros. *The Globe and Mail*. https://www.theglobeandmail.com/opinion/editorials/article-ai-regulation-tumbler-ridge/

Government of Canada, Ministry of Innovation, Science and Industry, & Ministry of Public Safety. (2022, May 19). Government of Canada announces intention to ban Huawei and ZTE from Canada's 5G networks [Policy statement]. https://www.canada.ca/en/innovation-science-economic-development/news/2022/05/statement-from-minister-champagne-on-telecommunications-security.html

Guan, J., Blanchard, T., Foerster, H., Jia, H., Huang, G., & Papernot, N. (2026). *AI agents enable adaptive computer worms* [Preprint]. arXiv. https://arxiv.org/abs/2606.03811

Huey, L., Ferguson, L., & Towns, Z. (2025). "Cops need doxxed": Releasing personal information of police officers as a tool of political harassment. *Crime & Delinquency*, *71*(3), 714–739. https://doi.org/10.1177/00111287241271069

Inuit Tapiriit Kanatami. (2018). *National Inuit strategy on research*. https://www.itk.ca/wp-content/uploads/2020/10/ITK-National-Inuit-Strategy-on-Research.pdf

International Association of Chiefs of Police (2026). IACP Connector. Washington, DC. https://www.theiacp.org/resources/iacp-connector

Internet Watch Foundation. (2026). Harm without limits: AI child sexual abuse material through the eyes of our analysts. Internet Watch Foundation. https://admin.iwf.org.uk/media/hl1nvdti/iwf_ai-CSAM-report-2026.pdf

Jacobs, G., van Houdt, F., & coons, g. (2024). Studying surveillance AI-cologies in public safety: How AI is in the world and the world in AI. *Surveillance & Society*, 22(2), 160–178. https://doi.org/10.24908/ss.v22i2.16104

Kasirzadeh, A. (2025). Two types of AI existential risk: Decisive and accumulative. *Philosophical Studies*, *182*(7), 1975–2003. https://doi.org/10.1007/s11098-025-02301-3

Ko v. Li, 2025 ONSC 2965 (CanLII). https://canlii.ca/t/kc6xx

Law Commission of Ontario, & Ontario Human Rights Commission. (2024). Human rights AI impact assessment (HRIA): Guide. https://lco-cdo.org/en/publications/human-rights-ai-impact-assessment-hria/

Lawrence, M., Homer-Dixon, T., Janzwood, S., Rockström, J., Renn, O., & Donges, J. F. (2024). Global polycrisis: The causal mechanisms of crisis entanglement. *Global Sustainability*, 7, Article e6, 1–16. https://doi.org/10.1017/sus.2024.1

Lockwood, C., Dobson, C., Phillips, C., & Pedersen, R. (2025, July 14). Artificial advocacy: How Canadian courts and legislators are responding to generative AI. Osler, Hoskin & Harcourt LLP.

https://www.osler.com/en/insights/updates/artificial-advocacy-how-canadian-courts-and-legislators-are-esponding-to-generative-ai/

Mehdi, T., & Morissette, R. (2024). Exposure to artificial intelligence in Canadian jobs: Experimental estimates (Analytical Studies Branch Research Paper Series, Catalogue No. 11F0019M). Statistics Canada. https://.www150.statcan.gc.ca/n1/pub/11f0019m2024005-eng.htm

Métis National Council. (2025). *Métis Nation data strategy*. https://www.metisnation.ca/wp-content/uploads/2025/12/Consolidated_Me_tis_Nation_Data_Strategy_Final_15_10_2025_3_56.pdf

Mosolova, D. (2026, February 27). How AI is supercharging Russia's online disinformation campaigns. BBC News. https://www.bbc.com/news/articles/cx2r7grrdwzo

MPS confirms the use of AI to spot potential misconduct. (2026, February 23). *Police Professional*. https://policeprofessional.com/?p=143282

Office of the Privacy Commissioner of Canada. (2021). PIPEDA Findings #2021-001: Joint investigation of Clearview AI, Inc. by the Office of the Privacy Commissioner of Canada, the Commission d'accès à l'information du Québec, the Information and Privacy Commissioner for British Columbia, and the Information Privacy Commissioner of Alberta. https://www.priv.gc.ca/en/opc-actions-and-decisions/investigations/investigations-into-businesses/2021/pipeda-2021-001/

Office of the United Nations High Commissioner for Human Rights (OHCHR), & Asia Pacific Forum of National Human Rights Institutions (APF). (2013). The United Nations Declaration on the Rights of Indigenous Peoples: A manual for national human rights institutions. United Nations. https://www.ohchr.org/sites/default/files/Documents/Issues/IPeoples/UNDRIPManualForNHRIs.pdf

Osler, Hoskin & Harcourt LLP (2026). *The Protecting Privacy and Consumer Data Act (Bill C-36): Key obligations and enforcement review.* June 2026. https://www.osler.com/wp-content/uploads/2026/07/Guide-to-the-PPCDA-Bill-C-36.pdf

Oxéus, G., Koplyay, T., & Gagnon, S. (2026). Chief Compliance Officers Navigating Ethical, Integrity, and Responsibility Challenges in Artificial Intelligence (AI) Projects and against Digital Disinformation. In *The Changing Role of the Chief Compliance Officer* (pp. 106-140). Routledge.

Parliament of Canada (2026). Bill C-36: An Act to enact the Protecting Privacy and Consumer Data Act, to amend the Personal Information Protection and Electronic Documents Act and to make amendments to other Acts. https://www.parl.ca/legisinfo/en/bill/45-1/c-36

Perrin, B. (2026). AI & policing: Roundtable report on the governance & use of artificial intelligence by police in Canada. UBC AI & Criminal Justice Initiative. https://canlii.ca/t/7nr4b

Perrin, B., Liew, G., & Sweeney, I. (2026). AI & policing: Research report on the governance & use of artificial intelligence by police in Canada (2026 CanLIIDocs 293). UBC AI & Criminal Justice Initiative. https://canlii.ca/t/7nrqc

R. v. Mohan, [1994] 2 S.C.R. 9. CanLII. https://www.canlii.org/en/ca/scc/doc/1994/1994canlii80/1994canlii80.html

R. v. Stinchcombe, [1991] 3 S.C.R. 326. CanLII. https://www.canlii.org/en/ca/scc/doc/1991/1991canlii45/1991canlii45.html

Sims, D. (2026, February 7). Waymo admits that its autopilot is often just guys from the Philippines. TechSpot. https://www.techspot.com/news/111233-waymo-admits-autopilot-often-guys-philippines.html

Standing Committee of the National People's Congress of the People's Republic of China. (2017, June 27). National Intelligence Law of the People's Republic of China (DiCarlo, Trans.). China Law Translate. https://www.chinalawtranslate.com/en/national-intelligence-law-of-the-peoples-republic-of-china-draft/

TechCabal. (2026, March 4). Kenyan workers say Meta Ray-Ban AI glasses expose intimate moments. https://techcabal.com/2026/03/04/kenya-meta-ray-ban-ai-intimate-moments/

Treasury Board of Canada Secretariat. (2023). Directive on automated decision-making. https://www.tbs-sct.canada.ca/pol/doc-eng.aspx?id=32592

Treasury Board of Canada Secretariat. (n.d.). Algorithmic impact assessment. Government of Canada. https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/automated-decision-making/algorithmic-impact-assessment.html

Truth and Reconciliation Commission of Canada. (2015). Truth and Reconciliation Commission of Canada: Calls to action. http://nctr.ca/assets/reports/Calls_to_Action_English2.pdf

Tuquero, L. (2026, September 2). AI agents are hacking systems without any input from humans. How did we get here? *PBS NewsHour*. https://www.pbs.org/newshour/science/ai-agents-are-hacking-systems-without-any-input-from-humans-how-did-we-get-here

United Nations Children's Fund. (2025). Artificial intelligence and child sexual abuse and exploitation: Emerging risks and policy considerations. UNICEF. https://www.unicef.org/media/178571/file/UNICEF%20AI%20CSEA%20Brief_2.pdf

United Nations Office on Drugs and Crime. (2025). Emerging threats: The intersection of criminal and technological innovation in the use of automation and AI. United Nations Office on Drugs and Crime. https://www.unodc.org/roseap/uploads/documents/Publications/2025/UNODC_Report_Emerging_threats_-_The_intersection_of_criminal_and_technological_innovation_in_the_use_of_automation_and_AI.pdf

Watson, S. (2023). Investigating the role of social media abuse in gender-based violence: The experiences of women police officers. *Criminology & Criminal Justice*, *23*(5), 829–844. https://doi.org/10.1177/17488958221087487

Watts, F., del Pozo, B., White, M.D. & Malm, A. (2025). Does automated feedback impact the acceptability of AI-generated police body worn camera review? An implementation science

natural experiment. *Journal of Experimental Criminology*. https://doi.org/10.1007/s11292-025-09711-7

Wischmeyer, T. (2019). Artificial intelligence and transparency: Opening the black box. In T. Wischmeyer & T. Rademacher (Eds.), *Regulating artificial intelligence*, Springer Nature (pp. 75-101). https://doi.org/10.1007/978-3-030-32361-5_4

World Economic Forum. (2025). The future of jobs report 2025. World Economic Forum. https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf

Acknowledgements

CSKA and the authors sincerely thank the distinguished panel of subject matter experts who generously reviewed and commented on an earlier version of this paper. Their candid insights and practical advice greatly strengthened the final report. Panel members took part in their personal capacities. Responsibility for the content rests with the authors.

We extend particular thanks to:

Alfonso Adamo	York Regional Police Service
Ron Anderson	Government of Alberta
Mara Cairo, P.Eng., PMP	Edmonton Police Service
Gideon Christian, PhD	University of Calgary
Christa Dicks, MBA, MEd	Cape Breton Regional Municipality
Nathan Dokis	Greater Sudbury Police Service
Gabrielle Jacobs, PhD	Erasmus University, School of Social and Behavioural Sciences, The Netherlands
Dale Johnstone	Saskatoon Police Service
Guy Leeson	Delta Police Department
Jamie Reynolds	Cape Breton Regional Police Service
Christiane Sadeler, MA	More Better Solutions
Howard Sapers	Canadian Civil Liberties Association
Trent Stevely	Regina Police Service
Tony Ventura, MA	Peel Regional Police Service

Report Authors

Janos Botschner, PhD

Janos is a social scientist who works at the intersection of issues related to community safety and well-being. He previously held senior positions in the broader public sector and academia, including as founding director of a community-based research institute, and chief researcher at a large branch of the Canadian Mental Health Association, where he was involved in the design and evaluation of policies and programs in the context of mental health transformation in Ontario. He has served on the editorial board of the Canadian Journal of Community Mental Health, and as a section chair of the Canadian Psychological Association.

He maintains an active practice related to the intersection of emerging technologies and community safety and well-being, with particular attention to cybersecurity and to the information environment. Over the past three decades, his mission has been to help build the capacity of communities and society to be safe, healthy, inclusive, supportive, and innovative.

Ritesh Kotak, JD, MBA

Ritesh is a technologist and legal advisor whose career uniquely bridges public safety, emerging technology, and innovation. He spent eight years in Canadian policing, where he co-developed a cybercrime unit and led public safety innovations, before transitioning to Big Tech to drive cybersecurity, IoT, and intelligence projects globally.

Currently practicing law, Ritesh advises clients on complex matters ranging from cyber breaches to corporate litigation. An accomplished subject matter expert, he has twice testified before Canadian Parliamentary Committees on digital competitiveness. His insights are frequently featured in mainstream global media, including *The Economist* and the *Financial Times*, and he has been recognized by IT World Canada as a leading cybersecurity specialist.

Ritesh holds a Juris Doctor from the University of Ottawa, an MBA with distinction from the University of Edinburgh, a Bachelor's of Business Administration from the University of Toronto, and completed the Emerging Leaders Executive Program at Harvard's Kennedy School.

Neil Dubord, PhD

Neil is a distinguished police and business leader with over 30 years of experience. He has served as Chief of the Delta Police Department, Chief of Metro Vancouver Transit Police, and Deputy Chief of the Edmonton Police Service. His exemplary service has earned him numerous honours, including the Queen's Platinum Jubilee Pin, multiple Police Exemplary Service Medals, and Officer of the Police Order of Merit.

A recognized thought leader, Neil has published widely on policing, leadership, and public trust. He holds a Ph.D. in Business Management, a Master's in Leadership and Training, and advanced certifications from the FBI National Academy and University of Virginia. Neil is an active member of the IACP, FBI National Academy, FBI-LEEDA, and the Canadian Association of Chiefs of Police.

Cal Corley, MBA

Cal is the CEO of the Community Safety Knowledge Alliance, a role he has held for the past 11 years. He is a former Assistant Commissioner of the Royal Canadian Mounted Police. From 2008–2014, he was head of the Canadian Police College and during that time served as the RCMP's Senior Envoy to Mexico and the Americas.

Over the course of his career, Cal gained extensive domestic and international experience in both operations and executive management, serving in such areas as community policing, major crime, drug enforcement, national security, criminal intelligence, strategic planning, human resources, and leading major reform initiatives. He also served on secondments at the Privy Council Office's Security and Intelligence Secretariat, and at Public Safety Canada.

About CSKA

The Community Safety Knowledge Alliance (CSKA) is a Canadian non-profit applied research and education organization dedicated to advancing community safety and well-being across Canada.

Our mission is to conduct independent, evidence-based research on multi-sectoral approaches to improving community safety, and to mobilize that knowledge so it strengthens the delivery and outcomes of community safety and human services for Canadians.

Independence and objectivity are central to CSKA's identity. Our work is built on upholding these values in every relationship with the governments, organizations, and communities we serve, allowing us to sustain a public-interest, evidence-based perspective alongside effective collaboration across sectors.

For more information on CSKA, please visit www.cskacanada.ca or contact CEO Cal Corley at (613) 297-6728 or by email at ccorley@cskacanada.ca.

AI Disclosure

In preparing this report, the authors used Perplexity for discovery, under human direction, to identify relevant literature. Claude was used during copy editing to identify typographical errors and inconsistencies and to flag opportunities to improve clarity.

The analysis, findings, and recommendations are the authors' own and reflect the review by CSKA's multidisciplinary expert panel. AI was not used to generate recommendations. The authors take full responsibility for the accuracy and integrity of this report.