EXECUTIVE INFOGRAPHIC · NOVEMBER 2026

AI and the Future of Canadian Policing

Implications and strategic considerations for chiefs of police, police boards and policymakers. A companion to the Executive Decision Report and the Comprehensive Report (October 2026).

Community Safety Knowledge Alliance
Research to Practice to Alignment

The bottom line

AI adoption is unavoidable: organized crime and other malicious actors are already using it. But AI is **not a neutral modernization step**. It reshapes what is visible, prioritized and acted upon. Treat AI as a **governed capability**: embedded within professional judgement, constrained by law and accountable to the communities police serve.

68% of Canadians say government must heavily regulate AI, even if it slows development

74% doubt any government can keep pace with the technology

21 recommendations across seven governance themes

5 Year-1 Core controls for every service, regardless of size

45 police services in Ontario alone, of vastly different capacity

3 evidence sources: literature synthesis, 7 expert interviews, expert panel review

1 WHY THIS IS A LEADERSHIP ISSUE NOW

Pressure to adopt
Staffing and fiscal pressure meet vendor promises. Start with the problem and ask "safety for whom?"

Public skepticism
Canadians doubt not only AI tools but institutions' capacity to oversee them.

Regulatory lag
Law and vendor auditing trail the technology; no mandatory platform reporting duty (cf. Tumbler Ridge, 2026).

Broader disruption
Job disruption amid a wider "polycrisis" may fuel insecurity and crime.

Without governance, harms accumulate quietly: erosion of trust · loss of professional judgement · disproportionate impact on racialized and equity-deserving communities · normalized surveillance · evidentiary failure in court · vendor dependence and loss of data sovereignty.

2 FOUR KEY FINDINGS

Benefits are real but conditional 1

- Transcription, drafting, triage, analytics, forecasting
- Value where AI augments judgement and solves a defined problem
- Vendor claims need independent validation

Ethical and practical risks are significant 2

- Bias reproduced from historical data (cf. Clearview AI)
- A human in the loop is not enough on its own
- Opacity, scope creep, hidden labour, workforce strain

New institutional and security risks 3

- Cyber exposure, shadow AI, vendor capture
- Foreign reach and data sovereignty (China, US)
- Decision shaping, practice drift, misaligned AI agents

AI-enabled threats are escalating 4

- Deepfake fraud; "crime-as-a-service"
- Synthetic CSAM, sextortion, online abuse
- Disinformation, hybrid threats, insider radicalization

3 TEN PRINCIPAL RISKS AND THE CONTROLS THAT ADDRESS THEM (# = Recommendation number; all 21 listed on page 2)

Rights and legitimacy

Bias and discriminatory impact
Bias audits; consultation; impact assessment; zero-option 4 8 18 21

Scope creep
Scoped approvals; re-assessment on change; public register 5 8

Loss of public trust
External panel; register; plain-language explanations; consultation 2 8 17 18 19

People and judgement

Automation bias
AI-literacy training; documented rationale; supervisory audits 14

Shadow AI
Acceptable-use policy with approved, secure alternatives 11

Insider radicalization
Screening indicators, set in consultation with associations and unions 16

Vendors and technology

Vendor capture and data sovereignty
Scored procurement criteria; ownership-change, cost and exit protections; knowledge transfer 9 10 12

Cyber compromise
Security criteria; minimum standards; joint sign-off before deployment 6 9 10

Evidence and capacity

Evidentiary and disclosure failure
Contract terms on disclosure and testimony; annual legal review; explanations 3 10 17

Capacity gaps and shared services
Vetted-tools registry; shared evaluation body; municipal agreements 7 13 20

4 LEGAL, RIGHTS AND ACCOUNTABILITY BOUNDARIES

Established obligations

- **Charter** ss. 7, 8, 9, 15: courts examine the technologies that shape decisions, not only outcomes
- **Privacy law**: vendor assurances do not relieve police of their own duties
- **Evidence**: AI outputs and prompts may be disclosable (*Stinchcombe*) and tested for reliability (*Mohan*). If vendors will not disclose, prosecution risk shifts to police and the Crown
- **s. 8**: predictive work that intrudes on privacy still needs judicial authorization; guard against drift

Benchmarks and evolving law

- Treasury Board *Directive on Automated Decision-Making*: a benchmark for municipal services
- Bill C-36 (June 2026): would replace Part 1 of PIPEDA, with a new enforcement commission
- Bill C-27 (AI and Data Act) and Bill C-26 (on cybersecurity) died on the Order Paper in 2025 but signal policy direction

Demonstrable public benefit

Efficiency alone is unlikely to justify AI that affects privacy. Leaders must explain how a tool advances safety and protects privacy and liberty. **Accountability remains with people, not systems.**

Continued overleaf: governance approach, roadmap and next steps →

From principle to practice: governing AI in policing

5 FIVE GOVERNING PRINCIPLES

01 Start with the problem, not the product
Documented problem statement and a "safety for whom?" test.

02 Match governance to risk
Administrative and public-facing uses are approved differently.

03 Approve within a defined scope
Any material change triggers re-assessment and approval.

04 Keep the zero-option real
Withdrawal is always possible, on conditions set in advance.

05 Be realistic
Controls actually implemented beat a framework that stalls.

6 THE YEAR-1 CORE: FOUNDATIONAL CONTROLS FOR EVERY SERVICE

1 Governance anchor
Named accountable executive, tool inventory, quarterly reporting.
Year 1 · Chief and board

5 Risk classification
Tiered scheme; board approval before high-risk procurement.
Year 1 priority · Chief with board

10 Standard contract
Ownership, disclosure, testimony, updates, cost and exit terms.
Year 1 · Legal counsel

11 Shadow AI policy
Prohibit unapproved tools; provide secure alternatives.
Within 90 days · IT security

14 AI literacy
Mandatory training on automation bias, with supervisory checks.
Within 12 months · Training

7 IMPLEMENTATION ROADMAP

Within 90 days
- Shadow AI policy 11

Within 6 months
- Evaluation-body working group 20

Within 12 months
- Year-1 Core 1 5 10 14
- Registry 7; feasibility 20

Next procurement
- Criteria 9; audits 4; pilot 21

Ongoing and annual
- Panel and legal reviews 2 3
- Training refresh 14; sequencing 15; engagement 19

8 ALL 21 RECOMMENDATIONS AT A GLANCE (# = Year-1 Core)

GOVERNANCE AND OVERSIGHT
1 Executive lead, tool inventory, quarterly reporting
2 External review panel for high-risk tools
3 Annual legal review of investigation protocols
4 Annual independent performance and bias audits

TRANSPARENCY
17 Plain-language explanation and complaint channel

DECISION GATES
5 Tiered risk classification; board approval for high risk
6 Joint legal, privacy and IT security sign-off
7 Join or help build a vetted-tools registry
8 Zero-option conditions and public register

ENGAGEMENT
18 Non-waivable consultation, medium and high risk
19 Budget line and lead for engagement

PROCUREMENT AND RISK
9 Problem statement; scored security and sovereignty criteria
10 Standard contract template
11 Acceptable-use policy prohibiting shadow AI
12 Vendor knowledge transfer
13 Agreements where IT or procurement is municipal

WORKFORCE
14 Mandatory AI literacy; automation-bias safeguards
15 Sequence high-risk deployments; board may pause
16 Radicalization indicators in insider screening

FRAGMENTED CAPACITY
20 Shared regional or provincial evaluation body
21 Pilot an impact assessment tool

9 CONTROLS SCALE WITH RISK

Every tool
Problem statement 9 · inventory and register 1 8 · vendor audits 4 · joint sign-off 6 · plain-language explanation 17

+ Medium risk
Structured community consultation 18 · impact assessment pilot 21. Consultation may raise a tool to high risk.

+ High risk
Board approval before procurement 5 · annual external panel review 2 · sequencing to capacity 15 · board may pause

10 SEVEN QUESTIONS TO ASK BEFORE ADOPTING ANY AI TOOL

A — Justification and necessity
What problem does it solve, and for whom? Could a non-AI option work?

B — Evidence and validation
Independently evaluated? Accurate across Canadian populations?

C — Human judgement
Can reviewers override? Who is accountable if it is wrong?

D — Legal and rights
Rights impact assessed? Will the vendor support disclosure and testimony?

E — Governance and transparency
Which approvals apply? What stops function creep?

F — Procurement and vendor risk
Who owns the vendor and the data? Exit, portability and OCAP® protections?

G — Monitoring and exit
What trust indicators? What conditions trigger pause or withdrawal?

11 SUGGESTED IMMEDIATE NEXT STEPS

1. Chief and board confirm the Year-1 Core and name the accountable executive
2. IT security begins the 90-day acceptable-use policy update
3. Chief, with the board, begins defining risk tiers and zero-option conditions
4. Legal counsel begins the standard contract template
5. Chief places a shared evaluation body on the next chiefs' association agenda
6. The service schedules an annual review of this guidance

For further study: Indigenous data sovereignty in AI; executive-level accountability for AI.

Public trust is built slowly and lost quickly. These harms rarely arrive as a single crisis; they build up quietly until they are hard to undo.

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

CSKA Community Safety Knowledge Alliance
Research to Practice to Alignment

Source: CSKA, *AI and the Future of Canadian Policing: Implications and Strategic Considerations*, Executive Decision Report (November 2026) and Comprehensive Report (October 2026). Authors: Janos Botschner, Ritesh Kotak, Neil Dubord, Cal Corley. Polling: Angus Reid Institute (June 2026).

cskacanada.ca
Cal Corley, CEO
(613) 297-6728 · ccorley@cskacanada.ca