COGNITIVE SECURITY
SÉCURITÉ COGNITIVE
CANADA

Perceive • Percevoir | Analyze • Analyser | Integrate • Intégrer

RESEARCH REPORT

Cognitive Security in Canada

Human-Systems Analysis, Civic Resilience, Public Trust, and the Digital Information Environment

Report number	COGSECCAN RR-002-EN
Publication date	May 2026
Prepared by	Cognitive Security Canada
Information class	Public-interest research

Perceive / Percevoir • Analyze / Analyser • Integrate / Intégrer

Document Control

Document type	Research Report
Report number	COGSECCAN RR-002-EN
Title	Cognitive Security in Canada: Human-Systems Analysis, Civic Resilience, Public Trust, and the Digital Information Environment
Publication status	Public Release
Prepared by	Cognitive Security Canada
Version	v1.0-EN
Notes	Research Report Master Template

Executive Summary

Canada's digital information environment is now a civic infrastructure. It shapes how residents receive public warnings, interpret political conflict, understand health and safety advice, access services, evaluate evidence, and decide whom to trust. Cognitive security is the public-interest practice of protecting the conditions that allow people and institutions to perceive, analyse, and integrate information with integrity. It is not a replacement for cybersecurity, public safety, media literacy, privacy law, or national security. It is a human-systems frame that connects those domains around a common question: how do information systems affect attention, interpretation, judgment, behaviour, and public trust?

The need for a national cognitive-security frame is visible across Canada. The Canadian Centre for Cyber Security's National Cyber Threat Assessment 2025-2026 states that Canada faces an expanding and complex cyber threat landscape, including state and non-state actors that target critical infrastructure and national security. The same assessment identifies combined cyber operations and online information campaigns as tools used to disrupt, divide, intimidate, and shape public opinion [1]. This makes the cognitive layer central to security: the intended effect is often not only data theft or system disruption, but a change in confidence, perception, and collective behaviour.

Foreign interference, domestic polarization, generative AI, synthetic media, platform recommender systems, targeted advertising, dark-pattern design, cybercrime, and institutional communication failures all affect how Canadians interpret reality. The Public Inquiry into Foreign Interference concluded that Canada's democratic institutions remain robust, while warning that disinformation is one of the most serious threats to democracy and public confidence [2]. That finding points to a practical policy gap. Canada has institutions for cyber defence, elections, privacy, policing, consumer protection, intelligence, and public communications, but it does not yet have a widely shared civic vocabulary for the human interpretive effects that connect them.

This report proposes cognitive security as a calm, non-partisan, civic and analytical framework for Canada. Its focus is not censorship, surveillance, or control of belief. Its focus is situational awareness, decision integrity, institutional transparency, and resilience against manipulation. A mature Canadian approach should protect freedom of expression and democratic disagreement while improving the public's ability to identify source, intent, authenticity, evidence quality, manipulative design, and information-system pressure.

The report recommends a national cognitive-security agenda built around five pillars: shared definitions, human-systems analysis, public-trust infrastructure, civic resilience, and accountable digital governance.

Implementation should be staged through public education, research protocols, institutional training, procurement standards, crisis communication practices, platform and AI transparency, and cross-sector partnerships. The objective is not to eliminate uncertainty. The objective is to improve Canada's capacity to recognize and manage uncertainty without degrading rights, trust, or public reasoning.

Key finding: Canada requires a national cognitive-security frame that links cybersecurity, foreign-interference awareness, AI governance, public communication, civic resilience, and human decision integrity without turning cognitive security into censorship, surveillance, or partisan messaging.

Content Map

Definition and Scope

For this report, cognitive security means the protection and strengthening of human perception, interpretation, judgment, and decision-making within complex information environments. It is concerned with how people and institutions understand signals, assign meaning, evaluate trust, and act under uncertainty. The concept includes protection against manipulation, but it also includes positive capacity-building: better public explanation, clearer institutional communication, stronger evidence habits, and more resilient civic sense-making.

A Canadian cognitive-security framework should be human-centered. It should treat residents, public servants, community leaders, journalists, researchers, and decision-makers as people embedded in systems of attention, incentives, design, emotion, culture, language, trust, history, and institutional power. The individual is not the weak link. The individual is the point at which systems become lived experience. A human-systems approach therefore studies the relationship between people and the systems that shape what they see, what they ignore, what they believe is normal, what they fear, and what they are able to decide.

This report is national in scope. It addresses Canada as a federation with federal, provincial, territorial, municipal, Indigenous, private-sector, academic, media, and civil-society actors. It does not assign cognitive-security responsibility to one department or one profession. The field sits between several existing domains: cybersecurity, national security, intelligence assessment, democratic resilience, AI governance, privacy, consumer protection, education, public health, emergency management, behavioral science, and public administration.

The report does not recommend policing thought, suppressing disagreement, or creating state authority over truth. Democratic life requires debate, dissent, advocacy, persuasion, satire, criticism, and uncertainty. The cognitive-security concern arises when information systems become opaque, asymmetrical,

manipulative, coercive, or so degraded that citizens cannot reasonably identify source, intent, evidence, authenticity, or risk. The public-interest objective is to preserve the conditions for independent judgment.

Research Questions and Method

This framing paper is guided by five research questions:

- What does cognitive security mean in a Canadian civic and institutional context?
- How do digital information systems affect perception, interpretation, trust, and decision-making?
- What risks emerge when cyber operations, disinformation, AI-generated content, platform design, and institutional communication failures interact?
- Which Canadian governance domains already address parts of cognitive security, and where are the gaps?
- What practical steps could Canada take to build civic resilience without weakening rights, pluralism, and open debate?

The method is analytic synthesis. The report draws on public Government of Canada cyber-threat material, the Public Inquiry into Foreign Interference, Canadian privacy and consumer-protection materials, Canadian AI governance documents, democratic-resilience literature, and selected international and academic work on information integrity and human-system risk. The report uses public sources only and does not make classified, investigative, medical, psychological, or legal findings.

The analysis is deliberately conservative. It distinguishes influence from manipulation, disagreement from disinformation, and vulnerability from incapacity. It also recognizes that trust cannot be demanded by institutions. Trust must be earned through accuracy, humility, responsiveness, transparency, and accountability. Cognitive security is therefore not only a threat framework. It is also a governance-quality framework.

Why Cognitive Security Matters in Canada

Canada's security environment is increasingly cognitive because many harms now operate through interpretation. A ransomware attack against a hospital, a data breach, a foreign influence campaign, a deepfake, or a confused emergency message can all produce technical, operational, and psychological effects at the same time. The public may ask: Is the information real? Who is responsible? Can I trust the institution? Am I safe? What should I do now? These questions are not secondary to security; they are part of security.

The Cyber Centre notes that online platforms and digital technologies mediate how Canadians work, shop, travel, socialize, get informed, and access critical services [1]. This means Canadians do not encounter information in a neutral space. They encounter it through feeds, search results, messaging apps, recommendation engines, targeted ads, platform incentives, workplace systems, service portals, influencers, and automated tools. These systems can improve access and efficiency, but they can also distort attention, amplify emotional content, segment publics, and make verification difficult.

Foreign interference adds a national-security dimension. The Public Inquiry into Foreign Interference found that foreign actors attempted to interfere in Canadian democratic processes and that disinformation can harm democratic confidence even when election outcomes are not changed [2]. This is important because the strategic target may be public trust itself. A hostile actor does not need to convince everyone of a falsehood. It may be enough to exhaust people, confuse evidence standards, intimidate communities, deepen suspicion, or make democratic institutions appear permanently compromised.

The domestic dimension is equally important. Cognitive insecurity can arise without a foreign actor. It can emerge from commercial manipulation, algorithmic amplification, poor institutional communication, low

media literacy, scam ecosystems, polarized identity conflict, opaque AI systems, or design practices that steer choices. The Government of Canada's Office of Consumer Affairs identifies dark patterns as design techniques that can mislead or pressure consumers online [4]. Such design issues show that cognitive security is also a daily-life issue, not only a national-security issue.

For Canada, cognitive security matters because it connects public trust to practical governance. A population that cannot distinguish credible evidence from synthetic manipulation is more vulnerable during emergencies. Institutions that communicate poorly create information vacuums. Platforms that optimize only engagement can reward outrage. Public agencies that deploy AI without explainability can weaken confidence even when tools are technically accurate. Cognitive security helps connect these problems within one human-systems frame.

Human-Systems Analysis

Human-systems analysis begins with a simple premise: people do not interpret information in isolation. They interpret it through memory, emotion, identity, prior experience, social trust, language, fatigue, fear, community norms, institutional credibility, and the design of the systems presenting the information. A cognitive-security approach therefore asks not only what information was provided, but how the surrounding environment shaped interpretation.

In a digital environment, the system is not only the message. It is the feed, the timing, the repetition, the ranking, the visual cue, the notification, the social proof, the missing context, the friction to verify, and the emotional state created by the interface. A person may believe they are freely choosing among visible options while a system has already narrowed what they see, made one option easier, made another harder, and attached social or emotional cues to the choice.

Human-systems analysis should not be used to excuse every decision or remove agency. Rather, it improves accountability by showing where individual choice meets environmental pressure. This is especially important for public policy. If a public communication fails because it assumes perfect attention, high literacy, stable trust, and unlimited time, the failure is partly systemic. If a platform design rewards sensationalism, the resulting public confusion is not simply a user weakness. If an AI tool produces confident but flawed explanations, the risk includes both technical error and human overreliance.

A useful human-systems model for cognitive security can be organized around six questions: What was made visible? What was hidden or made difficult to access? What emotional cues shaped interpretation? What social cues shaped perceived legitimacy? What choices were realistically available? What evidence would allow later review? These questions can be used by researchers, public communicators, investigators, procurement officers, educators, and policy analysts.

Analytical layer	Core question	Cognitive-security relevance
Perception	What information did the person actually encounter?	Determines the visible evidence environment.
Interpretation	What meaning was encouraged by framing, emotion, repetition, or identity cues?	Identifies how sense-making was shaped.
Choice architecture	Which actions were easy, difficult, hidden, or discouraged?	Shows whether decisions were meaningfully available.
Trust environment	Which institutions, messengers, or communities were treated as credible or suspect?	Links public confidence to information behavior.

Analytical layer	Core question	Cognitive-security relevance
Evidence trail	Can the exposure, design, source, and decision path be reviewed?	Enables accountability and learning.

Canada's Digital Information Environment

Canada's digital information environment is defined by high connectivity, platform dependence, data collection, AI-mediated content creation, cybercrime, and the speed of public controversy. The Cyber Centre's assessment states that Canada's state adversaries are becoming more aggressive in cyberspace and that state-sponsored cyber threat actors are almost certainly combining disruptive computer-network attacks with online information campaigns to intimidate and shape public opinion [1]. This is a direct bridge between cybersecurity and cognitive security.

Generative AI intensifies this environment by lowering the cost of producing persuasive text, images, audio, and video. Synthetic content can be used for creativity, accessibility, translation, and public education, but it can also support impersonation, fraud, harassment, fake evidence, and high-volume manipulation. AI does not create the need for cognitive security by itself; it accelerates existing vulnerabilities in attention, authentication, source verification, and trust.

Platform architecture also matters. Recommendation systems can make public life feel personalized while obscuring why particular content appears. Users may interpret repeated exposure as social reality, popularity, or evidence, even when it reflects engagement optimization or targeted distribution. The public-interest issue is not that ranking exists. Ranking is unavoidable in large information spaces. The issue is whether users, researchers, regulators, and institutions can understand and challenge high-impact ranking systems when they affect democratic participation, public safety, health behavior, or vulnerable communities.

Data-driven targeting adds another layer. Personal data, inferred preferences, emotional signals, location patterns, and social-graph information can be used to tailor messages. In ordinary contexts, personalization can be useful. In high-impact contexts, it can exploit vulnerability or fragment the public record. Different groups may receive different messages, different levels of fear, different promises, and different explanations. This makes common civic understanding harder to maintain.

Digital insecurity is also experienced through ordinary scams, fraud, ransomware, phishing, and identity theft. The Cyber Centre identifies cybercrime as a persistent and disruptive threat and notes that cybercrime-as-a-service supports a resilient global ecosystem [1]. These crimes are cognitive as well as technical: they use urgency, authority, fear, trust, shame, confusion, and social engineering to defeat judgment. A Canadian cognitive-security approach should therefore connect cyber hygiene with judgment hygiene.

Civic Resilience and Public Trust

Civic resilience is the ability of people and institutions to absorb information shocks, maintain lawful democratic disagreement, correct errors, and continue making decisions under uncertainty. It is not the absence of conflict. Democratic societies are supposed to contain disagreement. The resilience question is whether disagreement remains anchored to evidence, proportionality, human dignity, and institutional pathways for correction.

Public trust is not a communications product. It is a relationship. Institutions often attempt to restore trust by issuing more statements, but cognitive security requires attention to the conditions under which those statements are received. If communities have experienced neglect, discrimination, secrecy, contradiction, or administrative harm, then institutional messages may be interpreted through those histories. Effective communication must therefore be accurate, timely, humble, transparent about uncertainty, and responsive to public questions.

The foreign-interference inquiry is instructive because it highlights the relationship between threat communication and democratic confidence. Public warnings can inform citizens, but vague warnings can also create suspicion, stigma, or panic. Overstatement can damage trust; understatement can leave people vulnerable. Cognitive security requires calibrated communication that helps the public understand risk without manufacturing helplessness.

Civic resilience also requires local capacity. National institutions cannot interpret every local context. Municipal governments, Indigenous governments and organizations, schools, libraries, settlement organizations, community media, professional associations, and civil-society groups all help people interpret events. A national cognitive-security strategy should therefore support trusted intermediaries rather than centralizing all interpretation.

Media literacy remains important, but it is not enough by itself. Asking individuals to verify everything in a high-speed synthetic media environment can become unrealistic. Cognitive security should combine individual skills with platform accountability, institutional transparency, professional standards, accessible public data, and crisis communication design. The burden cannot rest solely on the user.

Threat / Risk Analysis

A cognitive-security risk is highest when an information mechanism is opaque, repeated, emotionally intense, identity-linked, difficult to verify, hard to refuse, and connected to a high-impact decision. The table below summarizes recurring risks for Canada.

Risk mechanism	Observable indicators	Potential public-interest harm	Governance concern
Cyber-enabled information operations	Data leaks paired with narratives; hack-and-leak activity; coordinated amplification.	Confusion, intimidation, distrust, and pressure on public institutions.	Cybersecurity, public communication, and democratic resilience must be linked.
Synthetic media and impersonation	Deepfake audio/video, cloned voices, fabricated screenshots, fake documents.	Fraud, reputational harm, false evidence, and crisis confusion.	Authentication, provenance, rapid correction, and public education.
Algorithmic amplification	Sudden virality, repeated outrage content, unexplained recommendations.	Distorted perception of public opinion and heightened polarization.	Platform transparency and research access.
Dark-pattern design	Obscured opt-outs, false urgency, confusing consent, forced continuity.	Degraded autonomy and consent in daily digital life.	Consumer protection and design accountability.
Institutional communication failure	Delayed, vague, defensive, or inconsistent public messaging.	Information vacuums and declining confidence.	Crisis communication standards and after-action review.
Targeted harassment or transnational repression	Diaspora intimidation, doxxing, surveillance, coordinated abuse.	Silencing of communities and reduced democratic participation.	Protection, reporting pathways, and community trust.

Risk analysis should avoid treating all influence as hostile. Public health guidance, emergency alerts, education campaigns, journalism, political advocacy, and community organizing all seek to influence behavior. The difference lies in legitimacy, transparency, evidence, consent, accountability, and proportionality. Cognitive security is concerned with influence that is hidden, coercive, exploitative, deceptive, or structurally harmful to autonomous judgment.

The most difficult challenge is evidence. Cognitive-security harms often involve cumulative exposure rather than a single event. A person may be influenced by hundreds of posts, messages, design cues, and social pressures. Legal and policy systems need evidence protocols that preserve content, timing, source indicators, interface design, user interactions, platform responses, and institutional decisions. Without evidence, cognitive-security claims can become either impossible to prove or too easy to exaggerate.

A Canadian approach should therefore combine humility with documentation. Not every perceived manipulation will be provable. Not every false claim will be strategic. Not every strategic influence operation will succeed. But systematic observation can reveal patterns: repeated narratives, coordinated accounts, targeted communities, design pressure, synthetic artifacts, and failures in institutional response.

Governance Landscape

Canada already has several governance tools relevant to cognitive security. Cybersecurity institutions address malicious cyber activity and resilience. Elections institutions protect electoral administration and public confidence. Privacy regulators address personal information and consent. Consumer-protection and competition bodies can address misleading practices and manipulative design. Public Safety and national-security agencies address foreign interference. ISED and Treasury Board policy tools address AI and digital government. The gap is not complete absence of governance; it is fragmentation.

The Artificial Intelligence and Data Act companion document described a proposed federal approach to high-impact AI systems, including obligations around risk mitigation, monitoring, transparency, and record keeping [3]. Regardless of legislative status, these policy categories are relevant to cognitive security because high-impact AI can affect how people are assessed, informed, persuaded, served, or excluded. Cognitive-security analysis should be included wherever AI systems shape human judgment or institutional decisions.

Consumer protection is also relevant. The Office of Consumer Affairs' public guidance on dark patterns identifies design practices that can steer, pressure, or mislead users [4]. This is a practical example of cognitive security in everyday life. A confusing cancellation flow, manipulative privacy prompt, or false urgency countdown may not look like national security, but it normalizes degraded consent and weakens public expectations for transparent digital design.

Foreign interference law and policy address a different part of the landscape. The Countering Foreign Interference Act received Royal Assent in 2024 and includes measures such as a foreign influence transparency framework [5]. This kind of tool may improve transparency around certain foreign-linked activities, but cognitive security remains broader. It also includes domestic manipulation, commercial design, AI-mediated persuasion, public communication, and institutional trust.

The governance challenge is to create interoperability among these domains without creating an overbroad concept. Cognitive security should not become a label applied to every controversial speech issue. It should be used where there is a recognizable human-systems concern: manipulation of perception, degradation of decision conditions, exploitative information architecture, targeted intimidation, or loss of public capacity to evaluate evidence.

Domain	Current contribution	Cognitive-security gap	Practical path
Cybersecurity	Threat assessment, incident response, critical-infrastructure protection.	Often focuses on systems more than public interpretation.	Integrate public-trust effects into cyber incident planning.

Domain	Current contribution	Cognitive-security gap	Practical path
Foreign interference	Threat awareness, legal tools, election safeguards.	Can be communicated too narrowly as an election-only issue.	Include diaspora protection, disinformation resilience, and local reporting pathways.
Privacy	Consent, accountability, data protection.	Data use may be lawful while influence effects remain opaque.	Treat sensitive inferences and persuasion profiles as higher-risk.
Consumer protection	Misleading practices and dark-pattern awareness.	Subtle manipulation may not always fit existing categories.	Develop enforceable standards for manipulative design.
AI governance	Risk mitigation, transparency, monitoring, record keeping.	Human judgment effects may be under-specified.	Add influence-impact and decision-integrity assessments.
Public administration	Communications, service design, procurement.	Trust impacts are often evaluated after harm occurs.	Build cognitive-security criteria into procurement and communications review.

Recommendations for Canada

Recommendation	Lead actors	Timeline	Complexity
Adopt a national working definition of cognitive security focused on perception, interpretation, judgment, decision integrity, and public trust.	Public Safety, Heritage, ISED, CSE/Cyber Centre, civil society, academic partners.	Short term	Medium
Create a human-systems assessment guide for public communications, AI tools, platform partnerships, and crisis response.	Treasury Board, federal departments, provinces, municipalities.	Short term	Medium
Add cognitive-security criteria to cyber incident response, including public-trust impact, misinformation risk, and clear correction channels.	CSE/Cyber Centre, Public Safety, critical infrastructure operators.	Short to medium term	Medium

Recommendation	Lead actors	Timeline	Complexity
Develop evidence protocols for information manipulation, including content preservation, provenance, timing, amplification, and design analysis.	Researchers, regulators, legal professionals, digital forensics experts.	Medium term	High
Strengthen public education around source, intent, authenticity, evidence, uncertainty, and manipulative design.	Schools, libraries, civic organizations, media, public agencies.	Short to medium term	Medium
Require high-impact AI and digital-service deployments to assess effects on human judgment, explanation, appeal, and trust.	ISED, Treasury Board, procurement authorities, regulators.	Medium term	High
Support trusted community intermediaries, especially for diaspora communities and groups targeted by intimidation or manipulation.	Federal/provincial governments, municipalities, community organizations.	Medium term	Medium
Create a Canadian cognitive-security research network linking law, psychology, public administration, cyber, AI, communications, and civic education.	Universities, think tanks, civil society, public-interest research bodies.	Medium term	Medium

These recommendations are designed to be incremental. Canada does not need to create a single cognitive-security super-agency. It needs a common vocabulary, shared assessment practices, and better coordination among existing institutions. The most important early step is definitional discipline: cognitive security should be framed as decision integrity and situational awareness, not as control of belief.

Public communications should be treated as critical infrastructure during crises. This means preparing plain-language explanations, uncertainty statements, correction mechanisms, community-specific outreach, and rapid myth-response capacity before incidents occur. A communication failure can intensify a technical incident by creating fear, speculation, and distrust.

Procurement is another practical lever. Governments purchasing AI tools, social media monitoring services, digital identity systems, education technologies, or public communication platforms can require cognitive-security criteria: explainability, audit logs, user control, accessibility, bias review, privacy protection, manipulation safeguards, and independent evaluation.

Implementation Timeline

Phase	Period	Priority actions	Expected output
Phase 1: Vocabulary and awareness	2026	Publish public primers; convene stakeholders; define cognitive security in Canadian civic terms.	Shared language and non-partisan public framing.
Phase 2: Human-systems tools	2026-2027	Develop assessment checklists for communication, AI, cyber incidents, and platform risk.	Practical tools for institutions and researchers.
Phase 3: Evidence and training	2027	Create documentation protocols; train public communicators, investigators, educators, and analysts.	Repeatable evidence and improved response quality.
Phase 4: Procurement and standards	2027-2028	Add cognitive-security criteria to digital procurement and high-impact AI review.	System-level accountability before deployment.
Phase 5: Institutionalization	2028 onward	Build durable research partnerships and public reporting mechanisms.	Sustained national capacity and civic resilience.

The timeline is staged because cognitive security depends on trust. Moving too quickly toward enforcement without shared definitions could create confusion or concern. Moving too slowly leaves Canada exposed to manipulative design, synthetic media, foreign interference, cyber-enabled intimidation, and institutional trust failures. The best path is disciplined, transparent, evidence-based development.

Research Gaps and Next Steps

Several research gaps require attention. First, Canada needs better evidence on how different communities experience information manipulation, institutional trust, and digital vulnerability. National averages can hide local realities. Second, Canada needs stronger methods for measuring the cognitive effects of cyber incidents and public communication failures. Third, more work is required on AI-mediated persuasion, especially where automated systems generate customized explanations or emotional appeals. Fourth, there is a need for practical evaluation tools that non-specialist organizations can use without adopting overly technical language.

A Canadian research agenda should include public-interest case studies, bilingual and multilingual public education, community consultation, platform-transparency research, procurement analysis, and ethics review for human-subjects research. The field must remain careful with claims about mental states. Cognitive security should study observable mechanisms and documented effects rather than speculative assertions.

For Cognitive Security Canada, the next step is to convert this framing paper into three companion products: a two-page public awareness brief, a human-systems assessment checklist, and a research-development roadmap for partnerships with academic, civic, public-sector, and community organizations. These products should preserve the same calm framing: perceive, analyze, integrate.

Priority next steps:

- Create a Canadian cognitive-security glossary for public and institutional use.
- Draft a human-systems assessment checklist for digital influence and public communication.
- Develop a bilingual public awareness brief on cognitive security and civic resilience.
- Map Canadian governance responsibilities across cyber, AI, privacy, consumer protection, elections, and public safety.
- Identify research partners for ethical, non-partisan study of digital information environments and public trust.

Conclusion

Cognitive security provides Canada with a practical way to understand the human layer of the digital information environment. It does not replace cybersecurity, privacy, media literacy, public safety, or democratic rights. It connects them through the lived reality of interpretation: what people see, what they believe is credible, how they make decisions, and whether institutions preserve the conditions for independent judgment.

Canada's challenge is not only hostile content. It is the interaction of platforms, AI systems, cyber operations, scams, institutional communication, social trust, and human vulnerability. A resilient Canada needs people who can evaluate uncertainty, institutions that communicate honestly, digital systems that are accountable, and communities that are supported rather than blamed.

A national cognitive-security framework should therefore be civic, transparent, rights-respecting, and evidence-based. Its purpose is to promote situational awareness, not to control belief. Its measure of success is not uniform agreement, but stronger public capacity to perceive, analyze, and integrate information under pressure.

References

[1] Canadian Centre for Cyber Security. National Cyber Threat Assessment 2025-2026. Government of Canada, 2024. https://www.cyber.gc.ca/en/guidance/national-cyber-threat-assessment-2025-2026

[2] Public Inquiry into Foreign Interference in Federal Electoral Processes and Democratic Institutions. Final Report. Government of Canada, 2025. See also Associated Press reporting on the final report, January 28, 2025.

[3] Innovation, Science and Economic Development Canada. The Artificial Intelligence and Data Act (AIDA) - Companion document. Government of Canada. https://ised-isde.canada.ca/site/innovation-better-canada/en/artificial-intelligence-and-data-act-aida-companion-document

[4] Government of Canada, Office of Consumer Affairs. Dark patterns. Innovation, Science and Economic Development Canada. https://ised-isde.canada.ca/site/office-consumer-affairs/en/dark-patterns

[5] Parliament of Canada. Bill C-70, Countering Foreign Interference Act, S.C. 2024, c. 16. Royal Assent, June 20, 2024.

[6] Office of the Privacy Commissioner of Canada. Privacy guidance, research, and policy materials on personal information, online influence, and digital trust. https://www.priv.gc.ca/

[7] Communications Security Establishment Canada. Cyber Threats to Canada's Democratic Process: 2023 Update. Government of Canada.

[8] Competition Bureau Canada. Deceptive marketing practices and Competition Act enforcement guidance. Government of Canada.

[9] Treasury Board of Canada Secretariat. Directive on Automated Decision-Making and Algorithmic Impact Assessment materials. Government of Canada.

[10] Canadian Radio-television and Telecommunications Commission and Government of Canada digital policy materials relevant to online platforms, communications, and public-interest regulation.

Suggested Citation

Cognitive Security Canada. (2026). Cognitive Security in Canada: Human-Systems Analysis, Civic Resilience, Public Trust, and the Digital Information Environment. Research Report COGSECCAN RR-002-EN.

Disclaimer

This report is intended for public-interest research and educational purposes. It does not constitute legal, medical, psychological, cybersecurity, procurement, financial, or professional advice. Findings are based on cited public sources and analytic synthesis at the time of publication.

Contact: info@cogsec.ca