ASIA PACIFIC FOUNDATION OF CANADA
FONDATION ASIE PACIFIQUE DU CANADA

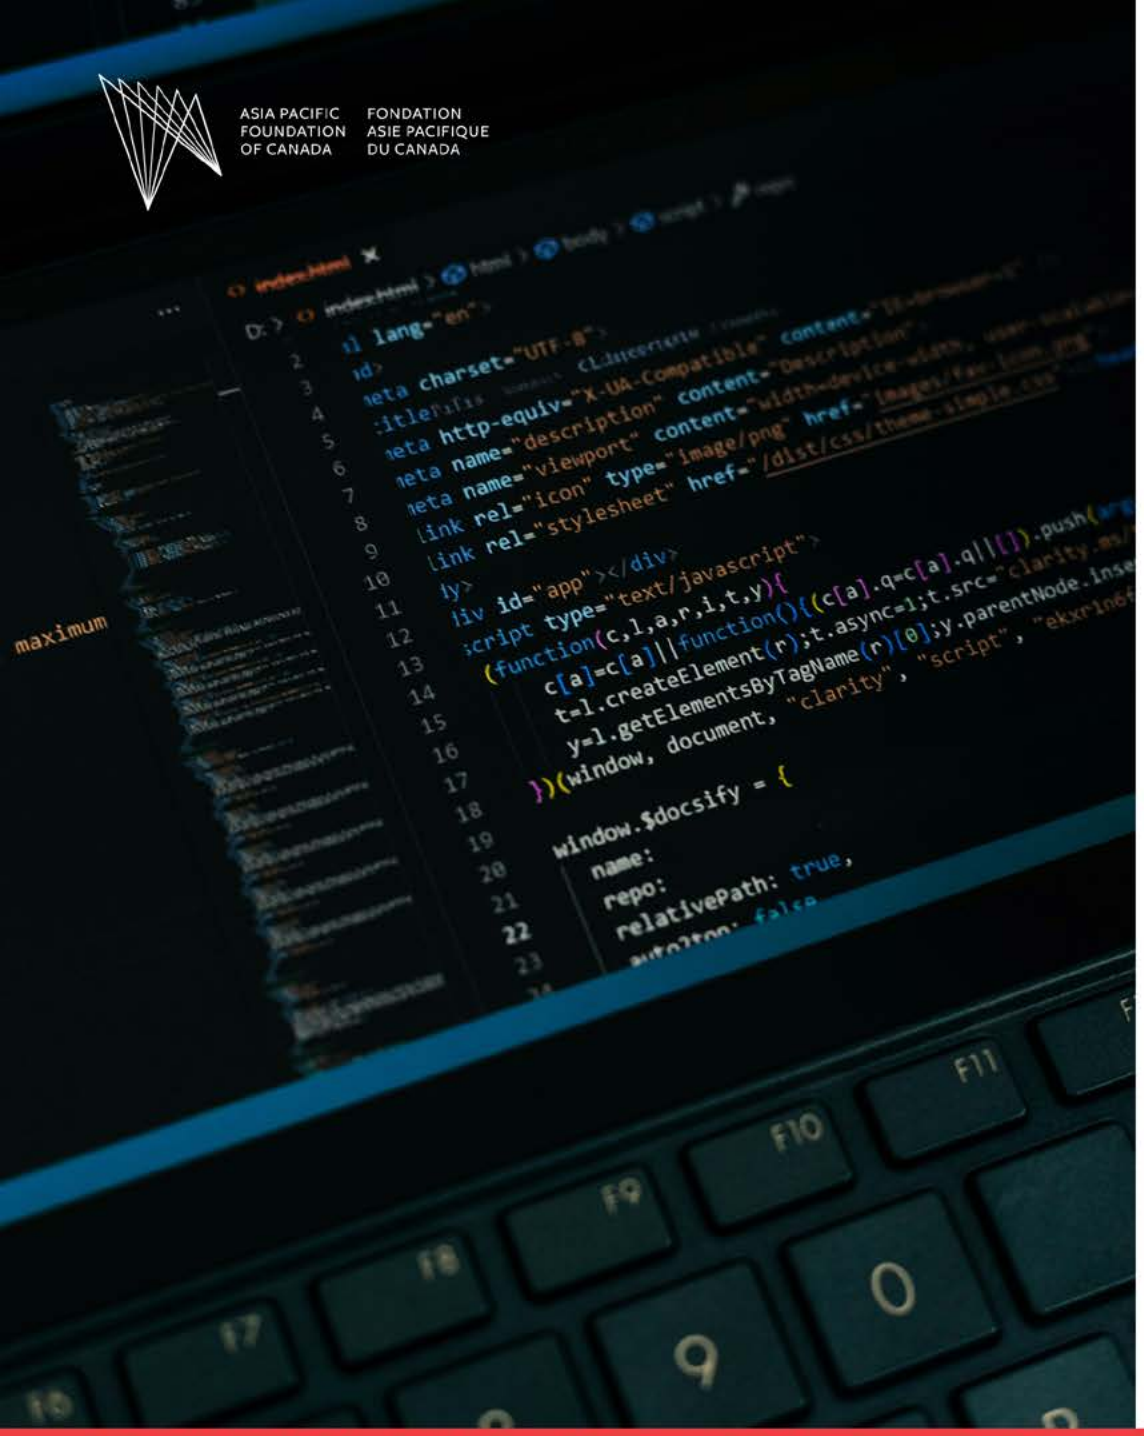

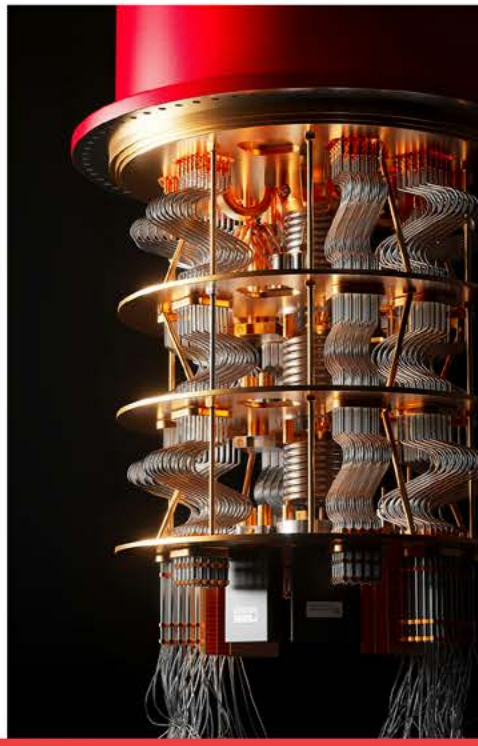

POLICY BRIEF 3: CANADA–SOUTH KOREA TECHNOLOGY SECURITY

Securing Open-Source Software Supply Chains Through Harmonization

BY TAE YEON EOM

Introduction

This series of three policy briefs examines where technology and security meet in the Canada–South Korea relationship. It follows that intersection across three domains: artificial intelligence, quantum technology, and open-source software supply chain security.

The briefs arrive at a pointed moment. On July 6, 2026, Canada selected Germany's ThyssenKrupp Marine Systems as the preferred supplier for its landmark 12-submarine procurement program, moving past South Korea's Hanwha Ocean. The decision opened the next chapter in Seoul's strategy toward the Canadian defence market. Spanning AI, quantum technology, and cybersecurity, the two countries' technology partnership already extends well beyond the Canadian Patrol Submarine Project. It rests on habits of co-operation that no single contract can define or take away.

The first brief defines military AI through the lens of cognitive sovereignty. Control over the software layer dictates battlefield decisions. It proposes concrete mechanisms: a non-U.S. AI chip development alliance, jointly mandated open-architecture standards, a bilateral AI security testing exchange, and an Arctic-Pacific physical AI laboratory. Canadian research strength and South Korean industrial scale combine through these channels into a strategic asset.

The second brief turns to quantum technology. Canada's foundational science and South Korea's commercial deployment capacity form a natural complement. Quantum computing will eventually break today's encryption, exposing defence and financial communications now in transit. The brief charts a path toward a formally recognized Canada–South Korea quantum security standard, anchored in defence industrial benefit obligations. Together, the two countries can become architects of the global quantum standard as the technology matures.

The third and concluding brief confronts the exposed foundations beneath both AI and quantum efforts. Open-source software supply chains form the invisible infrastructure of this entire emerging security architecture. This critical foundation is already under attack. In June 2026, a North Korea-linked group breached 144 code packages in an 88-minute window, highlighting how fast a single point of compromise can cascade through allied systems. The brief calls for a Canada–South Korea federated threat-intelligence bridge, harmonized software bill-of-materials standards, and a joint red-team testbed to close that gap before the next breach outpaces detection.

Together, the three briefs offer a working blueprint for Canada–South Korea technology co-operation. The author is grateful to the officials, researchers, and practitioners across both countries who shared their time and insight. This research was made possible by the Korea Foundation's Policy-Oriented Research Program, with the Asia Pacific Foundation of Canada providing institutional support and serving as the trusted Canadian platform for carrying this co-operation forward.

These pages stand as a lasting intellectual asset, binding the technology security of both nations for years to come.

Securing Open-Source Software Supply Chains Through Harmonization

Introduction: The Frontline Has Moved to the Code

In May and June 2026, two major cyber intrusions converged on the same target: the open-source software supply chain. One strike compromised Canadian government-linked infrastructure, while the other was tied to North Korea's weapons financing. This shared code infrastructure is what the world's AI and quantum technology now depends on, and it remains its most exposed point of failure.

Open-source software allows public downloading, inspection, modification, and redistribution without cost. This differs from proprietary, commercial products such as Microsoft Word or Adobe Acrobat. Prominent examples include AI models such as Meta's Llama, Mistral, DeepSeek, and Cohere's own North Mini Code, alongside foundational building blocks like the PyTorch machine-learning library and the Linux operating system kernel. This free, modifiable structure explains its rapid global adoption. It also explains why no single company is responsible for securing the entire stack.

On June 17, 2026, a hacking group tied to the North Korean state, tracked by Microsoft as Sapphire Sleet and linked to a similar attack on a popular web-development tool two months earlier, struck npm. As the world's largest public repository of pre-written software building blocks, npm is owned by GitHub (acquired by Microsoft in 2018), making this an attack on Microsoft's own platform. Millions of developers download and reuse these components to avoid coding from scratch. The process resembles a contractor drawing standard parts from a shared catalogue. Repositories such as npm and GitHub are not peripheral tools; they are the heart of the modern software supply chain. Their openness makes them invaluable and dangerous simultaneously. The accessibility that lets any developer reuse code instantly also allows a single malicious component to spread to thousands of downstream systems just as fast.

The accessibility that lets any developer reuse code instantly also allows a single malicious component to spread to thousands of downstream systems just as fast.

In a single 88-minute window, the attackers exploited precisely that openness. They secretly inserted malicious instructions into 144 building blocks inside Mastra AI, a widely used AI development toolkit. Anyone who installed one of these components during that window unknowingly handed over cloud passwords and digital access keys to a state-sponsored intelligence operation. Microsoft confirmed the attack two days later, matching the malware's code and infrastructure to previously

catalogued Sapphire Sleet campaigns with a high degree of probability. This is not an isolated case. Researchers and the United Nations Panel of Experts have traced these operations to North Korea's missile and nuclear weapons programs. This is a direct concern for South Korea and, through the alliance commitments examined in this brief, for Canada as well.

Canada has encountered a parallel pattern. On May 18, 2026, attackers compromised a GitHub employee's device using a tampered, malware-laced version of a popular coding tool, the Nx Console extension (built by Nrwl, a company headquartered in Toronto). Attackers used that foothold to copy roughly 3,800 internal GitHub repositories containing proprietary source code. It was the digital equivalent of breaking into a company's archive and photographing its blueprints. The Canadian Centre for Cyber Security (CCCS) issued a direct alert to Canadian IT professionals 11 days later, instructing them to change compromised passwords and inspect the automated systems that assemble their software.

This global vulnerability extended far beyond Canada. The CCCS became involved because Canadian critical infrastructure and financial systems heavily rely on these exact open-source packages. The significance for Canadians is structural: an attack on a developer tool built anywhere in the world can instantly compromise domestic banking or government networks that depend on it, irrespective of the attackers' specific geographic intent. Canadian financial institutions are especially exposed, as banks and payment processors increasingly build on the same open-source software components described in this brief.

While discussions of technology security often focus on export controls and hardware access, the AI and quantum policy specialists interviewed for this brief consistently named a different and more urgent vulnerability: the open-source software, including freely shared code libraries, frameworks, and components. AI models and quantum simulators are built from these tools, often without anyone in the chain verifying what is actually inside. A single tainted component can now reach thousands of downstream organizations within hours through routine, automated software updates alone.

The stakes are compounded by the dual-use nature of both technologies. The same open-source AI models that power benign commercial applications can be repurposed for offensive cyber operations, disinformation, or autonomous weapons targeting; the same open-source quantum simulators used for pharmaceutical or materials research are the tools a state would use to break current encryption standards. A supply chain compromise in either domain therefore carries consequences well beyond the immediate breach. Beyond mere data extraction, it grants adversaries a strategic foothold inside the tools shaping next-generation military and intelligence capability.

This brief examines two fronts where that risk is concentrated: the open-source infrastructure now driving frontier AI development, and the far less scrutinized open-source code beneath quantum computing research. It then sets out a co-operation agenda for Canada and South Korea, building on the defence and cyber framework the two governments put in place earlier this year.

Open Source and the New Politics of AI Sovereignty

Openly licensed models have moved from the margins of AI development to its mainstream. Drawing on Epoch AI's database of notable AI models, South Korea's Software Policy & Research Institute found that 47.3 per cent of all notable models released since 2018 have been open source, and the share keeps rising. The infrastructure supporting that shift is now substantial. GitHub hosted roughly 4.32 million AI-related projects as of 2024 and Hugging Face passed 2.25 million openly published models by the end of 2025. Across the enterprise sector, 89 per cent of companies now use open-source components in their AI development pipeline, citing innovation, lower cost, and de facto market standards.

Leadership of that innovation stack is shifting. Meta's Llama anchored the first wave of Western open models, while Chinese firms, notably DeepSeek and Alibaba's Qwen, have used aggressive licensing and strong price-performance to pull global developer attention since 2025. This shift carries serious security implications. When Western developers adopt high-performing Chinese

open models to cut costs, they inherit a dependency on code whose training data, safety tuning, and long-term maintenance are shaped by firms operating under Beijing's jurisdiction. This introduces a category of concentrated foreign control that poses a separate risk, independent of the software's licensing.

The strategic cost of dependence on centralized providers materialized decisively in mid-June 2026. The U.S. restricted foreign access to Anthropic's most capable models, Mythos and Fable, under an export-control directive. This abruptly cut off allied governments and enterprises that had built workflows around them. Cohere's chief executive told *Fortune* the episode demonstrated that centralized dependence on a single entity is a structural risk. He used a G7 leaders' gathering the same week to argue that democracies can no longer afford to rent their AI capability from centralized providers.

Canada has already begun acting on that logic. In February 2026, it signed a Joint Declaration of Intent on Artificial Intelligence with Germany and launched a Sovereign Technology Alliance to pool capability among trusted partners. Cohere's US$20-billion combination with Germany's Aleph Alpha has since turned this into a commercial reality, with Cohere leading the combined entity from Toronto. For South Korea, the lesson generalizes beyond any single vendor: dependence on AI capability that a foreign government can switch off at will is a supply chain exposure in its own right, regardless of licensing terms.

China's own posture is shifting in ways that could reshape this calculus. In July 2026, Reuters and the Financial Times reported that China's Ministry of Commerce had opened consultations with Alibaba, ByteDance, and Zhipu on export controls covering AI model weights, training data, and chip designs. Discussions reportedly include a tiered regime that could completely restrict foreign downloads of frontier open-weight models.

This move responds to the same dynamic driving Western sovereign-AI strategy. CNBC reported in July 2026 that the share of tokens used by U.S. companies on Chinese AI models via OpenRouter, a major AI model marketplace, has held above 30 per cent every week since February 2026, peaking at 46 per cent. That compares with an average of just 11 per cent over the prior twelve months, and 4.5 per cent in the first half of 2025.

TABLE 1

US$ per Million Tokens: Selected U.S. and Chinese AI Models

MODEL	DEVELOPER	COUNTRY	INPUT (US$)	OUTPUT (US$)
GPT-5.5	OpenAI	U.S.	$5.00	$30.00
Claude Opus 5	Anthropic	U.S.	$5.00	$25.00
Gemini 3.1 Pro	Google	U.S.	$2.00	$12.00
Gemini 3.5 Flash	Google	U.S.	$0.75	$4.50
Kimi K3	Moonshot AI	China	$3.00	$15.00
GLM-5.2	Z.ai (Zhipu AI)	China	$1.40	$4.40
Qwen3.5 397B	Alibaba	China	$0.60	$3.60
DeepSeek V4 Pro	DeepSeek	China	$0.44	$0.87
DeepSeek V4 Flash	DeepSeek	China	$0.14	$0.28

If Beijing follows through, North American startups that adopted Chinese open-weight models for cost reasons could face an abrupt loss of access. This scenario would mirror, in geographically inverted form, the disruption Anthropic's export controls inflicted on allied users of Mythos and Fable.

These two risks are distinct and compounding. The open-source structure lets state-backed attackers compromise code that lacks centralized oversight, as the npm incident showed. Dependence on foreign-controlled providers creates a separate exposure to access revocation by foreign policy directive, as the Anthropic export controls showed. Avoiding one does not neutralize the other.

Where Sovereignty Meets Its Limit

Against this backdrop, Cohere's own open-weight release illustrates both the appeal and the limits of the open-source route. On June 9, 2026, the company released North Mini Code, a 30-billion-parameter, Apache 2.0-licensed coding model built to run on a single Nvidia H100 GPU. This lets enterprises deploy agentic coding assistants entirely on premises, avoiding the need to send proprietary source code to a third-party API. The release is as much a security statement as a product launch. As coding agents move from writing snippets to autonomously browsing repositories and executing terminal commands, handing that level of access to a closed, foreign-hosted model becomes a critical supply chain decision. For South Korea, treating open-source AI as sovereign infrastructure is becoming a necessity, providing a natural point of convergence with Canada's approach.

Sovereignty achieved through open source is only real if the supply chain is defended.

This is where the sovereignty strategy meets its limit. The moment a government trades a closed foreign model for an open one, it inherits the open-source supply chain exposure itself. This is the exact supply chain that carried North Korea's malicious code into 144 AI packages in 88 minutes. Sovereignty achieved through open source is only real if the supply chain is defended.

This dynamic does not suggest South Korea should abandon open-source AI to build exclusively in-house systems. Homegrown code inherits the identical open-source supply chain risk the moment it relies on shared libraries and components. Whoever builds AI, domestically or abroad, must secure the open-source foundation beneath it. Self-reliance changes who controls the model, not whether the foundational code is safe. The more realistic path is building rigorous vetting mechanisms for open-source components and reserving fully domestic development for only the most critical strategic nodes, an approach that builds resilience without the cost of total technological autarky.

Why a Single Organization Can No Longer Defend Itself Alone

"Organization" here is used broadly. Modern software is assembled, not built. This applies equally to corporate, governmental, and open-source projects. A typical application pulls in dozens or hundreds of external libraries, plugins, and build tools, few of which are ever reviewed line by line before being trusted. Attackers have adapted accordingly. Instead of breaching a target directly, attackers compromise the credentials of an open-source maintainer or slip a malicious commit past review. They then let the package manager's own installation process carry the payload automatically to every downstream user. Incidents involving compromised GitHub Actions workflows and poisoned upstream packages show how a single corrupted dependency can expose confidential data and hand over system control across an entire software environment. The CCCS terms this a cascading supply chain attack.

The defensive consequence is concentration risk. Cloud platforms, shared databases, and third-party APIs connect institutions structurally. Drawing a perimeter around any single organization no longer reflects how systems actually work; treating each organization as an island leaves the connections undefended. Canada's reliance on hyperscale cloud providers (large-scale computing services such as Amazon Web Services, Microsoft Azure, or Google Cloud) sharpens this vulnerability. A single point of failure in a hyperscaler's authentication layer or Domain Name System (DNS) infrastructure could disable banking, government, and emergency services simultaneously.

AI and quantum computing form a dual-threat landscape: AI's open-source code faces active exploitation today, as the npm and GitHub incidents above show, while quantum computing's open-source code carries similar structural flaws that have not yet been exploited in a known attack. These are parallel software supply chain problems, differing only in whether the exploit has already happened.

The Quantum Blind Spot

Virtually every government, university, and technology company developing quantum algorithms today does so on classical computers running open-source quantum simulators, and until recently, almost no one had formally audited the security of that classical code.

A formal verification study, published in April 2026 by Broken Quantum, closed that gap. Applying formal verification methods, the study examined 45 open-source quantum simulation frameworks built by 22 organizations across 12 countries, including IBM, Google, Amazon, Xanadu, and Oak Ridge National Laboratory. Researchers identified 547 security findings, 40 of them rated critical. The vulnerabilities cluster into four classes: memory corruption in C++ backends such as IBM's Qiskit Aer; resource-exhaustion flaws in Python-based tools, including Canada's own Xanadu PennyLane; unsafe deserialization (a flaw that lets a booby-trapped data file execute hidden code) when external models are loaded; and a new attack class, QASM injection, in which unsanitized circuit code reaches quantum hardware parsers directly.

How vulnerabilities travel between projects is particularly consequential. Because open-source projects routinely reuse code, a flaw in IBM's commercial Qiskit Aer was found copied into the XACC framework used at a U.S. national laboratory. This is the first documented case of a vulnerability migrating from a commercial vendor into national security infrastructure through code reuse alone. A consistent threshold also emerged: simulators tend to fail once a calculation crosses roughly 32 qubits, a structural fragility that matters increasingly as governments scale up quantum capability.

2026: The Year the Software Bill of Materials Became Mandatory

Regulators on both sides of the Pacific have converged on the same response. They require organizations to know, and prove, exactly what is inside the software they run. A Software Bill of Materials (SBOM) is a structured inventory of every component in a piece of software. In plain terms, an SBOM is an ingredient list. Neither incident in this brief's introduction would have been prevented by one, as the attackers bypassed the list by compromising trusted credentials. What an SBOM does provide is rapid visibility: once a vulnerable ingredient is publicly disclosed, an organization with an SBOM can check within hours whether that ingredient is inside its own software, compressing a process that previously took days or weeks.

What is notable is not any single instrument but the sequence in which they have landed, each tightening the rules before it. Canada moved first. The Canadian Program for Cyber Security Certification (CPCSC), supported by Bill C-8, opened Level 1 self-assessment to defence suppliers on April 1, 2026, with cryptographically signed, verifiable SBOMs required at higher certification tiers.

Six weeks later, on May 12, 2026, the U.S. Cybersecurity and Infrastructure Security Agency (CISA) and G7 partners extended the same logic to AI with Software Bill of Materials for AI: Minimum Elements. This reached training datasets, model weights, and third-party plugins. Korea followed within six weeks, unveiling its SW Supply Chain Security Roadmap on June 24, 2026. This built on its 2024 guideline but embedded AI-specific defences and pursued

mutual recognition with allied regimes. South Korea is not a G7 member, yet it matched the G7's AI SBOM guidance inside a single fiscal quarter. This signifies that South Korea is shaping emerging standards alongside adopting them.

These mandates function as strict legal requirements. A defence supplier lacking a valid SBOM under CPCSC loses eligibility for Canadian defence contracts outright. South Korea's move signals it intends to be treated as a comparable, trusted partner in that same procurement system.

The European Union closes the loop with the most demanding deadline of all. Under the Cyber Resilience Act (CRA), manufacturers must report actively exploited vulnerabilities within 24 hours starting September 11, 2026. This is a legal requirement with real consequences: non-compliant products can be barred from the EU market entirely. Read together, these four measures compress a decade of gradual harmonization into a single year. SBOM disclosure is becoming the default precondition for selling software into any allied government's supply chain.

None of these instruments reach the quantum simulators described above. SBOM tooling catalogues known libraries and flags previously disclosed vulnerabilities in code that already exists. It has no mechanism for catching a malformed QASM string injected at runtime. This creates a critical regulatory gap: by September 2026, AI and classical software will be subject to some of the most detailed supply chain disclosure rules in the world, while the quantum simulators used to design next-generation cryptography remain largely unregulated. That gap is exactly what the joint quantum security testbed proposed later in this brief is designed to close, alongside the SBOM standards convergence proposed with it.

From Walls to Webs: Federated Threat Modelling

The common thread running through these regimes is a shift from defending perimeters to defending architectures. The most promising operational model for that shift is federated threat modelling, which lets organizations contribute to a shared defence without handing over sensitive data.

This approach lets each organization train threat-detection models on its own source code and network activity locally. They share only the resulting patterns, never the raw data, with a central system that builds a collective defence. A lightweight verification process authenticates every contribution and screens out poisoned updates. Once a vulnerability surfaces in one widely used library, every affected participant can be identified and contained immediately, avoiding the delay of institution-by-institution discovery. No government has yet deployed this operationally for open-source threat intelligence. The first pathway proposed below would position Canada and South Korea as early adopters of an emerging practice.

Federated threat modelling allows two allied governments to share top-tier threat intelligence without routing sensitive defence data through a hyperscale cloud platform headquartered in a third country. That makes it a data sovereignty tool as much as a technical one, and a natural complement to the sovereign-AI logic set out earlier in this brief. This mechanism drives the first of the three co-operation pathways detailed below.

Three Pathways for Canada–Korea Co-operation

Canada and South Korea do not need to build this co-operation from nothing. At the second Canada–Republic of Korea Foreign and Defence ('2+2') Ministerial in Ottawa on February 25, 2026, the two governments signed an Agreement on the Protection of Military and Defence Classified Information. They confirmed the inaugural Canada-Korea Cyber Policy Consultations would convene the following month, naming threat information sharing and critical infrastructure protection as priorities. That institutional foundation, layered onto the 2022 Comprehensive Strategic Partnership, supports the following three proposals.

First, a **Federated Threat Intelligence Bridge** linking the CCCS with South Korea's National Intelligence Service (NIS) and the Korea Internet & Security Agency (KISA). Built on the architecture described above, each country's sensitive data stays local while machine-learning signatures of emerging threats are exchanged in near real time. The primary objective is operational speed: signatures detected in Canadian code would automatically flag equivalent risks in Korean systems within hours, preempting independent discovery.

Second, **mutual recognition of AI and quantum SBOM standards.** Canada's CPCSC and South Korea's new Roadmap are converging on similar requirements. Formally harmonizing them, alongside a joint research effort on Cryptographic Bills of Materials, would let firms in both countries meet procurement requirements without duplicating compliance work, while jointly supporting the National Institute of Standards and Technology (NIST)-led migration to post-quantum cryptography. The objective here is administrative efficiency: a South Korean firm meeting domestic SBOM requirements would not need to repeat certification for the Canadian defence market. The Canada-Germany Sovereign Technology Alliance offers a usable template for this arrangement.

Third, a **joint red team and quantum security testbed.** The Canadian Quantum Technology Innovation Partnering Delegation travelling to South Korea from June 29 to July 4, 2026, provided an occasion to launch a standing exercise. Functioning similarly to an AI sandbox, a joint red team of Korean specialists would deliberately test Canadian quantum tools under controlled conditions. This shared environment would surface flaws, such as the 32-qubit failure threshold, while the consequences of finding them remain low.

Conclusion: Overcoming Implementation Barriers and the Strategic Role of APF Canada

Canada and South Korea are not at the same point in their regulatory cycles. Canada's certification regime is in its first, self-assessed tier, while South Korea's roadmap was unveiled only days ago. Canada acted as the first mover by launching CPCSC in April. Near-term harmonization will require South Korea to adapt to a Canadian framework that is already in motion. Threat-intelligence sharing between national-security institutions runs into classification barriers that take longer to resolve than any technical architecture. Furthermore, the small and mid-sized firms writing this open-source code lack the resources to absorb new compliance obligations without dedicated support.

That classification barrier is precisely why government channels cannot solve this alone. Direct threat-data sharing between the CCCS and South Korea's NIS will always be constrained by classification protocols not designed for the speed of an 88-minute supply chain attack. The code on the front line is written and maintained almost entirely outside government, by small firms and volunteer communities that neither agency can instruct directly. What is needed is a platform capable of carrying threat intelligence from classified sources to the unclassified developers who must act on it.

This partnership focuses on intelligence and technical co-operation, distinct from traditional technology transfers. Canada's contribution involves threat data from the CCCS and research tools like PennyLane; South Korea provides testbed capacity and its own threat intelligence from NIS and KISA. Both sides gain faster detection of compromises in open-source code. For example, if the CCCS detects a poisoned open-source package in a Canadian AI development environment, the resulting threat signature would be shared with KISA within hours, allowing South Korean firms using the same package to patch before exploitation. The reverse would hold if South Korea detects the threat first.

The Asia Pacific Foundation of Canada (APF Canada) is positioned to fill this gap. APF Canada, with South Korea's Science and Technology Policy Institute and the Canadian Embassy in Seoul, convened the inaugural Canada-Republic of Korea Track 1.5 Dialogue on Artificial Intelligence on July 9, 2025.

This event brought senior government, industry, and academic stakeholders into a single room outside formal diplomatic channels. APF Canada repeated that model from May 18 to 22, 2026, leading a Canada-Singapore

AI and quantum collaboration mission. That track record establishes APF Canada as a neutral and trusted platform capable of carrying the CCCS's threat data and KISA's testbed expertise to the private sector without waiting on declassification, and of transitioning the March 2026 Cyber Policy Consultations into a permanent standing partnership. The code supporting Canada and South Korea's AI and quantum futures is shared, largely unaudited, and thinly defended. Closing that gap is well within reach of the relationship the two countries established this year.

The code supporting Canada and South Korea's AI and quantum futures is shared, largely unaudited, and thinly defended.

ACKNOWLEDGMENTS

The author is grateful to the many practitioners, officials, researchers, and industry leaders across Canada and South Korea who generously gave their time and candour over the course of this research. Most requested anonymity; their perspectives and insights are present throughout these pages, even if their names are not. The analysis, conclusions, judgments, and any errors that remain are entirely those of the author.

This brief is in a series produced with the support of funding from the Korea Foundation.

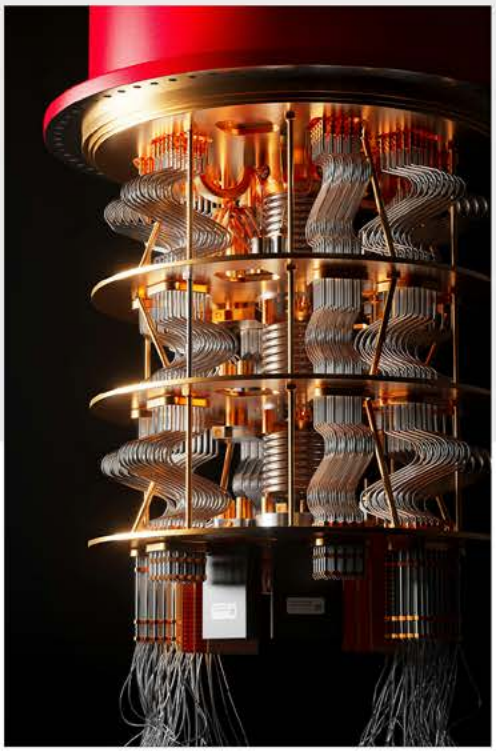

ASIA PACIFIC
FOUNDATION
OF CANADA

FONDATION
ASIE PACIFIQUE
DU CANADA