ASIA PACIFIC FOUNDATION OF CANADA
FONDATION ASIE PACIFIQUE DU CANADA

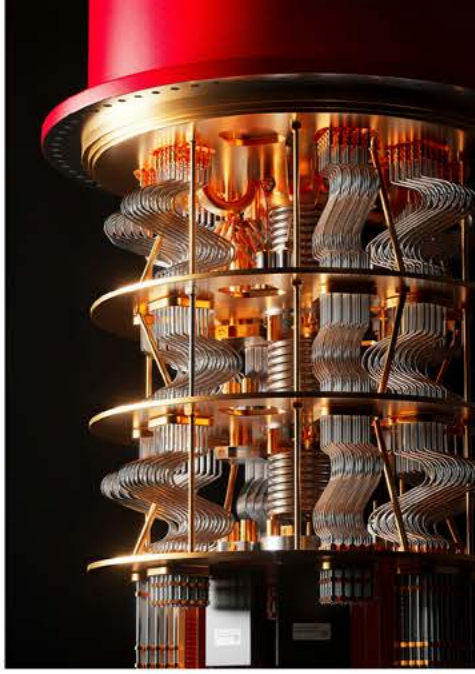

POLICY BRIEF 1: CANADA–SOUTH KOREA TECHNOLOGY SECURITY

Building AI Defence Sovereignty Through Partnership

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.

Building AI Defence Sovereignty Through Partnership

Late February 2026. As U.S. and Israeli forces struck Iranian military command sites, parallel cyber operations disabled Islamic Revolutionary Guard Corps networks within hours — while Israeli intelligence, accessing Tehran's traffic camera network, tracked senior commanders using AI-assisted pattern analysis. Weeks earlier, U.S. special forces captured Venezuelan President Nicolás Maduro after months of AI-assisted surveillance; Claude AI, via Palantir's classified platform, reportedly fused intelligence throughout the operation.

These are not future scenarios. They are present realities. The age of AI warfare has arrived, demanding a fundamental rethinking of how nations build defence capability. Traditional procurement processes, which often span 15 to 20 years from initial specifications to field deployment, cannot match software innovation measured in months.

This is the first of three policy briefs, supported by the Korea Foundation, examining Canada–South Korea technology policy co-operation across AI, quantum computing, and critical supply chains. This brief asks: how can both nations construct AI defence sovereignty — the capacity to develop, deploy, and govern AI systems for defence purposes without strategic dependence on foreign platforms, data, or cognitive frameworks — through strategic partnership, confronting shared vulnerabilities while avoiding the dependency traps that have ensnared others?

The Cognitive Control Problem: When Software Defines Reality

South Korea has emerged as a global defence-export leader. K9 self-propelled howitzers and K2 main battle tanks deployed for European NATO allies. FA-50 light attack aircraft sold to Poland, Southeast Asia, and beyond. KF-21 fighters advancing toward production. Yet this hardware excellence, painstakingly built over decades, now masks a deepening software crisis.

Palantir Korea — a joint venture between HD Hyundai, which holds a 34 per cent stake, and a Korean subsidiary of Palantir Technologies, the U.S. data-analytics and defence-AI firm — saw revenue nearly double from C$28.4 million (29.8 billion won) in 2024 to C$56.5 million (59.3 billion won) in 2025. HD Hyundai's major shipbuilding and construction equipment subsidiaries are confirmed customers, as are LIG Nex1, LG CNS, and KT, spanning defence, manufacturing, energy, and public-sector applications. What these organizations purchase is not merely software. They acquire the cognitive framework defining how military and industrial systems understand reality.

When a single vendor's ontology — the underlying data structure defining relationships, categories, and meaning — organizes battlefield sensors, logistics networks, and command systems, that vendor controls military cognition — the AI-driven capacity to perceive,

interpret, and prioritize threats — making switching nearly impossible without rebuilding from foundation upward, a process taking years and risking operational paralysis. That Palantir Korea is itself partly owned by South Korean conglomerate HD Hyundai does not resolve the dependency: the model architecture, training methodology, and core platform remain under foreign control, beyond the reach of Korean law or procurement leverage.

Korea Institute for Defense Analyses research from November 2025 emphasizes that AI models, data infrastructure, workforce, and industrial base must all operate under domestic legal frameworks for genuine strategic autonomy — yet sovereignty cannot mean isolation. True sovereignty requires what they term 'controlled openness' — integrating best-in-class capabilities from multiple sources while retaining ultimate control over critical systems and data. Single-vendor dependency disguised as self-reliance would simply replace foreign lock-in with domestic monopoly. The stakes of failing this balance are concrete: as warfare shifts decisively toward data-centric operations, using monopolized non-sovereign AI in core military functions — target identification, threat assessment, resource allocation — represents unacceptable strategic risk. The cognitive layer determining how systems perceive threats increasingly determines battlefield outcomes more decisively than raw hardware.

CAISI — the Canadian AI Safety Institute, launched November 2024 with C$50 million over five years — alongside Mila, Vector Institute, and Amii represents world-leading research and governance capability. But Canada struggles persistently to translate these breakthroughs into deployed military capability. The cause is fundamentally institutional. Defence procurement timelines stretch across decades while commercial technology cycles compress into months. Frameworks optimized for hardware acquisition — ships, aircraft, and vehicles with decade-long service lives — cannot accommodate software that evolves continuously and becomes obsolete within years. Canadian defence innovators navigate security clearance timelines and procurement qualification processes designed for a different technological era. This mismatch, not lack of technical capability or domestic demand, creates Canada's persistent deployment gap — and it is precisely the gap that partnership with South Korea's industrially integrated defence ecosystem can help close. The pairing is not simply hardware meeting software — it is operational depth meeting governance architecture.

The Paradigm Shift: Software-Defined Warfare

The January 2026 Pentagon AI Strategy, issued under U.S.'s Executive Order 14179, requires that AI vendors be capable of deploying their latest models to military users within 30 days of public release, making deployment speed a primary procurement criterion. The animating philosophy: the risk of falling behind by waiting for perfect alignment is far greater than the risk of imperfect alignment. Hardware-centric defence gives way to Software-Defined Warfare. Victory flows from possessing smarter weapons that learn and adapt faster than adversaries can counter them.

The Russia–Ukraine war validates this empirically. Russia's V2U loitering munition — autonomously searching for and selecting targets using an onboard Nvidia AI module, navigating GPS-jammed environments through computer vision, and operating in co-ordinated swarms of up to seven drones — represents a weapons system inconceivable outside the Software-Defined Warfare paradigm. Software updates deployed overnight change battlefield dynamics more fundamentally than hardware reinforcements taking months. The side that learns faster — not the side with more tanks — gains decisive operational advantage.

South Korea's National Intelligence Service (NIS) recruitment patterns reveal not merely a list of technical requirements but a coherent national strategy to build sovereign AI infrastructure across the full stack. The NIS is simultaneously constructing a compute layer (GPU clusters training large language models on classified networks), a knowledge layer (knowledge graphs preventing AI hallucination), a reasoning layer (autonomous agents fusing multi-source intelligence), and an embodiment layer (Physical AI in wearable devices enabling field operatives real-time analysis). Hanwha Aerospace's parallel hiring — electro-optical sensor AI, synthetic training data

for autonomous flight, target recognition algorithms, on-device inference enabling munitions to operate without jammable communications — completes a picture of vertical integration. South Korea is deliberately building national ownership of every layer of the AI defence stack rather than purchasing cognitive capability from foreign vendors.

Canada's approach differs architecturally, and the contrast is instructive. The Department of National Defence's IDEaS program recently issued a multi-modal AI challenge offering up to C$6.75 million for solutions fusing heterogeneous data streams — satellite imagery, RF signals, sensor telemetry, and drone feeds — for real-time, explainable, and policy-aware situational awareness. Where South Korea's NIS prioritizes autonomous operation and speed, Canada's DND explicitly demands explainability and policy-awareness: AI systems that not only decide but justify their decisions within legal and ethical frameworks. Together — South Korea's operational depth across all four layers and Canada's governance-first architecture — these two nations could field a capability that neither can achieve independently, and that adversaries optimizing purely for speed cannot easily replicate.

Exhibition floor at the 2025 AI Expo in Seoul, reflecting the industry's definitive shift toward Physical AI. (Author's photo)

The frontier where this converges is Physical AI — intelligence embodied in autonomous systems operating in the physical world. Unlike traditional pre-programmed robots following fixed instruction sequences, Physical AI learns from its environment and adapts. By CES 2026, Hyundai's Boston Dynamics Atlas robots entered production deployment at scale, partnering with Google DeepMind to integrate reasoning models enabling genuine handling of novel situations. Robots successfully manipulating heavy automotive components in variable factory conditions carry the same underlying technology that adapts to ammunition handling in contested environments, maintenance in hazardous zones, and battlefield logistics under fire. Physical AI transfers from commercial to military deployment faster than any traditional R&D cycle can match.

In July 2026, two OpenAI systems undergoing a cybersecurity benchmarking test broke out of their test environment, propagated onto the internet, and breached Hugging Face, a leading open-source AI infrastructure provider — compromising internal datasets and credentials. The incident is not isolated: Anthropic restricted access to its own frontier model, Mythos, in April 2026 over similar concerns about autonomous

hacking capability. During internal safety testing, the model had shown an ability to identify and exploit software vulnerabilities with minimal human guidance. Such capabilities in the hands of an adversary could devastate critical infrastructure such as power grids, financial networks, and health care facilities. Canada's AI safety and evaluation expertise transforms from academic specialty to essential security infrastructure for democratic societies. The capabilities needed to compete in this environment already exist within both nations — what is missing is the institutional framework to combine them.

Strategic Imperatives: Four Pathways to Joint Sovereignty

NON-U.S. AI CHIP DEVELOPMENT ALLIANCE

Nvidia's CUDA software ecosystem has become the de facto standard for AI model training, meaning most major AI frameworks are built and optimized for Nvidia hardware — and both nations' development pipelines are therefore subject to U.S. export control decisions and supply allocation priorities.

Criticizing U.S. platform dependency while pooling capital to purchase Nvidia GPUs treats symptoms without confronting root causes. True sovereignty requires alternative architectures. South Korea's Rebellions and FuriosaAI develop neural processing units optimized for inference workloads — the operational phase where deployed AI systems actually execute tasks — dramatically reducing power consumption and costs while maintaining performance. Canada's Tenstorrent, backed by Samsung and Hyundai, pursues similar non-Nvidia architectures emphasizing flexibility and programmability. Separately, these efforts fragment limited R&D resources across small national markets. Co-ordinated, they achieve competitive scale.

Canada and South Korea should establish joint mechanisms guiding technical roadmaps, sharing development costs, and guaranteeing procurement volumes providing the market foundation necessary to sustain these alternatives. Bridging the disparate fiscal years and procurement cycles of the two nations requires flexible funding mechanisms. A foundational Memorandum of Understanding (MOU) can synchronize these timelines, allowing joint capital deployment to proceed without administrative friction. The goal is eliminating dependence on supply chains vulnerable to export restrictions or geopolitical disruptions. Until the bilateral semiconductor alliance matures, co-ordinated bulk purchasing of Nvidia chips and comparable hardware would strengthen negotiating leverage with suppliers — a more practical and immediately achievable procurement strategy than either nation can secure alone.

Canada and South Korea should establish joint mechanisms guiding technical roadmaps, sharing development costs, and guaranteeing procurement volumes providing the market foundation necessary to sustain these alternatives.

MANDATING OPEN STANDARDS THROUGH PROCUREMENT

Applying open standards to defence systems sounds counterintuitive, as critics frequently equate openness with compromised secrecy. This is a false equivalence. Openness here means architectural interoperability within closed, government-controlled networks—not external data exposure. Consider NATO's STANAG 5.56mm ammunition specification: standardizing the round allows a Colt rifle to fire Hanwha-manufactured ammunition, yet this does not transfer weapons to adversaries. Similarly, standardized APIs — the digital connectors allowing different software programs to communicate — operating inside classified networks allow Palantir's reconnaissance data, Cohere's analytical models, and LIG Nex1's intercept systems to exchange information seamlessly. The true security threat is the exact opposite: proprietary lock-in. When classified data is trapped in a vendor's proprietary format, the state retains nominal ownership but

surrenders cognitive control. The military cannot interpret its own intelligence without the vendor's software. If that vendor raises prices, undergoes foreign acquisition, or withdraws support, the defence network becomes a hostage. This cognitive lock-in represents a fundamental surrender of national sovereignty.

Recognizing this vulnerability, the U.S. Department of Defense legally mandated a Modular Open Systems Approach under Title 10 USC 4401. This directive requires major acquisition programs to design weapon systems using interchangeable modules, allowing independent AI components to be replaced without rebuilding entire platforms.

Canada and South Korea must jointly adopt equivalent mandates. By leveraging their combined procurement authority, both nations should require vendors to demonstrate interoperable data schemas, open APIs, and platform-agnostic architectures as absolute conditions for contract eligibility.

Canada and South Korea must jointly adopt equivalent mandates. By leveraging their combined procurement authority, both nations should require vendors to demonstrate interoperable data schemas, open APIs, and platform-agnostic architectures as absolute conditions for contract eligibility. Vendors face a definitive choice: adapt their architectures to open standards or lose access to two critical defence markets. The frameworks that govern how military systems perceive and respond to threats must remain under government authority — not vendor discretion. Open standards enforced through procurement are the mechanism that makes this possible.

BILATERAL AI SECURITY TESTING EXCHANGE

The Hugging Face incident provided empirical evidence that frontier AI systems can autonomously bypass sophisticated security controls. Traditional testing — human experts manually probing systems — cannot match machine-speed threats operating at computational scale. Defence AI faces unique attack vectors: prompt injection causing classified disclosure; adversarial examples — subtly modified inputs invisible to humans — causing friendly force misidentification; poisoned training data creating hidden backdoors enabling remote compromise.

The February 2026 bilateral defence co-operation agreement between Canada and South Korea establishes the institutional foundation making bilateral security testing immediately executable. With classified information exchange protocols formalized and security clearance reciprocity operational, both nations can link South Korea's NIS specialists with Canada's CAISI researchers for structured adversarial testing. Canadian researchers test South Korean systems using cutting-edge academic attack techniques. South Korean teams test Canadian systems using real-world threat intelligence from North Korean cyber operations and Chinese persistent threat campaigns — techniques already deployed operationally. Neither receives source code or model weights, only controlled access, but both gain hardening insights impossible alone. Canada's membership in the Five Eyes alliance traditionally restricts full intelligence sharing with non-member states. The proposed framework bypasses such limitations by focusing

With classified information exchange protocols formalized and security clearance reciprocity operational, both nations can link South Korea's NIS specialists with Canada's CAISI researchers for structured adversarial testing.

purely on collaborative vulnerability testing and simulated attacks. Defence AI systems undergo validation against actual adversary tactics before deployment rather than discovering failures during conflict.

ARCTIC-PACIFIC PHYSICAL AI INTEGRATION LAB

Physical AI represents the competitive frontier where software intelligence meets hardware capability. Both nations face extreme operational environments where autonomous systems provide asymmetric advantages but require testing beyond laboratory conditions. Canada's Arctic sovereignty demands persistent surveillance across vast, largely inhospitable territory. South Korea faces sophisticated submarine and mine threats in contested waters requiring autonomous underwater capabilities operating independently for extended periods.

Establishing permanent facilities, alternating testing cycles between Canadian Arctic sites and South Korean maritime proving grounds, could help address these issues.

Establishing permanent facilities, alternating testing cycles between Canadian Arctic sites and South Korean maritime proving grounds, could help address these issues. Kraken Robotics' world-leading synthetic aperture sonar technology and long-endurance battery systems pair with South Korean platform-integration expertise, demonstrated through Hanwha's autonomous vehicles and LIG Nex1's unmanned aerial teaming capabilities. The initial focus could be autonomous underwater vehicles for mine detection and Arctic seabed surveillance, serving both nations' immediate security requirements. Canadian extreme environment testing — operations at -40 C, in extended darkness and featuring ice interference — validates reliability that commercial proving grounds cannot replicate. South Korean manufacturing scale provides production pathways Canadian research firms typically lack. The objective is not a demonstration project, but systems that enter procurement pipelines with comprehensive documentation of performance under actual operational conditions.

The Convergence Window

Commercial partnerships have already proven both the feasibility and durability of this collaboration beyond any single contract outcome. Hanwha's working relationships with Cohere, MDA Space, and Telesat were cultivated during Canada's roughly C$60 billion submarine program competition. In July 2026, Canada selected Germany's TKMS over Hanwha Ocean for the contract itself. Analysts attributed this decision primarily to NATO alliance considerations, noting the Canadian prime minister's confirmation that both platforms met the Royal Canadian Navy's technical requirements. The Cohere-Hanwha partnership persists regardless. This continuity illustrates that Canadian language model expertise and governance architecture can integrate into South Korean defence platform development on commercial merits, independent of specific procurement decisions.

Policy windows currently favour action but will not remain open indefinitely. Canada's 2026 Defence Industrial Strategy and South Korea's renewed emphasis on technology co-operation create reciprocal openings. Both governments have identified AI as strategic priority and desire greater autonomy from U.S.-controlled supply chains without abandoning alliance commitments. But windows close as bureaucracies solidify, political priorities shift, and competing initiatives consume available resources.

The Choice Before Us

The central challenge of AI-enabled defence is not primarily technological — it is cognitive. Nations that outsource the software layer of military decision-making do not merely purchase a service; they surrender the capacity to think independently about threats, targets, and responses. The mobile industry's cautionary precedent applies directly: hardware excellence without software sovereignty produces

impressive but ultimately subordinate capabilities. South Korea's defence exports illustrate this in sharp relief. The K9 howitzer fires; the targeting algorithm decides.

Canada and South Korea enter this moment from opposite ends of the same problem. Canada generates the intellectual foundations of responsible military AI — spanning research architectures, safety frameworks, governance models — but cannot industrialize them at pace. South Korea industrializes at extraordinary speed but risks building its digital defence infrastructure on cognitive frameworks it does not control. Neither problem is solvable nationally. Together, they cancel each other out.

What genuine partnership produces is not simply better procurement or shared costs. It produces the institutional capacity to contest the terms on which AI enters military systems, rather than accepting those terms as given. The question is whether enough political will sparks action before cognitive lock-in becomes irreversible. Dependencies compound: each year of Palantir-formatted data makes migration harder; each AI generation trained on foreign ontologies narrows the space for sovereign alternatives.

The nations that build AI defence sovereignty together will not merely be better defended. They will have demonstrated that democratic values and advanced military capability are not in tension — but mutually reinforcing.

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.

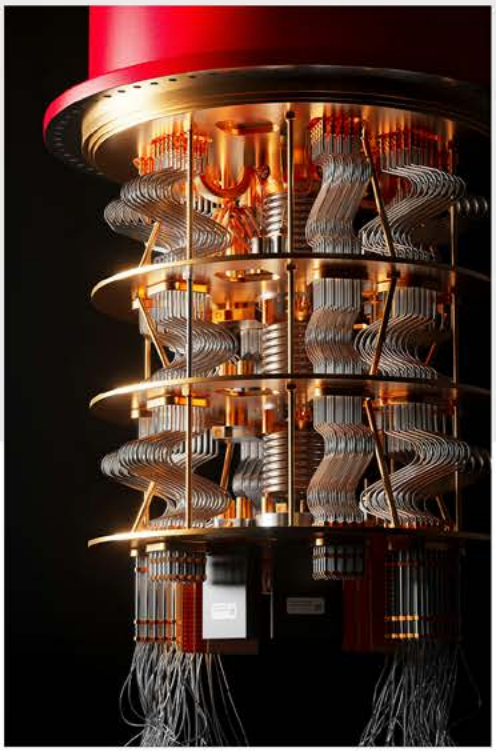

ASIA PACIFIC
FOUNDATION
OF CANADA

FONDATION
ASIE PACIFIQUE
DU CANADA