Innervation

G7 GovAI Grand Challenge funding recipient | Canada

Sovereign Multi-Agent Platform

Defence & National Security Capabilities

Governed Multi-Agent AI for Sovereign Defence Environments

Complete auditability · Canadian-built and hosted · Model-agnostic · PoC in weeks

G7 GovAI Grand Challenge Winner — Validated for responsible, production-grade AI in regulated public-sector environments.

The Strategic Problem

The gap is not in individual AI capability. It is in orchestration, governance, and auditability across coordinated AI workflows — where partial failure is not acceptable and every decision must be traceable. Adversaries operate at machine speed; the defence community's response has largely been siloed point solutions that do not coordinate, share context, or execute governed multi-step workflows.

✗ Point solutions	✗ General-purpose LLMs	✗ Internal builds
Siloed.	No guarantees. No audit trail.	12–18 month timelines.

The Innervation Platform

An orchestration layer coordinating specialised AI agents on complex, multi-step problems — with **mathematical workflow guarantees** (no race conditions, no partial failures, no silent failure mode), **complete observability** (every agent decision logged live), **sovereign infrastructure** (Canada/AWS, air-gap capable, configurable data residency), and **model-agnostic architecture** — run commercial, open-source, or classified models interchangeably.

Defence Capability Pillars

Pillar	Capabilities
01 Cyber Defence & Continuous Vulnerability Management	**Coordinated white-hat agent teams at machine speed.** Continuous red-team operations across the full attack surface; coordinated vulnerability discovery across app, infrastructure, and supply chain; automated blue-team response validation; compliance-mapped hardening with full audit trail.
02 Intelligence Analysis & Document Workflow Automation	**Multi-source synthesis and analyst augmentation.** Multi-source intelligence synthesis with confidence scoring; cross-reference against holdings and watchlists; structured report drafting (ACH, SWOT, Red Team); document bias and quality auditing; translation and summarisation with provenance tracking.
03 Sovereign AI Infrastructure & Supply Chain Governance	**Auditability and governance for AI in regulated defence environments.** Unified observability across multi-vendor, multi-model environments; supply chain AI risk assessment; policy-constrained workflow execution enforced at the orchestration layer (not prompt guardrails); procurement automation under Treasury Board frameworks; configurable human-in-the-loop escalation.

Technical Guarantees

Guarantee	Operational Meaning
No race conditions	Concurrency managed at the orchestration layer. Agents cannot produce conflicting outputs or corrupt shared state. Mathematical property, not a runtime check.
No partial failures	A workflow completes fully or fails explicitly — never a silent degraded state. Every failure surfaces with a full diagnostic trace.
Guaranteed run completion	Workflows terminate: successfully or with an explicit logged failure. No infinite loops, hanging processes, or orphaned agent states.
Complete audit trail	Every invocation, input, output, and decision logged in real time with timestamp, agent identity, model version, and full provenance. Immutable and exportable.
Policy-constrained execution	Operational boundaries enforced at the orchestration layer. A policy violation fails the workflow — it does not produce a non-compliant output.

Sovereignty & Compliance Posture

Infrastructure	Governance & Compliance
✓ Data residency configurable — stays within defined perimeter ✓ Supports air-gapped and restricted-network architectures ✓ Sovereign-configurable model hosting — Canadian or allied infrastructure ✓ No dependence on foreign LLM providers unless operationally chosen	✓ G7 GovAI Grand Challenge winner ✓ Audit trail designed for Treasury Board reporting requirements ✓ Human-in-the-loop decision gates at any workflow checkpoint ✓ Policy enforcement at the orchestration layer — not prompt guardrails ✓ Aligns with Directive on Automated Decision-Making

Engagement Model (Example)

01 Scoping & Threat Mapping	02 Proof-of-Concept Build	03 Evaluation & Validation	04 Production Deployment
Define the workflow, threat surface, or governance problem. Agree on success criteria. **1–2 weeks**	Working prototype on your actual data, in your actual environment — not a demo. **2–4 weeks**	Results against success criteria. Security review. Integration confirmed. Proceed or refine. **1–2 weeks**	Full deployment with operational SLAs, support model, and ongoing governance reporting. **4–8 weeks**

Why paid discovery: The PoC is a commercial arrangement — not a free evaluation. Both parties have skin in the game. You get a working system against real requirements. Programs that skip scoped evaluation spend more and wait longer.

Start with a Conversation

We scope every engagement individually. Bring your most complex workflow, your hardest governance constraint, or your most persistent threat surface.

contact@innervationai.com
innervationai.com

About Innervation AI

Canadian company building production-grade multi-agent AI orchestration for complex, regulated, and sovereign environments. G7 GovAI Grand Challenge Winner. Built and hosted in Canada on AWS.

Platform: Axon — multi-agent orchestration with mathematical workflow guarantees, complete observability, and embedded governance.

Unclassified — for preliminary discussion. All descriptions based on current production platform. Subject to applicable security and procurement requirements.