AI-Driven, Future-Ready

Enterprise AI maturity after the pilots end

What maturity actually looks like when the demos are over and the real transformation begins.



Celio Oliveira

Chief Data Officer | Government of Canada

Chair, CAIO/CDO Fall Boston 2026



Day 1 Opening Keynote • CAIO/CDO Fall Boston

CAIO/CDO FALL • BOSTON • OCTOBER 26–27, 2026

The pilot era is over

Four signals are converging in the same boardroom conversation.

PILOTS EVERYWHERE 01

Every enterprise can show a demo. Very few can point to a decision the demo changed. The pilot was the easy part.

AGENTS ARE ACTING 02

AI is moving from answering questions to taking actions inside real processes — at machine speed, and at machine scale.

BOARDS WANT BOTH 03

Value and risk now arrive in the same question, and directors expect evidence for each — not enthusiasm for one and hope for the other.

RULES ARE ARRIVING 04

Regulation, assurance, and audit expectations are landing while the building continues. Governance can no longer be the last slide.

In practice — at Finance Canada the question stopped being “which model?” months ago. It became “which decisions may it touch, and under what guardrails?”

Guardrails in practice

How a federal government deploys agents: inside the classification system — never around it.

PROTECTED B — AND NO FURTHER

DLP ensures agents read only up to Protected B. Confidential, Secret, or Top Secret material stops the flow: quarantined, logged, handed to a cleared human. “Proceed carefully” is not on the menu.

RASCI MATTERS MORE THAN EVER

When an agent acts, someone is Responsible and someone is Accountable — a person, with a name. The rest of the RASCI is designed in advance, not improvised after an incident.

ASSESS BEFORE YOU DEPLOY

Canada assesses first: the Algorithmic Impact Assessment (AIA) and the Privacy Impact Assessment (PIA) precede deployment — with transparency after, learning from the UK's ATRS.

A WORD OF ADVICE

Less risky: local agents, built by your own teams inside your own guardrails — over off-the-shelf agents whose data flows you cannot see. Build local where data is sensitive; buy carefully elsewhere.



PROTECTED B IS THE CEILING

From evidence-based decisions to decision architecture

A shift I have argued for a year: in the era of responsible AI, evidence alone is not enough. Decisions need an architecture.



GOAL	CONTEXT	SOURCE	EXPECTATION	VERIFICATION
The decision, and the outcome it serves	The situation, constraints, and mandate	Evidence that can be inspected	What good looks like — stated in advance	Checked against expectation, by a human able to overrule

The discipline behind it: ask AI for evidence, not confidence.

In practice — the Executive Copilot Lab at Finance Canada runs every executive question through these five moves. The habit scaled. The architecture is the habit, institutionalized.

The agentic leap

What actually changes when an agent goes from demo to production.

	BEFORE — THE DEMO	AFTER — IN PRODUCTION
WHO ACTS	A person, in a meeting, on the record.	An agent, between meetings, at machine speed.
WHAT SCALES	Judgment — slowly, and unevenly.	Judgment and error alike — instantly, and uniformly.
ACCOUNTABILITY	Attached to the signature at the bottom.	Designed in advance — or it attaches to no one.
RIGHTS	Protected by process, review, and delay.	Must survive automation. Canada: Charter s. 7 (security of the person), s. 8 (privacy), s. 15 (equality). The U.S.: due process, the Fourth Amendment, equal protection. Rights do not change when the decision-maker is software.

The failure mode writes its own headline: **“The agent did exactly what it was told to do.” That was the problem.**

Trust is infrastructure

A future-ready enterprise should not become an unaccountable one by default.

1 • LAWFUL AUTHORITY

Every system acts within a mandate someone can name. If no one can state the authority, the system does not run.

2 • HUMAN ACCOUNTABILITY

A human in the loop means a human with knowledge and authority — able to understand the recommendation and empowered to refuse it. A checkbox is not a human in the loop.

3 • EXIT BY DESIGN

AI with the ability to exit: portability of data, reversibility of decisions, and a tested path back to human process. Dependence without an exit is a risk position.

CONSENT

Collected for a purpose; used for that purpose.

INTEROPERABILITY

Systems that speak — without surrendering sovereignty.

COMMUNITIES AT THE TABLE

The people affected help design the guardrails.

STEWARDSHIP, IN WRITING

MOUs and shared agreements that outlast the project.

Trust is not a communications outcome. It is an architectural property.

Five moves for the future-ready enterprise

A practical roadmap for the year after the pilots.

1

Start with decisions, not models

Inventory the decisions that matter. Automate in that order.

2

Standardize meaning before scaling agents

Agents amplify ambiguity. Shared definitions are the real platform.

3

Design for reversibility

Every deployment ships with its exit: rollback, recourse, and a human path.

4

Ask for evidence, not confidence

Evaluate outputs against stated expectations — and publish the scorecard internally.

5

Name the accountable human

One owner per consequential system, with authority to pause it. No exceptions.

And count the footprint. Every model has one, and the North feels it first — the National Research Council's climate research in Canada's northern region is a standing reminder that efficiency is a governance metric too.

Maturity is not how much the machine decides.
It is how well we can answer for what it decided.

How do we build AI worthy of the institutions it serves? That is the work of the next two days.

Better decisions → **Better evidence** → **Better outcomes** → **Stronger institutions**

Interoperability without centralization • AI with the ability to exit • Accountability by design

Trust is not a communications outcome. It is an architectural property.

Thank you

Let us build AI worthy of the institutions it serves.

Celio Oliveira

Chief Data Officer, Government of Canada

linkedin.com/in/celioroliveira

CONNECT



Connect on LinkedIn

Scan to connect