

PUBLIC SECTOR INSIGHT


The Public Sector's Next Trust Boundary

Securing AI Agents, Machine Identities and the
Cryptography Beneath Them

Why responsible AI adoption now depends on
identity, policy, auditability and cryptographic
agility.

By Asim Khusro

Enterprise Technology & Cybersecurity | Securetron



Executive perspective

Public-sector AI adoption is entering a new phase. The challenge is no longer only whether an employee can safely use a generative AI tool. Increasingly, software agents can call APIs, retrieve records, invoke tools, make decisions and act across multiple systems. That changes the trust question from “who is the user?” to “who or what is acting, what is it allowed to do, and can we prove why it was allowed?”

Canada’s federal AI strategy explicitly aims to accelerate responsible AI adoption while strengthening governance, talent, transparency and value to Canadians. In 2026, the Government of Canada added specific guidance for agentic AI, reflecting the fact that autonomous systems create governance needs beyond conventional generative AI. At the same time, the Canadian Centre for Cyber Security is urging organizations to treat AI agents as distinct principals with cryptographically anchored identities, constrained privileges and continuous monitoring.[1][2][3]

A second transformation is happening underneath AI: cryptography itself is changing. The Government of Canada’s post-quantum cryptography roadmap requires departments to understand where public-key cryptography is used, build inventories, prioritize high-risk systems and progressively migrate to quantum-safe alternatives. For federal organizations, the initial departmental migration plan was due in April 2026, with high-priority systems targeted for completion by the end of 2031 and remaining systems by the end of 2035.[4]

The strategic implication: AI governance, machine identity and cryptographic modernization can no longer be treated as separate programs. They are converging into one digital trust problem.

The trust boundary has expanded

Public-sector security now has to govern people, software, machines and autonomous agents.




Figure 1. The public-sector trust boundary now includes autonomous software actors.

1. The identity problem is becoming a machine problem

For decades, identity programs were built primarily around people: employees, contractors, privileged administrators and citizens. Modern infrastructure expanded that model to service accounts, workloads,

devices, certificates and APIs. AI agents extend it again because they are not simply passive software. They can interpret context, choose tools, sequence actions and operate on behalf of a person or organization.

The Cyber Centre's 2026 guidance is unusually direct on this point. It recommends constructing each agent as a distinct principal with a cryptographically anchored identity, using strong authentication for agent-to-agent and agent-to-service calls, maintaining trusted registries, limiting permissions and continuously verifying identity and authorization at runtime.[3] NIST's 2026 concept work on software and AI-agent identity reaches the same set of questions: identification, authentication, authorization, delegation, auditability, non-repudiation and prompt-injection resilience.[6]

This matters in government because an agent may touch systems that carry citizen data, financial information, operational technology, case files, procurement records or internal policy. A shared API key is not enough. A public-sector security team should be able to answer: Which agent made this request? Which human or service authorized it? What role was it acting under? Which model and tool were used? What policy allowed the action? Can the credential be revoked immediately without disrupting unrelated workloads?

2. AI security needs an enforcement point, not only a policy document

Governance frameworks are essential, but AI traffic also needs technical enforcement. Prompt injection, sensitive-data disclosure, unapproved model use, credential leakage and over-privileged tool access occur at runtime. They cannot be managed only through acceptable-use policies or annual reviews.

The Government of Canada's guidance emphasizes risk assessment, safeguards, monitoring and responsible use. NIST's Generative AI Profile similarly frames trustworthiness as something organizations must manage across the AI lifecycle rather than bolt on after deployment.[1][5] For agentic systems, the Cyber Centre further recommends layered controls, human control points for higher-risk actions, strict privilege boundaries, runtime monitoring and auditable actions.[3]

The architectural response is to create a control point between users or agents and the models, data and tools they access. That control point should be able to identify the caller, authorize the request, inspect inputs and outputs, enforce policy, govern model selection and produce evidence for security and audit teams. This is the same logic that made gateways, identity providers and policy enforcement points foundational in earlier waves of enterprise architecture.

A unified digital trust architecture

The control plane sits between identities and the systems, data and models they can reach.



Figure 2. A unified control plane connects identity, AI governance, audit and cryptographic agility.

3. PKI becomes more important, not less, in the AI era

It is tempting to view PKI as legacy infrastructure and AI as a completely new security domain. In practice, the two are becoming more closely connected. Certificates and cryptographic keys already establish trust for servers, devices, workloads, APIs, code signing and encrypted communications. As AI agents become non-human operators, the same foundations can provide strong, revocable identities for software that acts autonomously.

That creates a lifecycle challenge. Every new machine or agent identity introduces credentials that must be issued, renewed, monitored, governed and revoked. Manual certificate operations become difficult to scale, especially across hybrid environments, multiple certificate authorities and cloud services. The certificate problem is therefore not simply “avoid expiry.” It is an identity-governance problem: know what exists, who owns it, what it protects, when it changes and whether it still complies with policy.

The quantum transition raises the stakes further. Canada’s PQC roadmap says departments need to identify where cryptography is used across applications, networks, endpoints, code pipelines and physical IT assets, then prioritize migration based on risk. The roadmap explicitly notes that discovery will require multiple methods and software tools, and that departments should begin with an incomplete inventory and iteratively improve it rather than waiting for perfect visibility.[4]

4. The public sector should manage three inventories as one trust problem

Most organizations already maintain multiple inventories, but they rarely line up. The AI team may know which models are approved. The PKI team may know which certificate authorities it operates. IAM may know the human identities. Infrastructure teams may know servers and cloud workloads. Security operations may see events without always knowing the business owner.

A stronger trust architecture connects at least three living inventories:

- **Cryptographic inventory:** certificates, keys, CAs, algorithms, owners, dependencies and expiry / lifecycle status.

- **AI inventory:** models, applications, copilots, agents, tools, data connections and approved use cases.
- **Identity and authorization inventory:** people, machines and agents, including credentials, roles, delegated authority and revocation state.

The value comes from the relationships between them. If an AI agent has access to a case-management API, security should be able to see the agent identity, the certificate or credential used, the policy governing the call, the data classification involved, the model selected and the resulting audit event. That is what turns “AI governance” from a collection of guidelines into an operational control system.

Where this matters first

Public-sector use case	Trust question	Priority controls
Employee copilots	What information can leave the organization, and which models are approved?	Data-loss controls, model registry, prompt/response inspection, identity-aware policy
Citizen-service AI	Can the organization explain and audit how AI was used in a public-facing service?	Governance, logging, human oversight, privacy controls, approved knowledge sources
AI agents with tools	What can the agent actually do, under whose authority and with what credentials?	Distinct agent identity, least privilege, tool-call authorization, mTLS / certificates, runtime audit
Machine & workload growth	Can every non-human identity be discovered, owned, renewed and revoked?	Certificate discovery, lifecycle automation, ownership, policy and revocation
PQC migration	Which systems still depend on quantum-vulnerable public-key cryptography?	Cryptographic inventory, risk prioritization, crypto-agility, vendor transition plans

5. A practical agenda: govern what exists before scaling what is new

Public-sector leaders do not need to solve every future AI and cryptography problem at once. They do need to establish visibility and control before autonomy and machine identity scale further. A practical starting agenda is:

- Discover the estate. Identify certificate authorities, certificates, machine identities, AI tools, models, agents and major data connections. Do not wait for a perfect inventory.
- Assign ownership. Every critical certificate, AI system and agent should have a business owner, technical owner and defined lifecycle.
- Separate identity from shared secrets. Where feasible, give agents and workloads distinct credentials that can be individually authenticated, authorized and revoked.
- Create a runtime policy layer. Govern prompts, responses, model access, agent permissions and tool calls at a technical enforcement point, not only in policy documents.
- Design for evidence. Security, privacy, audit and program owners need traceability without creating unnecessary new stores of sensitive content.
- Build cryptographic agility into procurement. New systems should be easier to migrate than the systems they replace, especially as PQC adoption progresses.

A practical 90-day agenda for public-sector leaders

Start with visibility and control. Modernization can then proceed in measured stages.




Figure 3. A 90-day agenda focused on visibility, control and evidence.

The leadership question has changed

The next phase of public-sector AI will not be defined only by better models. It will be defined by whether institutions can trust the identities, decisions and actions surrounding those models. That is why AI security, machine identity, PKI and post-quantum readiness are converging.

The most important question for a CIO or CISO is no longer simply, “Are we using AI securely?” A more useful question is: “Can we identify every human, machine and agent acting in our environment, constrain what each is allowed to do, inspect the paths where sensitive information moves, revoke trust when conditions change and produce evidence of every important decision?”

Organizations that can answer that question will be in a much stronger position to adopt AI faster, because trust will be engineered into the operating model rather than added after deployment.

A Securetron perspective

Securetron works at the intersection of these trust domains. PKI Trust Manager focuses on certificate discovery, lifecycle visibility, issuance, renewal, revocation, automation and Microsoft AD CS / multi-CA integration. AI-FW provides an AI security and governance gateway for identifying and authorizing callers, inspecting prompts and responses, enforcing policy, governing model access and maintaining auditable activity. The common design principle is simple: trust should be explicit, policy-driven, revocable and observable.[7][8]

About the author

Asim Khusro is an enterprise technology and cybersecurity sales leader with more than 24 years of experience across enterprise software, cloud, managed services, cybersecurity, digital trust, PKI and AI security. At Securetron, he works with organizations on PKI modernization, certificate lifecycle management, machine identity, post-quantum readiness and AI security and governance.

References and further reading

[1] Treasury Board of Canada Secretariat, AI Strategy for the Federal Public Service 2025-2027.

<https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/gc-ai-strategy-overview.html>

[2] Government of Canada, Progress on AI in government, including Guide on the Use of Agentic Artificial Intelligence (May 22, 2026).

<https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/advancing-ai/progress.html>

[3] Canadian Centre for Cyber Security, Careful adoption of agentic AI (2026).

<https://www.cyber.gc.ca/en/guidance/careful-adoption-agentic-ai>

[4] Canadian Centre for Cyber Security, Roadmap for the migration to post-quantum cryptography for the Government of Canada (ITSM.40.001), effective June 23, 2025.

<https://www.cyber.gc.ca/en/guidance/roadmap-migration-post-quantum-cryptography-government-canada-itsm40001>

[5] NIST, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1).

<https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence>

[6] NIST NCCoE, Accelerating the Adoption of Software and Artificial Intelligence Agent Identity and Authorization, Concept Paper (2026).

<https://csrc.nist.gov/pubs/other/2026/02/05/accelerating-the-adoption-of-software-and-ai-agent/ipd>

[7] Securetron, PKI Trust Manager and PKI product information.

<https://securetron.net/>

[8] Securetron AI-FW, AI firewall and governance gateway.

<https://aifw.io/>