




A new model for secure, scalable public sector automation

Visibility, control, and resilience for **agentic public systems**

Insights from Jean Le Bouthillier,
CEO and Chairman of Qohash



Government leaders agree that AI will fundamentally improve how the public sector serves constituents, unlock efficiencies, and fulfills new mandates.

Generative AI is already delivering measurable productivity gains in areas like routine tasks, case work, decision support, and citizen communication. Governments now have the opportunity to reimagine how services are delivered — with AI embedded at the core of public service delivery.

The next wave of transformation will be driven by agentic AI systems, creating a new operating model for government service delivery. These new agentic systems must access sensitive data broadly and continuously to function effectively.

To execute confidently, we need to take a thoughtful approach — and reimagine data security not just as a safeguard, but as the foundation and enabler of intelligent, sovereign public systems.

In this paper we tackle:



A new blueprint for AI-first government services



The data security foundation of visibility, control, and intelligence



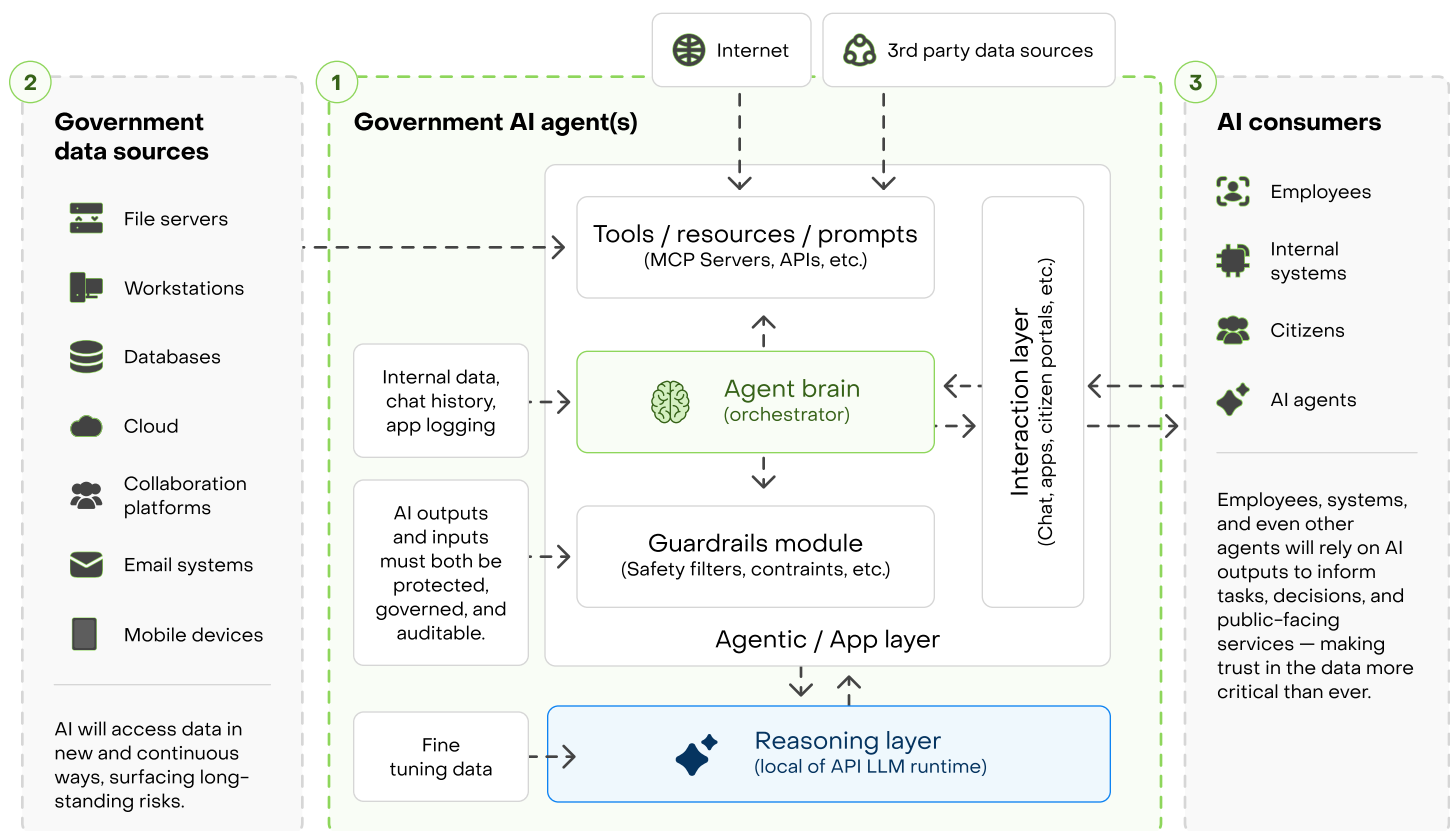
The five guiding principles for AI-ready data security

A new blueprint for AI-first government services

Governments around the world are developing blueprints for a new model of public service delivery powered by AI systems. At the core of this shift is the opportunity to create a wide range of ① **government AI agents** designed to improve how governments plan, operate, respond, and serve — at scale and cost-effectively.

These agents will need access to ② **context-rich data from internal government sources**, and will also generate sensitive outputs such as summaries, recommendations, and decisions that must be protected. This creates a new category of risk that didn't exist when services were delivered solely by humans.

At the same time, demand from ③ **AI consumers** such as employees, internal systems, citizens, and AI agents will accelerate dramatically. As these capabilities become available, expectations will shift. There will be no turning back. Governments must be ready to meet that demand by adopting systems that are trusted, controlled, and transparent.



A new operating model for public service: intelligent agents embedded across every layer of government. Tomorrow's government will rely on AI agents, autonomous systems that perceive, decide, and act to fulfill defined goals. These agents will support employees, enhance internal systems, and interact directly with citizens. They will operate at scale, making decisions and taking action based on context and data access, all powered by data that must be trusted, protected, and visible by design.


Foundation: visibility and control

AI thrives on context. To unlock its full potential, **governments will inevitably need to authorize access to key data sources for government AI agents**. Agents will summarize, decide, generate content, and act on behalf of the government.

To fulfill their purpose, agents will require access to sensitive information that will need protection. For example, Qohash is developing its own family of AI agents that continuously assess foundational data risks: (1) excessive retention, (2) internal overexposure, (3) external sharing, and (4) suspicious behaviours. The goal is to avoid unnecessary blocking, while still allowing enforcement where risks are high and clearly defined.



Emerging risk: unstructured data at rest



Governments have accumulated sensitive data for decades. This long tail of risk is especially acute in unstructured data, such as files at rest containing sensitive information. This latent risk will surface as AI begins accessing data in new, continuous ways. One dimension of this risk is unnecessary duplicate copies of the same sensitive data across systems.

Agentic AI systems require continuous, automated access to vast amounts of internal data to function effectively. Unlike humans, these systems:

- **Operate at an unprecedented scale and speed**, using more data, more often.
- **Lack context or intent**, making them prone to accessing data inappropriately if controls are missing.

Most government data environments such as file shares, collaboration platforms, and workstations were not built for machine access. This means:

- Data is often **unclassified, misclassified, or lacks meaningful labels**.
- Sensitive content is **stored in the wrong places or shared too broadly**.

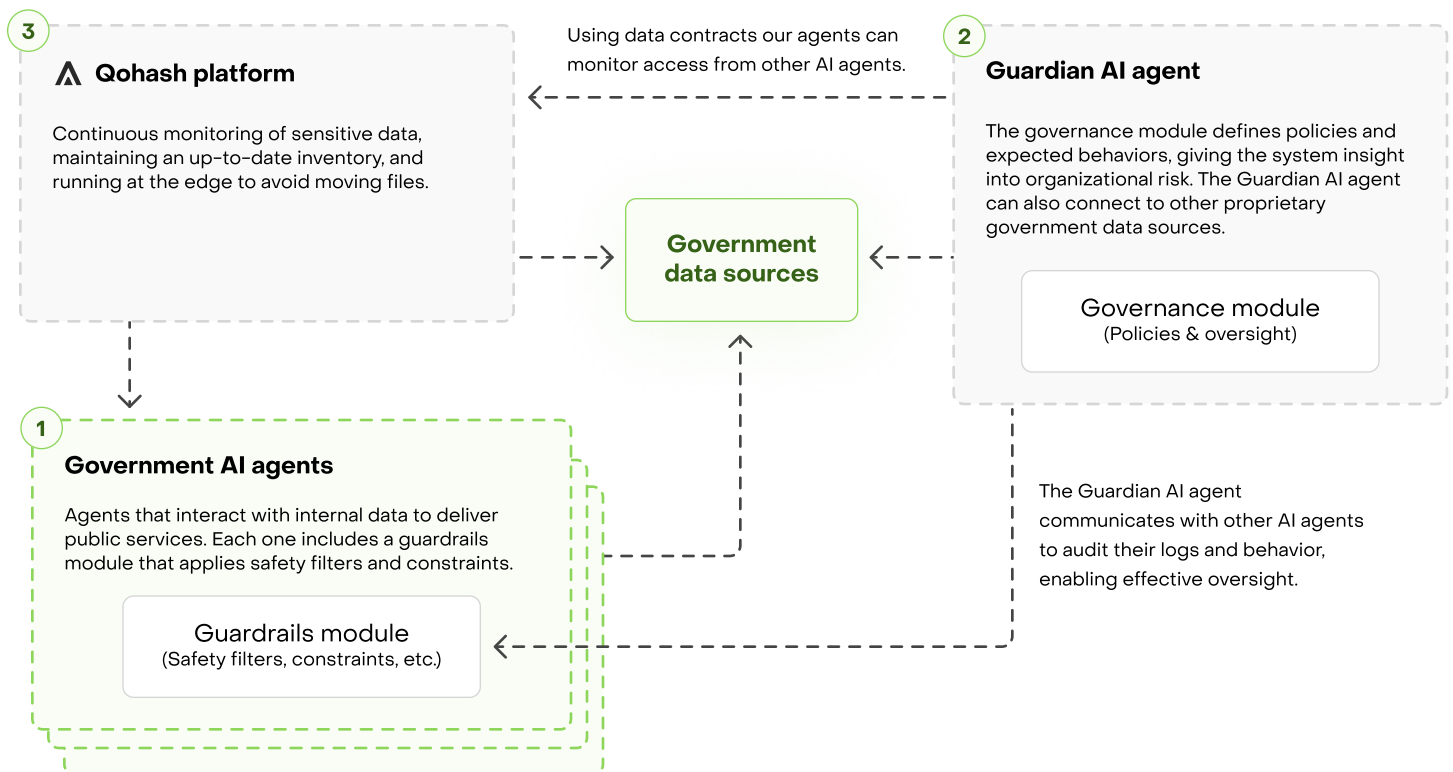


Governments need a solution that continuously discovers, classifies and labels data to confidently experiment with and adopt agentic AI for public service delivery. This intelligence must be available on demand to support AI guardrails and governance, and integrate seamlessly with the broader ecosystem.

Foundation: intelligence (at scale)

Once visibility and control are established, there is yet another challenge that must be addressed decisively: **scale**. That's because **government AI agents will operate continuously and generate millions of access events across vast data environments**.

To identify and mitigate misuse, it will be essential to supervise ① **Government AI agents** with an intelligent oversight layer: a ② **Guardian AI agent**. This oversight agent would integrate with the ③ **Qohash platform** to continuously assess which sensitive data is being accessed and surface other key data security insights.



To enable intelligent oversight at AI scale, such as with **Guardian AI agents**, systems must move beyond low-level APIs. Instead, they should rely on **structured data contracts** that define what to access, what to do, and what to analyze. This shift allows agents to express **intent through code**, enabling scalable, automated control.



The universal read layer

Retrieve any data using the Qohash model. Power dashboards, reporting, and AI agents with consistent and predictable access to Qohash data.



The universal execution layer

Manage actions such as deleting or moving files, notifying users, or revoking access through a consistent and secure execution layer.



Purpose-built data contracts

QXL enables the creation of custom views and intelligence pipelines that run continuously and adapt to emerging needs.

Appendix preview: See Appendix A for more details on Qohash's data contracts and how they enable visibility, control, and intelligence.

Beyond the foundations: putting principles into practice

Through our work with large enterprises and critical infrastructure organizations, we've seen firsthand what it takes to secure sensitive data in the AI era. Strategic partnerships were earned by understanding real-world problems and being willing to solve them head-on. Based on this experience, we are sharing five guiding principles to help government leaders secure data and unlock the benefits of AI.

1

AI readiness starts with data control

You can't safely use AI without controlling the sensitive data it sees. Visibility and policy enforcement must come first.

2

Focus on risk, not coverage

Coverage alone doesn't reduce risk. Use AI to assess context and assign confidence levels that generate appropriate actions or alerts.

3

Scan data where it lives

Never move files just to understand them. Scan in place with edge computing to reduce risk and analyze millions of files in parallel using your existing infrastructure.

4

Data grows, your bill shouldn't

Choose a technology architecture that scales with your data footprint. Avoid pricing models that tie cost to data volume, which limit coverage and create long-term risk.

5

Data security is national security

Ensure the code, infrastructure, and operations behind it are developed by trusted individuals, in trusted locations, under sovereign control.

More strategic data security insights

Qohash is a Canadian company, founded to deliver trusted, sovereign data security solutions for the nation's most vital organizations. We are innovative, pragmatic, and eager to roll up our sleeves to tackle emerging challenges with practical, forward-looking solutions.

I'd love to hear from you - please email me at jean@qohash.com.


v1.1 | Confidential © Qohash 2025




Appendix A – Qohash data contracts

We believe that one of the major changes essential to fully unlocking the value of AI agents is shifting how they interact with technology. To move beyond technical details and tedium, agents must operate through intent-driven communication. In response, we have taken proactive steps to address this shift by designing a data contract model that simplifies integration and increases reliability.

Through QQL, QAL, and QXL, we offer a predictable and structured way for data consumers – including software and AI agents – to access data, perform secure actions, and define custom materialized views based on specific needs.



The Qohash MCP server includes everything needed to translate natural language, intent-based queries into QQL, QAL, or QXL. To use it, simply connect your MCP client to the Qohash MCP server. The server will handle all secure communication with the Qohash platform.




QQL

“Qohash Query Language”
The universal read layer supporting **Visibility**

QQL allows you to formulate structured queries against the Qohash data model to retrieve specific entities or sets of entities that meet defined conditions. It provides a consistent and predictable interface for accessing Qohash data, enabling use cases such as dashboards, reporting, investigations, and AI-driven automation.

Who has access to more than 100 Canadian social insurance numbers?

```
Person: Count(info, info.type = "CanadianSSN") > 100
```




QAL

“Qohash Action Language”
The universal execution layer supporting **Control**

QAL defines executable actions such as delete, quarantine, notify, or revoke access, targeted at specific entities in the Qohash platform. Each QAL statement wraps a QQL query to precisely scope the set of entities the action applies to. This separation of retrieval (QQL) and execution (QAL) enables intent-driven automation across the platform.

I want to quarantine all files with sensitive data on a specific computer with MAC address .mac:44-38-39-ff-ef-57”

```
quarantine ( File: file.infoCount > 0 and datasource.id = "mac:44-38-39-ff-ef-57" )
```



QXL

Conceptual

Purpose-built data contracts supporting additional **Intelligence**

QXL enables users to define declarative data contracts that create normalized views of the Qohash data model for aggregation, time series analysis, or risk scoring. Once submitted, the platform continuously maintains the corresponding data layer, ensuring results are always available for workflows, dashboards, and integrations.

I want a daily updated dataset showing the weekly count of files containing PCI data on workstations, broken down by department.

```
contract PCI_Exposure_By_Dept {
  source:file
  where:info.type = "PCI" and
  datasource.type = "workstation"
  groupBy:file.department,
  week(file.discoveryDate)
  metrics:count(file.id) as fileCount
  refresh: daily }
```