The Rise of Agentic Data Governance

By Dr. Adi Hod, Founder & CEO, Velotix

Tel Aviv, Israel Sep 10, 2026 (Issuewire.com) - Yesterday, humans requested access to data and humans governed it. Tomorrow, humans and AI agents will both consume enterprise data - while governance agents continuously evaluate what data may be used, by whom, for what purpose, under what conditions, and for how long. That, to me, is the defining enterprise challenge of the next decade: governing data access at machine speed without losing policy, context, or accountability.

That single shift breaks thirty years of access control architecture. And it is why the next era of data security will not be won by more roles, more rules, or more approval workflows. It will be won by an agentic governance layer that understands intent, reasons over policy and context, and continuously governs how both humans and machines use data.

To see why, it helps to understand how we got here.

Three Eras of Access Control, One Fatal Assumption

When researchers at NIST formalized role-based access control in 1992, it was a genuine breakthrough. Instead of granting permissions user by user, you defined what an "analyst" could see, and every analyst inherited it. RBAC carried the enterprise for three decades - until data spread across 50 sources and three clouds, roles multiplied faster than employees, and every platform enforced its own local version of the truth.

Policy-based access control (PBAC) was the answer to that sprawl. Rather than asking "what role does this person have?", PBAC asks "what does our policy say about this person, this data, this purpose, right now?" Policy became a living statement of intent, evaluated in context and enforced centrally across every data source.

The third era added intelligence. A policy engine that merely stores rules still depends on humans to write every one of them - which relocates the bottleneck from IT ticketing queues to governance committees. A policy brain learns instead: AI-generated recommendations capture the judgment inside every access decision, and approved requests become candidate rules, so that yesterday's manual approval becomes tomorrow's automated policy. This approach is foundational enough to Velotix's platform that it is protected by US Patent 12,393,718, covering the system and method for managing data access requests. This evolution from static permissions toward intelligent, policy-driven governance is central to the approach we have taken at Velotix.

But all three eras quietly shared a more fundamental assumption: that access decisions would happen at human scale - with relatively stable identities, predictable purposes, and enough time for a human to intervene.

That assumption just expired.

The Agent Changed the Question

Enterprises are now deploying LLM copilots, autonomous agents, and agent-to-agent pipelines against their most sensitive data. A single agent can issue thousands of queries an hour, chain requests across systems, and act on behalf of users whose entitlements it may silently exceed. Non-human identities - service accounts, bots, and now agents - already outnumber human ones in most large organizations, and agentic workloads are multiplying them again.

Approval workflows designed for humans collapse at this scale. No committee can review ten thousand machine-initiated requests per day. No quarterly access certification can keep pace with agents that are created, cloned, and retired in minutes. And no role model can answer the question that now matters most: not just who is asking, but what is asking, on whose behalf, for what purpose, and under which conditions.

Much of the security industry has responded by racing to secure the agents themselves - authenticating

them, sandboxing them, watching their behavior. That work is necessary. It is also insufficient, because it treats the agent as the perimeter while leaving the deeper question unanswered: how should the use of enterprise data itself be governed, when the consumers of that data are humans, applications, and machines all at once?

What Agentic Data Governance Actually Means

Agentic data governance flips the frame. Instead of only securing agents, it puts agents to work governing data - a continuous, autonomous layer that evaluates every access, human or machine, against policy, context, sensitivity, and purpose, in real time.

Governance agents reason over identity, intent, purpose, data sensitivity, and regulatory context. They recommend and orchestrate the appropriate action, while a deterministic policy layer authorizes and enforces the final decision within enterprise-defined boundaries.

In practice, that layer behaves less like a rulebook and more like a colleague with perfect memory and infinite patience. Governance agents observe how data owners decide, identify patterns, and generalize those decisions into candidate policies. They can attach conditions and time limits, trigger the appropriate approval path, and maintain a complete audit trail for every decision - precisely the evidence that regulators and auditors will increasingly demand as AI-driven data use comes under scrutiny.

The building blocks are already in production. In its work with IBM watsonx, Velotix embedded governed intelligence into a conversational assistant that lets users request data in natural language, through the tools they already use, while policy enforcement runs silently underneath - collapsing access timelines from months to minutes. Velotix's inclusion as a representative vendor in the Gartner Guide for Data Security Platforms reflects how central this architecture has become to the data security conversation. The agentic layer is the natural continuation of the same design: governance that learns and enforces continuously, because nothing slower can survive contact with machine-speed consumption. For us, these production building blocks point toward governance that can operate at the same speed as the AI agents and applications consuming enterprise data.

The Arc Only Bends One Way

Step back and the trajectory is unmistakable:

RBAC gave us structure. PBAC gave us intent. The policy brain gave us learning. Agentic data governance gives us autonomy - governance that finally moves at the speed of the systems it governs.

Each transition happened for the same reason: the previous model could not keep pace with how data was actually being used. This one is no different, except in its urgency. Organizations will not get to choose whether agents consume their data. They will only choose whether that consumption is governed by an intelligent layer built for it, or by role hierarchies designed when the fax machine was cutting-edge collaboration technology.

The winners of the agentic era will not be the enterprises with the most rules. They will be the ones whose governance thinks - continuously, contextually, and autonomously - so that every human, every application, and every agent gets exactly the data it should, the moment it should, and not a byte more.

Access control had three eras. The fourth one governs itself.

See how an agentic policy brain works on your own data stack - book a demo.

About Dr. Adi Hod

Dr. Adi Hod is the Founder and CEO of Velotix, a data security platform that uses policy-based access control and AI to govern how humans, applications, and AI agents consume enterprise data. His work focuses on data access governance, the evolution of enterprise permission models, and the security challenges introduced by agentic AI.

Media Contact

citron

*********@gmail.com

Source : velotix

See on IssueWire