GIGAOM

GigaOm Perspective Brief

AI Agents Are Your New Blind Spot

A CISO's Framework for Governing Agentic AI at Enterprise Scale

Commissioned by TrustLogix

Whit Walters, Field CTO, GigaOm

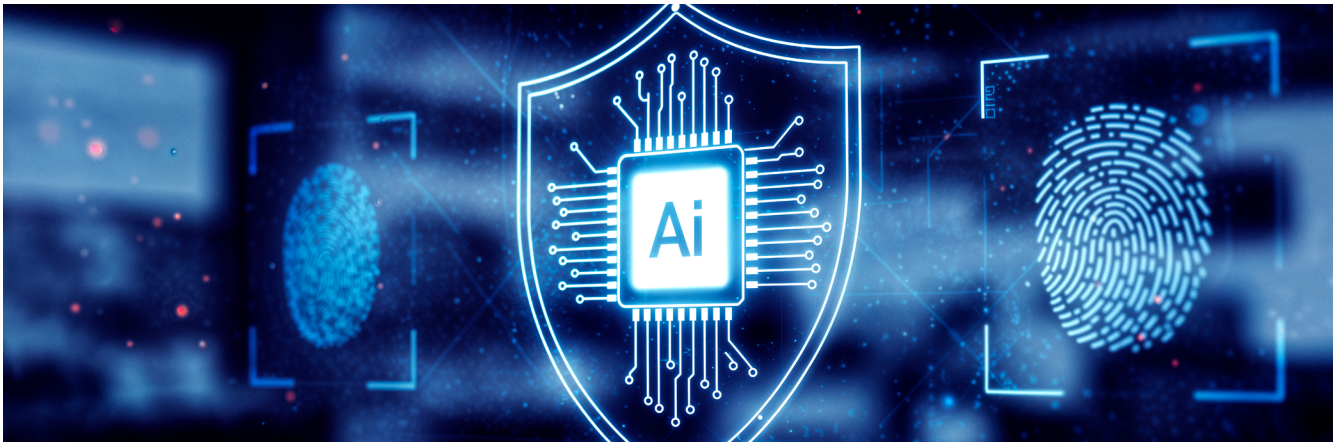

The Governance Question

If you are a CISO right now, you are fielding requests from every business unit to deploy AI agents. By AI agent, I mean any autonomous or semi-autonomous system that plans actions, calls tools or APIs, and queries data sources on behalf of a human user, whether that is a copilot, an orchestrated pipeline, or a fully autonomous agent. Sales wants them. Finance wants them. Operations wanted them yesterday. And you should be enabling them—the productivity gains are real.

But here is the question that should be on your whiteboard: Who is governing these agents once they are live?

I am not talking about prompt injection or LLM vulnerability scanning. Those are real problems, but they are different problems. What I am talking about is the layer where data, identity, and access control converge—where an autonomous agent decides which data to access, on whose behalf, and under what authority. That is the foundational governance layer. If you do not get that right, the security solutions sitting above it are protecting an unlocked house.

Your AI agent does not know it should not access executive compensation data. It does not understand that European customer records cannot leave the EU. It just executes, with whatever permissions it was granted on day one.

What You Are Actually Facing

Based on conversations with CISOs, CDOs, and CTOs across financial services, healthcare, and manufacturing, these patterns are consistent.

You have shadow AI. Business units are spinning up agents faster than your team can track them. Some are on approved platforms. Some are not. The service accounts those agents run on were provisioned months ago with broad access because someone needed to move fast. Developers share credentials across projects rather than requesting new accounts.

You have an accountability gap. When an AI agent queries your data warehouse, your logs show a service account, not the human who triggered the request. Your IAM tells you who a user is. It does not control what an agent acting on their behalf should see. You are missing the link between human identity and agent activity. Without auditing access versus entitlements and enforcing least privilege on the service account, you cannot verify the agent is only accessing what it's entitled to.

You have a privilege boundary you cannot enforce. Agents may hold permissions that their humans do not. The same service account may serve users with very different entitlements, with no control binding the agent's effective access to the human's. If you cannot audit that alignment, you cannot enforce least privilege.

You have compliance exposure you cannot quantify. PII is flowing into prompts. PHI is being summarized by models. Your legal team is asking questions you cannot answer: What data are these agents touching? Who authorized it? Can we prove purpose limitation? Regulators will not accept "we lack visibility into our agent data access" as an answer.

Your CDO has a specific problem. They have already built a data governance program: classifications, stewards, policies. But agents are accessing that data through a consumption pattern the program was never designed for. The policies that govern human access do not automatically extend to agents, and that gap widens every week.

FIVE QUESTIONS EVERY CISO SHOULD BE ASKING RIGHT NOW:

These five requirements extend the core principles your data governance program already enforces, including accountability, transparency, least privilege, and auditability, to the agentic AI surface. Use them as a vendor scorecard: if a solution cannot demonstrate each capability at your scale, it does not belong on your shortlist.

1. What are my AI agents actually accessing, and can I prove it to an auditor?

2. When an agent queries sensitive data, can I trace the request back to the human who triggered it?

3. Are my agents operating with least-privilege access, or running on overprivileged service accounts?

4. If something goes wrong—a data leak, a compliance violation—what is my forensic trail?

5. Is my security team enabling AI innovation, or has it become the place where agent projects go to die?

What Your Shortlist Should Require

Here is what separates adequate from effective in this space. This framework is based on production deployments and GigaOm's analysis of the data security platform market.

1. **Identity propagation across the full access chain *(accountability)*.** An agent runs under a service account with broad access. The human behind the request should not see everything the agent can see. You need a system that combines user identity, agent identity, and the policies for both, enforcing least privilege at the intersection. If Sarah in marketing cannot see raw customer SSNs, her AI assistant should not either. And the agent's own purpose-based policies must be enforced independently—an agent built for financial analysis should not query HR data, regardless of who triggered it.

   Require your vendor to demonstrate a complete audit record for every agent transaction: the human requestor, agent identity, tool invoked, dataset accessed, query executed, policy evaluated, decision rendered, and timestamp.

2. **Just-in-time access *(least privilege)*.** Agents operate continuously under standing credentials that never expire. Just-in-time access grants temporary, scoped entitlements tied to a specific task and automatically revokes them on completion. No orphaned service accounts. No credentials that outlive their purpose.

Verify that entitlements are scoped to a specific task and automatically revoked on completion, and ask the vendor to show you the lifecycle of a credential from grant to expiration with no manual intervention required.

3. **Enforcement, not just visibility *(policy enforcement)*.** Masking sensitive fields before data enters an AI context window. Blocking unauthorized queries at the data layer. Enforcing row-level security before data leaves the platform. Visibility is table stakes. Enforcement changes your risk posture.

   Test whether the platform blocks an unauthorized query before data is returned, not after. Ask for a live demonstration of field-level masking and row-level filtering applied before data enters an AI context window.

4. **Auditing at enterprise scale *(traceability)*.** Managing one agent with a few hundred users is almost a manual problem. A hundred agents with tens of thousands of users across multiple data platforms is logarithmic complexity. Test for this explicitly, not at pilot scale, but at a thousand agents in production.

   Require the vendor to demonstrate policy evaluation latency, conflict resolution across overlapping entitlements, and audit completeness at your projected agent count, not a sandbox demo with a handful of test users. Scaling agent governance is a multi-stakeholder exercise: define success criteria across security, data engineering, platform operations, and business units before greenlight.

5. **Continuous monitoring with actionable alerting *(transparency)*.** This means behavioral baselining for agent activity, with deviations routed into your existing SIEM and incident management workflow, not a separate dashboard your team does not have time to watch.

   Confirm that behavioral anomalies generate alerts routed directly into your existing SIEM and incident response workflows, and ask the vendor to show you a baseline deviation scenario from detection through ticket creation.

The Architecture Decision

Architecture has real implications for your deployment and your team's operational burden.

Evaluate architectural flexibility before you commit *(architectural versatility: mature AI governance deployments rarely fit a single enforcement model)*. Enterprise data environments are not uniform; the right enforcement point depends on the data flow, the agent architecture, and the access pattern in question. You did not move to Snowflake or Databricks to add latency, and a well-architected governance solution should not require that trade-off. Look for solutions that offer both a native enforcement path, where policy decisions happen at the data platform itself, and a gateway option that can broker and inspect agent interactions with MCP services where centralized control is warranted. A platform that forces you into one model will create gaps in the other. The goal is a unified policy control plane that can enforce consistently regardless of which path the data takes.

Demand platform coverage across your data estate *(unified: one policy control plane governing your entire data landscape, not siloed tools per platform)*. A point solution that covers Snowflake but not Databricks, or cloud databases but not on-prem, creates governance gaps. Evaluate for unified policy management across your full landscape.

Insist on metadata-only architecture *(managed-scope: govern access decisions without ever touching or transiting the sensitive data itself)*. Solutions that need to see your actual data to govern it create their own security problems. The right architecture governs access based on identity, classification, policy, and context, keeping sensitive data where it lives.

Where TrustLogix TrustAI Fits

TrustLogix is one of the vendors that addresses the framework above, and its TrustAI module is specifically designed for the agentic AI governance gap.

Their platform consists of three modules: TrustAccess for data access controls, TrustDSPM for data security posture management, and TrustAI for agent governance, all under a single license and unified policy control plane. This is not a bolt-on AI feature. TrustAI extends a platform already in production at Global 500 enterprises, securing petabyte-scale data environments across financial services, healthcare, semiconductor manufacturing, and telecommunications.

In GigaOm's Radar for Data Security Platforms, TrustLogix is positioned as a Leader and Fast Mover, with particular strength in access security, compliance reporting, and data anonymization. The Radar also identified areas for expansion—notably unstructured data and SaaS application coverage—worth factoring into a long-term evaluation.

The architecture is agentless and metadata-only, with enforcement options designed to fit your data environment. TrustAI provides dynamic authorization decisions through an architecture designed to meet enterprise environments as they are. Organizations can deploy the TrustLogix MCP Service for native, platform-level enforcement, where policy decisions are evaluated out-of-band and enforcement happens directly at the data platform. For workflows requiring centralized inspection and brokering of agent interactions, the TrustLogix MCP Security Gateway provides an active control point before queries execute. Both options evaluate requests against the agent's purpose-based entitlements and the human user's identity-based permissions, and both feed the same audit and monitoring pipeline.

Monitoring is near-real-time rather than inline: continuous activity tracking with alerting on a 30-to-40-minute cycle, feeding into existing security operations tools like Splunk, Microsoft Teams, and Slack. The platform integrates with Snowflake, Databricks, AWS, Azure, Oracle, SQL Server, and AI agent frameworks through MCP.

The Window Is Open. It Will Not Stay Open.

AI adoption is not slowing down. The question is whether you enable it securely or end up in reactive mode after an incident.

The organizations getting this right are treating AI governance as infrastructure—not a policy document, not an approval process, but an architectural layer that enforces controls automatically, at the data layer, at machine speed.

TrustLogix is one of the vendors that can deliver this. But regardless of which vendors you evaluate, the capabilities in this brief—identity propagation, just-in-time access, native enforcement, audit at scale, continuous monitoring—are what your shortlist should require.

Analyst's Take

TrustLogix is solving the right problem at the right layer. The data-identity-access boundary is where AI governance needs to happen, and its architectural versatility avoids the complexity trap that limits many security solutions at scale. The fact that TrustAI extends a production-proven platform gives it credibility that pure-play AI security startups cannot match.

The governance gap in agentic AI does not belong to the CISO alone. When non-human identities are querying classified, governed business data at machine speed, the CDO has as much at stake as the security team. The organizations getting this right are treating it as a shared infrastructure problem and evaluating platforms accordingly. For both personas: assess coverage gaps honestly, test for scale explicitly, and confirm that your enforcement architecture does not force a trade-off between security and performance. Start with visibility before enforcement. The organizations that figure out AI governance in 2026 will deploy agents at scale in 2027 without the security team being either the bottleneck or the liability.

Report Methodology

TRUSTLOGIX

THE GIGAOM PERSPECTIVE BRIEF provides an independent analyst framework for evaluating an emerging technology category, developed through vendor briefing, product documentation review, and analysis informed by GigaOm's ongoing research and market coverage including relevant Radar reports. The brief is buyer-centric by design, leading with the problem and evaluation criteria before addressing the featured solution, to ensure standalone value for enterprise decision-makers regardless of vendor selection.

ABOUT THE AUTHOR

Whit Walters is a Field CTO at GigaOm with 30 years of enterprise technology experience, including CTO roles at multiple startups, leadership positions within the Google Cloud partner ecosystem, and domain expertise in AI, machine learning, data architecture, and cloud security.

GIGAOM

GigaOm provides technical, operational, and business advice for IT's strategic digital enterprise and business initiatives. Enterprise business leaders, CIOs, and technology organizations partner with GigaOm for practical, actionable, strategic, and visionary advice for modernizing and transforming their business. GigaOm's advice empowers enterprises to successfully compete in an increasingly complicated business atmosphere that requires a solid understanding of constantly changing customer demands.

GigaOm works directly with enterprises both inside and outside of the IT organization to apply proven research and methodologies designed to avoid pitfalls and roadblocks while balancing risk and innovation. Research methodologies include but are not limited to adoption and benchmarking surveys, use cases, interviews, ROI/TCO, market landscapes, strategic trends, and technical benchmarks. Our analysts possess 20+ years of experience advising a spectrum of clients from early adopters to mainstream enterprises.

GigaOm's perspective is that of the unbiased enterprise practitioner. Through this perspective, GigaOm connects with engaged and loyal subscribers on a deep and meaningful level.