GREYREGENT FOUNDER'S NOTES, NO. 2

Shadow AI:

The Enterprise Within the Enterprise

Why banning artificial intelligence will not restore corporate control.

Geoffrey Gilroy · Founder & Principal, GreyRegent · November 2026 · 10 min read

We have been here before

I have spent most of my career in telecommunications, enterprise technology, technology finance and Enterprise Performance Management. The current debate about shadow AI is remarkably familiar.

Employees buying their own software. Departments procuring technology without IT. Business units building their own applications because the corporate alternative was inadequate or took eighteen months to arrive. We called it shadow IT, and the enterprise answered with procurement policies, approved vendor lists and governance committees. Yet here we are again.

This time the technology can interpret information, initiate transactions and delegate work to other systems. It is no longer an application running outside corporate control. It is an operational actor exercising authority the enterprise never granted.

The unit of ungoverned technology is no longer the tool. It is the decision.

The security industry is right, and misses the point

The cybersecurity industry has discovered shadow AI, and it tells a familiar story: employees are feeding corporate data into unapproved applications, the compliance exposure is substantial, and there is another platform to buy.

The figures support the alarm. Among the breached organisations IBM studied in 2025, one in five attributed a breach to shadow AI (IBM). Eight in ten employees surveyed by UpGuard use unauthorised AI tools. More striking still, 68% of the security leaders UpGuard surveyed build unauthorised AI into their daily work (UpGuard).

The people responsible for preventing the problem are taking part in it. When employees, senior managers and security professionals all route around a policy, the policy deserves as much scrutiny as the people breaking it.

Shadow AI is not the cause of a governance failure. It is the evidence of one. Where the enterprise has not provided a fast, governed route to intelligence, employees have built their own. That is a supply problem and an architecture problem, not only a security problem.

Shadow chat and shadow agents

There is a real difference between an employee pasting text into an unapproved chatbot and an autonomous agent holding credentials to corporate systems.

Shadow chat is employees using unsanctioned conversational AI: uploading documents, processing information and exposing corporate data. Most of today's research measures this.

Shadow agents are autonomous or semi-autonomous systems operating outside the enterprise's framework of identity, authority, financial control and accountability. Their building blocks are already spreading: agents, personal API keys, persistent integrations and the Model Context Protocol (MCP), which connects AI systems to enterprise tools. Their adoption is not yet measured at scale, so what follows about them is a forecast, not a finding.

The attack surface, however, is already being used. CrowdStrike documented adversaries exploiting legitimate generative AI tools at more than 90 organisations in 2025, and activity by AI-enabled adversaries rose 89% in a year (CrowdStrike). That is evidence about attackers, not about shadow agents. But it shows what ungoverned authority will be exposed to.

A shadow spreadsheet stored data. A shadow agent acts.

GOVERNANCE DIMENSION	SHADOW IT	SHADOW AI
Data	Uncontrolled storage and processing	Uncontrolled processing and disclosure
Identity	Unmanaged users and applications	Unmanaged human and non-human identities
Execution	Typically employee-driven	Potentially autonomous
Financial exposure	Unapproved expenditure	Variable consumption and delegated expenditure
Accountability	Unauthorised applications	Unauthorised decisions and actions
Auditability	Incomplete	Potentially absent across execution chains

Conventional software can also automate decisions. What sets agentic AI apart is the combination of adaptive reasoning, autonomous execution and delegated authority in one system.

A supply problem, not a discipline problem

The broader evidence points the same way.

- **Access is growing faster than governance.** Deloitte's 2026 survey of 3,235 leaders in 24 countries found that workforce access to sanctioned AI rose from under 40% to around 60% of workers in one year. Close to three-quarters of companies plan to deploy agentic AI within two years, yet only 21% of them report a mature agent governance model (Deloitte).
- **Shadow AI is costly.** Among IBM's 600 breached organisations, those with high levels of shadow AI faced breach costs averaging $670,000 more than those with little or none.
- **Prohibition is not holding.** In a Gartner survey of 302 cybersecurity leaders, 69% of organisations suspect or have evidence of employees using prohibited public generative AI. Gartner expects more than 40% of enterprises to suffer security or compliance incidents linked to unauthorised shadow AI by 2030 (Gartner).
- **Awareness is not enough.** UpGuard found that the employees who received AI safety training are among the most frequent users of unauthorised tools.

Demand for intelligence is effectively universal, at every level of seniority. The governed supply of it is scarce. Employees are not rejecting governance. They are choosing the only route that works.

Who gave the agent authority?

In The €50 Rule, I argued that access to intelligence must not be confused with authority to spend. The same principle applies to authority to act.

An employee does not gain the right to approve a purchase order because they can log into the ERP system. A controller does not gain unlimited spending authority because they can see the corporate bank account. An autonomous agent should be treated no differently.

The enterprise already knows how to do this. It has spent decades building approval hierarchies, financial mandates, access controls and accountability. AI does not invalidate those principles. It adds a new kind of actor to which they must apply.

CORPORATE GOVERNANCE QUESTION	PROPOSED CONTROL
Who is acting?	Agent identity and a named human owner
What may it access?	Data classification and access policy
What may it do?	Operational mandate
What may it spend?	Financial mandate
What did it actually do?	Execution audit trail and AI Charging Data Records (AI-CDRs)
Who accepted the result?	Independent validation
Was it economically justified?	Total Cost of Accepted Outcome (TCAO)

Both AI-CDRs and TCAO come from The €50 Rule. AI-CDRs borrow from telecom charging: an accountable record of consumption, execution and authority, whether or not an external invoice follows. TCAO measures the full cost of each independently accepted outcome, including failed attempts, infrastructure and human rework.

The framework is not new. Its subject is.

Fifty euros does not buy fifty thousand euros of authority

Consider a procurement agent asked to find alternative suppliers, analyse their terms and prepare a recommendation.

It receives a financial mandate of €50 for AI consumption, with a warning at €40 and a stop at €50 unless someone independent authorises more. Its operational mandate lets it read procurement data and research suppliers. It may not change supplier master data, approve contracts or issue purchase orders. Any agents it delegates to share the same €50. It cannot raise its own budget or widen its own permissions.

An agent authorised to spend €50 on intelligence has not been authorised to place a €50,000 order. Financial authority and operational authority are separate controls, and both must be enforced before execution, not reconstructed from an audit report afterwards.

Decentralised intelligence, centralised governance

The reflexive response to new technology is centralisation: one approved vendor, one approved application, one approved route. With AI, that is a commercial mistake. Model quality, capability and pricing change monthly, and an enterprise that ties its governance to one provider surrenders both leverage and independence.

The distinction that matters is between centralised governance and centralised execution. Policy, identity, financial authority and accountability belong to the enterprise. Execution can stay distributed across local models, external providers and specialist agents, provided each environment meets the enterprise's security and data requirements. Policy is defined centrally and enforced wherever the work runs. The provider may change; the mandate should not disappear with it.

This architecture has one obvious danger. The governance layer becomes the most powerful component in the enterprise: whatever can grant authority, change mandates or suppress audit records concentrates more risk than any single agent. It must therefore obey the rules it enforces:

- separation of administrative and operational duties;
- dual authorisation for material changes to sensitive mandates;
- independently controlled credentials and privileged access;
- tamper-resistant audit records and independent security reviews;
- a technical ban on any agent raising its own authority or limits;
- defined failure behaviour, suspending high-risk execution when authority cannot be verified.

Who governs the governors? An enterprise that cannot answer has not contained its risk. It has concentrated it.

Governance needs its own business case

Governance that costs more than the risk it removes is not governance. It is overhead. And governance that turns every AI request into a ticket will lose to the browser tab.

If an employee can get an answer from an unauthorised tool in thirty seconds, but needs three days and four approvals to get it through the sanctioned route, the outcome is predictable. The enterprise has created its own competitor.

Controls must be proportionate. A low-risk research task should not face the approvals required for an agent that changes financial records or handles sensitive customer data. Govern the exception. Automate the authorised.

Minimum viable governance in ninety days

This does not require a multi-year programme. A bounded first phase can be tested within a quarter:

PERIOD	OBJECTIVE	MINIMUM DELIVERABLE
Days 1–30	Visibility and containment	Initial inventory, named owners, spending baseline, immediate controls on high-risk agents
Days 31–60	Governed alternatives	Sanctioned access, operational mandates, enforceable financial limits, execution records
Days 61–90	Evaluation	Comparison of adoption, friction, exposure and the economics of accepted outcomes

Three existing sources give a practical start: identity and access records, AI-related spending on cards and expense claims, and network and SaaS audit logs.

Discovery can become an industry of its own, and by the time an inventory is finished the technology has moved on. Discovery and control must therefore run in parallel. From day one, any agent found with access to sensitive data, transactional systems or company money needs a named owner, restricted credentials, a defined mandate, an enforceable spending limit and a way to be suspended immediately. Where that cannot be put in place, high-risk execution stops until it can. An incomplete inventory is understandable. Unrestricted authority is not.

The aim is not to punish employees who exposed gaps in the corporate technology estate. It is to close those gaps without compromising security, and to give legitimate needs a safe route.

Measure the controls and their friction

Governance should be judged on two counts: whether exposure falls, and whether the sanctioned route gets easier.

MEASURE	WHAT IT SHOWS
Unsanctioned AI usage as a share of all AI usage	Whether observable exposure is falling
Time to sanctioned AI access	Whether barriers are unnecessary
Approval latency	How much friction governance adds
Policy violations	Whether controls are effective
Governance cost per task	What the control layer itself costs
TCAO	The full cost of each accepted business outcome

Read these with care. Network and proxy logs are a starting point; browser telemetry, SaaS audit records, AI gateways, card transactions and expense claims fill in more, within privacy and employment safeguards. Falling visible traffic may mean shadow AI has moved, not disappeared. Rising violations may mean better detection, not worse behaviour. Implementation and running costs belong in the assessment too.

If, after ninety days, the only result is a longer inventory and a larger governance department, the programme has failed.

Sovereignty over intelligence

In telecommunications I watched enterprises grow dependent on infrastructure that was indispensable yet largely outside their control. AI raises the same question: who controls the intelligence the enterprise now depends on?

Leaders are asking it. In Deloitte's survey, 83% call sovereign AI important to strategic planning and 77% weigh a vendor's country of origin when choosing. But data residency and jurisdiction, on their own, do not amount to sovereignty. Sovereignty means retaining authority, accountability, portability and control.

I propose a simple test.

The Provider Substitution Test. Replace the intelligence provider. What happens to the enterprise's identities, operational mandates, financial limits, audit history and acceptance criteria?

If they must be rebuilt each time the model provider changes, the enterprise does not own its governance. It is renting it.

A replacement model must still be qualified for security, capability and acceptance, and migration has real costs. The test is not whether a switch is free. It is whether the enterprise keeps its authority while it switches: identities and ownership stay under its control, mandates and spending limits remain enforceable, audit history stays accessible and reconcilable, and acceptance criteria remain its own.

Conclusion

Shadow AI is the enterprise within the enterprise: intelligence operating outside the structures through which the business normally delegates and exercises authority.

We have spent decades building financial controls, operational mandates, security policies and accountability. Artificial intelligence gives no reason to abandon them. Autonomous execution makes them more important.

Prohibition alone will not solve this. The organisations that succeed will make governed intelligence faster, better and cheaper to use than the ungoverned alternative.

In The €50 Rule, I asked who gave the agent authority to spend. The question here is broader: who gave it authority to act? Both come down to one thing: whether the enterprise remains in control of the technology it increasingly depends on.

Decentralised intelligence. Centralised governance. Enterprise sovereignty.

Disclosure and method

Disclosure. Geoffrey Gilroy is Founder & Principal of GreyRegent, an independent house of enterprise. GreyRegent develops budqi, a governed intelligence layer over the systems an organisation already uses, and other propositions in enterprise AI governance, intelligent orchestration and Delegation of Authority for autonomous systems, and so has a commercial interest in the adoption of the principles in this note. No cited vendor sponsored, commissioned or reviewed it.

The governance architecture, procurement scenario, ninety-day programme and Provider Substitution Test are proposals, not verified production implementations or demonstrated results. Survey findings show associations, not causes: IBM's figures apply to breached organisations only, UpGuard does not publish weighting or margins of error, and Gartner's 2030 figure is a forecast. Whether this approach reduces shadow AI at enterprise scale is a proposition to be tested.

Sources

All checked against the publishers' own releases on 26 November 2026.

1. Deloitte AI Institute, The State of AI in the Enterprise: The Untapped Edge, 21 January 2026. Survey of 3,235 business and IT leaders in 24 countries and six industries, August–September 2025.
2. IBM and Ponemon Institute, Cost of a Data Breach Report 2025, 30 July 2025. 600 organisations breached between March 2024 and February 2025.
3. UpGuard, State of Shadow AI, 10 November 2025. Surveys of 1,020 employees in the US and UK (Prolific, July–August 2025) and 542 security leaders (Dynata, August 2025).
4. Gartner, Gartner Identifies Critical GenAI Blind Spots That CIOs Must Urgently Address, 19 November 2025. Survey of 302 cybersecurity leaders, March–May 2025.
5. CrowdStrike, 2026 Global Threat Report, 24 February 2026. Cited as evidence of the AI attack surface, not of shadow-agent adoption.
6. Geoffrey Gilroy, *The €50 Rule: Why AI Agents Need a Spending Limit*, GreyRegent Founder's Notes, No. 1, November 2026. Introduces financial mandates, AI-CDRs and TCAO.