



ITF SECURITY MATTERS | LLOYD'S LIBRARY | 17 SEPT 2026

AI Governance: Tracking and Controlling Hidden Risks

Securing AI: visibility, control and shadow AI risk



Seetal R Patel

Cybersecurity Specialist, Cato Networks (SE – UK Majors)

WHO WE ARE



The Network and Security Cloud for the AI Era

THE PLATFORM

100+

PoPs / locations

60,000+

Sites & clouds

2.5M+

Remote ZTNA users

THE COMPANY

\$1B+

Total funding raised, \$4.8B
valuation

1,800+

Employees across global offices

4,800+

Enterprise customers worldwide

200+

Countries & territories served



Shlomo Kramer, CEO

Founder of Check Point & Imperva

THE PROBLEM

AI moves faster than the controls around it

These are the four risks we see most often.



DISCOVER

Shadow AI

Teams use AI tools without IT supervision, creating blind spots in data governance.

1 in 5

breaches linked to shadow AI

IBM Cost of a Data Breach, 2025



PROTECT

Data Leakage

Prompts and responses can expose sensitive data to AI providers or third parties.

86%

have no visibility into AI data flows

AIUC-1 Consortium & Stanford, 2026



GOVERN

Compliance & Governance

New regulations change accountability demands around AI use, outputs and auditability.

63%

lacked AI governance policies

IBM Cost of a Data Breach, 2025



DEFEND

AI Attacks & Vulnerabilities

Prompt injection, jailbreaks and attacks crafted for agentic systems are rising rapidly.

32%

of AI-app attacks leveraged prompts

Gartner, 2026



Shadow AI: the risk your insureds will tell you they don't have

47%

of enterprise gen-AI activity runs through personal accounts, outside corporate identity and logging

Akamai / Netskope, 2026

43%

of AI-related breaches now involve shadow AI (up from 20% a year earlier), averaging \$5.39m per breach

IBM Cost of a Data Breach, 2026

#1

data type uploaded to unsanctioned AI tools is source code, across 858k DLP events analysed

Verizon DBIR, 2026

Why it matters at the box: a proposal form that says "we don't use AI" is now almost certainly wrong. Browser extensions, Copilot-style features inside sanctioned SaaS, and personal ChatGPT accounts move client data, source code and PII out of the perimeter one prompt at a time, and if the insured can't see it, neither can the forensics team after a claim.



Governance frameworks exist on paper. Wordings and regulators are now testing them.

The London market

- 93% of LMA survey respondents have or are building AI governance frameworks. One in four still uses general third-party risk processes for AI suppliers. [LMA, April 2026](#)
- AI-related losses may be covered, limited, or excluded. Businesses should check their policy wording before an incident. [DAC Beachcroft, April 2026](#)
- Lloyd's underwriters now offer dedicated AI liability cover. Strong governance, testing, and human oversight can support an insurance submission. [Armilla/Chaucer, April 2025](#)

The regulatory clock


- **2 Aug 2026** EU AI Act Article 50 transparency duties live and enforceable (chatbots, AI-generated content)
- **2 Dec 2026** Gen-AI labelling grace period ends; new Article 5 prohibitions apply
- **Through 2026** UK Cyber Security & Resilience Bill widens mandatory incident reporting to MSPs, data centres and critical suppliers
- **2 Dec 2027** High-risk AI obligations (Annex III) apply, after the Digital Omnibus deferral




Five controls that make AI risk visible, and five questions worth asking at the box

Control	What good looks like	The underwriter's question
 Inventory	Continuous discovery of every AI tool, model, extension and MCP/agent endpoint, including features inside sanctioned SaaS	<i>"How many AI tools are in use, and how do you know?"</i>
 Identity-bound access	Zero Trust: AI reached only through corporate identity via SASE/SSE, personal-account access blocked or coached	<i>"Can staff reach ChatGPT on a personal login from a corporate device?"</i>
 Data controls in the prompt path	Inline DLP on prompts and uploads; classification decides what may leave; logs retained for forensics	<i>"What stops a claims file being pasted into a chatbot, and would you see it?"</i>
 Kill switch & playbook	A tested way to halt an AI system or agent mid-incident; AI scenarios included in resilience testing	<i>"Who can switch it off, and have you rehearsed it?"</i>
 Evidence trail	Risk-tiered register, owner per system, vendor AI addenda, training records: the LMA toolkit's five themes made auditable	<i>"Show me the register, not the policy."</i>

Full SASE platform, purpose built



The network becomes the AI control plane



Cato covers three AI surfaces

Each one works differently. Each one needs a different response.



Users

The AI your employees use

Unsanctioned usage. Sensitive data exposed in AI interactions with no visibility.

-
- Shadow AI discovery
 - Prompt & response inspection
 - Usage policy enforcement



Applications

The AI your organisation builds

Prompt injection and data leaks that traditional controls miss entirely.

-
- Prompt injection detection
 - PII masking & blocking
 - App-level audit trail



Agents

The AI that acts autonomously

Real actions across tools with no human in the loop and no visibility.

-
- Agent inventory & mapping
 - Tool call enforcement
 - Session-level audit

Cato closes the AI security gap

Inspect prompts, responses and tool calls to detect risk and enforce policy in real time.



EVERY SURFACE

Protect every AI interaction

Inspect every prompt and response inline: detect PII, jailbreaks and prompt injection before they cause damage.



CONVERGED

One solution across AI use cases

Enforce consistent policy across AI usage, applications and agents in one AI-native solution. No new appliances, no new agents.



AGENTS

Govern autonomous behaviour

Control agentic behaviour with runtime protection for prompts, tool usage and autonomous actions, across every surface.



THANK YOU!