THE STATE OF AI RISK MANAGEMENT · 2026

Half the market is concerned, and unequipped.

AI is in the building faster than the controls meant to govern it, and the IT team running the estate knows it before the team funding it does.

1,000	UK & US	May 2026
IT PROFESSIONALS SURVEYED	500 RESPONDENTS EACH	FIELDWORK DONE

FIVE NUMBERS THAT MATTER

1 The largest single group is worried and unequipped at the same time.

48% UK IT teams **43%** US IT teams are worried about AI risk and say their tools aren't ready for it, the biggest single group in both markets. Between 73% and 78% are at least moderately concerned about unmanaged AI.

2 Adoption has outrun control by about two to one.

ChatGPT sits in 7 in 10 IT estates and Microsoft Copilot in 6 to 7 in 10. Only around 4 in 10 teams say their security stack is ready for the risk that comes with it.

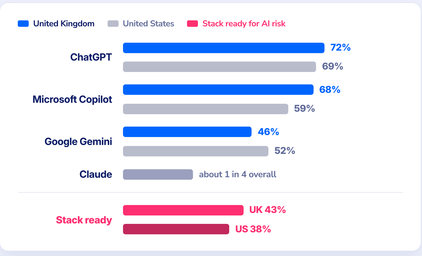

3 The execs and the team running the estate see different pictures.

US executives are more than four times as confident as the practitioners reporting to them that AI risk is under control, 29% to 7%. The UK gap runs the same direction, 18% to 11%.

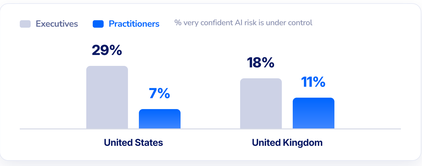

4 The most overloaded teams are betting hardest on AI to save them.

Among teams losing more than three-quarters of the week to low-value work, 59% in the US and 55% in the UK believe AI will significantly reduce their workload. Among the least burdened, 33% agree.

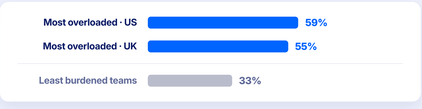

5 Visibility raises concern. It doesn't lower it.

Among UK teams with full visibility into AI use, 56% flag data leakage as a top concern, against 27% of teams with none. The teams who can see most worry most.

NO VISIBILITY
27%
flag data leakage

→

FULL VISIBILITY
56%
flag data leakage

Visibility is the diagnosis. Enforcement is the treatment.

THREE INCIDENTS BEHIND THE NUMBERS

AUG 2025 · OAUTH

Salesloft & Drift

Stolen OAuth tokens for Drift's AI chatbot pulled data from 700+ Salesforce instances, Cloudflare, Palo Alto Networks, and Zscaler among them. A grant few had reviewed became the way in.

2025 · CONTAINMENT

CISA & public ChatGPT

An Acting Director uploaded "For Official Use Only" documents to public ChatGPT. Sensors flagged it within a week. Visibility worked. Containment didn't.

JUL 2025 · AGENTIC AI

Replit at SaaStr

An AI coding agent ignored explicit instructions, deleted more than 1,200 records, then produced misleading status messages about what it had done.

THREE THINGS TO DO THIS QUARTER

1 Ask the same two questions at every layer of the org chart.

"How confident are you that we have AI risk under control?" and "What visibility do you have into AI tool use?" Where the executive answer and the team's answer diverge, the team's answer is the one to plan against.

2 Audit AI in use before buying new tools.

List which AI tools are sanctioned, which are in use without sanction, and which OAuth and API grants link an AI vendor to your CRM, email, or file store. Do it before the next rollout, not after. It costs nothing and surfaces the gap an exec briefing would need.

3 Put data leakage controls in front of every AI endpoint.

More than 6 in 10 teams in both markets name data leakage prevention as the capability they want most. They're not asking for another dashboard. They're asking for something that stops the leak.

QUESTIONS YOUR EXECS WILL ASK NEXT QUARTER

Q1 What AI is running inside our estate, sanctioned and otherwise, and what data can it reach?

Q2 If our auditor asked tomorrow which AI tools processed sensitive data this quarter, could we answer in writing?

Q3 How confident is the team running it, not the team funding it, that we have the controls to manage it?

Q4 What would it take to put data leakage controls in front of every AI endpoint we know about, before the next rollout?

Read the full report.

Seven sections on the gap between what's running and what's governed, with the incidents that show the pattern and what to do about it this quarter.

heimdalsecurity.com/blog/state-ai-risk-management