

The Agentic Enterprise Needs an Audit Trail

UK technology and security leaders will soon have to prove what their systems and agents did, sometimes within 24 hours. This briefing covers what the new reporting rules and the AI governance gap ask of your evidence.

24 hrs

to notify the regulator and NCSC of a significant incident under the Cyber Security and Resilience Bill¹

£1.9bn

estimated UK cost of the Jaguar Land Rover attack, the most damaging cyber event recorded in the UK so far³

92%

of organisations with an AI-related breach lacked proper AI access controls⁵

15×

growth in active agents in Microsoft 365 in one year, rising to 18× in large enterprises⁷

44%

of ransomware cases where the attacker was first to notify the victim⁴

01

The 24-hour reporting clock

The UK is taking different approaches to AI and cyber regulation, and both reach the same leadership team. The Cyber Security and Resilience Bill, now before the House of Lords, would require regulated organisations to notify their regulator and the NCSC of a significant incident within 24 hours and file a full report within 72.¹ Near misses would also be reportable, and serious breaches could draw fines of up to £17 million or 4% of worldwide turnover, whichever is higher.²

The NCSC handled 204 nationally significant attacks in the year to September 2025, up from 89 the year before.² The Cyber Monitoring Centre estimates the M&S and Co-op incidents cost £270–440m and puts the Jaguar Land Rover attack at roughly £1.9bn, the most damaging cyber event recorded in the UK so far.³

The Bill also brings managed service providers and data centres into regulation, with a duty to warn customers affected by an incident. For firms outside its scope, a supplier's notice could be the first indication of their own exposure.²

Meeting a 24-hour deadline depends on what your team can establish in the first day, starting with what happened and how far it spread. That's harder when someone else finds the intrusion first. Mandiant reports that in ransomware cases, the attacker notified the victim 44% of the time.⁴

02

What agents inherit from your data estate

The data estate panel describes AI as “the most unforgiving audit ever run”. An agent or copilot can read whatever its identity can reach, so permissions nobody has reviewed in years become live the day it's switched on. Microsoft reports 15 times more active agents in Microsoft 365 than a year earlier, and 18 times more in large enterprises.⁷

IBM's 2026 breach study found that 92% of organisations with an AI-related breach lacked proper AI access controls, and 68% of all breached organisations had no policy governing AI use.⁵ Among UK businesses using or adopting AI, 24% have practices in place to manage its cyber risk.⁶

When an agent does something unexpected, auditors will want a record of the data it accessed and the identity it acted under. That record exists only if the logs were collected and kept. For firms with EU exposure, the EU AI Act makes this explicit: high-risk systems, including those used for credit scoring and employment decisions, must log events automatically, and providers must keep those logs for at least six months.⁸

AI in the SOC has to show its reasoning

The SOC panel asks whether AI still adds capacity once a pilot ends. IBM found that organisations using AI and automation extensively in security operations reduced average breach costs by about \$1.93 million and contained breaches around 65 days sooner.⁵

Those gains depend on analysts accepting the output. If they can't see why an AI closed or escalated an alert, most will repeat the work by hand. In Microsoft's survey of 20,000 AI users, quality control of AI output ranked as the most important human skill as AI takes on more work.⁷

A confidence score on its own doesn't help an analyst verify a decision or an auditor accept it. When the AI names the detection rule that fired and the evidence behind it, a reviewer can check it quickly, and the same record can go into the 72-hour report.

Reporting to the board under the new regime

Among large UK businesses, 68% have a board member responsible for cyber security, and 69% identified a breach or attack in the past year. About one in four large businesses have no formal incident response plan.⁶

Activity metrics such as alerts triaged or rules deployed don't show a board whether your team could meet a 24-hour deadline. Measures closer to what a regulator will ask about include time to contain an incident and how much of the estate your team can see. It's also worth checking whether closed cases carry enough evidence to support a report.

SOURCES

1. UK Parliament, Hansard, House of Lords Second Reading, Cyber Security and Resilience (Network and Information Systems) Bill, 14 July 2026.
2. Travers Smith, "The UK's new Cyber Security and Resilience Bill", citing the NCSC Annual Review 2025. traverssmith.com
3. Cyber Monitoring Centre event assessments: M&S and Co-op (June 2025), and Jaguar Land Rover (October 2025).
4. Google Cloud / Mandiant, M-Trends 2026. cloud.google.com
5. IBM and Ponemon Institute, Cost of a Data Breach Report 2026 (July 2026). ibm.com/reports/data-breach
6. DSIT and Home Office, Cyber Security Breaches Survey 2025/2026 (April 2026). gov.uk
7. Microsoft, 2026 Work Trend Index Annual Report (May 2026). microsoft.com/worklab
8. Regulation (EU) 2024/1689 (EU AI Act), Articles 12 and 19, and Annex III.

Further reading from The Visibility Layer by Graylog on Medium: "What Responsible AI Requires Is a Log, Not a Statement", "The SOC Doesn't Need a Bigger Team. It Needs an Agentic One", "Your Board Wants Outcomes. Your SIEM Produces Activity Reports", and "Nobody Trusts AI Agents. Everyone Deploys Them Anyway". medium.com/the-visibility-layer

Questions to raise at the Think Tank

10:00 Britain's two-track bet

If our largest managed service provider reported an incident at 9am, how long before we'd know which of our systems it touched?

11:20 AI implementation in the SOC

When our AI closes an alert, can an analyst see why in under a minute?

12:10 What AI reveals about your data estate

Which of our agents can read data whose access no human has reviewed this year?

15:40 AI, risk and responsibility

Who signs the 24-hour notification, and what evidence will be in front of them when they do?

Where Graylog fits

AI is giving attackers more speed and scale, and it's also giving security teams new ways to keep up. Graylog, the sovereign SIEM and log management platform for the agentic era, covers both sides: it monitors and audits the AI agents in your environment, and it uses AI to help your analysts investigate and report, with its reasoning shown.

Agent and copilot activity shows up in cloud audit trails, identity logs, and application logs. Graylog collects those through prebuilt inputs for AWS, Azure, and Microsoft 365, and through APIs, syslog, or Kafka for anything else, so your agents' actions sit alongside the rest of your security data. When an auditor asks about a specific period or system, your team can search every storage tier from one place and retrieve only the logs that match. Data held in the data lake isn't charged until you pull it in, so keeping agent activity for the long term doesn't mean paying to process all of it.

When a security incident occurs, Graylog builds the investigation automatically with the related events attached, and AI Reporting drafts the incident report for an analyst to review. Graylog's own AI explains its reasoning, and you can connect your own LLM through MCP. Run Graylog on-premises or in your own cloud tenancy, and your logs stay in the jurisdiction you choose.

Talk to the Graylog team at the Think Tank, or visit graylog.com to book a demo.