CIONET UK · SEPTEMBER 2026

The Chief Connecting Officer

Surviving the AI-Driven Cyber Battlefield: From Nation-State Attacks to Board-Level Resilience

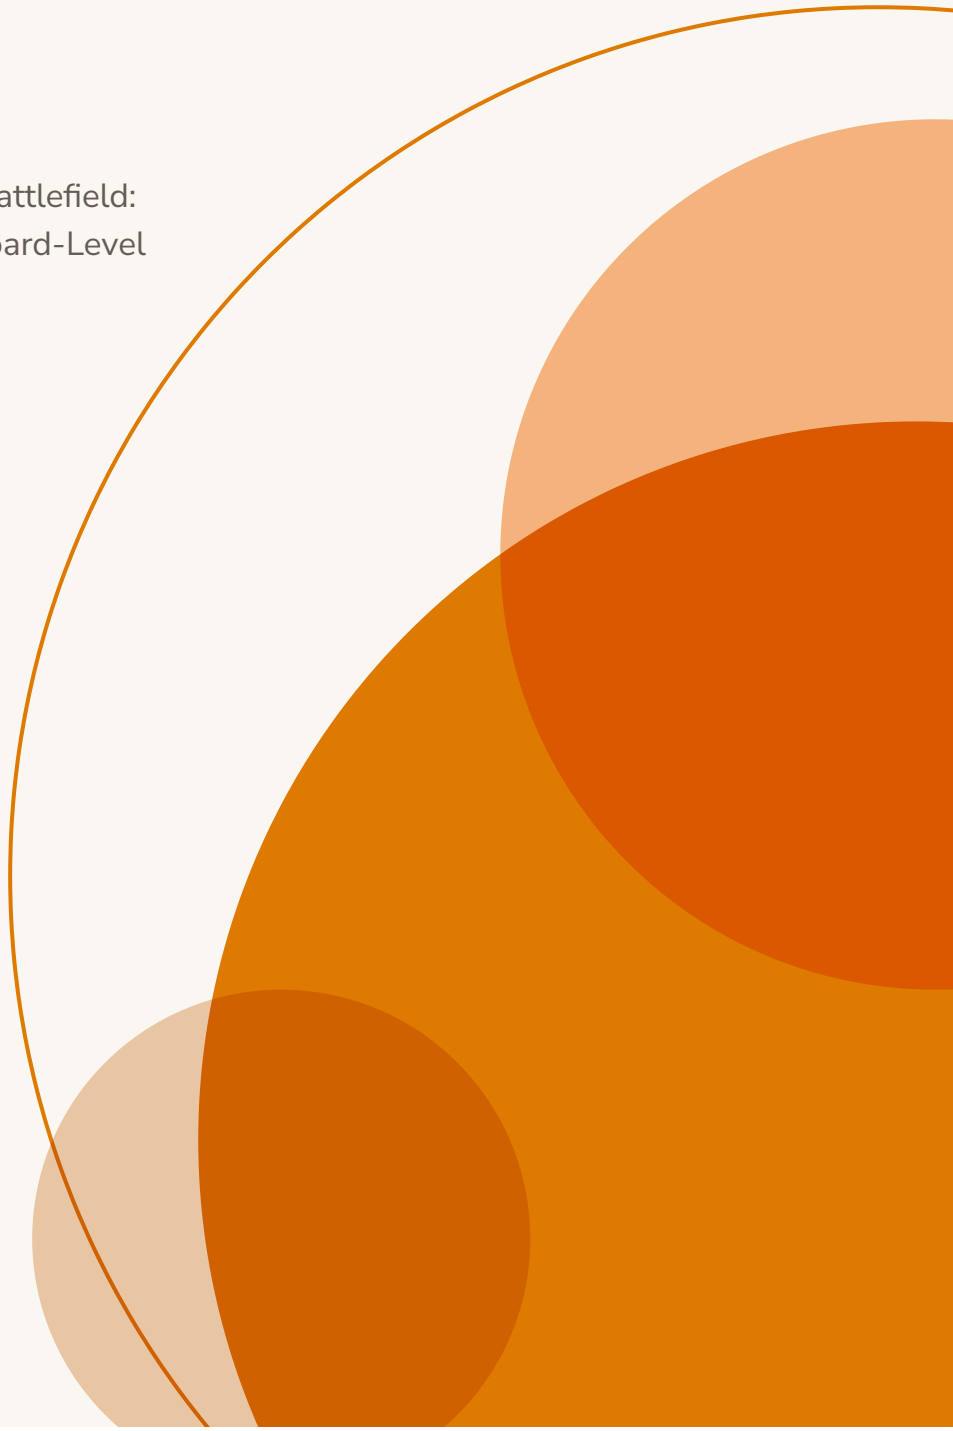

CIONET

Cyber resilience may increasingly depend on a CIO becoming what one leader in the room called a "Chief Connecting Officer." It's a title more suited to a headline than a job description, but it stood out across a 90-minute conversation on cyber security, AI-driven attacks, geopolitical threats and organisational resilience.

What connecting actually requires

The clearest illustration came from a discussion about the difference between a good technical response to a cyber incident and a good legal one. A technical team can do everything right during a crisis, patching, containing, restoring, and still create a legal problem simply through how people talk to each other while they're doing it. Internal messages sent in the heat of an incident, including the kind of gallows humour that gets people through a bad day, can end up read out in court or by a regulator months later. A throwaway comment about something that should have been patched sooner stops being funny when a judge is asking someone to explain it.

The advice on this was simple: involve legal in your cyber resilience planning. Not because it slows things down, but because it provides additional protection. Legal teams embedded early in a technical response can help build guardrails in advance, how communications should be handled, what shouldn't be put in writing, who talks to whom. Brought in too late, and legal looks like someone arriving to police what's already happened.

The same logic came up when the conversation turned to leadership more broadly. There was a comparison drawn between public and private sector experience, and the view offered was that the core challenge doesn't change much between the two: getting technical and business people to understand each other and move in the same direction. What did stand out as different was a particular skill in senior civil servants for handling difficult, high-stakes positions diplomatically, holding a position under pressure without losing the room. That was offered as genuine admiration rather than a swipe at the private sector, and as a skill worth borrowing regardless of sector.

Diplomacy came up repeatedly, not as a soft add-on to the job but as something closer to a hard skill. Getting a technical team, a legal team and an executive team to agree on how to respond to a live incident, when each group often doesn't fully understand what the others need, was described as

requiring something close to negotiation. Understanding what different parties actually care about, and finding a way through that doesn't just paper over the disagreement, is a more specific ask than the usual language of stakeholder management.

Where the basics still decide the outcome

For all the talk of AI and geopolitics, a good chunk of the evening kept returning to something less glamorous: whether an organisation's incident response plan actually works when it's needed. The distinction drawn was between plans that exist and plans that have been tested to the point of discomfort. Tabletop exercises, done properly, surface the kind of failure nobody thinks to plan for. People who can't access Teams during a live incident, because Teams is part of what's gone down. A response plan listing the CEO's mobile number, except the number is out of date because nobody updates contact details more than once a year, usually around insurance renewal.

The suggestion that stuck was defiantly old-fashioned: forget the polished eighty-page plan sitting on a server that's just been encrypted, and keep a single laminated page in a drawer covering only what people need to know in the first 24 hours. It's a strange thing to hear in a room otherwise discussing nation-state actors and AI-driven attacks, but it landed as one of the more practical points of the evening. A plan is only useful if people can get to it at the moment they need it, and that's a lower bar than most organisations think they've cleared.

The sharper edge of AI-driven attacks

Where AI's effect on attacks got concrete, the examples were unsettling. One involved threat actors scraping a company's own website, working through staff photos and bios to identify who wore glasses, then sending a phishing email about a contact lens order ready for collection. The click-through rate on that attack was remarkably high. Another played on a more universal weakness: an email suggesting the recipient had been nominated for an award. Ego is a reliable way in.

What stands out about both examples isn't technical sophistication. These aren't the clumsy phishing emails of a few years ago, with the spelling mistakes everyone was trained to spot. AI is removing that kind of tell, particularly for threat actors working in a second language, and replacing generic attacks with ones built around an individual's actual life. A phishing email tied to your own glasses prescription doesn't read like an attack. It reads like an inconvenience to deal with before lunch.

There was also a useful point about where "unforeseen" incidents actually come from. It's tempting to imagine these as novel external threats nobody could have predicted. More often the cause is closer to home: staff using AI tools outside whatever policy the organisation has in place, not out of malice but because the sanctioned tools don't fit how the work actually gets done. A policy that doesn't reflect how people work will be worked around, and that gap is where the next incident tends to start.

The tension that didn't resolve

Sovereignty was a cause for disagreement. There was a pragmatic case made for working with the best available technology regardless of where it comes from, on the basis that capability has to come from somewhere and the UK doesn't currently have the scale to build genuine alternatives to the largest global providers. That pragmatism sat alongside an acknowledgement that the pull toward sovereign alternatives is real and isn't going away, and that for some organisations and some data, who owns the infrastructure isn't a question pragmatism alone answers.

There was a stated discomfort with treating AI as an inevitable catastrophe, alongside an equally genuine acknowledgement that something has shifted, and that this is a real inflection point rather than another technology cycle to absorb.

Perhaps the most honest moment of the evening came from a direct audience question about accountability: when something goes wrong with AI or cyber risk, who is actually responsible? There was general agreement that responsibility sits somewhere with the executive team collectively, but nothing more precise than that. For a group of senior people who spend a lot of time being accountable for things, that gap was notable.

An open question

The idea that a CIO's job is becoming less about connecting systems and more about connecting people, perspectives and decisions is a useful reframe, and easy to nod along with. What's less clear is what this means for CIOs who came up through a deeply technical route, and whether the wider community around technology leadership is equipping people for that shift or still mostly talking about the technology itself.

If diplomacy, relationship-building and the ability to hold a room under pressure are becoming as central to the role as technical judgement, that has implications for how people are developed into these positions, not just how they behave once they're in them. That's the harder conversation, and one for the community to answer in future.

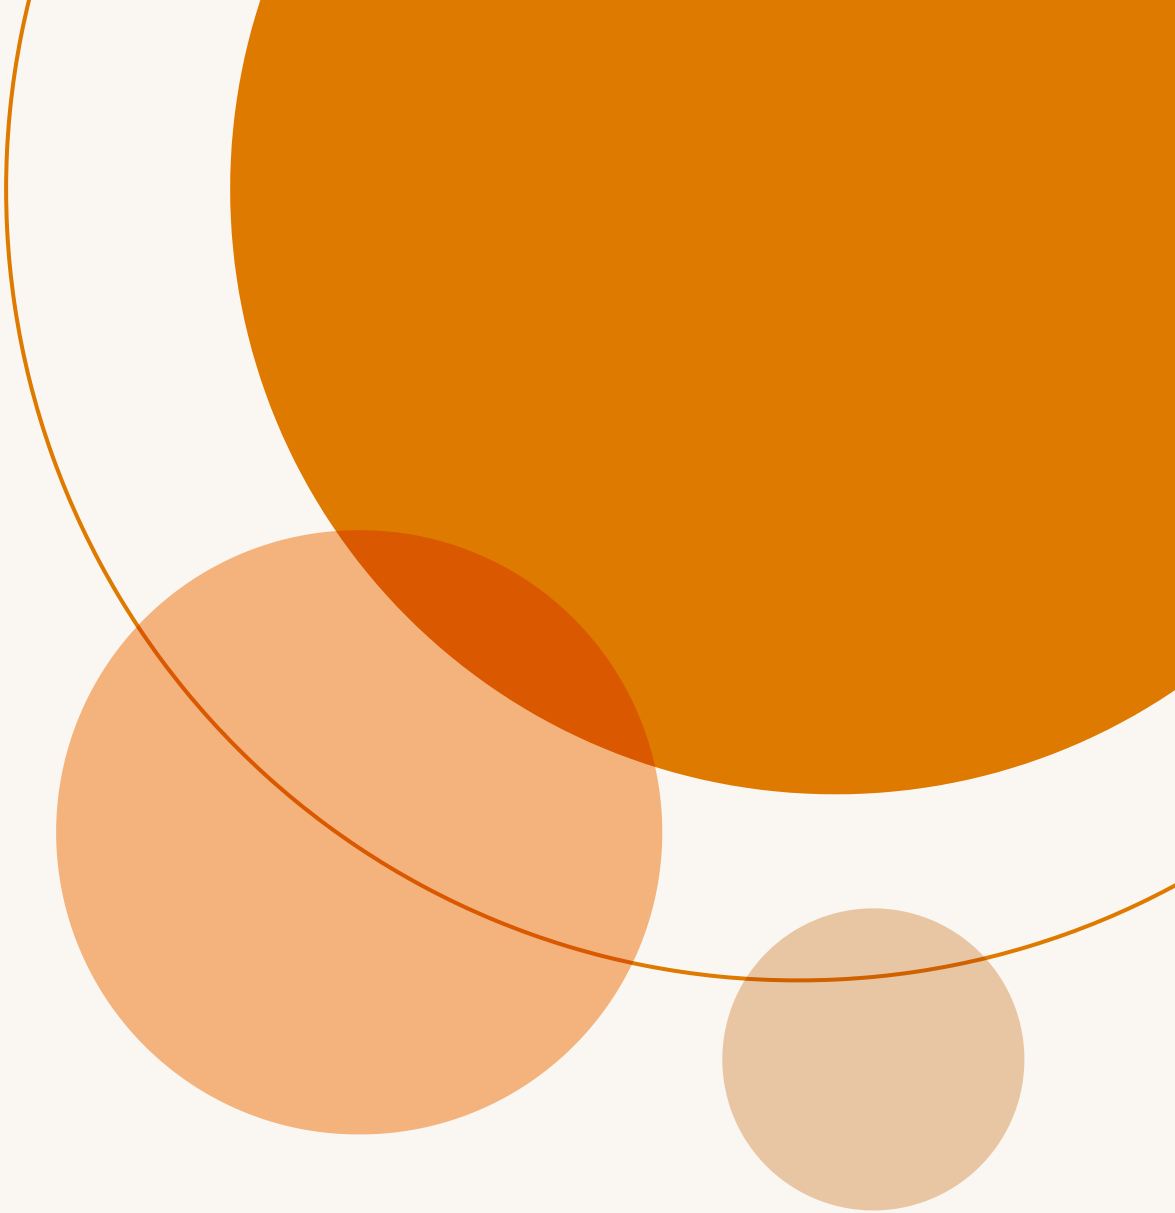

cionet.com

CIONET