

Post event report



Strategic Sponsors

DRATA

THREATLOCKER®
ZERO TRUST PLATFORM

Education Seminar Sponsors



FluidOne **CSA Cyber** **RISK LEDGER**

Networking Sponsor



Branding Sponsor

LastPass |

Inside this report:

- Sponsors
- Key themes
- Who attended?
- Speakers
- Agenda
- Education Seminars



Speakers

Matt Adams, Generative AI & Emerging Technology Security, **Citi**

Stephen Beckett, Global Security and Business Continuity Director, **Deloitte**

Haydn Brooks, CEO, Risk Ledger Phil Gough, CISO, **Pinsent Masons LLP**

Kevin Carr, Senior Manager, Solutions Engineering, **Drata**

Will Collinson, Technical Director – Cryptography, **HSBC**

Ian Dalby, Global Head of GRC, **A&O Shearman**

Rachel Dyges, Global Business Continuity Management Lead, **A&O Shearman**

Jai Ferguson, AI Regional Lead – Europe, **HSBC**

Jonathan Freedman, Director of Technology & Security, **Howard Kennedy**

Stephen Green, Regional Vice President of EMEA, **ConcentricAI**

Gayle Hedgecock, Business Continuity & Resilience Specialist, **Clifford Chance**

Amelia Hewitt, Director of Cyber Consulting, **Principle Defence**

Dale Hodgkinson, Former Head of Strategy and Architecture, **Slaughter and May**

Sam Hubery, BISO, **Fidelity International**

Oscar Javier Hernandez Rodriguez, Account Executive, **ThreatLocker**

Ellie Ludlam, Partner, **Pinsent Masons LLP**

Daniel Oxley, Senior Engineer, **Doppel**

Manit Sahib, Ethical Hacker & Former Head of Penetration Testing & Red Teaming, **Bank of England**

Tom Sampson, Head of Information Security, **Macfarlanes**

Samuel Scott, Cyber Risk and Advisory, **Marsh**

Dr Narayan Shiva, CTO and Enterprise Architect, **iBANK**

Geoffrey Taylor, Information Security Officer, **Nordea Asset Management**

Steve Velcev, Practice Lead for Offensive Security Engineering and Principal Red Team Consultant, **FluidOne**

Mark Walmsley, CISO, **Freshfields**

Philip Young, Co-founder and CEO, **Garfield AI**

Key themes

Identity, authority, and control for non-human actors

Securing algorithmic insiders


Data control when there is no perimeter

The power of automation

Integrity and the AI-enabled supply chain

Dealing with regulations

Who attended?



Cyber-security

We have a 15-year track record of producing the events cyber-security professionals take seriously



Risk Management

We attract senior risk officers with responsibility for information risk assessment and mitigation



Fraud, Audit, Compliance

We provide the go-to events for fraud prevention and compliance risk assessment at the world's key corporates



Data Protection & privacy

We are a key venue for decision-makers with budget and purchasing authority

Agenda			
08:00	Breakfast networking and registration		
08:50	Chair's welcome		
09:05	<p>Will the future law firm be indistinguishable from a tech company?</p> <p>Philip Young, Co-founder and CEO, Garfield AI</p> <ul style="list-style-type: none"> From people led to platform led delivery – how API integrations, automation, and 24/7 availability are reshaping client expectations, workflows, and scalability Regulation meets technology – the implications of running a regulated legal service through a software platform, including accountability, auditability, and risk New exposure, new responsibilities – how always-on, interconnected legal platforms change the firm's risk profile, including data security, resilience, and third-party dependencies 		
09:25	<p>Do I need a ROC as well as a SOC?</p> <p>Ian Dalby, Global Head of GRC, A&O Shearman</p> <ul style="list-style-type: none"> What if companies aren't optimising for operational security or resilience, but rather for protection against liability? How to distinguish real security from compliance-driven assurance and assess whether your organisation is truly resilient How compliance can quietly become a source of risk amplification rather than genuine risk reduction ... enter AI How to turn compliance into operational value ... welcome to The ROC! 		
09:45	<p>PANEL DISCUSSION The future of legal AI: Innovation with accountability</p> <p>Philip Young, Co-founder and CEO, Garfield AI; Amelia Hewitt, Director of Cyber Consulting, Principle Defence; Tom Sampson, Head of Information Security, Macfarlanes; Dale Hodgkinson, Former Head of Strategy and Architecture, Slaughter and May</p> <ul style="list-style-type: none"> What does 'good AI governance' actually look like inside a modern law firm? Unlike the EU AI Act, the UK has a principles-based approach. What does accountability look like for UK law firms right now? AI risk often sits across multiple silos. How should firms address the overlap between cyber, privacy, and AI governance? How can firms innovate with AI while preserving trust and meeting client expectations? How do we prevent AI from becoming a single point of failure in financial decision-making? 		
10:15	<p>Education Seminars Session 1</p> <table border="1"> <tr> <td> <p>Doppel</p> <p>Social engineering attack chains: Legal exposure, regulatory accountability, and organisational resilience in the AI era</p> <p>Daniel Oxley, Senior Engineer, Doppel</p> </td> <td> <p>FluidOne</p> <p>The AI-native attacker: How offensive AI is rewriting the playbook for breaching law firms</p> <p>Steve Velcev, Practice Lead for Offensive Security Engineering and Principal Red Team Consultant, FluidOne</p> </td> </tr> </table>	<p>Doppel</p> <p>Social engineering attack chains: Legal exposure, regulatory accountability, and organisational resilience in the AI era</p> <p>Daniel Oxley, Senior Engineer, Doppel</p>	<p>FluidOne</p> <p>The AI-native attacker: How offensive AI is rewriting the playbook for breaching law firms</p> <p>Steve Velcev, Practice Lead for Offensive Security Engineering and Principal Red Team Consultant, FluidOne</p>
<p>Doppel</p> <p>Social engineering attack chains: Legal exposure, regulatory accountability, and organisational resilience in the AI era</p> <p>Daniel Oxley, Senior Engineer, Doppel</p>	<p>FluidOne</p> <p>The AI-native attacker: How offensive AI is rewriting the playbook for breaching law firms</p> <p>Steve Velcev, Practice Lead for Offensive Security Engineering and Principal Red Team Consultant, FluidOne</p>		
10:55	Networking break		
11:30	<p>Actions speak louder than tokens: Treating frontier AI agents as insider threats</p> <p>Matt Adams, Generative AI & Emerging Technology Security, Citi</p> <ul style="list-style-type: none"> The alignment paradox: today's frontier models score well on macro-alignment – they reliably refuse explicit harmful requests – yet show poor micro-alignment, autonomously selecting dangerous methods in pursuit of legitimate goals A first formal framework adapting CERT's insider-threat dimensions to non-human actors – mapping motivation, opportunity, and capability onto optimisation objectives, tool access, and model capabilities – with a five-category STRIDE-derived taxonomy of agent threats Real-world validation from the March 2026 ROME incident, where a safety-trained agent autonomously mined cryptocurrency, opened SSH tunnels, and probed internal networks during RL training A structural playbook for financial services CISOs: stop assessing intent, monitor action-level telemetry, enforce least-privilege tool binding and ephemeral credentials, and fold AI agents into the insider-threat programs FSIs already run 		

Agenda

11:50	Quantum is coming. We can't afford to wait	
	<p>Will Collinson, Technical Director – Cryptography, HSBC</p> <ul style="list-style-type: none"> • Discover why the quantum threat to today's cryptography is closer and more disruptive than many realise • Hear what's at stake as quantum computing reshapes the cybersecurity landscape • Join the call for industry-wide collaboration to tackle one of cybersecurity's biggest ever challenges before the clock runs out • Learn what you can do today (or already should be doing) to reduce your risk 	
12:10	Zero Trust controls at the endpoint	
	<p>Oscar Javier Hernandez Rodriguez, Account Executive, ThreatLocker</p> <ul style="list-style-type: none"> • Discover how ThreatLocker applies Zero Trust at the endpoint, eliminating implicit trust by continuously verifying every application, executable, and action before authorisation • Learn how a deny-by-default, malware-proofing approach reduces ransomware risk, stopping unauthorised software and scripts even when other security layers are bypassed • Understand how least-privilege enforcement limits attacker capability, ensuring applications and users can perform only explicitly approved actions on enterprise devices • Explore how granular, policy-based endpoint control safeguards against modern threats, reducing enterprise exposure to ransomware and other advanced attacks 	
12:15	The new non-human insider: Governing the agents you can't see	
	<p>Kevin Carr, Senior Manager, Solutions Engineering, Drata</p> <ul style="list-style-type: none"> • Law firms are expected to protect client confidentiality, maintain operational resilience, and demonstrate sound governance every day, yet many still rely on point-in-time audits, static questionnaires, and fragmented tools. AI has widened that gap: clients now expect more current assurance, while firms are filling up with agents, APIs, copilots, and vendors that no one has fully classified as identities, each acting as a non-human insider with valid credentials and poorly understood access • This session argues that those agents should be treated as privileged actors, subject to the same scrutiny and control as human users. It explores why annual, manual approaches are no longer sufficient in a world shaped by client security requirements, professional duties around confidentiality and privilege, the EU AI Act, GDPR and UK GDPR, and the UK's incoming Cyber Security and Resilience Bill. It then looks at what firms need instead: continuous discovery of every agent, clear policy over what each one is allowed to do, detection of drift in behaviour or permissions, and evidence that can stand up to client, regulator, or board-level scrutiny on any given day • We'll also be candid about the limits firms face today: weak data quality, incomplete inventories, humans still on the critical path for approvals and exceptions, and open questions around how to govern the agents themselves. The session closes with a practical view of the next 18 months – what is realistic, what remains experimental, and where law firms should invest now 	
12:35	Education Seminars Session 2	
	<p>ConcentricAI AI is breaking data security... and fixing it: the new reality of ai-driven risk and how to stay ahead Stephen Green, Regional Vice President of EMEA, ConcentricAI</p>	<p>Risk Ledger Beyond the checkbox: When third-party risk becomes client disruption Haydn Brooks, CEO, Risk Ledger & Mark Walmsley, CISO, Freshfields</p>
13:15	Lunch networking break	
14:20	PANEL DISCUSSION Cyber-insurance for law firms: Protection, pitfalls, and practical use	
	<p>Ellie Ludlam, Partner, Pinsent Masons LLP (Moderator); Samuel Scott, Cyber Risk and Advisory, Marsh; Phil Gough, CISO, Pinsent Masons LLP</p> <ul style="list-style-type: none"> • The case for cyber-insurance: Exploring financial protection, specialist incident-response expertise, and the role insurance can play in strengthening governance and risk management • The limitations and challenges: Examining coverage gaps, rising premiums, policy exclusions, and the operational requirements that can affect claims and value • Making cyber-insurance work: Practical guidance on whether, when, and how to integrate cyber-insurance into a broader cyber-resilience and risk management strategy 	

Agenda			
14:45	In the age of AI, is security even possible?		
	<p>Jonathan Freedman, Director of Technology & Security, Howard Kennedy</p> <ul style="list-style-type: none"> • What AI-powered offensive capability really means in practice – from autonomous vulnerability discovery to agentic attack automation – and where the hype ends • Why foundational security controls remain the most effective defence against AI-enabled threats • How organisations can shift from preventing every attack to slowing, detecting, and responding to machine-speed compromise attempts before damage occurs 		
15:05	PANEL DISCUSSION	Business continuity in law firms: staying operational through cyber-disruption	
	<p>Jonathan Freedman, Director of Technology & Security, Howard Kennedy (Moderator); Gayle Hedgecock, Business Continuity & Resilience Specialist, Clifford Chance; Stephen Beckett, Global Security and Business Continuity Director, Dentons; Rachel Dyges, Global Business Continuity Management Lead, A&O Shearman</p> <ul style="list-style-type: none"> • When a cyber-incident hits, who actually makes the call – and is that genuinely clear in practice? • How does information flow in the first hour, who needs to know what, and how do you avoid confusion or bottlenecks? • Are we putting too much emphasis on backups and not enough on keeping the firm operational? • How do you handle client confidentiality and regulatory pressure while the situation is still unfolding? • When you've tested your plans, what's actually broken – and what caught you off guard? • And what are the headaches people don't usually plan for? 		
15:35	Conformity will not save you: AI risk beyond the EU AI Act		
	<p>Geoffrey Taylor, Information Security Officer, Nordea Asset Management</p> <p>Your assessment said low risk. Is it really?</p> <ul style="list-style-type: none"> • The EU AI Act requires organisations to classify their AI systems and demonstrate conformity. Conformity is similar to compliance – it is binary, a yes or a no at a point in time. It cannot calibrate impact when the unexpected occurs • On 24 April 2026, an AI agent deleted an entire company's production database in nine seconds. It was running the best model available, configured with explicit safety rules. When asked to explain itself, it produced a written confession: 'I violated every principle I was given.' • This session applies the Assume. Design. Test. framework to AI governance – shifting the question from 'are we compliant?' to 'how could we be impacted?' – and gives attendees a practical lens for assessing where their governance ends and their exposure begins 		
15:55	Networking break		
16:20	PANEL DISCUSSION	Customer data & AI: Control, exposure, and proof	
	<p>Simon Brady, Event Chairman; Sam Hubery, BISO, Fidelity International; Jai Ferguson, AI Regional Lead – Europe, HSBC; Dr Narayan Shiva, CTO and Enterprise Architect, iBANK</p> <ul style="list-style-type: none"> • As organisations adopt AI, where are you seeing customer data most commonly interact with this tool and how are you improving visibility over time? • What controls or approaches are proving most effective in practice for preventing customer data being exposed to AI tools – and where are you still seeing challenges? • Are you allowing any use of third-party or public AI tools (like ChatGPT) with customer data and what specific safeguards make that acceptable? • Can you demonstrate that customer data is properly controlled within AI systems? 		
16:50	Rise of autonomous attacks (Live Mythos-style hack)		
	<p>Manit Sahib, Ethical Hacker & Former Head of Penetration Testing & Red Teaming, Bank of England</p> <ul style="list-style-type: none"> • See how autonomous AI agents are now running the recon and exploitation phases of real-world attacks. and what that means for boards, CISOs, and red teams in 2026 • A first-hand look at how agentic offensive AI works in practice, driven by intent, not step-by-step instruction • See AI agent run reconnaissance against a controlled target, identify exploitable assets, and demonstrate the early stages of a kill chain in real time • A walk through real-world findings from recent engagements including critical vulnerabilities discovered by AI agents that automated scanners (Tenable, Qualys, Nessus) had missed for over 18 years • What defenders need to know: why traditional, control-based security models are structurally insufficient against goal-driven autonomous attackers, and the three specific actions every CISO should be taking before this becomes the default attacker model 		
17:10	Chairman's closing remarks	17:15	Drinks reception & networking
		18:30	End of conference

Education Seminars

ConcentricAI

AI is breaking data security... and fixing it: The new reality of AI-Driven risk and how to stay ahead

Stephen Green, Regional Vice President of EMEA, ConcentricAI

AI is rapidly becoming one of the biggest drivers of productivity and innovation in the enterprise – and one of the fastest-growing sources of data security risk. As copilots, assistants, and public AI tools become integrated into daily work, sensitive data is flowing into systems that most security teams can't fully see, understand, or control.

The problem is that traditional data security controls were never built for this. In fact, many organisations were already struggling to operationalise data security before AI accelerated the challenge. The good news? AI isn't just creating the problem – it's also enabling a smarter, more effective way to solve it.

Attendees will learn:

- Why AI has become one of the fastest-growing and least visible sources of enterprise risk
- How GenAI is creating new exposure points for sensitive data
- Why legacy data security tools have failed to keep up – and why AI is making those gaps harder to ignore
- How context-aware, AI-driven data security can deliver more accurate visibility, stronger controls, and real-time enforcement
- What organisations can do to enable AI innovation without expanding their risk surface
- Attendees will leave with a clearer understanding of how AI is reshaping data security – and how they can use that same technology to gain control, minimise exposure, and support safer AI adoption across the business

Doppel

Social engineering attack chains: Legal exposure, regulatory accountability, and organisational resilience in the AI era

Daniel Oxley, Senior Engineer, Doppel

Social engineering is no longer limited to isolated phishing emails. It has evolved into a sophisticated, AI-driven threat landscape that spans email, SMS, voice, collaboration platforms, social media, and synthetic media. As these attacks become more convincing and more scalable, they introduce significant legal, regulatory, governance, and operational risks that extend well beyond traditional cybersecurity controls.

Attendees will learn:

- Discover how threat actors are leveraging artificial intelligence to personalise and automate attacks across multiple channels, increasing their effectiveness while making attribution, evidence preservation, internal investigations, and legal defence significantly more challenging
- Gain insight into the legal and regulatory implications of modern social engineering campaigns, including data protection breaches, financial crime exposure, disclosure obligations, operational resilience requirements, third-party risk, contractual liability, and potential enforcement action
- Learn how organisations can evaluate and demonstrate their management of human-layer risk by identifying gaps between policy and practice, validating the effectiveness of controls, and evidencing reasonable and proportionate safeguards
- Explore human risk management as an emerging governance discipline that enables organisations to measure, monitor, and reduce human-targeted threats while strengthening compliance, auditability, and defensible decision-making
- Understand how legal, compliance, risk, security, and executive leadership teams can work together to build a unified, intelligence-led defence strategy that enhances regulatory readiness, strengthens incident response, improves legal preparedness, and drives long-term organisational resilience

Education Seminars

FluidOne

**The AI-native attacker:
How offensive AI is rewriting
the playbook for breaching
law firms**

Steve Velcev, Practice Lead for
Offensive Security Engineering
and Principal Red Team
Consultant, FluidOne

Your firm's most valuable asset – privileged client data, live-deal intelligence, litigation strategy no longer sits behind a firewall an attacker has to break through. It sits behind a single cloud login. This seminar takes you inside a real, end-to-end intrusion against a modern law firm, seen entirely through the attacker's eyes, and shows exactly where artificial intelligence now removes the friction, cost and skill that once stood between a criminal and your data. Led by an experienced working red teamer, it moves from AI-driven reconnaissance and MFA-bypass phishing to automated data theft from online services such as Microsoft 365 and then turns the page to the practical, achievable controls that actually break the chain. No hype, no vendor pitch: just what genuinely changed, what it means for your firm, and where to spend first.

Attendees will learn:

- How a modern breach actually unfolds – a step-by-step walk-through of the full attacker kill chain against a representative firm, from open-source reconnaissance to data exfiltration, with AI's role made explicit at every stage
- Why one stolen login now equals wholesale access – how attackers turn a single phished identity into the keys to the firm's most sensitive matters, and why this bypasses the controls most firms still rely on
- The threats you may not know are already mainstream – adversary-in-the-middle phishing that defeats most MFA, and ClickFix attacks that make your own people run the malware (now ~47% of tracked initial access) with no malicious file for filters or EDR to catch
- How AI has changed the economics of attacking you – why phishing-as-a-service, open-source AI tooling and machine-speed automation mean attacks are now faster, cheaper and more numerous, and what that demands of your defences
- What actually stops it – and where to start – a pragmatic, prioritised set of controls (phishing-resistant MFA, properly configured conditional access, Graph and data-layer monitoring) proven to disrupt this exact chain, framed as an order of operations for firms that can't do everything at once

Risk Ledger

**Beyond the checkbox:
When third-party risk
becomes client disruption**

Haydn Brooks, CEO,
Risk Ledger &
Mark Walmsley, CISO,
Freshfields

Third-party cyber-risk remains one of the biggest challenges facing security and legal teams. Recent industry research found that 75% of legal organisations say their biggest concern following a supplier incident is the impact on client service – from disrupted access to systems and data through to delays in delivering client work, while 80% say supplier audit rights are still difficult to enforce in practice.

Join Risk Ledger's CEO, Haydn Brooks and Mark Walmsley, CISO, Freshfields as they explore the gap between contractual best practise and operational reality – from how to respond effectively to vendor breaches, to navigating negotiations with large technology suppliers.

This panel discussion will examine how organisations can balance commercial priorities with cyber-risk and focus on the controls that meaningfully improve resilience.

Attendees will learn:

- How to respond effectively to vendor breaches
- How to navigate negotiations with large technology suppliers
- How organisations can balance commercial priorities with cyber-risk and focus on the controls that meaningfully improve resilience