AI IN LAW FIRMS

Data protection risks, opportunities and compliance

For law firm leaders, risk teams, compliance teams and fee earners

Kristy Gouldsmith

Spencer West Data Protection Partner

SPEAKER

Kristy Gouldsmith

kristy.gouldsmith@spencer-west.com 07545501200

Data protection law specialist

1. Advises organisations across multiple sectors on data protection, privacy and compliance.
2. Works with clients ranging from smaller businesses to large international companies.
3. Practical focus: helping firms balance innovation, client trust and regulatory risk.

Sector experience includes

SaaS platforms | Retailers | Manufacturers | Schools | Care homes | Financial services | Law firms

Today’s focus: responsible AI adoption, client confidentiality, privilege and UK GDPR compliance.

Why this matters now

AI use in law firms is no longer theoretical. The real question is whether use is visible, governed and safe.

Efficiency pressure

faster research, drafting and triage

Client expectation

speed, cost certainty and service quality

Risk pressure

confidentiality, privilege, accuracy and accountability

OUTCOMES

By the end, attendees should be able to make a defensible decision about AI use on a matter

1 Spot the risks

Identify when AI use involves personal data, confidential information, privilege or higher-risk processing.

2 Classify inputs

Decide what can go into which tool, and when anonymisation or approval is needed

3 Assess vendors

Ask the right questions about training, retention, security, transfers and sub-

4 Govern use

Put policies, training and matter-level sign-off around AI adoption.

AGENDA

- AI use cases in legal practice
- Confidentiality, privilege and data classification
- UK GDPR and data protection duties
- Vendor due diligence and contracts
- Governance, policies and implementation

USE CASES

Where AI is already appearing in law firms

The data protection analysis starts with the use case, not the brand name of the tool.

Research

case law summaries, horizon scanning, knowledge search

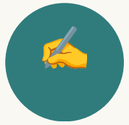

Drafting

first drafts, clause options, letters, bundles

Review

document review, extraction, due diligence, litigation disclosure

Operations

client intake, conflicts, billing narratives, HR and training

Key question: does the tool receive matter data, personal data, privileged material, or confidential strategy?

FRAMEWORK

The AI data risk map

Every AI workflow can be broken into four places where risk enters:

1 Input

What is submitted?

2 Processing

Where does it go?

3 Output

How is it used?

4 Reuse

Can it train or improve the tool?

If you cannot answer these four questions, the firm does not yet understand the processing

DATA CATEGORIES

Classify before you prompt

The tool decision should follow the data classification decision.

Public / open	Published law, public guidance, general drafting prompts
Firm internal	Policies, know-how and templates approved for AI use
Client confidential	Matter facts, client identity, commercial strategy
Privileged	Advice, litigation strategy, solicitor-client communications
Sensitive personal data	Health, biometrics, children, criminal offence data

Rule of thumb: the more matter-specific and sensitive the input, the stronger the approval, contract and technical controls must be.

PRACTICAL RULE SET

What can and cannot be entered into AI tools

Generally safer	Needs controls	Do not enter into public AI
Public law, generic prompts, approved templates, properly anonymised facts	Pseudonymised matter facts, limited confidential extracts, internal knowledge, approved enterprise AI	Client names, privileged communications, settlement strategy, witness evidence, special category or criminal offence data

The answer is rarely “AI or no AI”. It is “which tool, which data, which safeguards, which approval?”.

Confidentiality and privilege: the legal sector overlay

For law firms, the risk is not only a GDPR problem. It is also a client trust, privilege and professional conduct problem.

Confidentiality

Is information disclosed outside the permitted circle of confidence?

Legal privilege

Could use of a third-party tool undermine control over privileged material?

Client trust

Would the client expect this processing and understand the safeguard?

Mitigation: approved tools, confidentiality terms, no model training, access controls, audit logs, human supervision and clear client messaging.

Loss of legal professional privilege & confidentiality

Practical message for law firms using AI: privilege depends on confidentiality and public AI tools create a disclosure risk.

1 The legal foundation

Legal professional privilege protects confidential communications between a lawyer and their client, as well as between a lawyer and a third party created for the purpose of obtaining legal advice or information relevant to litigation.

2 The AI-specific warning

In *UK v Secretary of State for the Home Department* [2026] UKUT 81(IAC), the Upper Tribunal concluded that uploading client material into an open-source AI tool is to place that information on the internet in the public domain and thus to breach confidentiality and to waive privilege. By comparison, closed source AI tools that do not place information in the public domain were found to pose less risk.

3 The firm rule

Do not paste client confidential or privileged material into public AI tools.

LEGAL FRAMEWORK

UK GDPR principles through an AI lens

AI does not sit outside data protection law. The principles still apply.

Lawfulness, fairness, transparency

Can people understand what is happening?

Purpose limitation

Is AI use compatible with the original matter purpose?

Data minimisation

Are we submitting only what the tool actually needs?

Accuracy

Are outputs checked before they affect advice or decisions?

Storage limitation

How long are prompts, logs and outputs retained?

Security and accountability

Can the firm evidence the controls?

DATA ROLES

Who is the controller, processor or both?

Do not assume a vendor is always just a processor. Read the actual data use terms. Do not just rely on the marketing brochures.

Law firm

Controller for client and matter data

AI vendor

Processor, controller or mixed roles depending on terms

Sub-processors

Hosting, security, analytics, support and model infrastructure

Development & improvement

Separate controller purposes may appear here

Ask: What data is processed, for whose purposes, and on whose instructions?

Lawful basis, transparency and client expectations

1	**What is the processing?**	Prompt submission, document analysis, transcription, output storage or automated triage.
2	**What lawful basis applies?**	Legitimate interest. Not contract.
3	**Who needs to be told?**	Data subjects – clients, staff, candidates
4	**Does the privacy notice cover it?**	Privacy information should be clear enough to cover material AI use.

Use a Legitimate Interests Assessment where legitimate interests is relied on for higher-risk AI use.

DPIAs: make them useful, not ornamental

A DPIA is the bridge between “we want this tool” and “we can use it safely”.

Trigger examples: innovative tech, large-scale or sensitive data, systematic monitoring, profiling, significant automated decisions, or vulnerable individuals.

HIGHER-RISK MATTERS

Special category and criminal offence data

Law firms routinely encounter data that ordinary business tools were not designed to handle casually.

Employment

health, disability, grievances, equality c

Clinical negligence

medical records, experts, vulnerable cl

Family

children, safeguarding, financial vulner

Crime / investigations

criminal offence data, witnesses and al

Private client

capacity, health, beneficiaries and fam

For these matters, default to no public AI.

Use only approved environments with a DPIA, special safeguards and clear records.

International transfers and cloud AI chains

AI vendors often involve hosting, support, sub-processors and model infrastructure across jurisdictions.

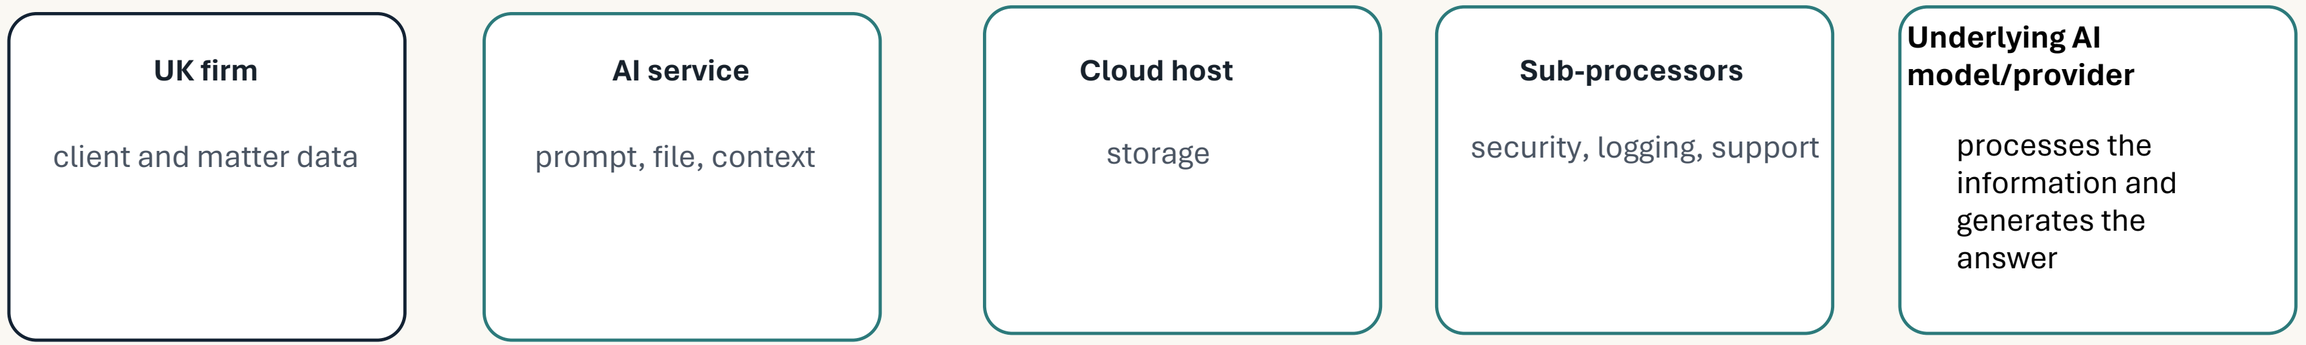

Identify countries where the data is processed and accessed, not only stored.

Check UK adequacy, use the IDTA or UK Addendum and do a transfer risk assessment.

Review the sub-processor contracts.

PROCUREMENT

AI vendor due diligence: the questions that matter

1	**Data use**	Is client data used to train models or improve the service? What do the sub-processors do with data?
2	**Retention**	How long are prompts, uploads, logs and outputs kept? The controller sets the retention period.
3	**Security**	Encryption, access controls, audit logs, segregation and incident response.
4	**Controls**	Can the firm configure no-training, retention, access and sharing?
5	**Evidence**	DPIA, security reports, sub-processor list, contractual terms, certifications

Procurement should not finish until Risk, IT/security and data protection have all signed off.

Contract protections to look for

The contract should match the technical promise. Marketing pages are not enough.

No training on client data

Prompts, uploads and outputs excluded from model training unless explicitly approved.

Processor terms

Clear instructions, confidentiality, assistance, deletion, audit and breach notice.

Sub-processors and transfers

Transparent list, change notice, objection route and transfer mechanism.

Retention and deletion

Configurable retention, prompt logs, file deletion and end-of-contract return/deletion.

Security and incidents

Access controls, audit trails, vulnerability management and short breach notice window.

Professional fit

Support for privilege, matter segregation, client commitments and regulator queries.

Accuracy, hallucination and professional supervision

Data protection risk does not stop when the answer is generated.

False confidence

AI may present wrong law, invented facts or false citations persuasively.

Personal data errors

Outputs can create or infer new personal data about clients or third parties.

Bias or unfairness

Automated triage or scoring can disadvantage individuals.

Over-reliance

A junior lawyer may treat an output as a finished product.

Control: no legal advice, court filing or client-facing document should be issued without competent human review.

CULTURE

Shadow AI: the risk behind the ban

A blanket ban often moves AI use out of sight. A governed pathway brings it back under control.

Why it happens

pressure, curiosity, no approved alternative

Why it is risky

unknown tools, unknown data use, no audit trail

What works

safe approved tools, clear rules, training and supervision

ACCOUNTABILITY

An AI governance model for law firms

AI governance needs named owners. Otherwise, everyone assumes someone else has checked it.

AI Sponsor

sets appetite and approves strategic use

Risk / COLP

professional duties, client commitments and supervision

DPO / Privacy

DPIAs, transparency, lawful basis, rights

IT / Security

access control, logging, security and integration

Practice Groups

matter use cases, output verification and training

Decision record: approved tool + approved use case + approved data categories + review date.

GOVERNANCE PACK

Policies and procedures every law firm should have

AI acceptable use policy

approved tools, prohibited data, examples and escalation

Tool register and approval process

owner, purpose, data categories, DPIA status and review date

Prompting and output guidance

anonymisation, verification, citations and client-facing safeguards

Vendor onboarding checklist

contract, security, transfers, sub-processors and deletion

Incident and breach playbook

what to report, who decides, evidence preservation

Training and supervision

role-based training for partners, juniors, support and IT

FROM IDEA TO OUTPUT

A matter-level workflow for safe AI use

1

Purpose

What are you asking AI to do?

2

Data

What will go in?

3

Tool

Is it approved for that data?

4

Controls

Have you done a DPIA?

5

Output

Who verifies before use?

Build this into matter opening, supervision notes, knowledge workflows or practice group guidance.

CLIENT TRUST

Client communications and terms of business

Clients do not need a technical manual, but they do need honest, accurate messaging about material AI use.

What to say	The firm may use approved technology, including AI-supported tools, to assist with legal services.
What to reassure	The firm remains responsible for advice and does not use uncontrolled public AI for confidential client data.
What to reserve	Some matter-specific AI uses may need express client instructions or separate approval.
What to warn about	Clients using public AI with firm advice or confidential documents may risk confidentiality and privilege.

Use wording that matches the firm’s actual controls. Do not overpromise.

RESPONSE

If something goes wrong: AI incident scenarios

Wrong tool

Fee earner uploads a client chronology into a public chatbot.

Wrong setting

A tool is configured to retain prompts longer than expected.

Wrong output

An AI-generated summary misstates a material fact in clie

Wrong recipient

A tool exposes matter data through permissions, links or s

Immediate steps: contain, preserve logs, identify data and tool terms, assess reporting duties, brief the client team, document the decision.

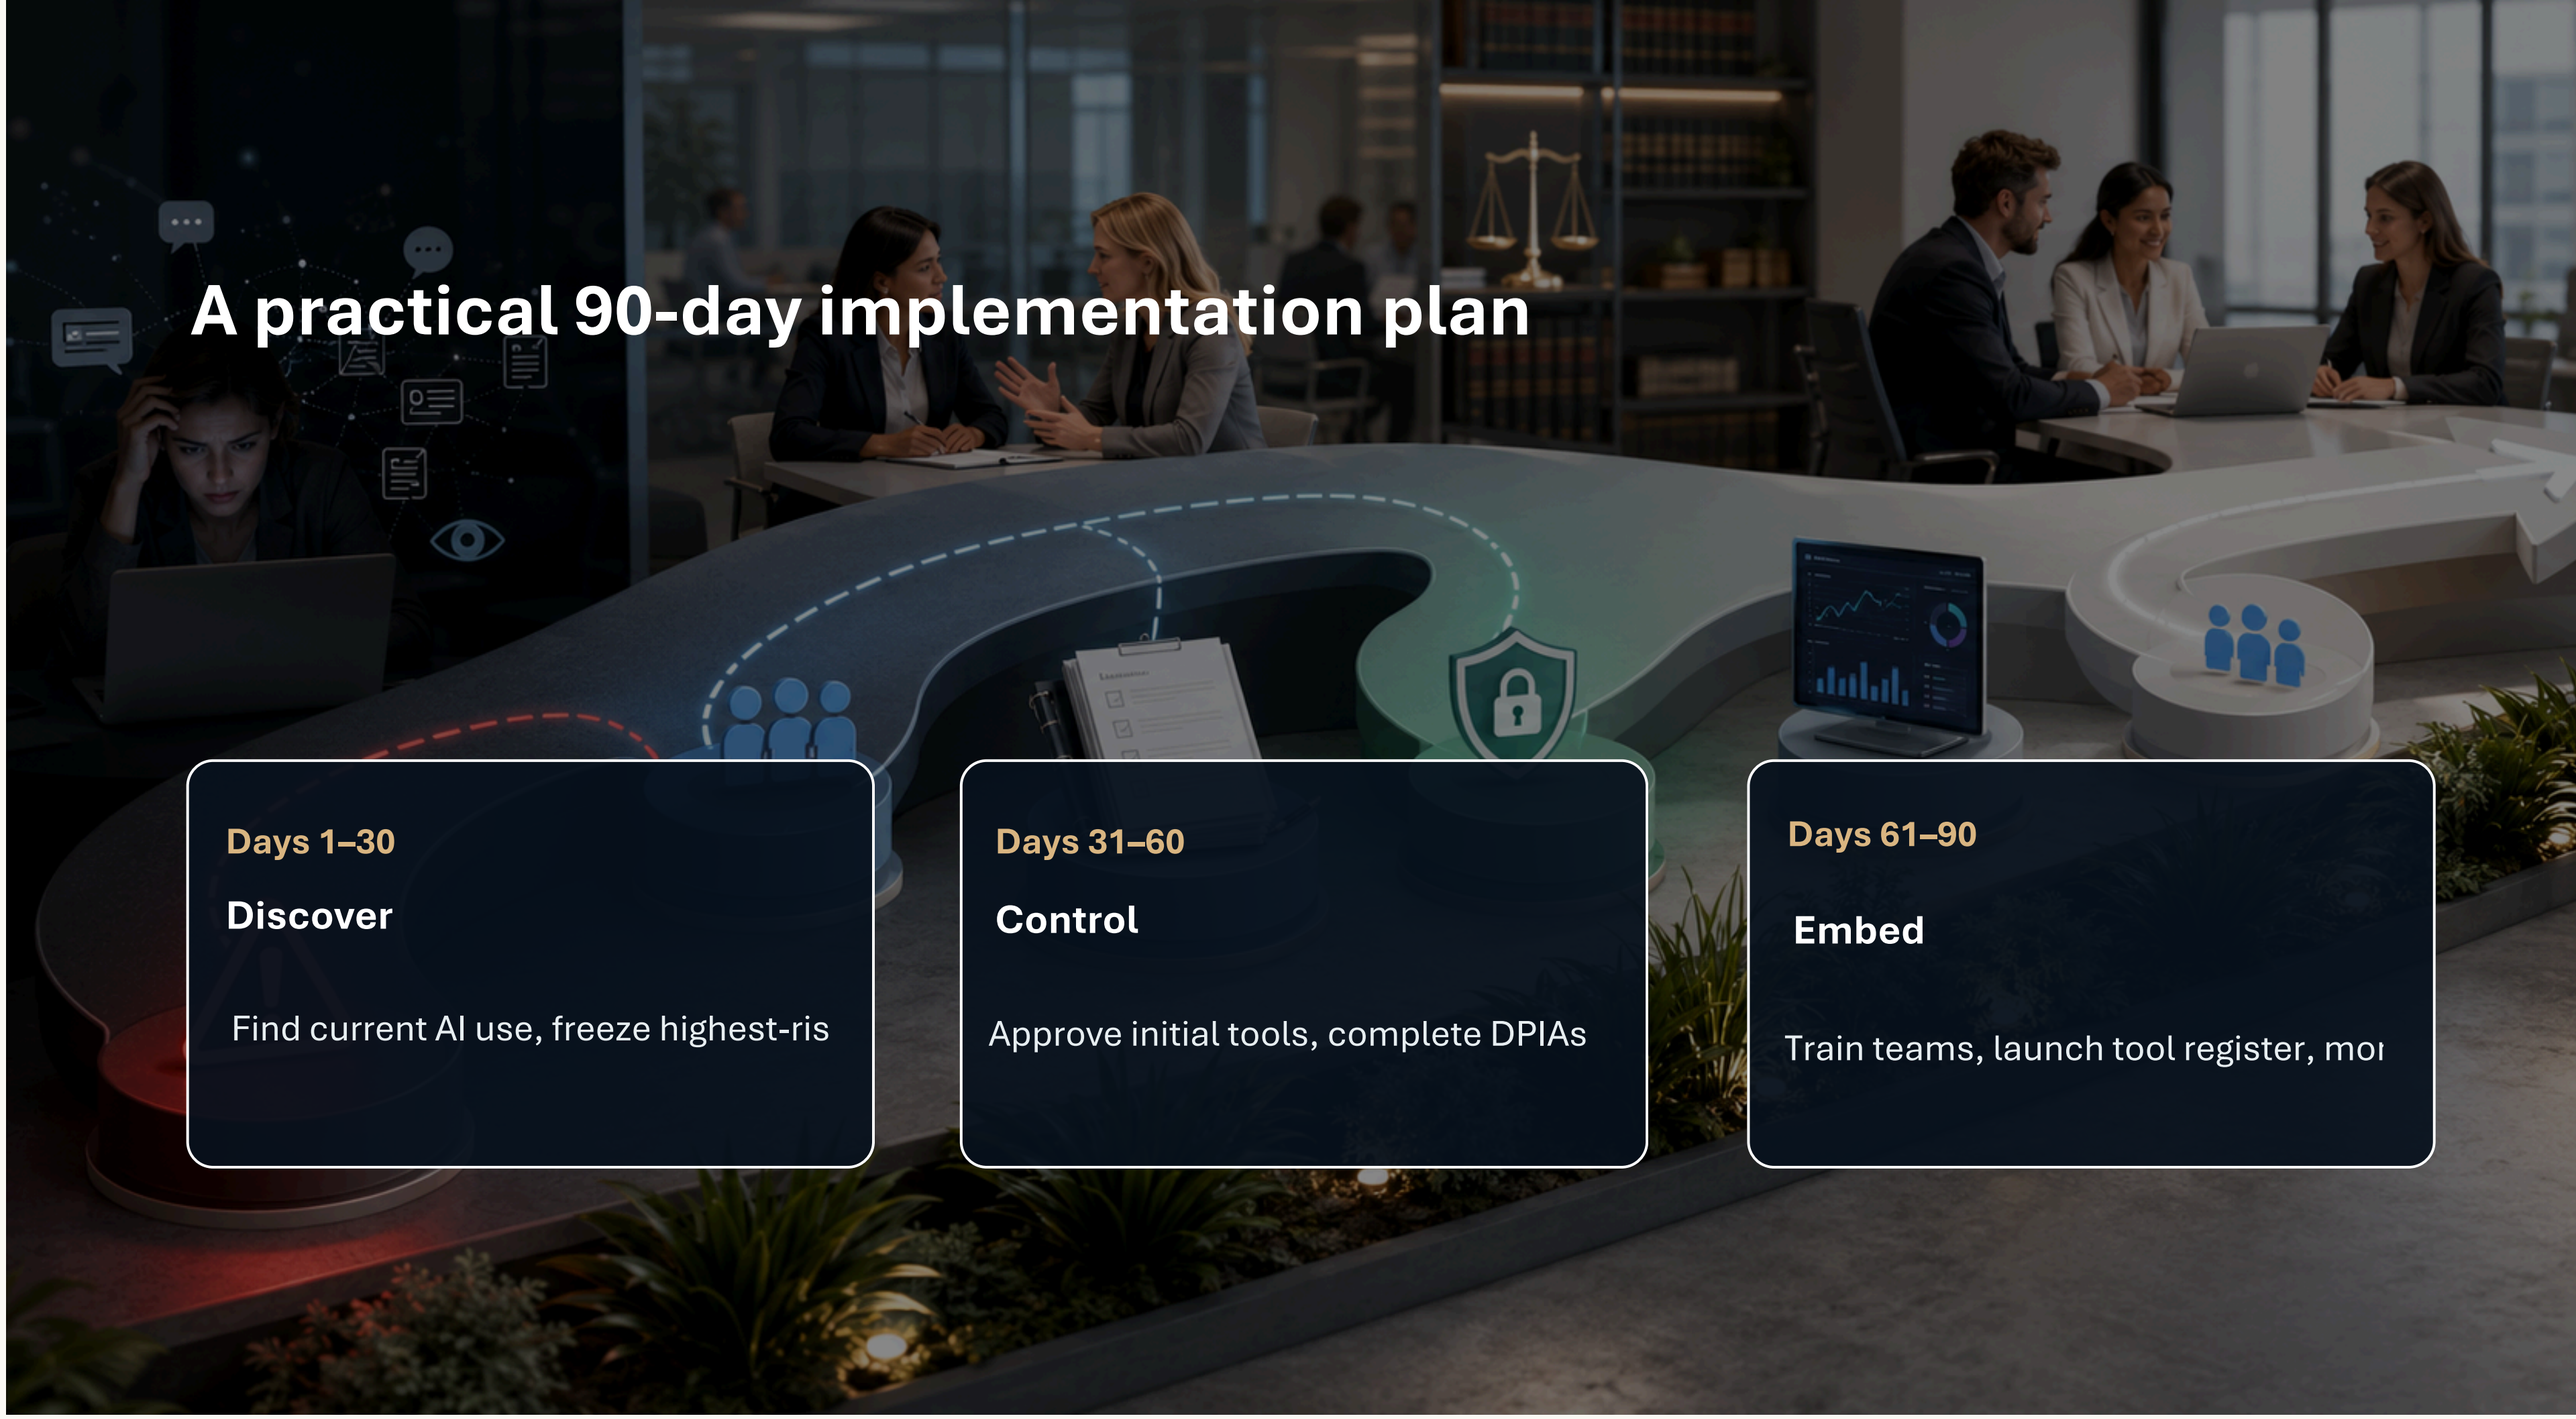

A practical 90-day implementation plan

Days 1–30

Discover

Find current AI use, freeze highest-ris

Days 31–60

Control

Approve initial tools, complete DPIAs

Days 61–90

Embed

Train teams, launch tool register, mor

TAKEAWAYS

The safe adoption checklist

- Know the use case and data categories before approving a tool.
- Keep client confidential and privileged material out of public AI.
- Complete DPIAs and vendor due diligence.
- Put contract terms around training, retention, security, transfers and deletion.
- Require human supervision and verification of AI outputs.
- Maintain policies, training, tool registers and review dates.

Responsible AI adoption is not about stopping innovation. It is about making use visible, controlled and defensible.

Questions