BEFORE YOU RELY ON IT

10 Questions Every Due Diligence Professional Should Ask About AI

A practical framework for compliance teams, KYC analysts, investigators, and AML professionals.

PRESENTED BY

[Presenter Name and Credentials]

[Organization / Event]

Replace the three bracketed lines above with your name, credentials, and event details; keep the rest unchanged. This guide is designed to be handed out at your seminar or shared directly with your clients.

How to Use This Guide

The compliance professionals who will win the next wave of client engagements are not the ones who avoided AI. They are the ones who can look a client in the eye and say: here is exactly how we use AI, here is where your subject data goes, and here is the evidence.

That question is already appearing in engagement letters, vendor questionnaires, and insurance renewals. The firms that can answer it confidently are keeping those relationships. The ones that cannot are finding out the hard way.

AI is already in your workflow, whether you introduced it deliberately or not. Colleagues are using it to screen names, summarize corporate structures, and draft adverse media narratives. The question is not whether AI belongs in due diligence. The question is whether your team is using it in a way that is defensible, auditable, and grounded in verified data.

Use this guide as a practical 20-minute checklist. If any answer is "Not yet," note who owns the gap and set a date to close it.

Part One: The Workflow

1. What specific due diligence task are you trying to accelerate?

AI performs very differently across due diligence tasks. Summarizing a publicly available company filing is a low-risk, high-value use. Running a PEP or adverse media screening with subject-identifying information attached is a different matter entirely: data exposure and in

terms of what happens when the output is wrong. Be specific about the task before you choose the tool.

> *Good starting points: translating foreign-language corporate records, summarizing company financials, or drafting narrative sections from structured findings. Riskier tasks: name matching, sanctions screening, and adverse media where subject identity is in the prompt.*

2. What data does this workflow actually touch?

Trace the data before you start. Does your AI workflow involve subject names, dates of birth, nationality, beneficial ownership structures, transaction data, or anything that could identify a natural person? Even partial information can be sufficient to constitute personal data under most privacy frameworks. The tool you use must be able to account for where that data goes.

> *Ask: Does this workflow involve personal data? If yes, where does it go when it leaves your system, who can access it, and how long is it retained? If you cannot answer those three questions, you are not ready to use that tool for this task.*
>
> *Rule of thumb: if it is in a client file, a watchlist workflow, or an investigation record, do not paste it into a consumer AI account. That rule applies regardless of which tool your colleagues are using.*

3. Is your AI output defensible as evidence of reasonable care?

Regulators, auditors, clients, and courts will ask how you reached your conclusions. If your AI tool cannot tell you which sources it used, what it excluded, or why it flagged something, the output is a starting point at best. A confident-sounding summary with no verifiable source trail is not a due diligence finding. It is a liability.

> *Ask: Could I explain this finding to a compliance officer, a regulator, or a judge? Could I reproduce it from primary sources if challenged? Can the tool output a source list with URLs or document IDs for each claim it makes? If not, the output needs verification before it enters any deliverable.*

4. Are you supplementing human judgment or replacing it?

AI is most powerful in due diligence as an accelerant, not a substitute. It handles volume, surfaces patterns, and reduces time on routine checks. The judgment call, whether a finding is material, whether a risk is acceptable, or whether a subject is the right individual, still belongs to the analyst. The workflow should make that boundary explicit and auditable.

> *The failure mode is not that AI gets things wrong. It is that analysts stop checking because AI sounds authoritative. Build verification into the workflow, not as an optional step but as a required one before any finding enters a report.*

> ***If you can answer Q1 to Q4, you have a defensible workflow foundation. If you cannot, keep AI use in low-risk tasks until these are answered.***

Part Two: Data, Security & Accuracy

5. Where does your subject data go when you run an AI query?

This is the question most due diligence teams have not fully answered. When you paste a subject name, a corporate registration, a beneficial ownership structure, or transaction details into an AI tool, that information is typically processed on vendor infrastructure outside your controlled environment, unless specifically architected otherwise. Depending on the tool and tier, it may be retained, logged, or used to train future models. You may be creating a record of your investigation outside your firm's systems without realizing it.

> *Ask the vendor directly and in writing: Where is data stored? In which jurisdiction? Who can access it? Is it used for model improvement or training? Ask for an architecture or data flow diagram and a DPA annex you can file as evidence. Written documentation beats verbal assurances in every audit.*

6. Does your AI use comply with your data handling obligations?

If you are handling personal data of EU subjects, GDPR applies regardless of where the AI tool is hosted. If you are screening for sanctions or financial crime, your data handling obligations under applicable AML and financial crime regulations apply too, including confidentiality, access control, retention, and auditability expectations. Most AI tools were not designed with these obligations in mind. The defensible position is a tool where you can demonstrate that subject data remained within a controlled boundary throughout the workflow.

> *The test: could you show a client or regulator exactly where their subject's data traveled during your AI-assisted screening process? If the answer is no, that is a gap worth closing before the next engagement.*

7. What happens when AI generates a false finding?

Treat any AI-generated finding as untrusted until verified against a primary source. No exceptions.

Hallucination in a due diligence context is not a minor inconvenience. If an AI tool generates a false adverse media finding, attributes a criminal record to the wrong individual, fabricates a sanctions link, or incorrectly connects a beneficial owner to a high-risk jurisdiction, and you act on it or share it with a client, you have a professional liability problem. If you deliver it in a report, you have a bigger one. AI systems that produce confident output with no source trail are the highest-risk category.

> *Mitigation: always verify AI-generated findings against primary sources before including them in a deliverable. This applies especially to adverse media narratives, beneficial ownership chains, and any finding that names a specific individual in a risk context.*

> ***If you can answer Q5 to Q7, you can evidence control over your data and your outputs.***

Part Three: The Questions Most Teams Have Not Asked Yet

8. Can you audit what your team is running through AI?

Compliance functions require audit trails. If a team member is using a personal AI account to run subject screening, those queries are invisible to your firm's oversight. That is a governance gap, and regulators are beginning to notice. The right infrastructure logs metadata: who ran the query, when, which tool, and which workflow. Query content stays governed under your own controls, not the vendor's.

> *Ask: If a regulator asked for your AI query logs from the past 90 days, what could you produce? If the answer is nothing, because those queries happened in personal accounts or consumer tools, that is not a defensible position.*

9. If a client asked how you handle their subject's data in AI workflows, could you answer today?

This question is already appearing in engagement letters, vendor due diligence questionnaires, and insurance renewals. The compliance teams that can answer it confidently will keep those client relationships. The ones that cannot will find themselves explaining, after the fact, why they did not think about it sooner. You do not need to have everything figured out. But you should be able to say: we have thought about this, we have a framework, and our AI use is auditable.

> *You do not need a perfect answer today. You need an honest one. "We are reviewing our AI governance framework and can walk you through our current approach" is a credible answer. "We use whatever tools the team prefers" is not.*

10. Is your AI infrastructure built for this work, or borrowed from somewhere else?

Most AI tools were designed for content creation, coding, or general productivity. They were not designed for environments where the data is a subject of a financial crime investigation, a PEP screening, a sanctions check, or a beneficial ownership inquiry in a high-risk jurisdiction. That design mismatch matters. The tool you use for compliance-sensitive due diligence should be one where data confidentiality is not a preference. It is an architectural control. The question is whether confidentiality is enforced by the vendor's promises, or by the design of the environment itself.

> *Ask: Is the confidentiality of subject data in this tool a function of architecture, or of the vendor's policy? Could that policy change with a terms update? Would your answer to that question satisfy a client audit or a regulatory examination?*

> ***If you can answer Q8 to Q10, you are ready for the client questions that are already coming.***

AI Readiness Scorecard for Due Diligence Teams

Use this with your compliance advisor as a working governance tool. If any item is Not Yet, name an Owner and set a Due date.

Question	Yes (approved)	Not Yet (gap)	Owner	Due
We have defined the specific due diligence tasks we are accelerating with AI	☐	☐		
We have traced what personal data our AI workflows actually touch	☐	☐		
Our AI outputs are verified against primary sources before entering any deliverable	☐	☐		

Question	Yes (approved)	Not Yet (gap)	Owner	Due
Our workflow makes the human judgment boundary explicit and auditable	☐	☐		
We know where subject data goes when we run a query and can document it	☐	☐		
Our AI use has been reviewed against applicable data handling and AML obligations	☐	☐		
We have a process for handling AI-generated false findings before they reach a client	☐	☐		
Our AI queries are logged and auditable by the firm, not just by the individual analyst	☐	☐		
We could explain our AI data practices to a client or regulator today	☐	☐		
Our AI infrastructure was chosen for compliance-sensitive data, not repurposed from general productivity tools	☐	☐		

Mostly Yes	A Mix	Mostly Not Yet
Strong position. Work with your advisor to formalize your AI governance framework and prepare a client-facing answer to the data handling question.	Normal for most teams in 2025. Use the Not Yet items as your preparation checklist before expanding AI use to higher-risk workflows.	Do not rush. Getting the foundations right protects your clients, your firm, and your professional standing. This guide is your starting point.

Prepared for you by your compliance advisor.

Your advisor: _____

If the questions in Part Three resonated, ask your advisor about AI infrastructure designed for compliance-sensitive workflows where subject data is designed to remain within your controlled environment.

TOTALLY PRIVATE AI

Air-gapped AI for regulated environments.

totallyprivate.ai