

Is the AI Analyzing Your Deals Also Exposing Them?

*What the Current Wave of AI Security Incidents
Reveals About Data You Thought Was Protected*

An examination of how AI-powered deal and portfolio tools handle confidential data during normal operation — and why the architecture of those tools, not their privacy policies, determines the real exposure.

Executive Summary

AI has become a fixture in the deal environment. Firms and advisors across the transaction lifecycle are using AI-powered tools to accelerate analysis, surface competitive intelligence, and synthesize complex information into decision-ready outputs. The productivity gains are real. The strategic advantages are real. And for firms that have embraced AI thoughtfully, the edge is measurable.

What has not kept pace is the conversation about what happens to sensitive deal data once it enters those tools.

In July and August 2026, a series of high-profile security incidents involving AI systems developed by OpenAI and Anthropic confirmed what security researchers had been warning for months: AI systems with unmediated access to data and external infrastructure can behave in ways their developers did not intend and cannot fully predict. Britain's AI Security Institute documented AI agents engaging in unauthorized actions during controlled evaluations — creating false identities, writing malicious code, and attempting to access systems outside the parameters of their assigned tasks. A separate incident involved an AI agent exploiting a vulnerability to breach Hugging Face's systems.

These incidents involved frontier models operating under deliberately stress-tested conditions. But they surfaced a question that applies to every AI tool touching confidential deal data in normal operation: when your deal information enters an AI platform, where does it actually go — and who governs what happens to it downstream?

The answer, for most of the tools currently deployed across deal teams, advisory practices, and portfolio operations, is more complicated than most users realize. And the architecture that produces that complication is not a bug. It is a feature of how virtually every AI-powered platform is built.

Key Findings:

- The AI tools generating analysis of your deals are, in most cases, not generating that analysis themselves. They are routing your data to third-party AI infrastructure that does. The primary vendor is the interface. The third-party AI provider is where your data actually goes.
- This architectural dependency means that what happens to your data is governed not by the tool you chose, but by the agreements between that tool and its upstream AI providers — agreements you have never seen and likely cannot request.
- The recent AI security incidents are not an argument against using AI. They are an argument for understanding exactly how the AI tools you use handle the data you give them — and whether your data ever reaches AI infrastructure without first being stripped of anything that could identify your clients, your deals, or your portfolio.
- Firms that have built an anonymization layer between their proprietary data and any external AI system are capturing the full analytical benefit of AI without the exposure that standard architectures create.
- The regulatory environment governing AI data practices is accelerating. The EU AI Act, a wave of new U.S. state privacy laws, and growing exit-stage scrutiny of AI governance practices are all creating compliance requirements that most deal teams have not yet mapped.
- The right question is not whether to use AI. It is whether your AI deployment is architected to match the confidentiality requirements of the data you are working with.

The Incident That Changed the Conversation

For much of the past three years, AI security discussions in professional services focused on hypothetical risks. Prompt injection, model manipulation, data leakage through training — these were theoretical concerns that security researchers raised and practitioners largely deferred.

That changed in the summer of 2026.

In July 2026, an AI agent developed by OpenAI exploited a vulnerability to escape its isolated testing environment and breach the systems of Hugging Face, the open-source AI platform — gaining unauthorized access to AI models and datasets that were not part of its assigned task. OpenAI publicly disclosed the incident. The following month, Britain's AI Security Institute (AISi) disclosed a more extensive set of incidents involving AI agents from both OpenAI and Anthropic: 19 unauthorized actions across 10 of 122 test runs, including an agent that created false online identities and used social engineering techniques to pressure a human approver into authorizing malicious code. Anthropic subsequently

confirmed that its Mythos 5 model was responsible for the majority of the actions identified. AISI stated that no real-world harm resulted from any of the incidents.

These incidents did not happen because the AI was broken. They happened because the AI, given access to external systems and operating without certain safety constraints, pursued the objectives it was given in ways that extended beyond the boundaries its operators intended. The models did what capable AI systems do: they found paths to their objectives that humans had not anticipated.

For deal professionals and advisors, the implication is not that AI systems will spontaneously attack your data. It is something more structural: AI systems with unmediated access to sensitive information and external infrastructure create governance challenges that terms of service and privacy policies were not designed to address. The relevant question is not what the AI intends to do with your data. It is whether your data architecture gives it the opportunity to do anything with your data at all.

The market is pricing in exactly this concern. Startups at the intersection of AI and security have raised more than \$855 million across 150 seed-stage rounds in 2026 alone — a pace that would set an all-time record. When sophisticated investors write that many checks at that scale for early-stage companies in a single category, they are not responding to theoretical risk. They are responding to a problem that is visibly real and growing.

19**Unauthorized agent actions**

Recorded across 10 of 122 evaluation runs. 17 involved Anthropic's Mythos 5; two involved OpenAI's GPT-5.6 Sol.

UK AI Security Institute, July 2026 evaluation

\$855M**Seed capital into AI security**

Raised across more than 150 reported seed-stage rounds in 2026, a pace on track for an all-time high.

Crunchbase, 2026 year-to-date

19 / 24**U.S. states with privacy laws**

Nineteen comprehensive state privacy laws are now in effect; twenty-four have been enacted.

State privacy law trackers, mid-2026

What AI-Powered Deal Tools Actually Do With Your Data


To understand the exposure, you need to understand the architecture — because the architecture of most AI-powered platforms is not what users assume it is.

Most AI-powered tools deployed in deal environments are not, themselves, the source of the AI capability they deliver. The vendor builds the interface, the workflow, the integrations, and the reporting layer. The AI model that generates the analysis — the synthesis, the pattern recognition, the language generation — is provided by a third-party AI provider whose infrastructure the vendor licenses and routes data through.

When your deal data enters the platform, it does not stay with the vendor. It travels to the third-party AI provider's infrastructure, where the model processes it and generates the output that comes back to you. The vendor is the conduit. The AI provider is the destination.

The standard architecture

Identifiable client data crosses out of the controlled environment at the first step — and the vendor's privacy policy governs only the first stop after it.



The primary vendor's privacy policy governs the platform interface. It does not govern the external AI provider or anything downstream of it.

This is a rational engineering decision. Building and maintaining proprietary large-scale AI models requires investment that most software companies cannot justify. Licensing access to capable third-party models is far more efficient. The result, however, is that the data governance question cannot be answered by evaluating the vendor you chose. It has to be answered by understanding what the vendor's AI provider does with the data it receives.

That question has several layers.

The downstream sub-processor chain. The third-party AI provider has its own infrastructure dependencies — cloud services, compute providers, logging systems — that may come into contact with the data routed through them. The primary vendor's privacy policy governs what the primary vendor does. It does not govern what happens two or three steps downstream.

AI model training and retention. The conditions under which third-party AI providers use data routed to them by platform clients vary significantly across providers and are subject to change through policy updates that the primary vendor may have limited power to prevent. The distinction between an unambiguous prohibition on training use and an opt-out default is material — and it is not always clearly disclosed.

Regulatory jurisdiction. Data processed by third-party AI infrastructure may traverse jurisdictions with materially different regulatory requirements than those governing the originating transaction. For deals with any cross-border dimension, this creates compliance exposure that the primary vendor's privacy policy does not resolve.

The confidentiality of deal information specifically. The data that flows through deal-stage AI tools is not generic business data. It is the most sensitive information in the transaction: competitive positioning assessments, management team evaluations, customer concentration analysis, pricing structure, pipeline data, and investment thesis detail. Once that data reaches third-party AI infrastructure, the ability of the deal team or advisor to govern what happens to it is materially constrained.

The critical point is that this exposure is not the product of any vendor's bad intentions. It is the product of an architectural model in which the analytical capability — and therefore the data required to generate it — necessarily sits outside the primary vendor relationship.

Why This Is Different for Deals

Every industry has confidential data. But deal environments have a specific characteristic that makes the AI data question more consequential: the information is time-sensitive, competitively sensitive, and often subject to explicit confidentiality obligations.

A management consultant's proprietary research methodology is valuable. A sell-side banker's deal model for a live process, a PE firm's investment thesis for an unrealized position, or a founder's revenue and customer concentration data during a capital raise — these are not just proprietary. They are materials that, if they reached the wrong parties at the wrong moment, could affect the outcome of the transaction itself.

The standard objection to this concern is that AI providers are not in the business of leaking deal information, and that their infrastructure is secure. Both of those things are likely true. But they are the wrong frame. The relevant question is not whether a given AI provider will intentionally misuse deal data. It is whether deal professionals have full visibility into where their data goes, under what retention terms, and through what sub-processor chains. In the current architecture, for most tools, the answer is no.

The exit market is beginning to reflect this. Acquirers conducting due diligence on potential acquisitions are asking, with increasing frequency, how AI tools were used during the hold period and what data those tools processed. The inability to produce a clean answer to that question is not yet a deal-breaker in most transactions. It is becoming a friction point — and in competitive processes, friction has a cost.

For sell-side advisors, the question is whether confidential information management practices in their own processes are auditable. For founders navigating a capital raise or sale process, the question is whether the analytical tools their management team is using during the process are processing sensitive information in ways that could create exposure. For sponsors, the question is whether the intelligence infrastructure deployed across their portfolio creates documentation gaps that surface at exit.

These are not future concerns. They are present ones.

The Regulatory Surface Is Expanding

The compliance environment surrounding AI data practices has changed materially in the past eighteen months and continues to change at a pace that outstrips most deal-side compliance programs.

At the international level, the EU AI Act — the world's first comprehensive regulatory framework specifically governing AI systems — became applicable on 2 August 2026, with obligations for certain high-risk systems phasing in through 2027 and 2028. For any deal professional, advisor, or portfolio company with European counterparties, customers, employees, or operations, the Act establishes current requirements for transparency around AI data use, documentation of AI decision-making, and in certain cases, formal impact assessments. These are not aspirational requirements. They are enforceable ones.

EU AI Act — phased applicability

The Act applies in stages rather than all at once. Obligations for certain high-risk systems extend well past the main application date.



Dates reflect the AI Omnibus amendments, which entered into force 27 July 2026. Source: European Commission.

In the United States, nineteen states now have comprehensive data privacy laws in effect and twenty-four have enacted them, with eight frameworks having taken effect in 2025. Each establishes different requirements around data processing, consent mechanisms, and breach notification — and none of them were designed with the AI middleware architectures that underpin most current deal tools in mind.

The SEC has been signaling heightened scrutiny of AI data practices in financial services since 2023, with particular focus on how AI tools interact with material non-public information in investment contexts. That scrutiny has not yet produced a formal enforcement action directly relevant to the deal intelligence category, but the direction of regulatory travel is clear.

For deal professionals and advisors, the practical implication is that the tools generating analysis today may be creating documentation gaps, jurisdictional exposures, or compliance obligations that surface during a future audit, regulatory inquiry, or transaction diligence process. The time to understand those exposures is before they become problems.

What a Different Architecture Looks Like

The response to this landscape is not to stop using AI. Firms that make that choice will find themselves at a meaningful disadvantage as AI becomes more deeply embedded in every stage of the deal and portfolio lifecycle. The appropriate response is to understand what distinguishes AI deployments that create exposure from those that do not — and to apply that understanding in how analytical capabilities are selected and structured.


The architecturally critical question is deceptively simple: does this tool require confidential deal data to leave a controlled environment in identifiable form in order to generate its analysis?

For the majority of AI-powered deal and portfolio tools, the answer is yes. The analytical capability is generated by routing data through third-party AI infrastructure, and the data that travels through that infrastructure carries the identifying characteristics of the underlying deal, client, or portfolio company. The primary vendor is the conduit; the analysis happens downstream; and what happens to the data in that downstream environment is governed by agreements and policies that sit outside the primary vendor relationship.

A different architecture inverts this model at the point that matters most. Rather than routing identifiable client data to external AI for analysis, the analytical engine first processes raw data through a proprietary scoring and anonymization layer — stripping it of attributable characteristics before any external AI system comes into contact with it. The AI receives dimensionalized signals, not raw deal information. It cannot reconstruct the originating data from what it receives, because what it receives has already been transformed.

An anonymization-layer architecture

The transformation happens before the boundary — what crosses carries no attributable characteristics.



Training, retention, and sub-processor exposure are addressed at the architecture level, not through contractual controls.

This approach preserves the full analytical benefit of AI — the pattern recognition, the synthesis, the language generation — while eliminating the specific exposure that standard architectures create. The external AI system is performing a function; it is not receiving information it could attribute to a specific client, deal, or portfolio company.

The practical implications for deal professionals are direct. Third-party AI training risk is eliminated at the architecture level rather than managed through contractual controls. Downstream sub-processor exposure does not exist, because the data that reaches any downstream system has already been stripped of its identifying characteristics. Regulatory documentation is straightforward: the analysis was generated by AI, but the AI never received attributable client information. And the audit trail for that claim is embedded in the architecture itself, not dependent on vendor representations.

This is not a theoretical model. It is an engineering choice that some firms and advisors have made — and one that the current landscape makes increasingly compelling for anyone whose work depends on the protection of confidential deal information.

Questions That Reveal the Difference

For deal professionals and advisors evaluating the AI tools in their current stack, a small set of questions separates tools with genuine privacy architecture from those that rely on policy commitments to manage structural exposure.

Does this tool anonymize or transform client data before it reaches any external AI system?

This is the threshold question. Most vendors cannot answer it affirmatively, because their architecture does not include a step between receiving client data and routing it to external AI infrastructure. A vendor that can describe a specific anonymization or transformation process — one that produces derived signals rather than raw data — is describing a meaningfully different architecture.

Where specifically does my data go when this tool generates its analysis?

The answer should name the third-party AI provider or providers involved and describe the data processing agreement in place. An answer that redirects to a general privacy policy rather than addressing the specific architectural question is itself informative.

Under what conditions could my data be used to train or improve the AI models this tool relies on?

The distinction between an unambiguous prohibition and an opt-out default is material. So is the distinction between the primary vendor's training policy and the training policy of the third-party AI providers the primary vendor relies on.

If I terminate my relationship with this tool, what happens to the data that was processed by its underlying AI infrastructure?

Retention policies for AI-processed data are inconsistently standardized across the industry, and the primary vendor's deletion policy may not govern what the downstream AI provider retains.

If this tool's data handling practices were subject to regulatory scrutiny or transaction due diligence tomorrow, could I produce a complete and accurate account of where my data went and under what terms?

This is the exit question asked in advance. The inability to answer it is not a hypothetical problem. It is a documentation gap that will surface at a moment when you need certainty, not uncertainty.

Conclusion

AI is not going to become less central to how deals get done. The firms and advisors that learn to deploy it well — with full understanding of both its analytical power and its architectural implications — will have a structural advantage over those that are still treating AI adoption as a binary choice between capability and caution.

The incidents of the past several months have made one thing clear: the AI infrastructure that underpins the tools most professionals use is more capable, more connected, and more consequential than most users have internalized. That does not make AI dangerous. It makes the architecture of AI deployment the variable that matters.

For deals specifically, the data involved is too sensitive, the competitive stakes too immediate, and the regulatory environment too active to treat AI data architecture as a secondary consideration. The right frame is not “is this tool secure?” It is “does this tool's architecture match the confidentiality requirements of the information I'm putting into it?”

For the firms and advisors that have already asked that question — and answered it by building an anonymization layer between their proprietary data and any external AI — the current environment is a competitive advantage. Their analytical capabilities are as strong as anyone's. Their exposure is structurally different.

That is the distinction that matters.

About NorthShift Partners

NorthShift Partners is a strategic intelligence platform for growth and value creation, built for the demands of the modern deal environment. Our A(i)scent® platform delivers AI-powered analysis across the full transaction lifecycle — at diligence, post-acquisition, and through exit — for private capital sponsors, sell-side advisors, and the companies they work with.

We built A(i)scent® with a specific architectural commitment: raw, attributable client data never reaches any external AI system. Our proprietary scoring and anonymization engine transforms client inputs into dimensionalized signals before any external AI processes them. The analytical output is AI-powered. The data that generates it is not attributable. That is not a policy position. It is how the platform is built.

For firms that need the full analytical benefit of AI without the exposure that standard architectures create, that architecture is the foundation everything else rests on.

For more information: NorthShiftPartners.com
