Confidencial

RESEARCH REPORT · MAY 2026

PHARMA & LIFE SCIENCES · DATA PROTECTION

The 2026 State of Data Protection in Pharma & Life Sciences

Third-Party Exposure, AI Governance, and the Chain of Custody Problem

ORGANIZATIONS ANALYZED

199 Pharma, Biotech & CRO Organizations

GEOGRAPHY

U.S., Canada, EU & UK

PUBLISHED

May 2026

THE PIPELINE HAS NO PERIMETER

The data leaving your walls represents decades of commercial value — and no one is governing where it goes.

Pharma and biotech companies don't sell physical product the way most industries do. They sell molecules, trial data, and the regulatory packages that turn one into the other. The data that moves through a sponsor's ecosystem isn't just sensitive — it represents the entire commercial value of the molecule, the regulatory credibility of the submission, and the trust that makes CRO relationships work. Lose control of that data, and you don't just have a security incident. You have a threat to your exclusivity window, your FDA relationship, and your ability to partner.

Maintaining that control has become structurally harder. By regulatory mandate, sponsors hold the most sensitive information their patients possess. By operational necessity, that information moves constantly — to CROs, to central labs, to imaging vendors, to regulators in every jurisdiction the trial touches, and now into AI tools that promise to compress development timelines by years. At every handoff, the Data Package leaves one environment and enters another. The governance doesn't follow it. The audit trail ends at the door.

> The data protection frameworks most pharma companies have built were designed for a world where clinical data stayed inside the four walls of the sponsor and its primary CRO. That world is gone.

This research, based on a structured analysis of 199 pharmaceutical, biotech, and clinical-research organizations, reveals a structural vulnerability that no clinical document platform currently solves natively: governance enforcement at the moment a document leaves the sponsor's environment. The findings are organized around one argument in four parts — data leaves without enforceable governance; it travels through a third-party network whose depth sponsors cannot fully see; AI has added an invisible layer at every link; and the audit trail that would prove what happened to your Data Package doesn't exist for most sponsors today.

29.6/40
Average Data Protection Exposure Score across 199 organizations

15%
Sponsors in the Critical exposure band (35–40)

$7.42M
Avg. cost of a healthcare data breach in 2025
IBM Cost of a Data Breach Report 2025

20%
Breaches in 2025 involving shadow AI — avg. cost premium $670K
IBM CODB 2025

82%
Sponsors scoring maximum on Therapeutic Sensitivity

52%
Sponsors scoring 4–5 on Jurisdictional Complexity

RESEARCH FRAMEWORK

The Eight-Dimension Model — Built on Open Source Data

The findings in this report are not based on surveys or sentiment. They are derived from a structured analysis of 199 pharmaceutical, biotech, and clinical-research organizations — drawing exclusively on open source data: company websites, published technology partnerships, clinical-trial registries, regulatory filings, breach notifications, and media records. Each dimension is scored 1–5, producing a maximum composite score of 40.

> Because this research relies on open source data, it represents a conservative floor of actual sector risk — especially for third-party network depth. Clinical trial registries and investor filings surface primary CRO relationships. Sub-CRO networks, central lab partnerships, and downstream AI tooling are rarely disclosed publicly. The actual third-party chain is longer than any external analysis can fully see. Read these scores as a floor, not a ceiling.

01
Therapeutic Sensitivity
Oncology, rare disease, pediatric, gene therapy, and genomics carry the highest regulatory and commercial sensitivity. Data from these programs cannot be recreated if compromised.

02
Regulated Data Exposure
Every sponsor handles HIPAA-regulated PHI, FDA clinical data, or EMA documentation. Most handle all three. Regulators expect end-to-end control, not just inside the sponsor's walls.

03
Jurisdictional Complexity
Cross-border trials, multi-region manufacturing, and global submissions multiply the regulatory surface. The molecule, not the headquarters, defines the exposure.

04
Clinical Trial Footprint
Trial count and phase mix proxy for data volume, partner count, and data supply chain breadth. More trials mean longer third-party chains and higher chain of custody risk.

05
Clinical Data Retention
Multi-year ePHI and IP retention under FDA 21 CFR Part 11 and EU Annex 11 means sensitive data accumulates and ages. The Data Package doesn't expire — it grows.

06
Technology Adoption
Breadth of the disclosed technology stack and resulting governance surface. Every new platform is a potential handoff point where chain of custody can break.

07
Industry Risk Baseline
Healthcare has been the most expensive sector for data breaches for 14 consecutive years. This is a structural floor applied uniformly. (IBM CODB 2025)

08
AI Adoption Depth
Confirmed enterprise AI plus breadth of LLM and ML platforms disclosed in the public stack — a proxy for shadow AI risk and ungoverned data movement into model pipelines.

Dimensions 1, 3, 5, and 8 are directly derived from structured enrichment of public data. Dimensions 2, 4, 6, and 7 are supplemental, derived from the same sources using a documented scoring rubric available on request. Risk bands: Critical 35–40 · High 28–34 · Elevated 18–27 · Moderate below 18. No organization in this cohort scored in the Moderate band.

FINDINGS — 199 ORGANIZATIONS

A sector average of 29.6 out of 40 — and the floor is structural.

Composite exposure scores produced a sector average of 29.6 out of a maximum of 40. Even the lowest-scoring organizations carry substantial inherent exposure — driven by the nature of pharmaceutical work itself, not by any specific governance failure. No organization scored below 19. The baseline risk of pharma data handling is structural and irreducible.

EXPOSURE BAND DISTRIBUTION — 199 ORGANIZATIONS

15% — **Critical — Score 35–40.** Large global sponsors with confirmed AI deployments, complex CRO networks, and highest-sensitivity therapeutic areas. Maximum exposure across most dimensions.

59% — **High — Score 28–34.** The majority of the sector. Elevated exposure across most dimensions, significant patient data risk, multi-jurisdictional trial footprint, and frequently underestimated AI exposure.

27% — **Elevated — Score 18–27.** Lower relative exposure, typically reflecting narrower therapeutic focus or pre-commercial stage. Still above the Moderate band — there is no low-risk pharma organization in this cohort.

HOT ZONES IN THE DATA

100% **Regulated data exposure is universal.**
Every sponsor in the cohort scored 4 or 5. All operate under HIPAA, FDA Part 11, EMA, PMDA, or equivalent. Regulators expect documentation control end-to-end — not just inside the sponsor's environment.

82% **Therapeutic sensitivity is concentrated at the top.**
Oncology, rare disease, pediatric, gene therapy, and genomics dominate the modern pipeline — all in the highest tier of regulatory and commercial sensitivity. The data these programs generate is irreplaceable. It cannot be recreated if compromised.

36% **Pharma is more cross-border than legal.**
36% of sponsors scored maximum on Jurisdictional Complexity — more than the comparable AmLaw 200 cohort at 30%. Pharma's data crosses more borders than legal's, and does so under more regimes. The molecule, not the headquarters, defines the exposure.

38% **AI adoption is already meaningful — and the disclosed number is the floor.**
38% of sponsors score 4–5 on AI Adoption. 19% have a frontier LLM (OpenAI, Claude, Gemini, Copilot, Azure OpenAI) confirmed in their public stack. The undisclosed floor is materially higher.

ARCHITECTURE GAP

Governance Ends at the Door

Every major clinical document platform in use across this cohort was designed to manage documents within a controlled environment. The moment a trial master file is exported to a CRO, an imaging dataset is shared with a central lab, a CMC document is submitted to a regulator, or a study protocol is uploaded to a partner portal, controls like access, retention, and classification become inoperative. From a governance standpoint, that data is now unmanaged. The document is still moving. The governance stopped at the door.

PLATFORM	ADOPTION	GOVERNANCE IMPLICATION
Veeva Vault (any module)	**43%**	Most-disclosed clinical content platform. Controls end at export.
OnCore CTMS	**23%**	Trial operations data exchanged with study sites and CROs.
Multi-cloud (AWS + Azure + GCP)	**22%**	Data residency complexity across three or more hyperscalers.
Medidata Rave / Clinical Cloud	**10%**	EDC and randomization data with multi-party access.
TrackWise / QMS systems	**7%**	Quality and deviation data with regulator visibility.

① SECURE (INTERNAL)

A trial lead exports the latest protocol amendment and informed-consent package from Veeva Vault eTMF, tagged under the sponsor's classification scheme.

② THE BREACH POINT

The package is sent to the CRO, three regional study coordinators, and the central lab. The moment it leaves the sponsor's environment, the classification tag becomes informational — not enforceable.

③ UNGOVERNED EXPOSURE

Recipients can forward the document, save to a personal drive, or drop it into a generative AI tool. The sponsor has no visibility, no mechanism to revoke access, and no audit trail.

> These platforms govern data inside the environment. None of them govern what happens after export. The document is still moving. The governance stopped at the door.

The ePHI blind spot. Structured EDC fields are well-governed. The blind spot is the PDFs, screenshots, source-document transfers, monitoring reports, and emailed amendments where PHI sits in the body of the file — unstructured, unmapped, and outside every policy written to protect it. The 2026 HIPAA Security Rule Final Rule explicitly requires asset inventories of every system that creates, receives, maintains, or transmits ePHI — including AI software. Most sponsors cannot produce that inventory today.

87% of healthcare and pharmaceutical companies have been negatively affected by a breach in their third-party ecosystem. *(Help Net Security, 2025)* · Healthcare is now the fourth most-targeted industry globally for ransomware, rising 4.8% year-over-year. *(CybelAngel, 2025)*

WHERE THE CHAIN BREAKS

The Chain Is Longer Than You Think

The conventional picture of pharma's data supply chain understates the risk by roughly half. Sponsors think about their primary CRO. They don't always think about what the primary CRO brings with it. A single Phase III oncology trial can involve a primary CRO, two or three sub-CROs handling specific geographies or functions, a central lab, an independent imaging vendor, a biomarker lab, a site management organization, and a pharmacovigilance partner. Each has its own technology stack, its own data handling practices, and its own staff making decisions about how to work with the sponsor's Data Package. The sponsor's governance covers none of them.

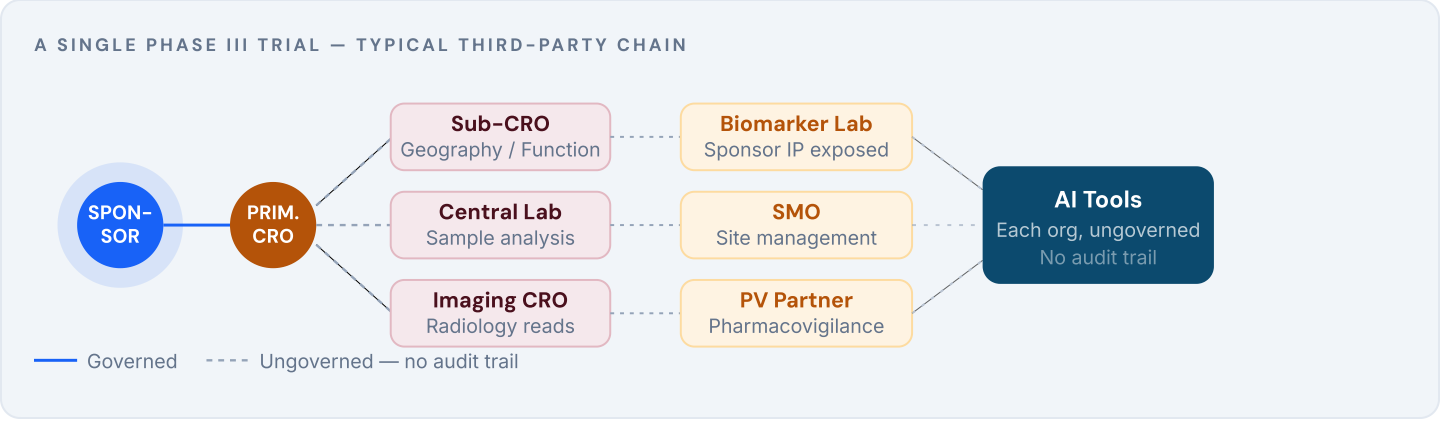

The Contract Gap

Data processing agreements define what organizations are permitted to do with sponsor data. They don't enforce it architecturally. A CRO statistician summarizing interim analysis in an AI tool isn't violating the letter of most agreements written before 2022 — because most agreements written before 2022 didn't anticipate that capability. The contract describes behavior. It does not control it.

CROs Can Learn Your IP Without Stealing It

A major pharma CIO recently flagged a situation that illustrates the aggregation problem precisely: a CRO their organization worked with was attempting to pool data across multiple sponsor clients — trial data, compound information, patient data — into a shared dataset. The CIO described it plainly: the CRO wanted to smush all of their clients' data together. Their CTO would have lost their mind. This isn't an edge case. CROs working across multiple sponsors have growing technical capability to identify patterns across datasets. A sponsor's IP doesn't have to be exfiltrated to be compromised. It can be learned. And it's nearly impossible to detect — because the sponsor has no audit trail once the file crosses the boundary. GxP principles require trial data not be used outside the scope of the agreement. Aggregation violates both data integrity requirements and those terms.

The Procurement Moment

Forward-looking sponsors are moving data protection requirements upstream — into procurement and contracting, before the first file transfer occurs. When protection is architectural rather than contractual, the chain of custody becomes enforceable regardless of where the file travels.

The sponsor retains visibility into every access event: who opened the file, when, from which environment, and whether any AI tool ingested the content. That audit trail persists across every hop — primary CRO, sub-CRO, central lab, partner AI pipeline. In a GxP context, this isn't a security capability. It's a data integrity record. It's what you produce when a regulator asks what happened to that file, or when a CRO's data handling practices are called into question.

Audit Trail Across Every Hop

Governance that travels with the data means the sponsor can answer: who touched this file, when, from where, and whether any model ingested it — at every link in the chain, permanently. That's not a security feature. That's GxP compliance infrastructure.

THE INVISIBLE LINK IN THE CHAIN

AI Has Added a Link to the Chain No One Governed For

Generative AI doesn't introduce a new category of risk to the Data Package. It extends the chain of custody problem into a new environment — one that is harder to see, harder to audit, and for most sponsors governed by nothing more than a policy statement and a data processing agreement written before ChatGPT existed. Every link in the third-party chain now has its own AI layer the sponsor didn't sanction and cannot audit.

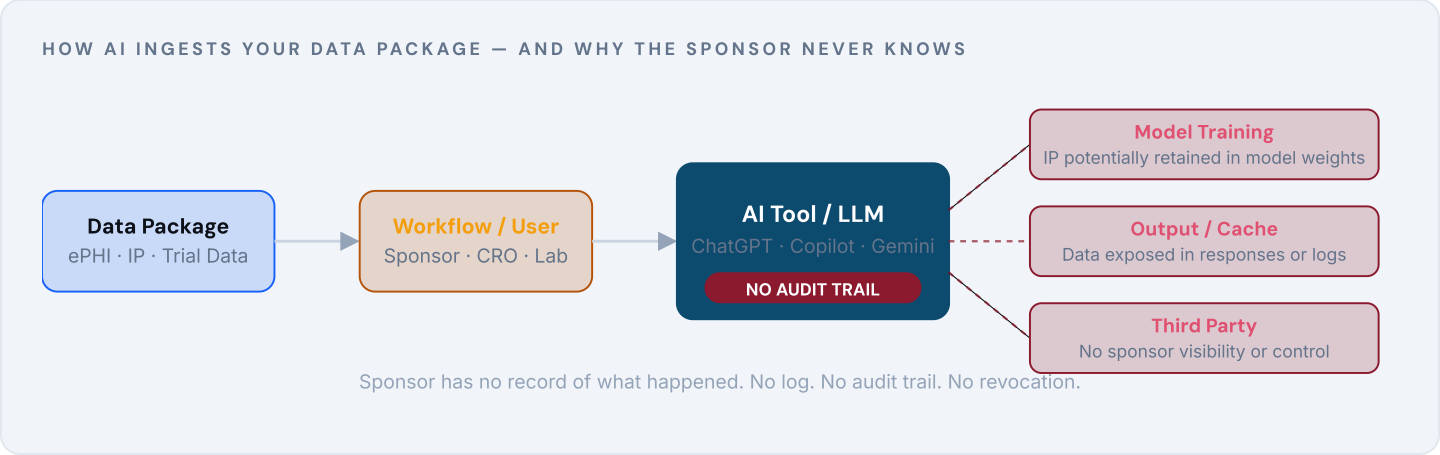

Why AI Has to Work in Pharma

Drug discovery: Rentosertib, the first end-to-end AI-discovered and AI-designed drug, is now in Phase II trials. Insilico Medicine nominated 22 AI-discovered preclinical candidates between 2021–2024, averaging 12–18 months to nomination vs. the industry standard of 4–6 years. *(Insilico / Nature Biotechnology, 2024)*

Clinical operations: McKinsey estimates AI's near-term opportunity in clinical development at $13–25B annually — protocol design, site selection, patient matching, and adverse-event detection. *(McKinsey Global Institute, 2024)*

Regulatory engagement: FDA has reviewed over 500 drug and biologic submissions with AI components since 2016. The competitive pressure to adopt is existential for mid-size sponsors.

Why It Hasn't Worked Yet

ePHI in unstructured files: Sending ePHI to a third-party LLM without protection is an unauthorized disclosure under HIPAA. The 2026 Security Rule overhaul makes encryption mandatory and explicitly includes AI in the asset inventory requirement.

IP leakage into foundation models: Compound structures, target rationales, and trial protocols sent into a foundation model can be retained or surfaced in training. A DPA is not an architectural control. Pharma data stays sensitive for 15–20 years.

CRO AI exposure: Many CROs are deploying their own AI on sponsor data under contracts that didn't anticipate generative AI. The third-party blind spot and the AI problem are the same problem — one link further down.

Regulatory uncertainty: FDA's January 2025 draft guidance introduces a credibility-assessment framework for AI in submissions. The EU AI Act treats AI in clinical trial recruitment as high-risk under Annex III — penalties up to €35M or 7% of global turnover.

Cross-jurisdictional residency: EU clinical trial data, US PHI, and Chinese patient data carry residency rules that conflict when the same model operates across all three jurisdictions.

THE SHADOW AI PROBLEM

19% Is the Floor, Not the Ceiling

19% of sponsors in the cohort have a frontier LLM confirmed in their public technology stack. The bigger story is what isn't disclosed. Shadow AI — unsanctioned tools used by individual employees — leaves no enterprise record, triggers no governance policy, and produces no audit trail.

20%
of 2025 breaches involved shadow AI — avg. cost premium $670K per incident
IBM CODB 2025

97%
of organizations with AI-related breaches lacked proper AI access controls
IBM CODB 2025

57%
of employees hid AI use from employers — only 40% of workplaces had any policy
KPMG 2025, 48K workers

46%
reported internal data leaks through GenAI applications
Cisco 2025

In pharma, none of these are policy violation problems. They are Data Package integrity problems. A clinical trial coordinator summarizing a protocol amendment in ChatGPT. A statistician uploading interim analysis to an AI tool. A medical writer pasting patient narratives into a consumer model. None captured in enterprise deployment data. All ePHI or IP leaving the sponsor's environment without governance, without logging, and without any audit trail to produce when a regulator asks what happened to that data.

WHAT MOST SPONSORS HAVE TODAY

A Policy Document

"Do not use unsanctioned AI tools."
No enforcement. No audit. No record.

WHAT ACTUALLY CONTROLS IT

Architecture

Encryption travels with the file.
Every access logged. Every hop audited.

THE BOTTOM LINE

A policy that says "do not use unsanctioned AI tools" is not a control. It is a wish.

Compound IP, clinical trial data, and ePHI are moving through AI tools right now — inside sponsor walls, inside CRO environments, at every link in the chain. The organizations that govern it architecturally will know exactly what happened and when. The ones that don't will find out later.

JURISDICTIONAL RISK

The Data Follows the Molecule

Regulatory exposure is not defined by where a sponsor's headquarters is located. It's defined by where patients enrolled, where labs ran samples, and where data crossed a border. 52% of assessed sponsors scored 4 or 5 on Jurisdictional Complexity. 36% scored maximum. A US-headquartered sponsor running a Phase III oncology trial with sites in the EU, UK, China, and Japan is simultaneously operating under HIPAA, GDPR, the UK Data Protection Act, PIPL, and APPI. Every protocol amendment, source data transfer, and pharmacovigilance report flowing between sites crosses every one of those regimes.

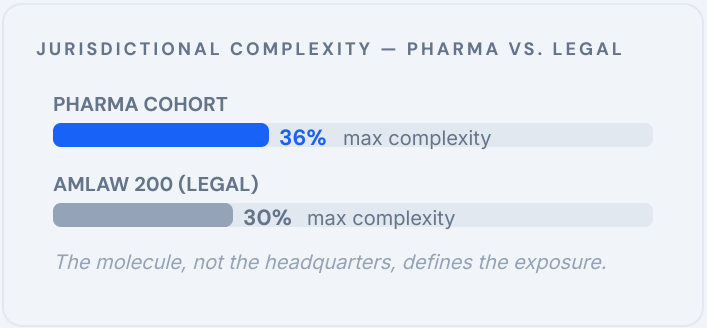

CLOSING THE GAP

Governance That Travels With the Data

Historically, data protection in pharma was a defensive IT function. Today it is a strategic differentiator. Sponsors that build strong chain of custody capabilities gain a decisive advantage in CRO negotiations, regulator interactions, and safe AI adoption — protecting the commercial value of their Data Package regardless of where it goes. Closing the gap requires governance that travels with the data rather than guarding the environment the data leaves behind.

1. **Field-Level Encryption at the Source**
   Sensitive elements — ePHI, patient identifiers, compound structures, target data, trial documentation — are automatically discovered and selectively encrypted at the field level. The rest of the document remains fully structured and usable. Trial monitors, statisticians, and AI workflows continue normally. Sensitive fields are protected without disrupting surrounding context.

2. **Protection That Travels With the File**
   Encryption is bound to a governance policy that travels with the document to any destination — the CRO's eTMF, a regulator portal, a central lab's pipeline, or a generative AI tool. Protection remains attached even after the file leaves the sponsor's environment, across every hop.

3. **Governance Enforced Wherever the File Lands**
   Data-level access is enforced regardless of where the document resides. Different recipients see different portions based on role and permission. The sponsor retains the ability to modify or revoke access at any time — even after sharing with a CRO, sub-CRO, partner lab, or regulator.

4. **A Persistent Audit Trail Across Every Hop**
   Every access event is logged — who opened the file, when, from which environment, and whether any AI tool ingested the content. That audit trail persists regardless of where the file travels. Across a primary CRO, sub-CRO, central lab, and partner AI pipeline — the sponsor can produce an evidence record at every link. In a GxP context, this is data integrity infrastructure, not a security feature. It's what you produce when a regulator asks what happened to that file.

5. **Sensitive Data Never Reaches the Model in Cleartext**
   Before content enters a model — RAG, summarization, or analysis — sensitive fields are encrypted or masked while surrounding context remains intact. ePHI, compound IP, and patient identifiers never reach the model in cleartext. AI Guard governs AI ingestion the same way it governs human access.

On post-quantum cryptography: NIST finalized PQC standards in 2024. Pharma data stays sensitive for 15–20 years. Governance controls that live at the data layer can adapt as cryptographic standards advance — without rebuilding underlying infrastructure. The Data Package stays protected not just for today's threat environment, but for the full length of the exclusivity window it represents.

QUANTIFY YOUR EXPOSURE

Stop Guessing.
Quantify Your Exposure.

This report presents industry-wide findings. It identifies where the sector's exposure is concentrated, why it is structural rather than incidental, and what closing the governance gap requires architecturally. What it cannot do is tell a specific sponsor where its own exposure sits — or how long its third-party chain actually runs. Protecting your Data Package starts with understanding your baseline.

OPTION 01

Request Your Free Custom Sponsor Assessment

We have analyzed the open source data footprints of 199 pharma and life sciences organizations. We can generate a named, sponsor-specific Data Protection Risk Assessment — a personalized 8-dimension scorecard benchmarking your therapeutic mix, CRO footprint, AI deployment, and jurisdictional spread against your peers.

TIME COMMITMENT: 15-MINUTE EXECUTIVE BRIEFING

OPTION 02

The Free Data Protection Risk Scan

For sponsors ready to move beyond public benchmarks, Confidencial offers a read-only scan of your environment to discover unprotected ePHI, IP, and regulated trial data in unstructured files. You receive a detailed Risk Report quantifying your financial exposure, highlighting over-permissioned access, and identifying where sensitive data is vulnerable to AI ingestion or ungoverned external sharing.

TIME COMMITMENT: MINIMAL IT LIFT · RESULTS WITHIN DAYS

Ready to see your organization's exposure score?
Initiate either assessment at the link below. Takes less than 15 minutes to get started.

confidencial.io/contact →

APPENDIX A

Financial Exposure Worksheet

The following uses verified benchmarks to frame order-of-magnitude exposure for a sponsor-specific assessment. Every figure is sourced. This is an illustrative model intended to frame exposure, not predict a specific outcome.

Stage 1 — Industry Baseline (Verified)

VARIABLE	VALUE	SOURCE
Avg. cost of a healthcare data breach (2025)	**$7.42M**	IBM CODB 2025
Healthcare consecutive years as most expensive sector	**14 years**	IBM CODB 2025
Average breach lifecycle (healthcare)	**279 days**	IBM CODB 2025
Average US breach cost (all sectors, 2025)	**$10.22M**	IBM CODB 2025
Shadow AI premium per breach	**+$670K**	IBM CODB 2025
Average Phase III trial cost (completed 2024)	**$36.58M**	Tufts CSDD via MedPath, 2024

Stage 2 — Pipeline Scale Multiplier

VARIABLE	INPUT LOGIC
Active clinical trials (all phases)	___ trials. Each trial = data flow to one or more CROs, labs, and partners.
Phase III trials specifically	___ Phase III. $36.58M avg. per Phase III. Breach disruption calculated against trial cost.
Stage 2 Base Exposure	$7.42M × (trials/10). Trial count proxies for data-supply-chain scale.

Stage 3 — Jurisdictional Notification Cost

VARIABLE	INPUT LOGIC
Primary regulatory jurisdictions	___ jurisdictions. Count distinct regimes: HIPAA, GDPR, UK DPA, PIPL, APPI, LGPD.
Multi-country breach notification cost	+~10% of base per regime. Multi-country complexity is a top breach-cost amplifier. (IBM CODB 2024)

Stage 4 — Statutory Penalty Ceilings (Verified)

PENALTY	AMOUNT	SOURCE
HIPAA Tier 1 (lack of knowledge), per violation	**$145–$73,011**	HHS Federal Register, Jan 28, 2026
HIPAA Tier 4 (willful neglect, uncorrected), per violation	**$73,011–$2,190,294**	HHS Federal Register, Jan 28, 2026
HIPAA annual cap per identical provision (Tier 4)	**$2,190,294**	HHS Federal Register, Jan 28, 2026
EU AI Act maximum penalty	**€35M or 7% global turnover**	EU Regulation 2024/1689
GDPR maximum penalty	**€20M or 4% global turnover**	GDPR Article 83

Precedent Reference: Merck's 2017 NotPetya incident — characterized by the White House as the most destructive and costly cyberattack in history — resulted in initial SEC filing estimates of $870M in damage and a $1.4B insurance settlement after a multi-year dispute. Manufacturing was disrupted globally for weeks. The regulatory penalty environment has changed materially since — GDPR, the 2026 HIPAA overhaul, and the EU AI Act all post-date this incident. A comparable event today would carry substantially higher statutory exposure. *(Merck SEC filings 2017–2018; Security Magazine, Jan 2022; White House statement on NotPetya, Feb 2018)*

APPENDIX B

Methodology & Sources

This research evaluates 199 pharmaceutical, biotech, and clinical-research organizations across the U.S., Canada, EU, and UK. The cohort was derived from a base list of 243 organizations associated with the American Society of Clinical Oncology (ASCO) ecosystem and filtered to 199 by removing patient advocacy nonprofits, professional associations, IRBs, distributors, agencies, hospice organizations, and entities where open source data was insufficient for scoring.

Each organization was scored 1–5 across eight dimensions. The four directly-derived dimensions — Therapeutic Sensitivity, Jurisdictional Complexity, Clinical Data Retention, and AI Adoption — came from a structured enrichment process drawing on company websites, regulatory filings, and clinical-trial registries. The four supplemental dimensions — Regulated Data Exposure, Clinical Trial Footprint, Technology Adoption Breadth, and Industry Risk Baseline — were derived from the same source data using a documented scoring rubric available on request. Risk-band thresholds were calibrated against the equivalent framework used in Confidencial's 2026 Legal Sector report to enable cross-sector comparison. No organization in the pharma cohort scored in the Moderate band (below 18).

Sources Cited

SOURCE
IBM. *Cost of a Data Breach Report 2025.* ibm.com/reports/data-breach
McKinsey Global Institute. "Generative AI in the pharmaceutical industry: Moving from hype to reality." January 9, 2024.
Insilico Medicine. Developmental candidate benchmarks and *Nature Biotechnology* publication, 2024–2025.
U.S. FDA. Draft Guidance: "Considerations for the Use of Artificial Intelligence To Support Regulatory Decision-Making for Drug and Biological Products." January 6, 2025.
HHS Office for Civil Rights. HIPAA Security Rule Final Rule effective 2026. Civil Monetary Penalty adjustment, Federal Register, January 28, 2026.
EU Regulation 2024/1689 (AI Act). Phased applicability through August 2, 2027 for high-risk AI in regulated products.
NIST. Post-Quantum Cryptography Standards (FIPS 203, 204, 205). Finalized August 2024.
KPMG. "Trust in artificial intelligence: A global study." 2025. (48,000+ respondents, 47 countries.)
Cisco. 2025 Data Privacy Benchmark Study.
Gartner. 2025 cybersecurity leaders survey on shadow AI prevalence.
Tufts Center for the Study of Drug Development, via MedPath. Phase III trial cost analysis, 2024.
CybelAngel. "Reviewing Pharmaceutical Threats in 2025." 2025.
Help Net Security. "Attackers are coming for drug formulas and patient data." September 2025.
Merck & Co. SEC filings 2017–2018. Security Magazine. "Merck wins $1.4B lawsuit over NotPetya attack." January 2022. White House. Statement on NotPetya attribution. February 2018.

Published by Confidencial
Data-layer encryption for pharma, legal, and enterprise.
confidencial.io