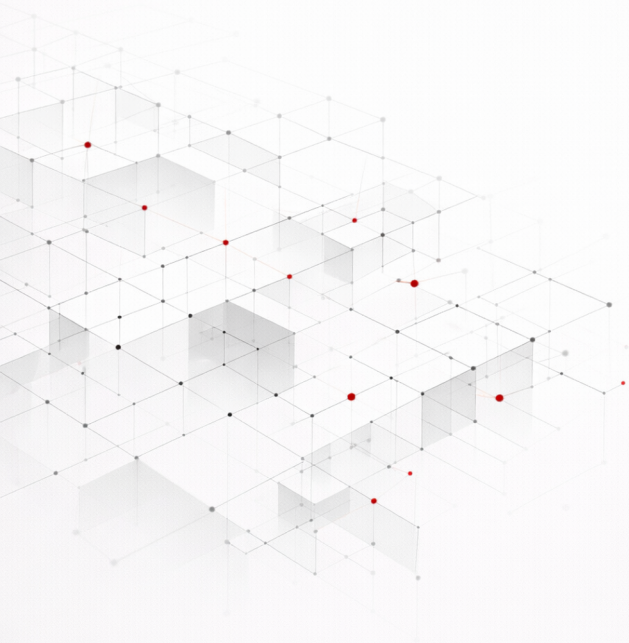

APOLLO Data Auditor

WHITE PAPER · RISK EXPOSURE

# How Much Is Your Data Really Costing You?

Measuring Financial Exposure Before the Breach

# How Much Is Your Data Really Costing You?

## MEASURING FINANCIAL EXPOSURE BEFORE THE BREACH

**Risk Exposure Module — April 2026**

## EXECUTIVE SUMMARY

443 data breach notifications per day. That is the current pace in Europe — a 22% increase in a single year.

Most of those organizations discovered the scope of the problem after the incident, not before. They had no map of their sensitive data. No financial estimate of their regulatory exposure. No way to answer the only question that matters to a board, an insurer, or a regulator: how much?

This paper examines why that question remains unanswered in most SMBs, why the tools that could answer it are out of reach, and what a different approach looks like.

## 1. YOU DON'T KNOW WHAT YOU STORE

A mid-size law firm in the United Kingdom was breached through a legacy service account. The password was unknown to the IT team. No multi-factor authentication. The attacker moved laterally across the network and exfiltrated 32.4 GB of judicial case files — court documents, photographs, personal data of 791 individuals. All published on the dark web.

The firm did not identify the incident as a reportable breach for 43 days. The GDPR obligation is 72 hours.

The supervisory authority fined them £60,000.

The compromised account had excessive privileges and had never been audited. No one knew it existed. This is what shadow access looks like in practice: credentials and data that sit outside everyone's visibility, until an attacker finds them first.

This pattern repeats across sectors and geographies.

In New York, an accounting firm experienced two successive incidents exposing Social Security numbers, financial accounts, and medical benefits data of 4,700 individuals. The firm took over a year to notify victims. Settlement with the Attorney General: $60,000, plus mandatory security overhaul.

In Estonia, a pharmacy loyalty program had accumulated six years of purchase history — pregnancy tests, blood pressure monitors, intimate hygiene products — for 750,000 individuals. No MFA, no access logs, no defined roles. The data had never been audited. Fine: €3,000,000, the largest ever imposed in Estonia.

None of these were Fortune 500 companies. They were organizations with 50 to 500 employees that had no inventory of their sensitive data.

---

**The data confirms this is not anecdotal:**

- 35% of all data breaches involve shadow data — data organizations don't know they have. These breaches cost 16% more and take 26% longer to detect (IBM 2024).
- 59% of security leaders admit they cannot maintain a detailed data inventory (Forrester/Cyera 2024).
- 48% of ransomware victims in France are SMBs (ANSSI 2025). 88% of breaches affecting SMBs involve ransomware (Verizon DBIR 2025).
- The average time to detect a breach is 241 days. Supply chain breaches: approximately 9 months (IBM 2025).
- GDPR fines have reached €7.1 billion cumulative. In Spain, 70% of 2024 enforcement actions targeted SMBs (AEPD).

---

## APOLLO IN ACTION — THREE REAL CASES SCORED

Each incident above followed the same pattern: no one knew what data they had, where it was, or how exposed they were — until an attacker or regulator found it first. Below is what an APOLLO Data Auditor™ scan would have surfaced for each organization, run before the incident.

---

## Case 1 — UK Law Firm, 791 individuals, £60,000 fine

Scan profile: 8,400 files — judicial case files, court photographs, personal correspondence. Legacy service account with full network access, unknown to IT. No MFA. No processing register.

| Score | Value | Status |
|---|---|---|
| **Risk Exposure — S013 Global** | **19 / 100** | CRITICAL |
| Compliance GDPR | 11 / 100 | Grade F |
| Data Protection | 14 / 100 | CRITICAL |
| **Estimated financial exposure** | **€ 340,000** | — |

| GDPR Article | Score | Finding |
|---|---|---|
| Art. 9 — Sensitive data | 2 / 100 | Judicial files classified special category — 100% unprotected |
| Art. 32 — Security | 0 / 100 | No MFA. Shadow service account with full read access. No encryption. |
| Art. 30 — Documentation | 10 / 100 | No processing register. Shadow account not referenced. |
| Art. 5 — Retention | 40 / 100 | Files dating 8+ years. No retention policy in force. |

| Priority | Action | Estimated impact |
|---|---|---|
| P1 | Audit and remove all service accounts not listed in LDAP | – €190,000 exposure |
| P1 | Enable MFA on all access points (files + network shares) | Art. 32: from 0 to partial compliance |
| P1 | Create processing register for judicial case file types | Required — Art. 30 currently missing |

The ICO fine was £60,000. APOLLO estimated €340,000 potential exposure. The gap: the ICO applied a significant cooperation discount. Had the processing register existed, the Article 30 violation — and a portion of the fine — would not have been applicable.

## Case 2 — Estonian Pharmacy Loyalty Program, 750,000 individuals, €3,000,000 fine

Scan profile: loyalty database, 2.1M records, 6 years of purchase history. Product categories include pregnancy tests, blood pressure devices, intimate hygiene — classified as health-adjacent under GDPR Article 9. No MFA. No access logs. No defined roles. No retention policy.

| Score | Value | Status |
|---|---|---|
| **Risk Exposure — S013 Global** | **8 / 100** | CRITICAL |
| Compliance GDPR | 0 / 100 | Grade F |
| Data Protection | 7 / 100 | CRITICAL |
| **Estimated financial exposure** | **€ 4,100,000** | — |

| GDPR Article | Score | Finding |
|---|---|---|
| Art. 9 — Sensitive data | 0 / 100 | 100% of loyalty records contain health-adjacent product data |
| Art. 32 — Security | 0 / 100 | No MFA. No access logs. No role-based access control. |
| Art. 30 — Documentation | 0 / 100 | No processing register. No controller documentation. |
| Art. 5 — Retention | 0 / 100 | 6 years retention with no identified legal basis beyond 12 months. |

| Priority | Action | Estimated impact |
|---|---|---|
| P1 | Purge all records older than 12 months — no retention basis identified | – €2,400,000 exposure |
| P1 | Define role-based access — currently full read access for all staff | Art. 32: from 0 to partial |
| P1 | Document processing basis for loyalty program under Art. 6 | Required — no register exists |

The fine was €3,000,000 — the largest ever imposed in Estonia. APOLLO's model: €4.1M. The difference: the DPA applied the 4% revenue cap below the calculated ceiling. The data had never been audited. Had APOLLO run one year prior, the purge alone would have reduced fine exposure by an estimated 58%.

## Case 3 — New York Accounting Firm, 4,700 individuals, $60,000 settlement

Scan profile: mixed file and database estate. PII types detected: Social Security Numbers, financial account numbers, medical benefits records. Two successive incidents within 24 months. No structured access control. Notification delayed over 12 months.

| Score | Value | Status |
|---|---|---|
| **Risk Exposure — S013 Global** | **24 / 100** | CRITICAL |
| Compliance CCPA | Grade F | CRITICAL |
| Data Protection | 21 / 100 | CRITICAL |
| **Estimated financial exposure** | **$ 420,000** | — |

| Toxic Combination | Classification | Multiplier applied |
|---|---|---|
| Social Security Number + Medical benefits | Tier 1 — Maximum sensitivity | × 3.0 |
| Financial account number + SSN | Tier 1 — Maximum sensitivity | × 2.5 |

| Priority | Action | Estimated impact |
|---|---|---|
| P1 | Segment SSN files from financial files — co-located in 3 directories | Eliminates Tier 1 toxic combination |
| P1 | Delete medical benefits files past 3-year retention period | – $180,000 CCPA exposure |
| P2 | Implement access logging on all directories containing SSN | Enables breach detection (current: 241-day avg detection time) |

The AG settlement was $60,000. APOLLO's exposure model: $420,000. The gap reflects a negotiated resolution — and excludes the cost of two separate incidents, mandatory notifications to 4,700 individuals, and the required security overhaul. All three could have been avoided by a $2,999/year scan that would have flagged the toxic combination on day one.

## 2. WHY EXISTING SOLUTIONS DON'T ANSWER THE QUESTION

If you are a CFO preparing for a board meeting and you ask "how much are we exposed to under GDPR?", today no one in your organization can answer with a number. Your DPO has a processing register built from declarations. Your IT team knows which databases exist but not what they contain. Your CISO has security controls in place but no financial impact model.

The market offers several categories of tools. None of them answer the question.

| What you need | Enterprise data security platforms | Compliance automation | Privacy governance tools | SMB data discovery | Consulting firms |
|---|---|---|---|---|---|
| **Scans actual data** | ✅ Yes | ❌ Declarations | ❌ Questionnaires | ⚠️ Windows files only | ❌ Interviews |
| **Calculates exposure in € / $** | ❌ Risk scores, not amounts | ⚠️ Generic models | ❌ High/medium/low | ❌ No | ⚠️ One-time estimate |
| **Simulates breach impact** | ❌ Posture, not impact | ❌ Rare | ❌ No | ❌ No | ⚠️ Manual, billed separately |
| **Typical price** | $50K–$500K/yr | $10K–$200K/yr | €5K–€50K/yr | €1K–€15K/yr | €5K–€50K per engagement |
| **Time to deploy** | Weeks to months | Weeks | Days | Minutes | 4–8 weeks |

Enterprise platforms scan data but don't translate findings into regulatory fine amounts. Compliance tools model risk abstractly but never look at actual data. Privacy tools manage processes declaratively. Consulting firms provide a one-time snapshot based on interviews — not repeatable, not scalable, not affordable for most SMBs.

To cover all the dimensions of the problem — scanning, quantification, and simulation — an SMB would need to combine three or four of these categories. Multiple vendors, multiple contracts, multiple interfaces, at a combined cost that exceeds €100K per year. That is not realistic for a 50–to-500–employee organization.

**The gap:** no existing category under €50K/year combines real data scanning + automated financial exposure calculation + breach impact simulation in a single tool.

## 3. WHAT WOULD A BREACH COST YOU TOMORROW?

Imagine your cyber insurer asks you to quantify your PII exposure. You have 47,000 files across local servers, two databases, a SharePoint instance, and an Active Directory with 230 accounts. Today, you cannot answer.

What you need is not another compliance checklist. You need five things:

**A PII map.** A complete inventory of which files, tables, and cloud documents contain which types of personal data. Without it, everything else is guesswork.

**Financial exposure in € and $.** Computed from the PII actually detected, the source types involved, and the regulatory penalty formulas — GDPR Article 83 (up to 4% of revenue or €20M) and CCPA ($2,663 per standard violation, $7,988 per intentional violation). Not an abstract score. A number your CFO can present to the board and your insurer can use to price a policy.

**Breach impact simulation.** What would happen if this data were breached tomorrow? What fine, what notification cost, what remediation cost, what business interruption? Modeled from the PII types and volumes actually found in your infrastructure — not a theoretical scenario.

**Toxic combination detection.** An IBAN alone is sensitive. A Social Security number alone is sensitive. Both in the same file is a regulatory catastrophe. These co-locations multiply exposure and must be identified before an incident reveals them.

**Risk zone identification.** Which directories, databases, or cloud containers have the highest concentration of unprotected PII? This tells you where to focus first — and where the damage would be greatest.

---

## 4. HOW APOLLO DATA AUDITOR ANSWERS THE QUESTION

The five capabilities described above exist today, in a single tool, at a price point built for SMBs. Here is how it works.

**Your data stays yours.** A native agent runs locally on your infrastructure — Windows, Linux, and macOS (arm64). It scans 11 source types: PostgreSQL, MySQL, MongoDB, SQL Server, OneDrive, SharePoint, Active Directory/LDAP, ERP (Pennylane), local files, NFS/SMB shares, and infrastructure. Only metadata and counters transit to the cloud dashboard. Raw data never leaves your perimeter. The agent is open source (BSL 1.1) — every line of code is verifiable on GitHub. This is not a claim. It is auditable.

**You get a financial answer.** The dashboard shows your GDPR and CCPA exposure in euros and dollars — article by article, source by source. It maps every PII type detected, flags toxic combinations, identifies your highest-risk zones, and simulates the financial impact of a breach scenario based on your actual data profile.

**You get a remediation plan.** Prioritized actions (P1/P2/P3) with estimated financial impact for each. Not a generic best-practices list — actions tied to the specific risks found in your scan.

**What this looks like in practice:** a cyber insurance underwriter asks a mid-market manufacturer for a data risk assessment. The manufacturer has 47,000 files, two PostgreSQL databases, a SharePoint instance, and 230 Active Directory accounts. The scan runs in under 48 hours. Result: €2.3M of potential GDPR exposure concentrated in three directories, 12 toxic PII combinations, and 1,400 files containing health-related data that no one had classified. The manufacturer now has a number for the insurer, a remediation priority list, and a baseline to measure improvement.

**44 PII types detected** across EU and US regulations — IBAN, SSN, NIR, PESEL, BSN, NIE, NISS, codice fiscale, passport, email, phone, health data, and 32 more. **129 auditable scores with published formulas.** Every score is transparent, reproducible, and verifiable. No black box. Up to **1.16M rows per second.**

---

### THE PRICE COMPARISON

| | Enterprise platforms | Consulting firms | APOLLO Data Auditor |
|---|---|---|---|
| **Annual cost** | $50,000 – $500,000 | €5,000 – €15,000 per audit | **€2,999 / year** (Starter) |
| **Deployment** | Weeks to months | 4–8 weeks | **Under 48 hours** |
| **Repeatable** | Yes | No (one-time) | **Yes — every scan** |
| **Calculates exposure in €/$** | No | Partial (manual) | **Yes — automated** |
| **Breach impact simulation** | No | Manual (billed separately) | **Yes — included** |

---

## 5. ONE SCAN. ONE ANSWER. FOUR DIMENSIONS.

Most SMBs would need to combine three or four vendors to cover what a single data audit should deliver: knowing what data you have, what it costs you, whether it is protected, and whether it is ready for what comes next. That is not a viable strategy for a 200-employee company with no dedicated security team.

APOLLO Data Auditor was built to solve exactly this. One scan, one tool, four dimensions of your data risk — at a price that reflects the reality of SMB budgets, not enterprise procurement cycles.

→ **Risk Exposure** — what this paper covers. Financial quantification in € and $, breach impact simulation, PII mapping, toxic combinations, risk zones.

→ **Compliance** — GDPR scored by article (Art. 5, 9, 30, 32), CCPA, NIS2, SOC2, DORA. Grades A through F based on your actual data — not declarations. Prioritized remediation plan (P1/P2/P3).

→ **Data Protection** — encryption coverage, backup resilience, ransomware readiness simulation, access control hygiene, dormant account detection. The technical posture that determines whether a breach stays contained or becomes catastrophic.

→ **Intelligence** — AI Readiness scoring, data quality metrics, AI Act pre-compliance (Article 10 data governance, Article 15 cybersecurity posture). Whether your data is ready for AI — or whether it will block deployment.

No other tool under €5,000/year covers all four. That is not a marketing claim — it is the result of a market where enterprise DSPM platforms start at $50K, GRC tools don't scan data, and consulting firms deliver one-time reports that are outdated before the invoice is paid.

**APOLLO™ Data Auditor** exists because SMBs deserve the same visibility into their data risk that a Fortune 500 company gets — without the Fortune 500 budget.

---

Sources: IBM Cost of a Data Breach Report 2024 & 2025 · Verizon 2025 DBIR SMB Snapshot · DLA Piper GDPR Fines & Data Breach Survey 2026 · ANSSI Panorama de la cybermenace 2025 · CMS GDPR Enforcement Tracker 2024/2025 · CPPA CCPA Fines 2025 · ICO UK Enforcement Actions · Linklaters AEPD FY24 · Forrester/Cyera Data Security Study 2024 · FBI IC3 Annual Report 2024

---

**APOLLO™ Data Auditor** Every file is a risk. Measure it.

→ https://apollo.aiia-tech.com → GitHub: https://ggabrie2025.github.io/apollo_data_auditor/ → contact@aiia-tech.com