PRIVAINI · CASE STUDY

A top-3 global broker makes privacy the new placement edge.

How one of the world's three largest cyber brokers built privacy risk assessment into application prep — and turned a market gap into an advisory advantage.

INDUSTRY	PRACTICE	ASSESSMENT TIME	OUTCOME
Insurance brokerage	Global cyber risk & insurance	Minutes per client	Placement differentiation

EXECUTIVE SUMMARY

One of the world's three largest cyber brokers watched the privacy claims curve bend upward before most of the market — 78% litigation growth over four years, settlements averaging $11.5 million, underwriters responding with pointed questions about applicant privacy posture — and drew a practical conclusion: the brokers who could assess and communicate client privacy risk would win placements, and those who couldn't would cede ground.

This case study describes how the broker's global cyber practice integrated Privaini into client advisory and application preparation; the statutory landscape behind the underwriting questions its clients were suddenly facing; and the results — an evidence-backed advisory capability most competitors cannot match, deployed in minutes per client with zero client burden.

It is also, candidly, a story about positioning. The broker's cyber leadership frames privacy as 'the new MFA' — the control-point that shifts from differentiator to table stakes within a market cycle. Firms that build the capability early set the standard the rest of the market is measured against.

MARKET CONTEXT

The enforcement economy.

Privacy violations now cost organizations more than **$1.7 billion in regulatory fines every year** — and fines are the smaller half of the story. US data privacy litigation surged **78% between 2020 and 2024**, average settlements under the Biometric Information Privacy Act and the Video Privacy Protection Act reached **$11.5 million**, and wrongful-collection claims now average **$1.9 million**. Behind those numbers sits a plaintiff's bar that has industrialized privacy litigation: automated scanning of corporate websites for tracking technologies, template complaints keyed to statutory damages provisions, and demand letters that arrive weeks after a new marketing tag goes live.

What makes this risk category structurally different from breach risk is **where the evidence lives**. A data breach requires an attacker, a vulnerability, and an incident. A privacy violation requires none of those things — it is committed in public, continuously, by the company's own digital properties. The trackers firing before consent, the session-replay script recording keystrokes, the video player sharing viewing data with an advertising pixel, the privacy policy that describes none of it: all of it is observable by anyone who looks. Regulators look. Plaintiff's attorneys look. Insurers are learning to look.

This is why mature privacy programs still generate enforcement actions. Internal tools — consent management platforms, data mapping software, assessment workflows — manage what the organization *says* and *documents*. Enforcement is triggered by what the organization *does*, as experienced from the outside. The gap between the two is where regulatory actions originate, class actions are filed, and insurance claims are generated. Most organizations have never measured that gap, because none of their tools can see it.

INDUSTRY EXPOSURE

The broker's dilemma: advising on an invisible risk

A cyber broker's value concentrates at two moments: preparing a client to face the underwriting market, and negotiating that market on the client's behalf. Both moments now turn partly on privacy posture. Underwriters ask about consent practices, session-replay tooling, video-pixel pairings, biometric features, and state-law compliance — because those are the fact patterns generating their claims. A client surprised by those questions mid-placement is a client whose broker did not prepare them.

The standard toolkit gave brokers nothing to prepare them with. Security ratings measure breach likelihood, not litigation exposure. Client questionnaires reproduce the client's own documentation — the same documentation whose gaps create the risk. Engaging outside counsel for a privacy audit per placement is economically absurd at brokerage scale. The result was a structural advisory blind spot precisely where underwriter scrutiny was intensifying fastest.

For a top-three broker, the blind spot cut both ways: as risk, in placements complicated by unaddressed privacy exposure; and as opportunity, because a credible privacy assessment capability — deployable across a global book, without burdening clients — would be a genuine differentiator in a market where differentiation is scarce.

Timing sharpened the calculus. Capabilities of this kind confer the most advantage in the window between a risk becoming material and its assessment becoming standard practice. Arrive too early and clients don't yet feel the pain; arrive late and the capability is table stakes commanding no premium. The practice's leadership judged — correctly, as underwriter behavior has since confirmed — that privacy risk was entering that window: claims data made it material, carrier guideline changes were making it priced, and no incumbent advisory toolkit yet addressed it. The firm that moved first would spend the window winning business with a capability competitors would eventually need just to compete.

THE CHALLENGE

Underwriters moved first. Brokers had no tool.

The broker's cyber risk advisory team saw the privacy claims trend accelerating before most of the market. Their concern was practical: brokers had no tool to assess or communicate client privacy posture during the application and renewal process. Privacy risk was becoming a differentiating factor in placement — and the standard broker toolkit had a gap.

Security ratings and questionnaires didn't fill it. Ratings measure breach risk, not privacy litigation exposure. Questionnaires reflect what clients say about their internal programs, not what is observable from the outside. Neither gives an advisor the ability to walk into an application prep conversation with a credible, evidence-backed assessment of the client's privacy posture — or to flag, before an underwriter does, that a client's marketing stack contains the exact fact pattern behind the market's fastest-growing claim category.

The competitive implication was clear, and the practice's leadership stated it internally in exactly those terms: brokers who could offer privacy risk assessment as part of application preparation would differentiate their offering and improve placement outcomes. Brokers who couldn't would find their clients surprised in underwriting meetings — and their renewals increasingly contested by firms that could.

THE REGULATORY LANDSCAPE

The questions underwriters now ask.

Underwriter scrutiny maps directly to the statutes driving claims. A broker preparing a client application today needs credible, evidence-backed answers on four fronts — each corresponding to a regime with statutory damages or regulator enforcement attached:

Video Privacy Protection Act (VPPA)

The VPPA is a 1988 federal statute born from an analog scandal: a newspaper obtained and published Supreme Court nominee Robert Bork's video rental history during his confirmation hearings. Congress responded by prohibiting 'video tape service providers' from knowingly disclosing information that identifies a person as having requested or obtained specific video materials, and set **liquidated damages at $2,500 per person** — with no requirement to prove actual harm.

Thirty-five years later, that analog statute has become the digital plaintiff bar's favorite instrument. Courts have applied the VPPA to streaming platforms and to ordinary websites that embed video content, and the modern fact pattern is devastatingly simple: a site hosts video, an advertising or analytics pixel transmits the viewer's identity together with the video title to a third party, and every affected visitor becomes a class member at $2,500 each. A publisher with a few hundred thousand video viewers is looking at nine-figure theoretical exposure from a single misconfigured tag. The wave of VPPA class actions filed since 2022 has targeted media companies, publishers, sports leagues, and healthcare systems — anyone whose properties pair video with tracking.

The critical operational fact: VPPA exposure is created and cured at the level of observable site behavior. Whether a pixel transmits video-viewing data is not a question about policy language — it is a question about what the page actually does when a user presses play.

Biometric Information Privacy Act (BIPA)

Illinois' BIPA, enacted in 2008, requires informed written consent before collecting biometric identifiers — fingerprints, voiceprints, face geometry — along with published retention schedules and destruction timelines. Its private right of action grants **$1,000 per negligent violation and $5,000 per reckless or intentional violation**, and the Illinois Supreme Court held in *Rosenbach v. Six Flags* (2019) that a plaintiff need not show any actual injury: the statutory violation itself is the harm.

The arithmetic is what makes BIPA the largest source of privacy settlements on record. Violations accrue per person, and courts have held they can accrue per collection event; although 2024 amendments tempered the most extreme per-scan accrual theories, exposure still compounds across

a workforce or user base with brutal speed. Voice interfaces, photo-tagging features, identity-verification selfies, and employee time clocks have all generated eight- and nine-figure BIPA outcomes — and copycat statutes with private rights of action continue to advance in other state legislatures.

Like the VPPA, BIPA liability turns on observable behavior: what a digital experience actually collects, and whether the consent flow the user encounters satisfies the statute — not what internal documentation says should happen.

State wiretapping statutes and CIPA

A second front opened when plaintiff's attorneys began applying all-party-consent wiretapping statutes — written for telephone taps — to modern website tooling. The California Invasion of Privacy Act (CIPA) is the leading vehicle: its provisions against intercepting communications without all parties' consent carry **$5,000 in statutory damages per violation**, and since 2022 they have been deployed against session-replay tools that record keystrokes and mouse movements, third-party chatbots that route conversations through vendors' servers, and tracking pixels characterized as unlawful 'pen registers' under CIPA's device provisions.

Massachusetts, Pennsylvania, Florida, and other all-party-consent states have seen parallel filings under their own statutes. The pattern is consistent: a routine analytics or support tool, deployed by marketing without legal review, silently creates a per-visitor statutory damages theory. Because these tools are visible in the page source and network traffic of the live site, plaintiffs can — and do — identify defendants by scanning the web at scale. The defense posture is equally observable: companies that can demonstrate exactly what their properties transmit, when, and under what consent state are the ones that can dispose of these claims early.

The US state privacy law patchwork

Twenty states have now enacted comprehensive consumer privacy statutes, and no two are identical. California's CCPA as amended by the CPRA created a dedicated enforcement agency — the California Privacy Protection Agency — alongside the Attorney General, with civil penalties of **$2,500 per violation and $7,500 per intentional violation or violation involving minors**, counted per consumer. Other states vest enforcement in their Attorneys General with penalty structures of similar magnitude.

The obligations that matter most are the ones visible from the outside: honoring opt-out preference signals like the Global Privacy Control — the basis of California's landmark Sephora settlement — presenting compliant notice at collection, operating functional 'do not sell or share' mechanisms, and ensuring that stated practices match observed data flows. Obligations attach based on whose data a company processes, not where the company is headquartered, which means a single national website is simultaneously subject to twenty overlapping regimes. A configuration compliant in one state can be a violation in the next — and regulators have shown they will test sites directly rather than take a company's paperwork at its word.

THE CLAIM MACHINE

Anatomy of a modern privacy claim.

How a modern privacy claim is born

Understanding the exposure requires understanding the machinery. A modern privacy class action rarely begins with an aggrieved user. It begins with software: plaintiff-side firms and their vendors continuously scan the public web for actionable fact patterns — video players paired with advertising pixels, session-replay scripts, chatbots routing through third-party servers, consent banners that load trackers before anyone clicks. A hit becomes a target list entry; a target list entry becomes a demand letter citing statutory damages; and because the underlying conduct is observable and continuous, every day of delay grows the putative class.

The economics then do the work. At $2,500 per person under the VPPA or $5,000 per violation under CIPA, a mid-sized website's monthly traffic generates theoretical exposure large enough that even a strong defense carries unacceptable variance — which is why so many of these matters settle, and why average settlements have climbed to eight figures. The defendant's practical leverage comes almost entirely from evidence: an organization that can demonstrate precisely what its properties transmitted, when, and under what consent state can distinguish defensible facts from indefensible ones early — and fix the indefensible ones before the scan finds them.

That is the strategic meaning of outside-in monitoring: it runs the plaintiff's playbook against your own estate, first.

THE SOLUTION

Privacy assessment inside application prep.

The broker deployed Privaini to integrate privacy risk assessment into its client advisory and application preparation workflow. Because Privaini works from a company name and domain — no IT involvement, no questionnaires, no client cooperation required — advisors could generate an outside-in privacy posture assessment for any client or prospect in minutes, at any point in the placement cycle, without adding a single task to the client's plate.

The assessment arrives in the form advisory work needs: findings mapped to the statutes underwriters ask about, severity-ranked so remediation conversations have an agenda, and reproducible so a client's technology and legal teams can verify every claim independently. The same capability runs pre-renewal, surfacing emerging exposure while there is still time to remediate before the market sees it — and it runs on prospects, giving new-business conversations an opening that competitors cannot easily replicate.

Key capabilities used

→ Client privacy risk assessment during application prep — without requiring client participation

→ Outside-in posture report: consent flows, tracker behavior, policy-practice gaps, jurisdiction exposure

→ Litigation exposure signals: VPPA, BIPA, wiretapping, wrongful collection

→ Pre-renewal risk assessment to identify emerging exposure before renewal conversations

→ Competitive differentiation: demonstrating advisory sophistication to clients

How outside-in monitoring works

("

Observation, not attestation

Privaini approaches privacy risk from the same vantage point as the parties who create it: outside the firewall. Starting from nothing more than a company name and a domain, the platform observes the organization's public digital footprint — websites, subdomains, consent experiences, embedded technologies, policies, and disclosures — exactly as a regulator, plaintiff's attorney, or underwriter would encounter them. No questionnaires. No internal data feeds. No IT integration. Nothing to install and no one's calendar to book.

That independence is methodological, not incidental. Self-reported assessments inherit every blind spot of the organization's own documentation; an outside-in assessment inherits none of them. If a tag manager deployed a new pixel last Tuesday, the outside view shows it — whether or not any internal record does.

", "

One hundred signals, mapped to law

The platform evaluates **100+ proprietary signals** spanning tracking technologies, consent flow behavior, cookie practices, policy language, AI data-collection indicators, dark-pattern markers, and disclosure consistency — and then does the step that turns detection into intelligence: mapping each observation to the specific regimes it implicates across **144 countries and 20 US states**. A raw finding like '47 trackers detected' is operationally useless. Privaini's jurisdiction mapping converts it into '12 violate LGPD in Brazil; 8 create wiretapping exposure in California; 3 fire before consent for UK visitors in breach of PECR.' Risk becomes addressable the moment it becomes specific.

", "

Policy-practice gap analysis

The single most predictive indicator of enforcement and claim risk is the distance between what a company's privacy policy says and what its properties observably do. Privaini reads the policy the way opposing counsel would — as a set of testable claims — and verifies each against observed behavior: data collected but not disclosed, sharing relationships not described, retention promises contradicted by practice, consent representations the user experience does not honor. Every gap is a ready-made allegation; closing them in order of statutory exposure is the highest-leverage remediation work a privacy team can do.

", "

Regulator-grade evidence

Every finding is reproducible, timestamped, and jurisdiction-specific — evidence engineered to the standard a regulator applies, not an AI's unverifiable inference. This is the same observation standard Privaini applied in a live UK ICO regulatory response, where findings had to withstand the scrutiny of the regulator that prompted them. The distinction matters everywhere evidence is consumed: a privacy team prioritizing remediation, a board consuming assurance reporting, an underwriter pricing a submission, or counsel responding to a demand letter can all act on findings that can be independently re-verified.

", '

Continuous, because exposure is

A point-in-time audit describes a website that no longer exists: marketing ships new tags weekly, vendors update SDKs silently, and regulations amend on their own calendars. Privaini re-observes continuously, so posture is a living record rather than an annual snapshot — new exposure surfaces when it appears, remediation is verified when it lands, and the organization holds an auditable timeline of its observable behavior. In enforcement, in litigation, and at renewal, the timeline itself becomes an asset: evidence of a program that monitors and corrects.

')

Implementation: the first thirty days

Day zero to steady state

Deployment inverts the usual enterprise-software timeline because there is nothing to deploy. Day zero is a domain name: Privaini begins observing the public estate immediately, and first findings arrive in about thirty minutes. Within the first week, the working team — privacy, legal, and whoever owns the digital properties — reviews a jurisdiction-mapped, severity-ranked findings briefing: what fires before consent and where that is unlawful, which disclosures fail against observed behavior, which properties carry litigation-pattern signals.

Weeks two through four convert findings into a remediation queue sequenced by statutory exposure: per-person statutory-damages patterns first, regulator-penalty categories second, experience and friction issues third. Because every finding is reproducible, each fix is verified by re-observation rather than by assertion — and the before-and-after record becomes part of the organization's defensibility file.

Steady state is a cadence, not a project: continuous re-observation across the estate, alerting when new exposure appears, and a standing evidence trail that answers — for any date — the question every regulator, plaintiff, and underwriter eventually asks: what did your properties actually do?

RESULTS

An advisory capability most competitors can't match.

The practical change was immediate: advisors stopped entering application prep conversations dependent on client self-reporting. With an independent outside-in view of each client's privacy posture, the team could identify material gaps ahead of underwriting review — the pre-consent trackers, undisclosed session tooling, and video-pixel pairings that would otherwise surface as underwriter objections or, worse, as claims — and convert them into proactive remediation conversations. Clients entered the market with cleaner submissions; placements moved with fewer surprises; outcomes improved.

The relationship effects compounded the transactional ones. A broker who arrives with independent evidence about the client's own digital estate — findings the client's teams can verify, mapped to the statutes their general counsel worries about — is performing visibly senior advisory work. At renewal, the pre-renewal assessment turned what is often a defensive conversation into a demonstration of vigilance: exposure identified, remediation sequenced, posture documented before any underwriter asked.

Strategically, the practice positioned itself ahead of the privacy curve it had forecast — consistent with its own thesis that privacy is the next MFA-level dimension in cyber. The capability now functions as both shield and sword: protecting existing relationships by eliminating placement surprises, and winning new business from incumbents who still walk into underwriting meetings blind on the market's fastest-growing risk.

The working playbook that emerged is repeatable across any book of business. New client or prospect: run the assessment before the first strategy meeting, so the relationship begins with the broker knowing something material the client doesn't. Application prep: review findings with the client sixty to ninety days ahead of market submission, sequence quick remediations, and prepare evidence-backed answers for the underwriting questions the findings predict. Renewal: re-assess in advance, brief the client on posture changes, and walk into the market with a documented improvement story rather than a defense. Each step costs the client nothing, costs the advisor minutes — and compounds into exactly the kind of advisory intimacy that renewals are retained with.

By the numbers.

Ahead of the curve on the fastest-growing dimension of cyber risk.

Min	**0**	**78%**	**Top 3**
Per-client assessment, from a domain	Client questionnaires or IT lift	The litigation surge the practice is ahead of	Global cyber broker

“

Privacy is the new MFA. We are watching privacy claims heat up.
Use Privaini to compete for new business.

CYBER RISK LEADERSHIP · TOP-3 GLOBAL BROKER

WHY IT MATTERS

The broker's 'privacy is the new MFA' framing is worth taking literally. Multifactor authentication traveled a predictable arc: from sophisticated differentiator, to underwriter expectation, to a binary condition of insurability — in roughly three market cycles. Privacy posture is on the same arc. Underwriters are already asking; carriers are already updating guidelines; the only open question is which advisors will be fluent when fluency becomes mandatory.

For the brokerage market broadly, the lesson generalizes: advisory differentiation now lives in risk dimensions that are real, measurable, and not yet commoditized. Privacy risk qualifies on all three counts — and the firms that build evidence-backed capability early don't just win the placements in front of them. They define the standard of advice the rest of the market must eventually meet.

THE ECONOMICS

The build-vs-buy math

The traditional route to an outside view is manual: outside counsel and consultants reviewing properties jurisdiction by jurisdiction. Market rates for a single-market regulatory assessment routinely run into six figures — organizations have cited internal estimates of **$200,000+ per market** — and the deliverable is a snapshot that begins aging immediately, across an estate that changes weekly. Scaling that model to twenty US states and every international market a company touches is not a budget line; it is a fantasy.

Automated outside-in monitoring changes the unit economics by an order of magnitude: every property, every jurisdiction, continuously, for less than the cost of assessing a single market manually — with findings that are reproducible rather than opinion-based. The comparison to the cost of being wrong is starker still: a single VPPA class action's defense costs alone typically exceed years of monitoring, before any settlement. The honest framing is not whether an organization can afford continuous outside-in verification; it is that the alternative was never actually available at any price.

THE ROAD AHEAD

The direction of travel: enforcement is compounding

Every structural indicator points the same way. California's dedicated privacy regulator has moved from rulemaking to active enforcement sweeps, and state attorneys general from Texas to Connecticut have opened privacy dockets of their own. The FTC has revived decades-old authorities to police tracking and children's data, with per-violation penalties that turn systematic conduct into existential math. European regulators continue to publish enforcement decisions grounded in direct

observation of live properties, and the first wave of AI-specific enforcement — data collection for model training, undisclosed AI features processing user content — has already produced multimillion-euro outcomes.

Each new regime adds statutes; none retires the old ones. The organizations that navigate this decade well will not be those with the thickest policy binders — they will be those that can answer, quickly and with evidence, the only question every enforcement theory shares: *what do your properties actually do?* That question does not get easier with time. The estate grows, the tooling multiplies, and the scanners on the other side never stop running.

EVALUATION CRITERIA

What to demand from privacy risk intelligence

For organizations evaluating this category — whether as an enterprise verifying its own program, an insurer pricing a book, or an advisor preparing clients — five criteria separate intelligence from noise. **Independence:** findings must derive from outside observation, not self-reported inputs, or they inherit the blind spots they exist to eliminate. **Evidence standard:** every finding should be reproducible, timestamped, and specific enough to hand to counsel — an unverifiable score is an opinion, not intelligence. **Jurisdictional mapping:** raw detections only become decisions when tied to the statutes they implicate, in every market that matters. **Continuity:** a snapshot describes a website that no longer exists; only a cadence describes a company. **Zero operational burden:** any methodology requiring questionnaires, integrations, or the target's cooperation cannot scale — and quietly reintroduces self-reporting through the side door.

These criteria are not abstractions; they are the difference between intelligence that survives contact with a regulator, a courtroom, or an underwriting committee — and intelligence that doesn't. They are also, not coincidentally, a description of how Privaini was engineered.

FAQ

Frequently asked questions

How is this different from a security rating? Security ratings model breach likelihood — patching, exposed services, credential hygiene. Privaini models privacy liability: what your properties observably collect, share, and disclose, mapped to the statutes that turn those behaviors into claims. The two signals are orthogonal; strong security scores routinely coexist with severe privacy exposure.

Do you need access to our systems? No. The assessment is built entirely from publicly observable signals — the same evidence available to a regulator or plaintiff's attorney. No questionnaires, no agents, no integrations, no internal data.

How current are the findings? Continuous. Properties are re-observed on an ongoing cadence, so new trackers, tags, and consent changes surface as they ship — not at the next annual audit.

Are the findings legal advice? No — they are evidence. Privaini produces reproducible, timestamped, jurisdiction-mapped observations of what your properties do. Your counsel decides what the law requires; Privaini shows what the facts are.

Will the findings hold up with regulators, courts, and underwriters? They are engineered to. Every finding can be independently re-verified, carries its observation timestamp, and cites the jurisdictional rule it implicates — the same evidentiary standard applied in a live UK ICO regulatory response.

How fast is the first assessment? From a company name and domain to first findings in about thirty minutes. A full jurisdiction-mapped briefing follows within days, not quarters.

Can findings be shared with third parties — regulators, insurers, acquirers, clients? Yes, and that is much of their value. Because findings are observations rather than privileged analysis, organizations routinely share them with underwriters at submission, with counsel during a regulatory response, with boards for assurance reporting, and with acquirers during diligence. The reproducibility is what makes them portable: any recipient can verify any finding independently.

What happens after remediation? Verification, then vigilance. Fixed findings are confirmed by re-observation and recorded, producing a defensibility file that documents both the exposure and its cure. Monitoring then continues on cadence — because the next marketing tag, vendor update, or acquired microsite is always coming, and posture is only as good as its most recent observation.

A PRACTICAL STARTING POINT

What would a regulator see on your properties today?

Twelve questions any organization can ask about its observable posture this quarter — each answerable from the outside, without a single internal meeting:

1. Which trackers fire on our properties *before* a user interacts with the consent banner — and in which jurisdictions is that sequence unlawful?
2. Does any embedded video player transmit viewer identity alongside video titles to an advertising or analytics endpoint?
3. Are session-replay, chatbot, or heat-mapping tools capturing user interactions — and do our consent flows disclose them by name and purpose?
4. Does our privacy policy disclose every third party that observably receives data from our pages?
5. Do our properties honor the Global Privacy Control and state opt-out signals, verifiably?
6. What does our consent experience look like to a user in London? In São Paulo? In Sacramento?
7. Which of our subdomains and microsites carry tags our privacy team has never reviewed?
8. Are AI or model-training data flows present on our properties that no current disclosure describes?
9. Could we produce timestamped evidence of our consent behavior for any date in the past year?

10. If a demand letter cited CIPA tomorrow, could we show exactly what our chat and replay tools transmit?
11. Does any property serving minors or students carry advertising or analytics tracking?
12. When did we last verify any of the above — and would the answer survive a regulator doing the same check today?

ABOUT PRIVAINI

Privaini is the privacy risk intelligence platform for enterprises, insurers, and risk advisors. From a company name and domain, Privaini builds a regulator-grade, outside-in view of privacy posture — 100+ signals mapped across 144 countries and 20 US states — and monitors it continuously. Two of the top three global cyber brokers and leading cyber insurers use Privaini to see privacy risk before it becomes loss. From your domain name to findings in 30 minutes: **privaini.com**.

Your exposure is already visible. See it before they do.

No questionnaires · No installation · Domain name only · Results in 30 min

privaini.com/request-a-demo