Vault JS

VAULT JS RESEARCH
OCTOBER 2026

RESEARCH REPORT • CIPA · 318 DECISIONS · SB 690

CIPA in 2026: what a claim is really *worth*.

An analysis of 318 court decisions in internet-tracking litigation under the California Invasion of Privacy Act, decided since *Javier v. Assurance IQ* (9th Cir., 2022). **Not one has reached a trial verdict.** Defendants win on a motion or they pay. The number in play is the settlement.

318
Decisions coded and analyzed, May 2022 to June 2026

0
Trial verdicts for either party in the analyzed record

31.1%
Of decisions end in a win for the defense

$3.7M
Median disclosed settlement paid by a brand


Karel Kubicek
Lead Privacy & AI Researcher


Julie Oberweis
COO & Co-Founder

VAULT JS RESEARCH
Released 5 October 2026

IN THIS REPORT

What the record says, what it *costs*, what to do.

Every internet-tracking decision under CIPA handed down between *Javier v. Assurance IQ* (May 2022) and June 2026, coded for outcome, claim type and disposition, and analyzed as data. Conclusions follow the numbers.

00 **Executive summary.** p. 03
Five findings that hold, and one widely believed theory that does not

01 **The landscape.** p. 05
A permanent operating condition · who gets sued · what gets pleaded

02 **Where the risk sits.** p. 09
Vendor category · deciding grounds · consent posture · plaintiffs · sensitivity · sector

03 **Eight measures that reduce exposure.** p. 18
Ordered by the lever each one acts on: selection, exposure, odds

04 **What the moves cost.** p. 20
What this record can price · the one number a company controls

05 **If a letter arrives.** p. 22
Three axes · two demand archetypes · the best defense on each axis

06 **What it costs, and what we cannot see.** p. 25
42 disclosed settlements · the median · the disclosure gap

07 **What changes next: SB 690.** p. 29
The pen-register route, and what removing it does and does not do

08 **The first ninety days.** p. 32
Measure, gate and prove the gate, then procure and operate

A **Appendix A: scope and method.** p. 34
What is analyzed, what is set aside, and the corrections behind the rates

B **Appendix B: what this report does not claim.** p. 39
The five limits on every number in these pages

AT A GLANCE

318
Decisions since *Javier*, against 134 in the five decades before

266
Distinct companies sued, 238 of them exactly once

16.7%
Defense win rate where an opt-out was recorded, then ignored

$275.3M
Total disclosed settlement funds paid by brands

THE FINDINGS

Five things hold. One widely believed theory *does not*.

00

The findings below come from the decided record. We pulled every internet-tracking decision under CIPA handed down between *Javier v. Assurance IQ* (May 2022) and June 2026, coded each for outcome, claim type and disposition, and analyzed the coded set as data. Conclusions follow the numbers, not our reading of any single opinion. Across 318 decisions, five things hold and one widely believed theory does not.

01 **In cases to date, the exposure is priced, not adjudicated.**
Not one of the 318 decisions is a trial verdict, and the handful that reach the merits do so on summary judgment about what the vendor did, as in *Gutierrez v. Converse Inc.* (9th Cir. 2025). The recorded outcomes are 99 defense wins, 89 mixed, 62 settled, 42 rulings for the plaintiff and 26 undecided. The commercial question is what a claim costs and whether it can be ended on a motion, not whether the conduct was lawful (see Sections 1 and 6).

02 **The defense wins about three times in ten, and no company is exempt.**
The defense prevails in 31.1% of the analyzed decisions, and in 31.9% of the 254 in which the brand running the website was the defendant. Across the analyzed decisions, 266 distinct companies appear as defendants, 238 of them exactly once (see Sections 1 and 2.7).

03 **A category's win rate partly reflects who chose to sue over it.**
Live chat appears to defend far better than pixels: 40.4% against 24.5%. But a small group of serial filers concentrates on particular technologies, so a category rate is never only a fact about the technology. What a vendor does matters. Who sued matters more (see Section 2.5).

04 **Sensitive data lowers the defense's odds and raises the price.**
The defense wins 35.6% of decisions where no sensitive data is alleged and 22.4% where health data is. The same marketing pixel runs at 29.6% without sensitive data and 17.5% with it. Health data also brings claims CIPA does not have, so the matter is worth more: healthcare accounts for 29 of the 35 disclosed settlement funds paid by brands. Financial data is the exception and sits near the baseline (see Sections 2.6 and 6).

05 **The one number a company controls is the third-party host count on page load.** A demand letter's first multiplier is the number of distinct third parties a claimant can enumerate from a visit to the site. It can be measured in-house first and reduced before anyone else counts it (see Sections 3 and 4.3).

THE THEORY THAT DOES NOT HOLD

A vendor's business model does not decide the *case*.

Where the plaintiff alleges the third party used the data for itself, the defense wins 32.5%. Where a court accepts the vendor as a service provider, 37.0%. The baseline is 31.9%. The three numbers are statistically indistinguishable (see Section 2.2).

32.5%	37.0%	31.9%	0
Third party alleged to use the data for its own purposes (160 decisions)	Vendor accepted as a service provider acting on instruction (54 decisions)	The baseline across all 254 brand-defendant decisions	Statistically significant difference between the three

A note on language and method

Throughout, a “win” is a win for the defense. A rate is defense wins divided by all decisions in the stated group; mixed, settled, plaintiff and undecided decisions stay in the denominator, and ‘undecided’ means the ruling did not end the case. Rates are quoted on brand defendants (the company running the site, excluding platforms and tracking vendors) unless stated otherwise. Appendix A defines the analyzed set, explains how records are counted, and shows why the server-side and live-chat rates overstate their effect.

A PERMANENT OPERATING CONDITION

Not a wave that *crests*.

01

CIPA has been on the books since 1967. For five decades it produced a trickle of decisions, most of them about telephone wiretaps and call recording. *Javier v. Assurance IQ* put §631(a) squarely on the web in May 2022. Since then, internet-tracking decisions under it have climbed from 8 in 2022 to 103 in 2025: 318 in just over four years, against 134 in the five decades before. That is the shape of a standing operating condition, not a wave that crests. In August 2026 the California Assembly and Senate approved SB 690, and the Governor signed it on September 30. The bill, discussed in Section 7, closes one pleading route; it does not change that shape.

The 134 are every corpus record decided before *Javier*, across all three set-aside groups in Figure 11: 40 internet-tracking, 87 telephone or call-centre recording and 7 with no recorded surface; 12 records carry no decision date and fall in neither count. This report analyzes those 318 decisions. They are court decisions, each carrying a CIPA count, not lawsuits filed. A decision is the only record that shows how the defense fared, and that is what the rates in this report measure. The full 493-record dataset, which includes cases outside websites and mobile apps, is described in Appendix A.

> ▸ NOTE
>
> **Published filing counts are larger and measure something else. The Fisher Phillips digital wiretapping litigation tracker lists 4,235 matters (2 October 2026): complaints filed under any privacy statute in any state. It counts how often companies are sued over tracking technology, not how often under CIPA, nor how those suits end.**

The outcomes are almost entirely motion practice: 99 defense wins, 89 mixed, 62 settled, 42 rulings for the plaintiff, 26 undecided at the recorded posture. A defense win on a motion is not always the end of the matter: in 42 of the 99 the ruling dismissed at least one claim with leave to amend, and no later ruling is on record. Trial verdicts are absent and merits rulings are rare, and the latter turn on evidence about the vendor rather than on the lawfulness of the tracking. That is the single most important commercial fact about this statute: these cases end on a motion or on a check.

BY THE NUMBERS · THE ANALYZED RECORD

318 decisions, and not one of them a trial *verdict*.

99	89	62	42 + 26
Defense wins	Mixed rulings	Settled	Plaintiff rulings, plus undecided at the recorded posture

Two baselines carry through the report. The defense wins 31.1% of the 318 analyzed decisions. Where the defendant is the brand whose website was involved, rather than a platform or a tracking vendor, it wins 31.9% over 254 decisions. The split matters because the same technology lands on opposite sides of the baseline depending on which type of company was sued: ad-tech tags run 41.7% where the brand was the defendant and 17.6% where the ad-tech company itself was (Appendix A.3). Across all 64 decisions against a platform or vendor, whatever the technology, the rate is 28.1% (Figure 01).

Forum matters for one reason. 252 of the 318 decisions were in federal district courts, 11 in federal appellate courts, 39 in California state courts and 9 in other state courts, and 7 in another or unrecorded forum. The Article III standing defense exists only in federal court, and Section 2.2 shows it is the most reliable defense in the record. A claim filed in state court takes that option away.

▸ AND NO COMPANY IS EXEMPT

266 distinct defendants appear in the analyzed set, 238 of them exactly once. The most-sued names are the platforms, Meta (10) and Google LLC (9); no other company appears more than 4 times. The rest range from national retailers to single-site stores, and most were sued once. Recent defendants include a hotel chain, a coffee company, a theme-park operator and a record label. Every company is a potential target.

FIGURE 01 · WHO WAS SUED

N = 318

318 decisions, two kinds of defendant

Every analyzed decision is counted once, by whether the company sued ran the website or made the tracking technology.



The brand running the website
254 decisions (80%) · 31.9% defense win rate

The platform or tracking vendor itself
64 decisions (20%) · 28.1% defense win rate

266 distinct companies sued across the 318 decisions
238 of them exactly once. The most any one company faced: 10 (Meta).

The 318 analyzed decisions by who was sued: the brand running the website, or the platform or tracking vendor itself. Across the 318 decisions, 266 distinct companies were sued, 238 of them exactly once.

What gets sued has also changed. In 2023, 62% of the decisions against brands alleged a live-chat widget and 17% a marketing pixel. In 2026 the shares are 2% and 57%. The early docket was built on a strict reading of the wiretap statute, chat transcripts and session recordings, where the content of a communication was easy to plead. The current docket is pixels, analytics and ad-exchange tags, increasingly pleaded under §638.51, where there is no content element to contest (Section 7). A compliance program designed against the 2023 docket is defending the wrong thing.

FIGURE 02 · WHAT GETS SUED

BRAND DEFENDANTS

2023 was chat. 2025 and 2026 are pixels.

Share of each year's decisions against brands that alleged each tracker category. Categories overlap, so a column does not sum to 100%.

	2023 n=42	2024 n=78	2025 n=81	2026 to Jun · n=49
Live chat widget	62%	31%	6%	2%
Session replay	26%	13%	5%	2%
Marketing pixel	17%	29%	59%	57%
Analytics	5%	15%	20%	37%
Ad and data-broker tags	19%	17%	17%	27%

Share of each year's decisions against brands that alleged each tracker category. Categories overlap, so a column does not sum to 100%; 2022 is omitted (four decisions) and 2026 runs to June.

▸ DEFENDING THE WRONG DOCKET

The early docket was chat transcripts and session recordings, where the content of a communication was easy to plead. The current docket is pixels, analytics and ad-exchange tags, pleaded under §638.51, where there is no content element to contest. A compliance program designed against the 2023 docket is defending the wrong thing.

VENDORS, GROUNDS AND PLAINTIFFS

What actually decides the *outcome*.

02

The intuitive model of CIPA risk is that the more a vendor takes for itself, the worse the exposure. The record supports only part of that model, and not the part most compliance programs are built on.

2.1 Win rate by vendor category

On the numbers alone, the ranking appears to confirm a common assumption: vendors acting on the brand's behalf fare better than vendors collecting data for their own use. Live chat and session replay defend at 40.4% and 41.4%. Marketing pixels and analytics, which are Meta and Google collecting for themselves, sit at 24.5% and 22.4%. Two things complicate it. The ad-tag row is not one population but two, and its average hides that. More importantly, these are only the cases plaintiffs chose to file, and Section 2.5 shows why who files matters as much as what they sued over.

FIGURE 03 · TRACKER CATEGORY

BASELINE 31.9%

Pixels and analytics lose most often. Chat and replay lose least.

Defense win rate by alleged tracker category, brand defendants (n=254). Least defensible at the top. Categories overlap: one decision can allege several products.

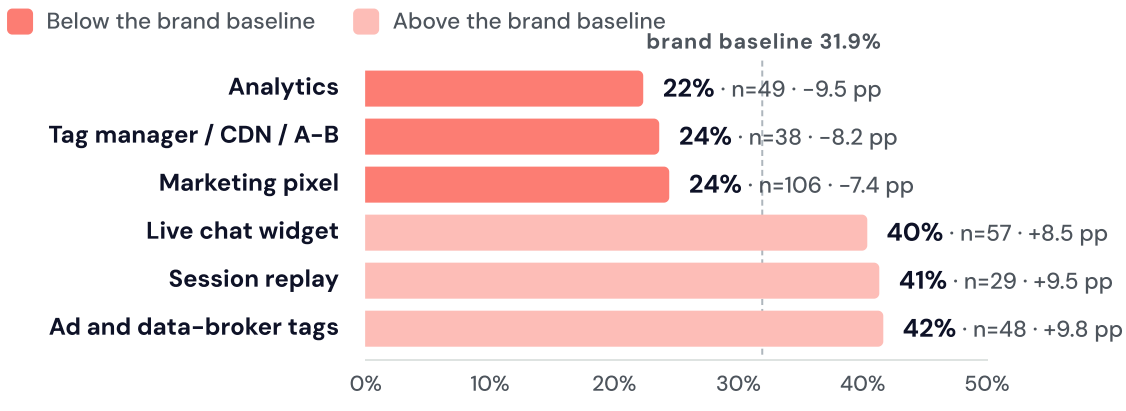

Defense win rate by alleged tracker category, excluding platforms and tracking vendors (n=254), against the 31.9% baseline for that group. Categories overlap: one decision can allege several products. A high rate means the defense won more often; it is not a measure of how much data the technology captures, which runs the other way (Section 2.3).

▸ WHY AD TAGS READ ABOVE THE BASELINE

Ad and data-broker tags read 41.7%, the opposite of the service-provider story. Two populations sit inside that number. 16 of the 48 decisions also allege chat or session replay, most of them brought by repeat filers, and run at 62.5%. The other 32 allege ad tags alone and run at 31.2%, within a point of the 31.9% baseline. That second group is new: 26 of the 32 were decided in 2025 or 2026, 12 were filed by one firm (Bursor & Fisher), and 11 are publishers sued over ad-exchange trackers. It is a young docket and its rate will move.

2.2 What decides the outcome: the legal ground, not the vendor

The record does not support the story that a vendor's business model decides the case. Where a plaintiff alleges the third party used the data for its own purposes, on 160 decisions, the defense wins 32.5%. Where a court accepts the vendor as a service provider acting only on instruction, on 54 decisions, 37.0%. The baseline is 31.9%. There is no statistically significant difference between the three numbers.

What separates outcomes is the ground the case is decided on, and three grounds carry most of it.

- **Standing is the earliest win and the most reliable.** 30 decisions turn on it and the defense takes 80.0% of them. It is available only in federal court, because it is the Article III test of *TransUnion LLC v. Ramirez* (2021). In *Valenzuela v. Keurig Green Mountain* (N.D. Cal. 2023) a chat-interception claim was dismissed because the complaint never identified any private content or resulting harm. In *Maghoney v. Dotdash Media Inc.* (S.D. Cal. 2026) health-related search queries shared with third parties were held not to be a concrete injury because they were never tied to the plaintiff's identity. Neither case decided whether the tracking violated CIPA; both asked, under *TransUnion*, whether the alleged disclosure was a concrete injury. The Ninth Circuit set the current bar in *Popa v. Microsoft Corp.* (9th Cir. 2025): session-replay capture of a visit to a pet-supply site, pleaded under Pennsylvania's wiretap statute, was not a concrete injury. District courts in California now apply it to CIPA claims, *Maghoney* among them, though they remain split on how far it reaches.
- **The party exception is the specific defense to content capture.** 17 decisions, 64.7% to the defense. The argument is that the vendor was the brand's own recording tool rather than a third party listening in. *Williams v. What If Holdings* (N.D. Cal. 2022) is the clean statement of it: a form-capture product was held to be a tool, not an eavesdropper. *Pena v. GameStop, Inc.* (S.D. Cal. 2023) applied it to a website chat. It is the deciding ground in 7 of 57 chat decisions and 3 of 106 pixel decisions. It is one doctrinal reason chat and pixels behave differently; Section 2.5 has the other.
- **Element failure is the largest bucket and the least predictable.** 77 decisions at 50.6%: no interception in transit, no contents, no confidential communication. Most of the doctrinal disagreement between judges sits here.

The ground a company can build in advance ranks last. Consent disclosure is the primary ground in 6 of 318 decisions, rare enough that it cannot be the plan (Section 2.4).

THE DECIDING GROUNDS · WHERE CASES ARE WON

Three grounds carry most of the *record*.

80.0%	64.7%	50.6%	6
Standing · 30 decisions · federal court only	Party exception · 17 decisions · content capture	Element failure · 77 decisions · the largest bucket	Decisions out of 318 where consent disclosure was the primary ground

2.3 The same technology, two opposite rulings

The deciding grounds are doctrinal rather than technical, so the same product can produce opposite outcomes on similar facts. A brand cannot read its own risk off the vendor's name. Two examples from the record make the point.

Two chat rulings, two answers

CHAT WIDGET · AIDING AND ABETTING 2024–25

In *Rodriguez v. Ford Motor Company* (S.D. Cal. 2024) the court let an aiding-and-abetting claim proceed at the pleading stage because the chat provider allegedly had the capability to use the transcripts for its own purposes and the brand was plausibly aware of it. In *Gutierrez v. Converse Inc.* (9th Cir. 2025) the Ninth Circuit affirmed summary judgment for the brand on the same theory against a different chat provider, because the plaintiff had no evidence the vendor actually read the messages.

WHAT SEPARATED THEM: **THE STAGE, AND THE EVIDENCE OF THE VENDOR'S CONDUCT**

Session replay defends well despite capturing the most

SESSION REPLAY · §631(A) 2025

Replay defends comparatively well, and it is not safe. In *Mikulsky v. Bloomingdale's, LLC* (9th Cir. 2025) the Ninth Circuit held a §631(a) claim based on session-replay interception adequately pleaded, and the matter settled. Better odds on a motion do not make a matter cheaper: replay and chat capture transcripts, form contents before submission and keystrokes, the richest material any plaintiff in this record has to work with.

RECORDED OUTCOME: **SETTLED**

▸ EASIER TO DEFEND IS NOT SAFER TO RUN

Session replay and live chat are easier to defend than a marketing pixel, and they capture far more: keystrokes, form contents and whole transcripts. A risk assessment has to weigh the odds on the motion and the size of the settlement together.

2.4 Consent posture: what attracts the claim, not what decides it

FIGURE 04 · CONSENT POSTURE

BASELINE 31.9%

Only separate written consent clearly beats the baseline. A banner does not.

Defense win rate by alleged consent posture, brand defendants (n=254). Least defensible at the top. Postures overlap.

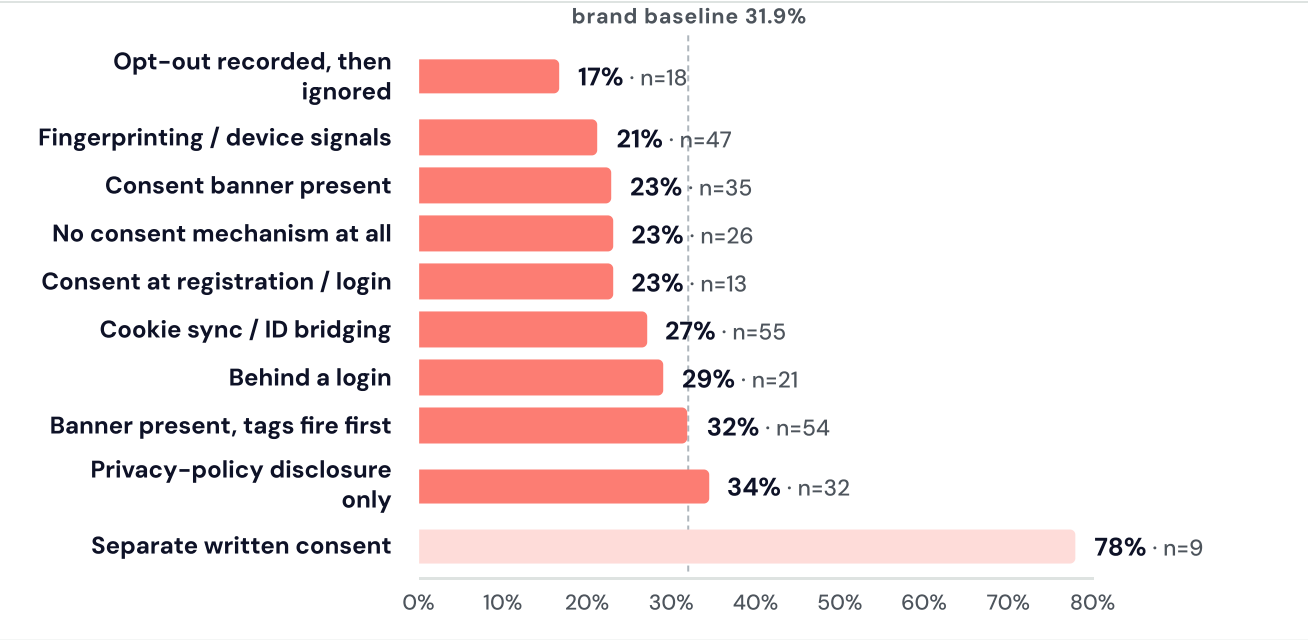

Defense win rate by alleged consent posture, excluding platforms and tracking vendors (n=254). Postures overlap: a banner that fires late is both 'banner present' and 'tags fire first.' Separate written consent, the only posture clearly above the baseline, stands on 9 decisions.

The litigated population is almost entirely partial consent. The worst posture in the figure is the one that looks like compliance from the inside: an opt-out that is recorded and then ignored, at 16.7% on 18 decisions. It is also the newest: 1 brand decision alleged it in 2024, 7 in 2025 and 10 in the first half of 2026, one in five of that year's. *Casillas v. Six Flags Entertainment Corporation* (C.D. Cal. 2025) is the archetype. The complaint alleges that a theme-park site kept transmitting browsing and interaction data to advertising and data-broker services after visitors chose “Deny Non-Essential.” A banner was present. A consent platform was deployed. The claim turned on whether the tags obeyed the choice. The lesson is the one European privacy teams learned under the ePrivacy Directive and

GDPR: configuration is the work, third-party chains drift, and the consent state has to be re-measured rather than assumed.

Two rows invite a wrong reading. ‘Privacy-policy disclosure only’ (34.4%) sits above ‘consent banner present’ (22.9%), as if a banner hurt. It does not. The two rows share 1 decision and differ by three wins (11 of 32 against 8 of 35), and the two flags enter the record from different parties for different reasons. A banner is recorded mostly because the plaintiff says it failed: 15 of the 35 banner decisions also allege an ignored opt-out and 13 allege tags firing before any choice. A policy-only disclosure is recorded mostly because the defendant offered it as consent and the court weighed it. The figure therefore carries one finding, not ten: every alleged consent posture except separate written consent runs at or below the baseline (policy-only disclosure sits within a win above it), because a posture appears in a pleading when the plaintiff has a story about how it failed. Decisions with no consent posture recorded at all run at 34.3% (n=140), above every row in the figure except separate written consent and policy-only disclosure.

One row is deliberately absent. ‘Banner present and respected’ appears on six decisions; we reviewed all of them. One is a criminal §632 prosecution that is not a website matter at all, two are mixed, and three are defense wins: a writ petition with the caption reversed, a dismissal for lack of standing, and a ruling on consent to the platform's own account terms rather than to a site's banner. The flag is asserted only by courts and by defendants, never by a plaintiff. It is what a court says when it rejects a consent theory, not a measurement of a working platform.

Where consent does decide a case, it is because the assent can be proven. In *Snyder v. G6 Hospitality LLC* (C.D. Cal. 2025) a guest had to tick a box reading “I agree to rate and room policies of this reservation” before the booking would go through; the asterisk led to the policies, the privacy policy among them. The court held that box to be assent, found the policy described the sharing, and dismissed without leave to amend. There was no cookie banner in the analysis: what won was a checkbox the user could not skip and a policy that matched the conduct.

▸ WHAT TO TAKE FROM THIS

Consent posture is what attracts the claim, almost never what *decides* it.

Getting the banner right changes what a claimant finds; it does not win the motion. The weakest posture in the record is the one that looks compliant from the inside: a consent platform deployed, a choice recorded, and tags that keep firing anyway.

16.7%	34.3%	15 of 35	6
Opt-out recorded, then ignored (18 decisions). The weakest posture on file	Decisions with no consent posture recorded at all (140 decisions)	Banner decisions that also allege an ignored opt-out	Decisions flagged 'banner present and respected'. None is a working platform

2.5 Repeat plaintiffs and the firms behind them: read the name first

A small number of serial plaintiffs send thousands of templated demand letters under CIPA. For a company that receives one, the name on the letter is the most informative fact about the matter, and it is available before anything else is.

CIPA claims arrive from a spectrum, and the two ends are different businesses. At one end is the templated demand aimed at a settlement below the cost of responding. On July 20, 2026, in *Vivek Shah v. Crain Communications, Inc.*, No. 2:26-cv-03070-RGK-CTS, the Central District of California declared one such filer a vexatious litigant. The court counted at least 29 proceedings since 2021, all but one filed pro se, and seven near-identical CIPA complaints against seven different defendants in the seven months before that suit. It described template claims calibrated to clear the amount-in-controversy threshold and voluntarily dismissed once a defendant moved to dismiss. The only recovery identified in the order was a $3,000 arbitration award. The revenue in that model is not in judgments. It is in the letters answered before anyone reaches a judge. The prefiling order is narrow: it requires leave of court for new CIPA and digital-privacy filings in that district only, and does not reach other federal districts, California state court, other claim types, or the filer's pending matters. The designation constrains one forum rather than ending the practice, and a defendant has to know to cite it. The high-volume demand-letter firms, among them Pacific Trial Attorneys, Tauler Smith and Manning Law, appear in none of the settlements with a disclosed amount behind Section 6. That is what the model looks like from outside: it resolves privately, and rarely produces a public number.

At the other end are class-action firms for whom a single matter is the product. Simmons Hanly Conroy has appeared for plaintiffs in multiple CIPA-related settlements with hospital systems,

totaling more than $40 million.

The name on the letter, and the counsel behind it, tell a defense team which end it is holding, and the two call for different first moves. A templated demand is priced to settle below the cost of a motion. A class action is built to survive one, and the exposure is the class rather than the letter.

The claimant's purpose is also becoming a defense in its own right. In *Rodriguez v. Autotrader.com Inc.* (C.D. Cal. 2025) the court held that a privacy tester who visited a site expecting an invasion of privacy had no reasonable expectation of privacy and therefore no Article III injury, and dismissed the §631 claim with prejudice. That route runs through standing, so it is available only in federal court, as Section 2.2 sets out.

THE VEXATIOUS-LITIGANT ORDER · SHAH V. CRAIN COMMUNICATIONS

What a templated-demand practice looks like from the *docket*.

29	7	$3,000	1
Proceedings since 2021, all but one filed pro se	Near-identical CIPA complaints in seven months, against seven defendants	The only recovery identified anywhere in the court's order	District the prefiling order covers. It does not reach any other forum

▸ READ THE NAME FIRST

Before assessing the technology, check the claimant and their counsel against their filing history. A quick search of the public docket usually shows whether this is a templated demand or the opening of a class action, and the two call for different responses.

SOURCES

IAPP, 'You've got mail: From a serial privacy plaintiff: A privacy pro's CIPA playbook,' September 10, 2026. *Vivek Shah v. Crain Communications, Inc.*, No. 2:26-cv-03070-RGK-CTS (C.D. Cal. July 20, 2026) (Klausner, J.), reported in Duane Morris LLP and Pillsbury Winthrop Shaw Pittman client alerts, July 2026. Settlement amounts for Sutter Health, Partners HealthCare, BJC HealthCare and MarinHealth Medical Center are drawn from the ledger described in Section 6 and Appendix A; one of the four is recorded at a ceiling rather than a paid amount, so the total is stated below the recorded sum. Firm identifications follow ArentFox Schiff, 'Class Action Year in Review: Wiretapping Update,' and Barnes & Thornburg, 'CIPA/ECPA Website-Tracking Privacy Litigation in 2026.'

2.6 Sensitive data lowers the win odds and raises the price

On paper, CIPA requires no proof of actual damages, so the sensitivity of the data should not change whether a claim survives. In the record it does. Across all 318 analyzed decisions, rather than the

brand subset the rest of Section 2 uses, the defense wins 35.6% of decisions where no sensitive data is alleged and 25.4% where any is. Health and intimate-health data run at 22.4%, children's data at 22.2%. The same marketing pixel runs at 29.6% without sensitive data and 17.5% with it. Financial data is the exception at 33.3%, near the baseline, and only when it travels without health data: alongside health data it runs at 25.0% (n=24), on its own at 42.9% (n=21). The gradient runs through health, intimate health, and children's data specifically. And the other side of the split is not comfortable either: with no sensitive data on the page, the defense still wins only 35.6% of the time.

Two mechanisms are visible in the record and a third is a hypothesis; they compound. First, sensitive data brings claims CIPA does not have (the Confidentiality of Medical Information Act, the constitutional privacy right, ECPA), which changes the ground the case is decided on and the price it settles at. Second, selection: the pattern is consistent with plaintiffs targeting sites that hold sensitive data because those defendants settle. Sensitive-data decisions end in settlement more often (22% against 17%) and in an outright ruling for the plaintiff less often (10% against 16%). That is the profile of a defendant who pays rather than litigates, and it is why the selection bias runs against health-related brands. Third, and not measurable in this record: a court may read harm into a pixel on a symptom checker that it does not read into the same pixel on a checkout page, whatever the statute says. We record that as a hypothesis.

FIGURE 05 · DATA SENSITIVITY

BASELINE 31.1%

Sensitive data lowers the defense win rate

Defense win rate by the sensitivity of the data alleged to have been transmitted (n=318). Least defensible at the top.

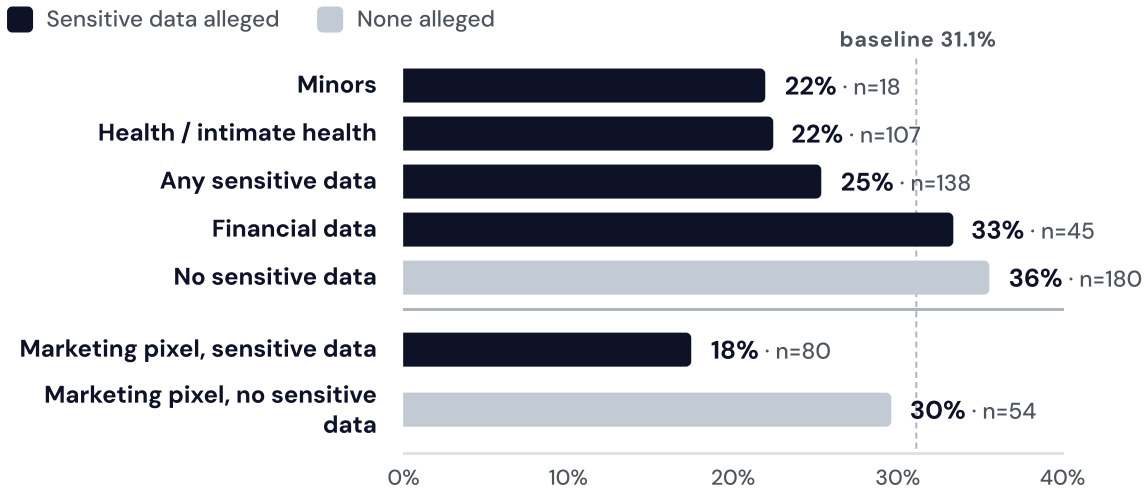

Defense win rate by the sensitivity of the data alleged to have been transmitted. Color marks whether sensitive data was alleged at all; the two marketing-pixel rows below the rule are the same technology cut by that question.

The mechanism is visible in the pleadings. *Pattison v. Teladoc Health, Inc.* (S.D.N.Y. 2025) is the pattern: a telehealth provider's use of a marketing pixel and its server-side counterpart to transmit

patient information supported federal wiretap and state-law claims together, and the court let most of them proceed. The same tag on a checkout page would have carried the CIPA count alone.

> ▸ THE LINE TO REMEMBER
>
> The same pixel on a checkout page or on a symptom-checker page is the same transmission, with worse odds for the defense and a higher price.

2.7 Sector

Retail is the most-sued industry and the best-defended. Healthcare is the second most-sued and among the worst-defended. That pairing is the sector story. Industry here is a hand-coded judgment about the defendant's business rather than an extracted fact.

FIGURE 06 · INDUSTRY

BASELINE 31.9%

Retail is sued most and defended best. Healthcare is the reverse.

Defense win rate by the defendant brand's industry, brand defendants only (n=254). Industries with fewer than ten decisions are omitted.

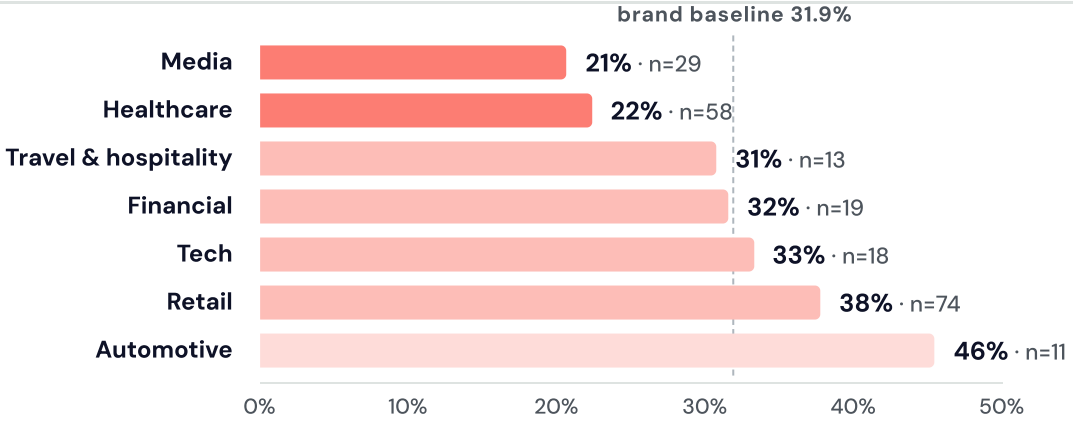

Defense win rate by the defendant brand's industry, excluding platforms and tracking vendors. Industries with fewer than ten decisions are omitted.

> ▸ WHAT SECTION 2 ADDS UP TO
>
> What a vendor does matters. Who sued matters more, and what the data is matters most of all. The technology on the page sets which claim arrives; the plaintiff sets how hard it is to end; the sensitivity of the data sets both the odds on the motion and the price of the settlement.

SELECTION, EXPOSURE, ODDS

Eight measures that reduce *exposure*.

03

Eight moves, ordered by the lever each acts on rather than by compliance logic. **Selection** is the probability a claimant picks the company at all, decided by what a website user can see in browser developer tools in ten minutes. **Exposure** is the size of the number once picked, driven by the third-party host count on page load and the sensitivity of the data. **Odds** is the probability of winning the one motion that decides it.

01 Nothing third-party fires before consent

LOWERS SELECTION AND EXPOSURE

Removes the two commonest alleged patterns at once and shrinks the multiplier in a demand letter, which is counted from exactly those requests.

02 Verify the consent platform is obeyed, continuously

LOWERS SELECTION AND EXPOSURE

A banner is alleged on 35 decisions and an ignored opt-out on 18 (16.7% defense wins, the weakest posture on the table). This is a monitoring problem, not a procurement one.

03 Real prior consent where the data is sensitive

LOWERS EXPOSURE, RAISES THE ODDS

The only consent posture clearly above the baseline, 77.8% on 9 decisions (directional), and sensitive data is where both the odds and the price are set.

04 No implied consent, pre-checked boxes or delayed choice

RAISES THE ODDS

Javier, where this analysis starts, rejects consent obtained only after the collection; implied consent and pre-checked boxes were not decided there and need their own support.

05 Separate essential functionality from tracking; replace vendors that cannot sit behind a consent gate

LOWERS SELECTION AND EXPOSURE

The vendors that cannot be gated generate the alleged pattern in the first place.

06 Do not treat server-side tagging as remediation

DOES NOT CHANGE THE ODDS

No decision in the record pleads server-side forwarding alone. All 19 plead it beside a marketing pixel, and the pixel's odds do not change in any way the record can distinguish: 23.5% with a server-side allegation against 16.7% without, on the same sensitive-data pixel cases (4 wins in 17 against 7 in 42). Server-side hides the request from a browser capture, so it may lower the chance of being picked, since both demand archetypes start from a network log. It does not change what was sent once a claim is filed, and *Pattison v. Teladoc Health, Inc.* (S.D.N.Y. 2025) is what it looks like when it is pleaded alongside the pixel rather than instead of it (see Appendix A.4).

07 Do not rely on the privacy policy

DOES NOT CHANGE THE ODDS

Disclosure-only consent runs 34.4% on 32 decisions, level with the baseline rather than above it.

08 Inventory the estate repeatedly, and measure what the vendors load

LOWERS SELECTION AND EXPOSURE

Building an accurate vendor inventory is difficult enough. The greater challenge is understanding where those vendors' request chains reach and what additional services they introduce. At scale, no single team has a complete view.

WHAT CAN BE PRICED

What the moves cost, and what can be *measured*.

04

Every rate here is conditioned on a suit having already been filed. It can move the odds in litigation. It says nothing directly about the chance of being selected, and Section 2.5 shows that who brings a claim matters as much as the technology it targets. Anyone who says that adding consent makes a company X% more likely to win is misreading outcome data. We hold no data on defense costs and publish no cost-of-defense figure. Two rates in Section 2 are inflated by specific factors: server-side tagging by its association with Meta Pixel and live chat by unusually frequent federal court litigation. See Appendix A.4.

4.2 The cost side

Four costs recur, and none of them is legal spend.

- **Measurement and attribution loss from consent-gating.** For most brands a measurement loss; for an ad-funded publisher a direct revenue loss, because unconsented inventory sells for less. Its size is knowable in advance from the company's own consent rate.
- **Vendor replacement.** The expensive one, and the only one that removes an alleged pattern rather than documenting it.
- **Engineering time across the estate.** The unit is properties, not tags, and the property count is always higher than the room believes.
- **Continuous monitoring.** An operating cost, not a project cost.

4.3 The measurable number: third-party hosts on page load

One of the demand letter's multipliers is the count of distinct third parties a claimant can enumerate from a visit to the site. That is measurable, and it is measurable in-house before anybody else does it. Start with a cold load: one visit, no interaction, no consent given, counting requests, distinct third-party hosts, and cookies set before any consent decision. Then repeat it the way a claimant does: type into a search box or a form without submitting it, to see whether a pixel captures the keystrokes or the email address, and visit with Global Privacy Control (GPC) set, to see whether anything honors it. Then count twice: once as the site runs today, and once with nothing third-party firing before consent. The difference is how many third parties move 1 takes out of the letter's count.

THE ONE NUMBER A COMPANY CONTROLS

Count the third-party hosts before a claimant *does*.

A cold load with no interaction and no consent given. Then the interactions a claimant tries: text typed into a form without submitting, and a visit with Global Privacy Control set. Then count twice: once as the site runs today, and once with nothing third-party firing before consent. The difference is how many third parties move 1 takes out of the letter's count.

4.4 Why statutory structure dominates what a company controls

The Video Privacy Protection Act (VPPA) is the closest benchmark: same pixels, same era, same plaintiff firms. It has real statutory gates ('consumer,' 'personally identifiable information,' 'video tape service provider'), and defendants win on them: in the 186 VPPA records we hold, a separately assembled set, the defense wins 42.5%, against 31.1% under CIPA. CIPA has almost none of those gates, which is why it is the vehicle of choice. The comparison does not isolate statute from conduct, but statutory structure does more to decide defensibility than conduct does, and that is also the right lens for reading SB 690 (Section 7).

BENCHMARK · VPPA AGAINST CIPA

Statutory gates are what a defendant actually wins *on*.

42.5%	31.1%	3
Defense win rate across 186 VPPA records	Defense win rate across 318 CIPA decisions	Statutory gates the VPPA has and CIPA effectively does not

THREE AXES

If a letter *arrives*.

05

A letter has three axes: who sent it, which provision it invokes, and what technology it read. Every available defense follows from those three answers, and the two archetypes on file sit at opposite ends of all three.

Archetype A

STACKED-MULTIPLIER ARBITRATION DEMAND

Demand

The arithmetic is built from a single network capture: two confirmed interceptions plus every other third party found on the page, multiplied by the statutory figure, then asserted twice, once under §631(a) and again under §638.51 as expressly cumulative, plus a federal statutory claim, and delivered through the company's own pre-arbitration notice provision.

OPENS AT **100% OF COMPUTED EXPOSURE**

Archetype B

BARE PEN-REGISTER COMPLAINT

Complaint

No calculation at all: third-party trackers on page load are pen registers, an IP address is addressing information, no court order was obtained. No contents, no sensitive data, no harm and no pixel required. It is the maximally broad theory, and Section 7 shows how fast it grew.

STATES **NO NUMBER**. A FILED COMPLAINT SETS THE ANCHOR

Several things follow from Archetype A. There are two multipliers, the number of third parties and the number of transmissions per third party, and both can be contested: a letter's 'sixty third parties' typically counts vendors acting on the brand's own behalf, several endpoints of one vendor, and hosts that intercept nothing, while the transmission count grows with every interaction and with the size of a putative class. §638.51 doubles the same conduct; SB 690 removes that second count, and a wiretap-only letter has to plead interception in transit, which can be put to proof (Section 7). Only the first one or two interceptions are evidenced; the rest is inferred. And the arbitration clause is the delivery mechanism.

Archetype B's prayer for relief is, in effect, a technical remediation checklist written by the adversary: block third-party requests until consent, no implied or pre-checked or delayed choice, segregate essential from tracking, replace vendors that cannot comply, submit to independent audits. Section 3 is that list.

5.4 The best defense, by axis

— **The plaintiff.** Standing is the primary ground in 30 decisions and the earliest place to win. Whether the claimant is a repeat filer, and which firm's roster they belong to, is the most

informative single fact about a new matter (Section 2.5); check the name against the docket before anything else. A vexatious-litigant designation exists in one court: in *Shah v. Crain Communications, Inc.*, No. 2:26-cv-03070-RGK-CTS (C.D. Cal. July 20, 2026), the court entered a prefiling order barring the plaintiff from filing any new CIPA or related digital-privacy claim in the Central District without leave.

- **The statute.** The provisions are not interchangeable, and the rates overlap because one decision pleads several theories: §631(a) wiretapping 156 decisions (35.3% defense wins), §631(a) in-transit reading 141 decisions (28.4%), §631(a) aiding and abetting 141 decisions (31.9%), §638.51 pen register / trap and trace 103 decisions (30.1%). 'In transit' belongs only to §631. A letter stacking both provisions on one set of facts is asserting two theories and should be answered as two. Figure 07 puts the four exclusive combinations side by side: pleading §638.51 instead of §631 does not change the odds.

- **The technology.** Per-third-party counting is the weakest link in the chain. The claimant's evidence is a network capture of their own visit, but the claimant was never a customer: the site was found in a crawl and loaded once, under conditions the claimant chose. *Gutierrez v. Converse Inc.* (9th Cir. 2025) is the reason this matters: summary judgment for the brand because the plaintiff could not show the vendor actually read anything. A continuous internal record of what fired, when, and against which consent state is the only thing that contradicts the claimant's capture, and it has to exist before the letter arrives.

FIGURE 07 · PROVISIONS PLEADED

BASELINE 31.1%

Pleading §638.51 instead of §631 does not change the odds

Defense win rate by the CIPA provisions pleaded, all analyzed decisions (n=318). The rows are exclusive: each decision is counted once.

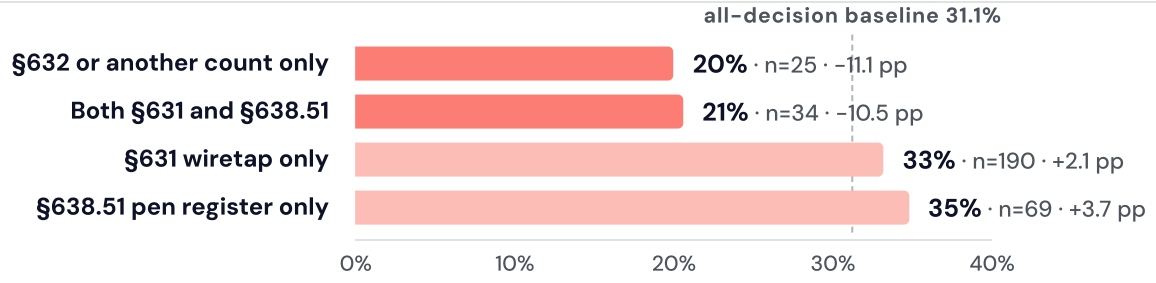

Defense win rate by the CIPA provisions pleaded, all analyzed decisions, against the 31.1% baseline for the whole set. The rows are exclusive, so they can be read against each other; the two small rows are directional. Where both provisions are pleaded, half the decisions are mixed: one count dismissed, the other allowed to proceed, which leaves fewer clean defense wins to record.

This describes patterns in documents we hold. It is not legal advice and not a litigation strategy.

5.6 What the stack does to the price

A demand letter stacks claims the way a complaint does. Co-pleaded claims run intrusion upon seclusion (common law) in 100, the California constitutional privacy right in 98, California's Unfair Competition Law (Business and Professions Code §17200) in 76 decisions. CIPA is the hook; the stacked claim sets the number, which is why sensitivity shows up in the price as well as in the odds. *Ahringer v. LoanDepot, Inc.* (C.D. Cal. 2024) is the shape of it: the wiretap theory against the brand itself failed, the aiding-and-abetting theory and a California constitutional privacy claim over sensitive financial information survived, and the matter resolved without a disclosed figure.

THE STACK · CLAIMS PLEADED ALONGSIDE CIPA

CIPA is the hook. The stacked claim sets the *number*.

100	98	76
Decisions co-pleading intrusion upon seclusion (common law)	Decisions co-pleading the California constitutional privacy right	Decisions co-pleading California's Unfair Competition Law

▸ THE FIRST HOUR, IN ORDER

Read the name. Read the provision. Then read the *capture*.

Check the claimant and their counsel against the public docket before assessing anything technical. Answer a stacked letter as two theories rather than one. And put the per-third-party count to proof: the claimant's evidence is a single network capture taken by an adverse party under conditions they chose.

30	2	1	0
Decisions decided on standing, the earliest place to win. Federal court only	Theories in a letter that stacks §631 and §638.51 on one set of facts	Visit behind the claimant's evidence, under conditions the claimant chose	Of it answerable without a dated internal record of what fired, and when

THE DISCLOSED RECORD

What it costs, and what we *cannot see*.

06

Publicly disclosed settlement figures exist for 42 internet-tracking matters; 33 further settlements have no public number at all. 35 of the disclosed figures were agreed by a brand for tracking on its own website or app, the population every rate in this report is quoted on. Across those: median $3.7M, quartiles $2.3M to $9.2M, largest $47.5M (Kaiser Permanente), total $275.3M. Every figure was checked against the settlement agreement, the approval order or the official settlement site before it was printed.

How a claims-made settlement is counted

9 of the 35 brand figures are not a stated fund. The agreement sets a payment per claimant and caps the fees, or adds a claims-made supplement to a fund, so there is no single number. Those are counted at the most the defendant agreed to pay: the fund plus the supplement, or every class member paid at the per-claim rate plus the fee and award caps. That is what was allocated. What is actually paid out is rarely disclosed and always lower; in the one matter where it is, Legacy Health redeemed $2.0M of a $20.7M ceiling.

The platform funds are a different case

The other 7 disclosed figures, $363.2M between them, were paid by a platform or vendor sued over its own product rather than by a brand over third-party tags: Oracle America, Inc. $115.0M, ByteDance, Inc. $92.0M, Zoom Video Communications Inc. $85.0M, Google LLC $48.0M, PowerSchool Group Holdings Inc. $17.2M, Flurry Inc. $3.5M, Nebuad, Inc. $2.4M. Their median is $48.0M, and the two largest funds on file are among them. They are drawn separately in the figures below and are counted in the totals, not in the brand median.

FIGURE 08 · WHAT A SETTLEMENT COSTS

42 DISCLOSED

The brand distribution, and the platform funds beside it

35 disclosed funds agreed by brands over tracking on their own properties, on a log scale, with the 7 platform and vendor funds shown separately.

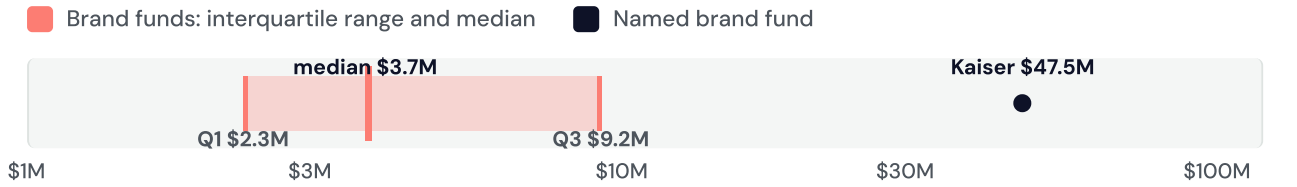

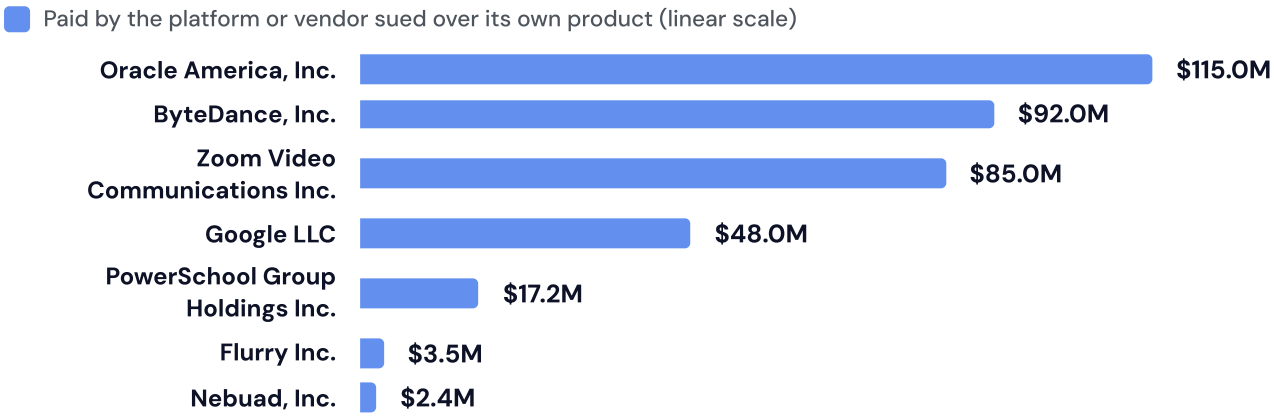

0.04%	2.00%	4.66%	12.76%
Kaiser Permanente	**Boston Globe Media**	**Flo Health, Inc.**	**Choose Your Horizon**
$47.5M on $116B revenue	$4.0M on $200M revenue	$8.0M on $172M revenue	$0.4M on $3.13M revenue

The dollar spread across the disclosed settlements is three orders of magnitude, and most of it is the size of the defendant rather than the seriousness of the conduct. The bottom row shows each named settlement as a share of the payer's annual revenue: the largest fund on file is the smallest share. Three payers carry no usable revenue figure (Call-On-Doc, Inc. and Nebuad, Inc.), and for Mount Nittany Health System the revenue we hold ($0.8M against a $1.8M settlement) points to the wrong legal entity.

Settlement dollars by industry

FIGURE 09 · INDUSTRY

$638M TOTAL

Healthcare pays most often: 29 of 35 brand settlements

Disclosed settlement funds by industry (42 settlements, $638M in total), split by who paid. The largest funds were paid by platforms.

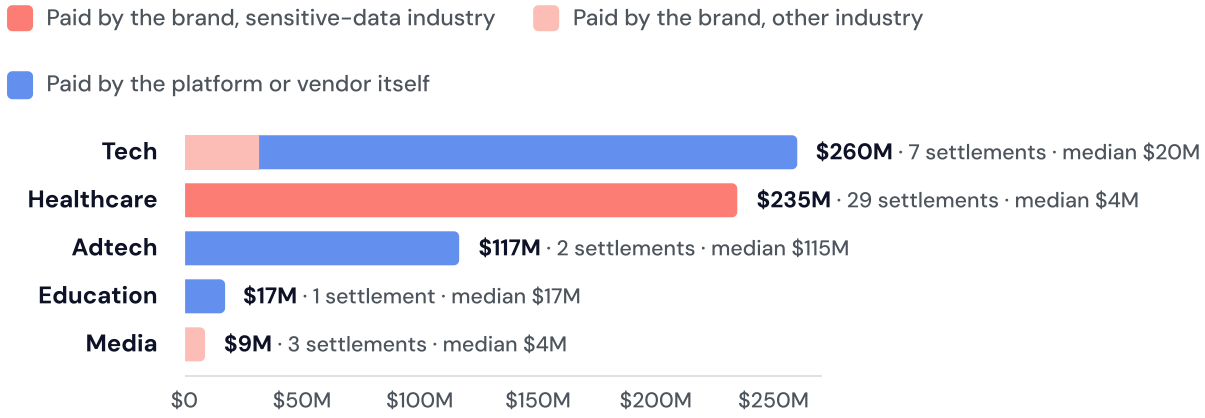

Disclosed settlement funds by the payer's industry, with the count and median beside each bar, split by who paid: the brand running the site (salmon where the industry's data is sensitive, light salmon otherwise) or the platform or vendor itself (blue). Streaming brands are counted under tech.

The defensible statement is that sensitive data concentrates both the filings and the dollars.

Read this section with more caution than the rest

The gap is not in the amounts. It is in the disclosure. Counted by matter, 42 settlements have a disclosed amount and 33 do not. Counted by decision record, the corpus holds 78 settled records, 62 of them among the decisions analyzed here, and only 21 of the 78 carry a public number. A settlement notice rarely enumerates the statutory counts that decide corpus membership, so the disclosed figures and the decisions analyzed here overlap only partly: 22 of the valued figures belong to matters in the corpus (one record carries two settlements) and 19 to matters among the analyzed decisions. The distribution above is the visible tail of a larger one, and it is biased upward, because a private resolution is not filed anywhere and a pre-suit resolution is not even docketed. The volume of undisclosed settlements is large, and this report cannot size it.

What the check changed

Every figure was checked against the settlement agreement, the final-approval order or the official settlement site, and seven recorded amounts did not survive. The largest of the seven is the widely reported $32.0M GoodRx settlement: the court denied preliminary approval twice, most recently in January 2026, so there is no fund to report.

What is counted apart

The same search confirmed 7 large settlements, $153.6M between them, that are call-recording and voice-assistant matters rather than internet tracking. They are excluded from the figures above, for the same reason Appendix A sets that population aside: the largest of them ($95.0M) is bigger than any brand-paid fund above, and the price of a voice-recording claim is not the price of a pixel claim.

▸ WHAT THE DISCLOSED RECORD DOES AND DOES NOT SHOW

The visible tail of a larger distribution, biased *upward*.

A private resolution is not filed anywhere, and a pre-suit resolution is not even docketed. The volume of undisclosed settlements is large, and this report cannot size it.

42 / 33	21 of 78	7	$2.0M
Settlements with a disclosed amount, against those with no public number	Settled records in the corpus that carry a public figure	Recorded amounts that did not survive the check against the agreement or order	Actually redeemed by Legacy Health against a $20.7M ceiling, the one matter where payout is known

WHAT CHANGES NEXT

SB 690 and the pen-register *route*.

07

Everything in this section is stated as of 1 October 2026. SB 690 passed both houses of the California Legislature on August 28, 2026, and the Governor signed it on September 30. It takes effect on January 1, 2027. The final bill is far narrower than the one introduced. It leaves §631 untouched and changes one thing: it removes the private right of action (Penal Code §637.2) for a §638.51 pen-register or trap-and-trace claim arising from a website or app, leaving that claim to the Attorney General. That limitation reaches back to pending claims filed within two years of the operative date (January 1, 2025).

The record says this matters more than the commentary suggests, and differently. The pen-register theory is not a bolt-on that doubles a §631 claim. It is mostly a substitute for one. Of the 103 analyzed decisions carrying a §638.51 count, 69 plead it without any §631 count at all, against 34 that plead both. It is also new: it appears in 1 in 2023, 22 in 2024, 48 in 2025, 32 in 2026. The route opened with *Greenley v. Kochava Inc.* (S.D. Cal. 2023), which held that software collecting identifiers and browsing data could plausibly be a pen register. *Moody v. C2 Educational Systems Inc.* (C.D. Cal. 2024) then let a TikTok pixel through on the same theory, and *Shah v. Fandom, Inc.* (N.D. Cal. 2024) extended it to ordinary third-party scripts. The figure shows the result: wiretap-only decisions fell from 68 in 2024 to 45 in 2025 while pen-register decisions rose, and the two are pleaded together in a growing minority (32 of the 34 both-pleaded decisions date from 2025 and 2026).

FIGURE 10 · PROVISIONS BY YEAR

N = 318

The pen-register route grew while wiretap-only decisions fell

Analyzed decisions by year and by the CIPA provisions pleaded (n=318). 2026 runs to June.

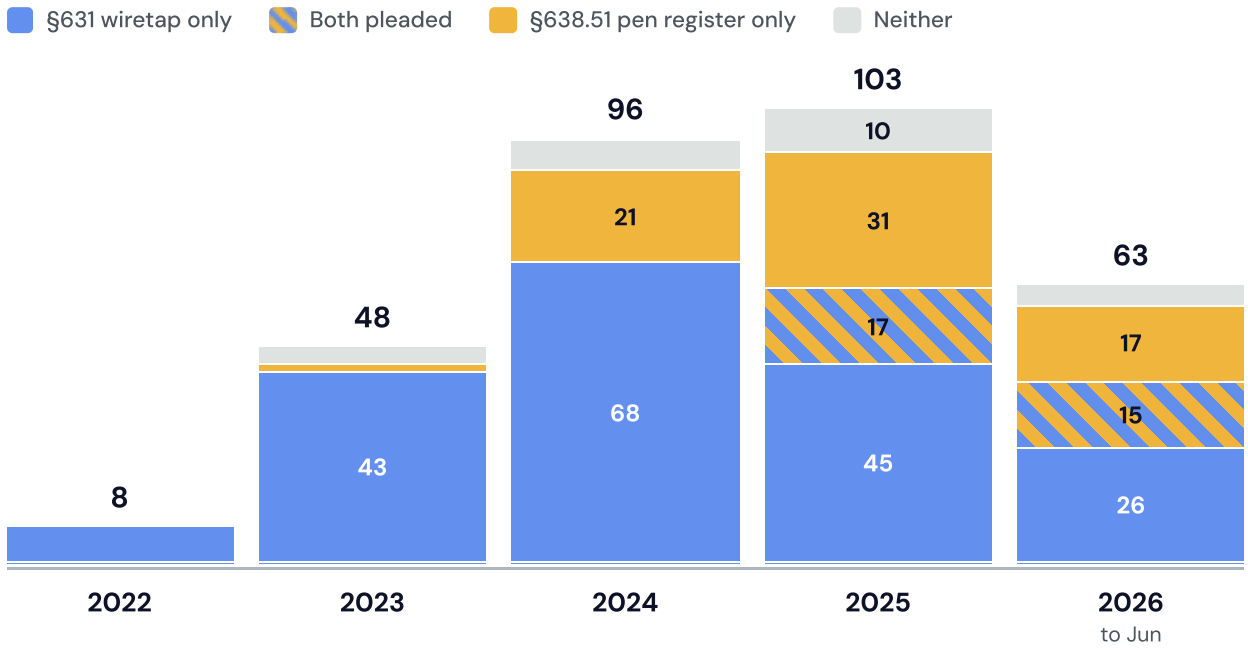

Analyzed decisions by year and by the CIPA provisions pleaded: §631 only, §638.51 only, both, or neither (a §632 or other count alone). 2026 runs to June. The §631 route did not disappear; it was joined and then partly displaced.

Why plaintiffs took it is visible in the grounds. §631 decisions turn on elements: interception in transit, contents, a confidential communication. Element failure is the largest single ground in the record (77 decisions). A pen-register claim needs none of that: a third-party request on page load and an IP address are the whole theory. TikTok's software is the clearest example. It is named in 29 analyzed decisions, 21 of them under §638.51 and only 6 under §631, although a pixel that transmits a page URL would fit the wiretap statute if the plaintiff could plead interception in transit. Under §638.51 the request itself is the device, so nobody has to. It was the path of least friction, not the path to better odds. The defense wins 34.8% of the pen-register-only decisions and 33.2% of the §631-only decisions, which is the same number (Figure 07).

Where both are pleaded, the defense wins 20.6% on 34 decisions. That is a small and recent cell, and no single ruling explains it: 32 of the 34 fall in 2025 and 2026, which is when plaintiffs began pleading the two provisions together. Much of the gap is arithmetic rather than doctrine. 17 of the 34 are recorded as mixed (50.0%), against 32.1% of the §631-only decisions and 10.1% of the pen-register-only ones, because a court that dismisses one count and lets the other proceed produces a mixed record by definition. Pleading two counts makes a clean defence win harder to record, whatever the merits. We record the cell rather than build on it.

A view often heard from defense practitioners runs the other way: §638.51 is hard to defend, §631 can be fought, because 'in transit' and 'contents' are elements a plaintiff needs a rare kind of expert to prove, so SB 690 should thin the docket. The record supports the second half of that and not the first. Decisions that plead only §638.51 end in a defense win 34.8% of the time, decisions that plead only §631 33.2%, and where in-transit reading is pleaded specifically the defense wins 28.4% on 141 decisions, below the 31.1% baseline. The elements are contested, and element failure is the largest ground in the record, but they are contested in both directions and do not produce a better rate. What §631 changes is cost, not odds: a claim that is harder to plead is also more expensive to defend, and a company that would rather settle than fund that defense pays either way. As long as it does, the letters keep coming.

▸ SO WHAT ACTUALLY CHANGES

SB 690 removes a pleading route and a damages count, not a *probability*.

The same fact patterns can still be pleaded under §631, with §631's elements to prove. Plan for the same fact patterns arriving as wiretap claims, and for a two-year retroactive window that reaches into matters already on file.

§631	69 of 103	2 yrs	Jan 2027
Left untouched by the final bill. Private claims under it stand	Pen-register decisions that plead no §631 count at all	Retroactive reach, to claims filed within two years of 1 January 2025	Effective date. Signed by the Governor on 30 September 2026

Two other things do not change. Vexatious-litigant designations bind one court and redistribute filings rather than ending them. And a strictly configured analytics-only deployment, with no content and no identifiers, is genuinely better placed. It is also a minority position on the estates we measure.

A MEASUREMENT PROGRAM

The first ninety *days*.

08

The work is a measurement program, not a policy exercise, and it is harder than it sounds. The object being measured, what third parties actually load in each state a visitor can be in, is not visible from any system of record. Nothing in a tag manager, a consent platform's configuration or a vendor contract records what actually fired. Only a browser does.

Days 0–30: measure

- **Enumerate the properties.** The unit is domains and apps, including the ones marketing bought without telling anybody: campaign microsites, regional storefronts, an acquired brand's estate still on its own stack. The property count is always higher than the room believes, and an unowned property is where the letter comes from.
- **Load each property the way a claimant does, and count.** First a cold visit, no interaction, no consent given: distinct third-party hosts, requests, and cookies or storage written before any consent decision. Then the interactions a claimant tries: text typed into a form or a search box without submitting, and a visit with GPC set. That host count is the multiplier in Archetype A and the only number in this report the company controls.
- **Inventory by interaction type, not just by page.** A pixel on a listing page and the same pixel on a form are different exposures. Enumerate the interactions that carry content: chat widgets, search boxes, symptom checkers, quote forms, account areas, anything that takes free text, and measure what fires during them, not just on load. Section 2.3 is why: content capture is what the party exception and the element arguments are fought over.
- **Classify what the data is, not only where it goes.** The sensitivity gradient in Section 2.6 is both the odds gradient and the price gradient. Health, intimate health and children's data change what is stacked alongside the CIPA count.

Days 30–60: gate, and prove the gate

- **Measure every consent state separately.** Before any choice, after acceptance, after rejection, after a granular partial choice, after a stored choice on a return visit, and while authenticated. A platform that is obeyed in the first state and ignored in the fourth is the 18-decision fact pattern with the worst outcomes in the record (16.7% defense wins).
- **Chase the second hop.** The vendor a tag loads is rarely the only vendor that ends up on the page. Piggybacked calls, ID-bridging partners and CDN-hosted scripts are outside any contract

the company signed, and no single team owns the answer.

- **Fix the disclosure to match the measurement.** The consent notice and the privacy policy should name the vendors and categories the measurement actually found. Disclosure alone does not win (34.4% on 32 decisions, level with the baseline), and a disclosure contradicted by a network capture is worse than a thin one. *Snyder v. G6 Hospitality LLC* (C.D. Cal. 2025) is what a disclosure that matches the conduct is worth.
- **Keep the evidence, dated.** The claimant's proof is a capture taken by an adverse party under conditions they chose. A dated internal record of what fired, when, and against which consent state is the only thing that answers it, and *Gutierrez v. Converse Inc.* (9th Cir. 2025) is what having that evidence is worth.

Days 60–90: procure and operate

- **Replace what cannot sit behind a gate.** A vendor that will not function after a rejection is a vendor whose risk the company is carrying permanently. This is the expensive move and the only one that removes a pattern rather than documenting it.
- **Make re-measurement continuous.** Third-party chains drift, marketing ships tags weekly, and a quarterly audit measures a configuration that no longer exists. This is an operating cost.
- **Watch the docket as well as the estate.** Check an incoming claimant's name against their filing history before anything else. Section 2.5 makes it the single most informative fact about a new matter.

> ***Most privacy teams can specify this program. Few can execute it across a real estate without dedicated measurement infrastructure, because every step above is a browser-level observation repeated across properties, consent states and time.***
>
> THAT IS THE WORK WE DO

THE CORPUS AND THE CODING

What is analyzed, and what is set *aside*.

CIPA is older than the web, and a fifth of what we hold is telephone and call-center recording, the part of the statute with the longest adjudicated history and an entirely different fact pattern. Reading it alongside pixel litigation would distort both. The report therefore analyzes internet-tracking decisions handed down on or after *Javier v. Assurance IQ* (9th Cir., May 31, 2022) whose surface is a website, a mobile app or software. The three groups set aside are reported with their own win rates, so the exclusion can be audited: each runs below the analyzed set, which means leaving them in would have made every number in this report look worse rather than better. 13 of the set-aside decisions are themselves post-*Javier* telephone and call-recording matters, so the volume of CIPA litigation since 2022 is larger than the 318 analyzed here. Voice AI and call-recording products are alleged in 104 records and cut across the groups rather than forming one of them: 92 are in the telephony group and 4 among the analyzed decisions.

FIGURE 11 · THE CORPUS

N = 493

What this report analyzes, and what it sets aside

Every record on file falls in exactly one of these four groups.



- **Internet tracking, decided since *Javier***
  318 records (65%) · 31.1% defense win · analyzed in this report
- **Internet tracking, decided before *Javier***
  51 records (10%) · 25.5% defense win
- **Telephone and call-center recording**
  100 records (20%) · 25.0% defense win
- **Surface not recorded**
  24 records (5%) · 20.8% defense win

The 493 records on file in four mutually exclusive groups: 318 internet-tracking decisions since Javier (analyzed), 51 internet decisions outside the analyzed period, 40 of them dated before Javier and the rest undated, 100 telephone or call-center recording matters, and 24 with no recorded surface.

A.2 A record is a decision, not a matter

One dispute can produce several decisions. A district ruling and the appeal that follows it are kept as two records on purpose: collapsing them would either discard the district court's reasoning or let the appellate outcome overwrite it, and both are evidence. Counts throughout are records; where the distinction changes a number, we say so. Duplicate rows for the same decision were merged before this draft: 39 clusters, 42 rows, and one settlement whose $21.5M had been sitting on four separate records.

A.3 Why every rate is split by who was sued

The same technology produces opposite outcomes depending on whether the brand running the website or the tracking vendor was the defendant, and every row moves the same way, so any rate published unsplit is the average of two opposite facts. Ad-tech and data-broker matters read 35.4% unsplit: 41.7% where the brand was sued and 17.6% where the ad-tech company was. That is why every conditional rate in Section 2 is quoted on brand defendants only, and why the sector figure in Section 2.7 excludes platform defendants.

FIGURE 12 · SPLIT BY DEFENDANT

BRAND VS VENDOR

The same technology, two opposite answers

Defense win rate by tracker category, split by who was sued. A row needs at least three decisions on each side to appear.

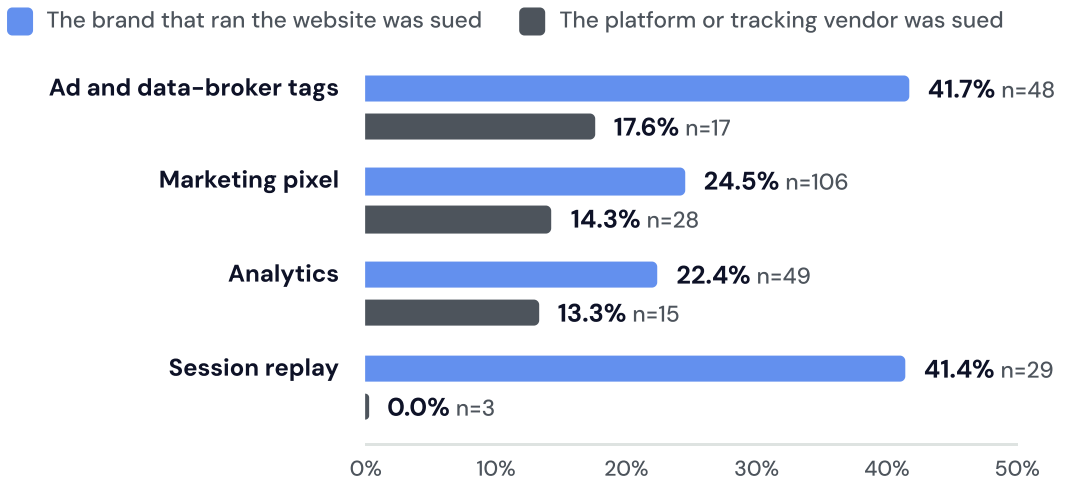

Defense win rate by tracker category, split by who was sued. A row needs at least three decisions on each side to appear. Any rate published unsplit is the average of two opposite facts.

A.4 Two rates that overstate their effect

— **Server-side tagging.** 21.1% on 19 decisions reads as an 11-point penalty. It is an artifact of what server-side tagging appears alongside. Server-side forwarding is a delivery technique, not a

vendor category, and no decision pleads it alone: all 19 plead it beside a marketing pixel, and 17 of them involve sensitive data. On the same sensitive-data pixel cases, the defense wins 23.5% where server-side is also alleged and 16.7% where it is not. The row has therefore left the category figure. Server-side hides the request from a browser capture; it does not change what was sent, and it plausibly does lower the chance of being selected, because both demand archetypes begin with somebody reading a network log.

- **Live chat, by forum.** The chat advantage is +14 points in federal court on 44 decisions and -7 in California state court on 12. Forum and plaintiff identity together account for most of it.
- **Consent postures.** Only separate written consent stands clear of the baseline. The other rows sit within a few wins of it and of each other, and their order moved when record outcomes were re-read from the latest ruling (see Method); Section 2.4 explains why two adjacent rows cannot be compared at all. Quote the top of the ladder, not the decimals.

A.5 Where these numbers come from

The corpus is 493 records of CIPA and related wiretapping litigation, assembled from court documents and public reporting. Membership is decided by what an extraction found in the document (a §631, §632 or §638 count in the claim mapping) rather than by a statute name in free text, because the latter is written by a model and whether it says the act or the code is not a fact about the case. Every decision named in this report was selected by query against the same strata as the statistics, not by hand, so an example cannot illustrate a claim the numbers do not support.

The four biases, stated in the direction they push

- **Public record only.** Settlements mostly do not reach a docket and demand letters never do, so volume is understated and outcomes are mispriced in a knowable direction.
- **Pleading-stage facts.** Most of what we know about any defendant's tracking is what a plaintiff alleged and a court assumed at the pleading stage, not what was proven.
- **Settlement coverage.** Only 19 of 42 valued settlements are among the decisions analyzed here, and a search found more that we do not hold; see Section 6.
- **Coded judgment.** Defendant industry and vendor role are hand-coded best-effort labels, not extracted facts, and are marked as such wherever they appear.

Method

Outcomes are classified from the court's own disposition language into five values, and a rate is defense wins divided by all decisions in the group. A decision that leaves the case undecided at the recorded posture is one of those five values and stays in the denominator; 26 of the 318 are of that kind. Where a record holds several rulings, its outcome is that of the latest ruling on the merits (a motion to dismiss, demurrer, summary judgment or appeal), read document by document; a later

sealing, scheduling or class-certification order does not replace it, and a settlement after the ruling does. Against the previous release this moved 23 of the 318 records to another outcome and left the headline rate at 31.1%. A ruling that dismisses with leave to amend counts as a defense win when no later ruling is on record; 42 of the 99 defense wins are of that kind, and the record does not show whether the plaintiff amended. A record with no disposition at all is excluded, not counted as a loss. The coding taxonomy was developed in collaboration with an independent legal academic, who reviewed a sample of the coded data and found the results trustworthy. Before publication every conditional rate was re-computed inside each stratum of six possible confounders (defendant role, surface, technology, era, sensitivity and forum), and any rate whose direction reversed between strata was either split or dropped. Plaintiff identity has been added as a seventh. Four findings in earlier drafts did not survive these tests and have been rewritten; the most recent is Section 2.6, where the data reversed an earlier claim that sensitivity changes the price but not the odds.

Rates on fewer than about thirty decisions are directional. Several rows in Section 2 sit below that line and are marked with their n wherever they appear; the ones that carry an argument in the text (the party exception, separate written consent, the tester finding) are quoted with their counts in the sentence rather than only in the figure. Settlement dollars come from publicly disclosed amounts, collected and checked by hand; counts and dates come from the court record; the two are never mixed inside one figure.

A.6 Decisions cited by name in this report

Every decision named anywhere in this report appears here. Those in the analyzed set were selected by query and pinned, so that a later change to the corpus fails the build rather than leaving a stale citation in the prose. Four are outside the analyzed set and are marked as such: two decided before *Javier*, one that carries no decision date in the corpus, and one decided under another state's wiretap statute.

- ***Javier v. Assurance IQ, LLC*** · 9th Cir., 2022-05-31. No. 21-16351 (9th Cir. mem. op.). The decision that opens the analyzed period. PLAINTIFF
- ***Williams v. What If Holdings*** · N.D. Cal., 2022-12-22. DEFENDANT
- ***Pena v. GameStop, Inc.*** · S.D. Cal., 2023-04-27. DEFENDANT
- ***Valenzuela v. Keurig Green Mountain*** · N.D. Cal., 2023-10-10. DEFENDANT
- ***Maghoney v. Dotdash Media Inc.*** · S.D. Cal., 2026-02-23. DEFENDANT
- ***Rodriguez v. Autotrader.com Inc.*** · C.D. Cal., 2025-03-14. DEFENDANT
- ***Rodriguez v. Ford Motor Company*** · S.D. Cal., 2024-12-03. MIXED

- ***Gutierrez v. Converse Inc.*** · 9th Cir., 2025-07-09. No. 24-4797 (9th Cir. mem. op.), D.C. No. 2:23-cv-06547-KK-MAR. DEFENDANT
- ***Snyder v. G6 Hospitality LLC*** · C.D. Cal., 2025-08-18. DEFENDANT
- ***Mikulsky v. Bloomingdale's, LLC*** · 9th Cir., 2025-06-20. No. 24-3564 / No. 24-3837 (9th Cir. mem. op.) / D.C. No. 3:23-cv-00425-L-VET. SETTLED
- ***Casillas v. Six Flags Entertainment Corporation*** · C.D. Cal., 2025-12-15. SETTLED
- ***Pattison v. Teladoc Health, Inc.*** · S.D.N.Y., 2025-06-25. Pattison et al. v. Teladoc Health, Inc., 23-cv-11305 (S.D.N.Y.). MIXED
- ***Ahringer v. LoanDepot, Inc.*** · C.D. Cal., 2024-02-06. SETTLED
- ***Greenley v. Kochava Inc.*** · S.D. Cal., 2023-07-27. SETTLED
- ***Moody v. C2 Educational Systems Inc.*** · C.D. Cal., 2024-07-25. PLAINTIFF
- ***Shah v. Fandom, Inc.*** · N.D. Cal., 2024-10-21. No. 3:24-cv-01062-RFL (N.D. Cal.). PLAINTIFF
- ***Popa v. Microsoft Corp.*** · 9th Cir., 2025. 153 F.4th 784. Outside the analyzed set: a session-replay claim under Pennsylvania's wiretap statute, not CIPA, which the corpus does not hold. Cited in Section 2.2 for the Ninth Circuit's application of *TransUnion* to website tracking. Outcome: dismissal for lack of standing, affirmed. DEFENDANT
- ***In re Facebook, Inc. Internet Tracking Litigation*** · 9th Cir., 2020-04-09. Outside the analyzed set: decided before *Javier*, and listed here because it is the reason the report dates the analyzed period from *Javier* rather than claiming *Javier* invented the theory. The Ninth Circuit had already held that §631(a) claims over internet tracking were plausibly pleaded. MIXED
- ***Shah v. Crain Communications, Inc.*** · C.D. Cal., 2026-07-20. No. 2:26-cv-03070-RGK-CTS. Outside the analyzed set: the corpus record carries no decision date. Cited for the prefiling order in Sections 2.5 and 5.4.

THE LIMITS ON EVERY NUMBER HERE

What this report does not *claim*.

B

Five limits apply to every number in these pages. They are stated here so that a reader can hold the findings at the strength the record supports and no further.

01 **That any of these rates predicts the chance of being sued.**
Every one is conditioned on litigation having started.

02 **That correlation is causation.**
A category's defense win rate reflects which claims plaintiffs bring against it, and which plaintiffs bring them, as much as how defensible it is. Section 2.5 is a worked example of exactly that.

03 **That we can price a defense.**
We hold no cost data and publish none.

04 **That the settlement distribution is complete.**
Section 6 says how incomplete, and in which direction.

05 **That any of this is legal advice.**
It is a description of a documented record.

TALK TO VAULT JS

Count the hosts before a claimant *does*.

Vault JS measures what actually fires on your web, mobile and CTV properties: every third-party host, in every consent state, dated and evidenced. That record is the one thing that answers a claimant's network capture, and it has to exist before the letter arrives.

To learn more, visit vaultjs.com →

© 2026 Vault JS · This report describes a documented record. It is not legal advice.