ADCRIME / THE PLAYBOOK / V1.0

The Affiliate Brand-Bidding Fraud Playbook

Detection, recovery, and prevention for affiliate managers at $10M–$500M brands.

Written by Mario Vaher, Founder, AdCrime

Published April 2026 · Staromeda OÜ, Estonia

Length ~15,000 words · 12 sections · Free, ungated

BEFORE YOU START

A note from the founder.

I wrote this Playbook because the version I needed did not exist.

I run AdCrime. We operate a global scanner network that watches affiliate programs for unauthorized brand bidding — the specific, expensive, well-documented fraud that happens when one of your affiliates decides your brand name is a free keyword they can rent from Google. Every program we have ever scanned has contained unauthorized bidders. Every single one. That is not an exaggeration for marketing purposes. That is the number.

So I sat down to write a single document an affiliate manager could read in one sitting and walk away with everything they need — the mechanics of the fraud, the economics of why it happens, how to detect it, how to claw back commissions, how to build a program that makes the fraud harder in the first place, and what the actual legal landscape looks like in 2026 versus what the networks pretend it looks like.

I am not a lawyer. I am a founder who has spent the last year looking at affiliate programs under a microscope and I have seen things that would make the finance team at any mid-market DTC brand lose sleep. This is what I wish somebody had handed me on day one.

One more thing. I run a company that sells a solution in this space — but the Playbook does not require you to buy anything from me. Every tactic here can be executed with tools you already own. If you get to the bottom and decide you want help, you know where to find me. If you handle it yourself, that is also a win — because the only outcome I actually care about is one fewer fraudster getting away with it.

Let's go.

— Mario

FOUNDER, ADCRIME

CONTENTS

What's inside.

SECTION 1

The size and shape of the problem

Affiliate fraud is not a rounding error. It is the largest unmonitored leak in the performance marketing stack.

The CHEQ and University of Baltimore study puts the cost of affiliate fraud at roughly $3.4 billion per year in the United States alone, up from $1.4 billion in 2020. That is a 143% increase in 24 months. Juniper Research's broader ad fraud estimate landed at $84 billion for 2023, projected to hit $172 billion by 2028. Those are headline numbers, and I have the same instinct you do when you see them — "vendor research, take it with salt." But the trend direction is not in dispute. Fraud is up. Fraud is up faster than brand spend. And the specific subcategory of brand-bidding fraud is up faster than the rest of the fraud.

BforeAI's March 2025 analysis of a 33,000-ad sample found that 6.9% of ads on branded queries in their dataset were from unauthorized affiliates, and that 31% of the brands in the sample had at least one affiliate bidding their terms. That is the more interesting number for you, because it is per-brand. It means roughly one in three affiliate programs has a live infection at any given moment, whether the brand knows it or not.

At AdCrime we operate a global scanner network with regional hubs across the Americas, Europe, and Asia-Pacific, plus the ability to reach into any specific market a client's customers actually come from. Our detection rates on programs that have never been audited before are consistently higher than the BforeAI number. The reason is geographic coverage — if you only scan from one country, you miss the fraudsters running geo-targeted campaigns designed to hide from the brand's local IP. The more vantage points you scan from, the more fraud you surface. This is not complicated. It is operationally expensive, which is why most in-house teams never build it and why most vendors cut corners on it.

Why brand-bidding fraud has accelerated since 2020

Four things happened between 2020 and 2024 that made this problem structurally worse.

One — Google changed its trademark enforcement policy. On July 24, 2023, Google shifted away from blanket trademark protections to a complaint-specific enforcement model. The practical effect is that unless you file a specific trademark complaint against a specific advertiser, Google will let the bidding happen. The burden shifted from Google's automated filters to your legal team. Most brands did not notice. Fraudsters noticed immediately.

Two — Third-party cookies started dying. Chrome's cookie deprecation timeline slipped repeatedly, but the strategic signal was received in 2022. Marketing budgets started moving from cross-site retargeting toward observable last-click channels, and affiliate is the most observable last-click channel there is. More budget flowed into affiliate. More attention from people who wanted to capture a percentage of it.

Three — Coupon behavior metastasized. Industry surveys consistently show that a majority of U.S. consumers search for a discount code before completing a checkout. That is not a fraud statistic in itself. It is the atmospheric condition that makes coupon-intercept fraud profitable. Every consumer who opens a new tab to search "[brand] coupon" is a potential attribution hijack, and the coupon site industry has been optimizing itself around that moment for 15 years.

Four — AI content generators made fraud infrastructure 100x cheaper to build. In 2020, spinning up 50 cloaked landing pages required a small team and three weeks. In 2025, one person with a ChatGPT API key and a $20 WordPress theme can spin up 5,000 pages in an afternoon. The QuoIntelligence report from March 2025 documented a NordVPN-targeted fraud ring running AI-generated content across more than 1,000 subdomains on seven primary domains, all tracked back to a single affiliate ID. One operator. Thousands of pages. Millions of dollars in disputed commissions.

These four forces are still active. None of them are reversing. The volume of fraud will keep climbing until the cost of executing it rises, and the cost will only rise when detection improves. That is the thesis. That is why I built AdCrime. That is why this Playbook exists.

The verticals getting hit hardest

Every vertical is a target, but the economics favor fraudsters in a handful of specific categories. If you run affiliate in any of these, your prior should be that you already have an active infection:

- **VPN, hosting, and SaaS with recurring revenue** — because the per-sale commission is often $50–$100 and the CPCs on branded queries are under $5. The arbitrage math is obvious.
- **Gambling and sports betting** — because the affiliate commissions are a percentage of lifetime value, which is lucrative, and the regulatory environment already tolerates ambiguous behavior.
- **Personal finance and lending** — because one qualified lead can pay $200+. CHEQ data has flagged personal finance as one of the highest-fraud verticals, meaningfully above the cross-industry baseline.
- **DTC fashion, beauty, and wellness** — because the coupon-search behavior is near-universal in these categories. Nordstrom lost approximately $1.4 million to a single cash-back portal exploiting a system error — the two brothers involved pled guilty to wire fraud (FBI, 2012).
- **Travel and booking** — one unnamed travel site has been reported as spending $7 million per month on affiliate acquisition, a non-trivial portion of which industry observers believe was being siphoned through attribution fraud.

If you are running a program in any of these categories and you have not audited it in the last 90 days, you are almost certainly paying for commissions you should not be paying for. The only question is how much.

$3.4B

U.S. affiliate fraud losses, 2022 (CHEQ / University of Baltimore)

1 in 3

Brands with at least one affiliate bidding their terms (BforeAI, March 2025)

100%

Affiliate programs AdCrime has scanned that contained unauthorized brand bidders

SECTION 2

The five fraud patterns

Brand-bidding fraud is not one thing. It is five distinct operational patterns, each with its own mechanics, its own detection signature, and its own recovery approach. Understanding which pattern you are fighting is the first step to fighting it correctly.

Pattern 1 — Direct brand-term bidding

What it is. An affiliate buys Google Ads on your exact brand name — "nordvpn," "monday.com," "asana login" — and collects a commission on every user who clicks through and converts. In the purest version, the affiliate is not adding any value whatsoever. They are inserting themselves between a customer who was already going to buy from you and your checkout page.

How it works. The affiliate's ad shows up above or next to your own organic listing. The ad copy often does not even disguise what is happening — it just shows your brand name and a discount promise. The user clicks, hits an affiliate-tagged redirect, and lands on your site. The sale completes. The affiliate earns 10–30%. You paid the Google CPC for them to intercept your own customer.

The evidentiary fingerprint. This is the easiest pattern to detect. Run a Google search for your brand name from a clean browser. If you see any paid ad above your own organic listing that is not your own ad, you are looking at direct brand-term bidding. Cross-reference the destination URL with your affiliate network's click tracking and you have the publisher ID.

Real case. In the Sage Financial Software incident documented in industry trade press, a single affiliate was found to be running 89,000 brand infringements across the Sage program. Once Sage identified and shut down

the activity, their own branded CPC dropped 75% in five days. That 75% number is the cost that was being invisibly imposed on Sage by a single bad actor.

Pattern 2 — Trademark-plus-modifier bidding

What it is. A sneakier version of Pattern 1. Instead of bidding your exact brand name (which is easier for you to catch), the affiliate bids on your brand plus a modifier — "[brand] coupon," "[brand] discount," "[brand] review," "[brand] login," "[brand] vs [competitor]." These queries are often technically allowed under program terms that only prohibit exact-match bidding on the brand name, and fraudsters have gotten very good at finding the loophole.

How it works. The affiliate bids on "[brand] coupon" or "[brand] review." Their ad shows up. The user is in late-funnel, discount-seeking mode. They click. They land on a thin review page or a coupon page with a "Reveal code" button. The button fires the affiliate cookie. The user goes back to your site to buy. Commission earned, nothing of value created.

The evidentiary fingerprint. Run branded searches with the common modifiers appended — "review," "coupon," "discount," "promo code," "login," "alternative," "vs," "pricing." Note which ads appear. Pull the destination URLs. Cross-reference with affiliate tracking. The affiliate with the highest click-through on modified queries relative to raw branded queries is the one bidding modifiers.

Why this pattern is harder to shut down. Courts have gotten progressively less friendly to pure trademark-plus-modifier claims. *1-800 Contacts v. JAND* (2d Cir., October 8, 2024) held that keyword purchase alone is "permissible and standard." *Lerner & Rowe v. Brown Engstrand* (9th Cir., October 22, 2024) found that 236 confusion instances out of 109,322 impressions — 0.216% — was de minimis and not actionable. These decisions do not kill your case, but they mean you cannot rely on generic trademark law. Your leverage has to come from your program terms (see Section 9) and from the FTC's material connection disclosure rules (see Section 7).

Pattern 3 — Typosquatting and lookalike domains

What it is. The affiliate registers a domain that looks almost like yours — "amaz0n-deals.com," "nord-vpn-discount.net," "mondaycom.app" — and runs ads on branded queries that lead to the lookalike domain first, where they drop their cookie, then redirect to your real site.

How it works. The user searches "nordvpn discount," clicks an ad that says "NordVPN Official Discount," lands on nord-vpn-discount.net, sees a "Continue to NordVPN" button, clicks it, and lands on your real site with the affiliate cookie set. To the casual user the entire flow looks like it was your brand. To your attribution system it looks like a legitimate affiliate referral.

The evidentiary fingerprint. Run your brand through a typosquat detection tool. Check WHOIS on any suspicious domain — registration in the last 60 days, privacy-protected contact, registered in an offshore jurisdiction are all strong signals. Use `archive.org/wayback` to check the domain's history. Typosquat domains almost always have no history. The most reliable signal of all is a domain that resolves to an affiliate-style landing page you did not approve.

Recovery option. Typosquat domains are the one area where legal enforcement is fast and reliable. File a UDRP complaint with WIPO or the National Arbitration Forum (see Section 7). Costs around $1,500, takes 45–60 days, almost always succeeds when the domain is objectively a typosquat.

Pattern 4 — Cloaking

What it is. The fraudster shows a clean landing page to Google's review bots and the brand's compliance team, but a completely different page to real users — the real page being either a hard-sell discount harvester, a fake coupon wall, or a straight attribution hijack.

How it works. The cloaking infrastructure identifies the visitor based on IP, user agent, cookie state, click source, geo, time of day, and referrer. A visit from a Google crawler sees Version A — clean, compliant, boring. A visit from an affiliate manager's IP range sees Version A too. A visit from a real consumer in a target market during peak shopping hours sees Version B — the one that drops the cookie and harvests the sale. Detection systems that check the URL once and move on will miss it 100% of the time.

The evidentiary fingerprint. Cloaking requires geographic detection to catch. You need to fetch the landing page from multiple countries, from multiple user-agent profiles, during multiple times of day, and compare the results. mFilterIt research has reported that roughly 45% of the ad fraud they observe involves some form of cloaking. This is the single biggest reason single-region scanning gives brands false confidence — your compliance team runs a check from the headquarters office and sees nothing, because the fraudster cloaked them.

Why this matters for your scanning strategy. If your current detection is a marketing ops person running a manual SERP check from a Mountain View IP address once a week, your detection is effectively zero. You will never see the ads that are only shown to consumers in Bangkok on Friday evenings. You will never see the ads that only show between 8pm and midnight on payday weeks. The fraudsters know this and they design around it.

Pattern 5 — Attribution hijacking via browser extensions and toolbars

What it is. The fraudster does not buy Google Ads at all. Instead, they run a browser extension, toolbar, or coupon-finder plugin that users install voluntarily. When the user navigates to any merchant site, the extension quietly fires an affiliate cookie, overwriting whatever attribution was previously in place.

How it works. The user installs "Honey" or a competitor. User visits your site through whatever channel — Google organic, email, social ad, direct type-in, whatever. At checkout, the user clicks "apply coupon" in the extension. The extension fires an affiliate cookie. The sale attributes to the extension's affiliate publisher ID. The actual channel that drove the sale — the channel you paid for — gets no credit. The brand pays the affiliate commission on top of whatever it paid to acquire the customer the first time.

The Honey/PayPal scandal. This pattern blew up publicly in December 2024 when the MegaLag YouTube investigation (13 million views and counting) documented that Honey — now owned by PayPal — was systematically replacing last-click attribution on merchant sites where Honey held an affiliate relationship. The reported economics were stark: Honey was collecting $35 commissions on transactions where the user had received a $0.89 reward, or in many cases no reward at all. As of the In Re PayPal Honey case (5:24-cv-09470, N.D. Cal.), the litigation is active. Judge Beth Labson Freeman dismissed the first amended complaint in late 2025; a

second amended complaint was filed in January 2026 with merchant contracts and specific commission examples. As of April 2026, the case remains active.

What changed. In March 2025, Google updated Chrome Web Store policy to require clearer disclosure of affiliate relationships in browser extensions. Rakuten and several other extension operators have adjusted their practices. This category of fraud is not solved, but it is under more scrutiny than it has been in a decade.

The evidentiary fingerprint. This pattern is hard to catch from SERP scanning alone because there is no ad on the SERP — the extension is doing the work client-side. The signal shows up in your attribution data: one or two affiliate publishers with impossibly high conversion rates, shorter-than-usual click-to-conversion windows, and a disproportionate share of commissions from users who also touched your brand through non-affiliate channels in the same session. Cross-reference your affiliate attribution with your server-side analytics. Anomalies there are the tell.

SECTION 3

Why the networks don't catch it

There is a question you should ask yourself at this point, and the answer is important for everything that comes after.

If brand-bidding fraud is this widespread, this documented, and this profitable to run, why don't the affiliate networks just stop it?

The networks have the data. They own the click logs, the publisher IDs, the conversion records. CJ, Impact, Awin, ShareASale, PartnerStack — each of them could, in theory, build a real-time detection layer that flagged unauthorized brand bidding the moment it happened. None of them have done so in any serious way. Why?

The honest answer is that the networks have structured their business models around not noticing.

The revenue model

Every major affiliate network is paid as a percentage of commission volume flowing through its platform. CJ takes an override of roughly 20–30% on every commission processed. Awin charges a tracking fee plus a percentage. PartnerStack and Impact have their own variations on the theme. The economics are simple: **every fraudulent commission that flows through the network is also a commission the network is taking its cut on.** Detection costs money. Letting the commission through costs nothing and generates revenue. This is not a conspiracy theory, it is just how the incentives are pointed.

Ben Edelman — the Harvard researcher who has been studying affiliate compliance longer than anyone — has documented this structural conflict across multiple papers. The most pointed observation, which I am

paraphrasing because I cannot reproduce the original at length: for one large client, he personally proved that nine out of the client's top ten affiliates were breaking program rules, and the network had not flagged any of them. Nine out of ten. Top ten by volume. That is not a detection failure. That is a detection refusal.

The terms of service

Read any major affiliate network's publisher agreement carefully and you will find language that shifts the responsibility for compliance enforcement almost entirely onto the brand. Awin's Clause 5.7 is the most aggressive on record: it effectively limits advertiser recovery rights for commissions paid on approved transactions. CJ, Impact, and ShareASale have slightly softer versions of the same idea. The pattern is consistent: the network collects the fee, and the brand carries the risk.

When you push a network on this during a dispute, the language in their reply will always come back to program terms. "Our platform merely facilitates the relationship. Enforcement of your program rules is your responsibility." This is legally defensible and operationally convenient. It is also the single biggest reason brand-bidding fraud has been allowed to metastasize.

The visibility argument the networks will make

In fairness to the networks, they will argue — correctly, up to a point — that they cannot see what happens upstream of the affiliate click. They cannot see the Google Ads campaign the affiliate is running. They cannot see the cloaked landing page. They cannot see the browser extension firing client-side. All they see is a click and a conversion. Everything before the click is invisible to them.

This is true. But it is also beside the point. The networks have two signals they could act on even without upstream visibility:

1. **The conversion rate anomaly.** A publisher with a 40% conversion rate on branded queries is almost certainly either typosquatting, cloaking, or attribution-hijacking. The networks can see this.
2. **The explicit complaint.** When a brand files a dispute with evidence, the network can process it and terminate the partner. They can do this today.

Most of them drag their feet because processing the dispute costs them the commission they would have earned.

The structural point is that the networks have built a business where they get paid regardless of whether the commission is legitimate, and where processing disputes costs them money. You cannot fix an incentive problem by hoping the counterparty rediscovers their ethics. You can only fix it by making the detection and the disputes your own responsibility.

That is the thesis of this Playbook, and it is not a comfortable thesis. But it is the one that matches reality.

> **"I once proved that nine of a large client's top ten biggest affiliates were breaking its rules."**

SECTION 4

How the math destroys the brand

Here is the part every CFO needs to read.

The direct cost of affiliate brand-bidding fraud — the commissions you pay to fraudsters — is only the tip of the iceberg. The hidden costs compound across at least four different line items, and in aggregate they are larger than the commission leakage by a factor of three to five.

Let me walk through a worked example. This is a generic $30M ARR DTC brand scenario, not a specific customer, but the shape of the math holds for any program in the $10M–$500M range.

Assumption set. One fraudster affiliate. They are running direct brand-term bidding plus a trademark-plus-modifier layer. They generate $20,000 per month in affiliate commissions at a 15% commission rate. Total flow: $133,000 per month in attributed gross merchandise value routed through this one affiliate.

Cost one — Commission leakage. The $20,000 you paid in commissions this month. This is the only number most finance teams see. In this example, the cost per year is $240,000.

Cost two — Organic cannibalization. Of the $133,000 GMV this affiliate claims credit for, how much would have converted anyway through organic search, direct type-in, or branded campaigns you were already running? In my experience, the answer is between 85% and 95%. The customers were going to buy from you regardless. Let us say 90%. That means $120,000 of the $133,000 is revenue you would have captured anyway, free. The $20,000 in commissions is not buying you $133,000 of incremental revenue — it is buying you at most $13,000. Your real ROI on this affiliate is negative $7,000 per month. Annualized: negative $84,000.

Cost three — Inflated paid-search CPCs. This is the cost most finance teams never catch. When the fraudster bids on your branded keywords, they drive up the auction price for those keywords. Your own branded search campaign now pays higher CPCs to maintain position one. The Sage Financial Software case documented a 75% CPC drop after the fraudster was shut down, meaning the fraudster had been inflating Sage's own branded search costs by 300% during the infection. In our $30M example, assume the brand was spending $15,000/month on branded paid search at inflated CPCs, and would have spent $8,000/month at natural CPCs. That is $7,000/month in additional paid media waste. Annualized: $84,000.

Cost four — Attribution distortion and bad decisions. The fraudster's fake "high performance" numbers now show up in your mix modeling. Your marketing team sees affiliate as a high-ROI channel and shifts budget into it. That budget shift funds more fraud. The feedback loop is vicious and invisible. The opportunity cost of budget misallocation is hard to quantify but easy to observe after the fact, because every brand that has cleaned up its affiliate fraud has simultaneously discovered that some of its "best" channels were lies. Kevin Frisch, formerly of Uber, has publicly reported that Uber turned off two thirds of its ad spend — $100 million out of $150 million — and observed "basically no change" in rider app installs. That is a $100 million attribution lie.

Total annualized cost from one fraudster affiliate in this example: $240K direct + $84K cannibalization + $84K CPC inflation + unknown attribution distortion = **approximately $408,000 per year, before you count the attribution damage.**

This is why I cannot take "affiliate is only a small part of our spend" seriously as an argument for not auditing. The impact is not proportional to the affiliate budget. It is proportional to the fraction of your organic and branded revenue the fraudster is inserting themselves into. A fraudster can extract more money from a brand that spends $0 on affiliate than the brand spends on the affiliate channel itself.

Multiply by the typical infection rate — eight to twelve active infringing affiliates on a mid-market program — and the seven-figure annual loss becomes the median case, not the worst case.

$408,000

Estimated annualized cost to a $30M ARR brand from a single infringing affiliate, fully loaded

SECTION 5

Detection: the evidentiary fingerprint

Every fraud pattern in Section 2 leaves a signature in data you already have access to. You do not need a specialized tool to start — you need to know where to look. This section is the translation key.

For Pattern 1 (Direct brand-term bidding): - A paid ad on your exact brand query that is not yours - Destination URL contains an affiliate tracking parameter - Auction Insights shows an unfamiliar domain consistently in the top five - Your own branded CPC has drifted upward over the last 60 days with no explanatory change in your bid strategy

For Pattern 2 (Trademark-plus-modifier): - Branded queries with "review," "coupon," "discount," "alternative," "vs" modifiers show affiliate ads - One affiliate is driving disproportionate commissions from queries that imply late-funnel intent - Conversion rate for that affiliate is 3–5x your program average

For Pattern 3 (Typosquatting): - WHOIS lookup on any unfamiliar domain shows recent registration - `archive.org/wayback` returns no history - Redirect chain goes lookalike domain → your real domain → affiliate cookie set - The lookalike domain runs ads on the same queries as Pattern 1 or 2

For Pattern 4 (Cloaking): - You cannot reproduce the fraudulent ad from your own IP, but your customers are reporting seeing it - SERP scans from different geographic regions return inconsistent results - The landing page you see from a headquarters IP looks clean; the landing page a customer reports seeing is a coupon hard-sell or a branded-lookalike page

For Pattern 5 (Attribution hijacking): - One or two affiliate publishers have conversion rates above 25% - Click-to-conversion windows are suspiciously short (under 2 minutes for a considered purchase) - These same

publishers overlap heavily with customers who also touched your brand through non-affiliate channels in the same session - Commissions from these publishers spike after you run any non-affiliate brand campaign

The single most powerful detection signal across all five patterns is **geographic inconsistency**. A clean SERP from your home country combined with fraud hits from other countries means you have cloaking or geo-targeted bidding. If you only scan from one region you will never see the pattern. At AdCrime we operate a global scanner network spanning the Americas, Europe, and Asia-Pacific, with the ability to reach into any specific market a client cares about — if 40% of your revenue comes from Canada or Australia, that is where your next scan runs from. You can replicate the approach manually with VPN rotation if you have the discipline to do it every week. Most teams do not have the discipline. That is fine — it is a thing I built software to solve because it was the thing I kept seeing teams fail at.

SECTION 6

The 30-day detection playbook

You have read enough theory. Here is the runbook.

This playbook assumes you run a mid-market program with three to six affiliate networks and that you have not done a serious audit in the last quarter. Execute it over 30 days. Do it once as a cleanup exercise, then fold the weekly cadence into your prevention program in Section 9.

Week 1 — Define the rules

Day 1. Pull your program terms document from every network you run. Read them. Highlight every clause related to keyword bidding, ad copy, landing pages, and trademark use. If your program terms are vague or missing these clauses, stop and go to Section 9 before continuing. You cannot enforce a contract that does not exist.

Day 2. Build your prohibited keywords list. Your brand name in exact, plural, possessive, and misspelled forms. Your brand plus every modifier you care about — "review," "coupon," "discount," "promo code," "login," "alternative," "vs," "versus," "competitor," "price," "cost," "refund," "cancel," "support," "official." Your product names. Your domain names. Put all of this in a spreadsheet.

Day 3. Define your geographic coverage. At minimum, you need three regions — your home market, one European market, and one Asia-Pacific market. More is better, and the right number is determined by where your actual customers are. If 40% of your revenue comes from Canada or Australia, you need to be scanning from those markets too. Expand the list until you are covering every geography that matters to your business.

Day 4. Set up your scanning environment. This can be as simple as a rotating VPN connection, a private browsing window, and a spreadsheet for results. It can be as elaborate as a cloud-based scraping stack. For the first audit, simple is fine.

Day 5. Run your first baseline scan. Go through your full prohibited keyword list across all geographic regions. Screenshot anything that looks off. Do not investigate yet. Just collect.

Week 2 — Map the evidence

Day 6. Go back through the screenshots from Day 5. For every affiliate ad you did not approve, capture the destination URL by right-clicking the ad and selecting "copy link address." Save these URLs.

Day 7. For each destination URL, trace the redirect chain. Paste the URL into `httpstatus.io` or use browser developer tools. Record the full chain from first click to landing page, including any affiliate tracking parameters.

Day 8. Cross-reference each redirect with your affiliate network dashboards. Look for the publisher ID that matches the tracking parameter. If the URL contains `aff_id=12345`, that is the publisher you are looking for. Record the matches.

Day 9. Pull the last 60 days of transaction data for each identified publisher. Note their commission volume, conversion rate, and click-to-conversion windows. Compare against the program average.

Day 10. Build your evidence package for each confirmed violator. Use the ten-item template in Section 8. Do not file disputes yet — you are still building the case.

Week 3 — Enforcement

Day 11. File dispute tickets on every confirmed violator, using the template in Section 8. Do not batch these — file each one individually so they cannot be bulk-dismissed.

Day 12. Email your network account managers with a short heads-up: "We filed X disputes this week, expect to see them in your queue, here is the pattern, we would appreciate expedited review." Copy the dispute ticket numbers.

Day 13. For any publisher driving significant volume, send a direct demand letter in parallel with the network dispute. See Section 8 for the content.

Day 14. Update your affiliate program terms if Week 1 revealed gaps. You will file the updated terms with every network and require acceptance for continued participation. This is the single highest-leverage action in the whole playbook.

Day 15. Document everything. Write a one-page memo for your CFO summarizing what you found, what you filed, and what the recoverable amount looks like. This memo is what gets you the budget to build the prevention program.

Week 4 — Close the loop

Day 16. Follow up on any network disputes that have not been acknowledged within 72 hours.

Day 17. For any network that is slow-walking the disputes, escalate to the account manager or compliance team.

Day 18. Begin standing up the weekly cadence from Section 9 — Monday SERP scan, Tuesday Auction Insights, Wednesday transaction anomalies, Thursday dispute status, Friday finance summary. The first time through takes a full day. By week three it will take 90 minutes.

Day 19. Reconcile commission reversals with finance. Make sure the money actually came back to your account — some networks will mark a dispute "resolved" without actually processing the reversal, and it is on you to catch it.

Day 20. Write the Prevention Program Review — the single-page document that closes out this audit and sets the baseline for next quarter.

Everything after Day 20 is maintenance. The audit is done. You know who was stealing from you, you have filed the disputes, you have updated the terms, and you have the weekly cadence running. This is the state every affiliate program should be in and almost none of them are.

SECTION 7

The legal landscape in 2026

A short tour of the rules you can lean on. This is not legal advice and you should talk to real counsel before you file anything. But it is the context you need to decide when to escalate.

Federal trademark law — the Lanham Act

The Lanham Act (15 U.S.C. § 1125) is the main federal vehicle for trademark enforcement in the United States. Under the Act, a trademark owner can sue for infringement if another party's use of the mark is likely to cause consumer confusion. This is the legal theory underneath most affiliate brand-bidding cases.

The current state of the case law. The last two years have been rough for brands trying to bring pure keyword-purchase cases. *1-800 Contacts v. JAND* (Second Circuit, October 8, 2024) held that keyword purchase is "permissible and standard," echoing earlier decisions. *Lerner & Rowe v. Brown Engstrand* (Ninth Circuit, October 22, 2024) held that 236 confusion instances out of 109,322 impressions was de minimis and not actionable. The Second and Ninth Circuits together cover a huge share of the country, and their signal is unambiguous: you cannot win a Lanham Act case on the theory that a competitor bought your keyword.

What still works. Brand-bidding fraud is almost never *pure* keyword purchase. It involves trademark use in ad copy, display URLs designed to confuse, cloaked landing pages that impersonate the merchant, typosquat domains, or attribution hijacking. Each of those pushes the case back toward classical likelihood-of-confusion and away from the permissive precedents. If you have evidence that the affiliate used your brand name in ad copy without

authorization, or ran a landing page that mimicked your site, or typosquatted your domain, you have a real case. If all you have is "they bought our keyword," you have a harder case than you had two years ago.

The FTC and the Endorsement Guides

The Federal Trade Commission updated its Endorsement Guides in July 2023 to require clearer disclosure of material connections between promoters and merchants. The Guides are not a private right of action — you cannot sue an affiliate under them — but they create an enforcement lever the FTC can pull, and they create a documentation standard you can use in your disputes and demand letters.

Recent FTC enforcement activity you should know about. *FTC v. Traffic and Funnels* (June 2024) resulted in a $1 million settlement for affiliate marketing deception. *FTC v. Total Wealth Academy* (August 2024) established individual liability for promoters of deceptive affiliate schemes. The FTC has also updated the Fake Reviews Rule (August 2024) and is actively enforcing against affiliate coupon sites that fail to disclose their relationships. The trend line is in your favor if you are a brand being defrauded by an affiliate running hidden schemes.

State-level consumer protection

Most U.S. states have unfair and deceptive practices statutes that reach further than federal trademark law. California's Business and Professions Code § 17200 is the most well-known, and it is explicitly broad — "any unlawful, unfair, or fraudulent business act or practice" is actionable. Illinois, New York, Florida, and Washington all have analogous statutes. State claims are often more useful than federal claims for brand-bidding fraud because the bar for "deceptive" is lower than the bar for "likelihood of confusion."

State claims also open the door to small claims court for modest amounts — typically $5,000 to $15,000 depending on the state. If the affiliate is a U.S. LLC or individual and the amount they stole is under the small claims limit, you can file yourself without a lawyer. Most affiliates do not show up for small claims hearings and you win by default. (Limits range from $2,500 to $25,000 depending on the state.)

The Honey / PayPal litigation as a bellwether

The In Re PayPal Honey class action (5:24-cv-09470, N.D. Cal.) is the case everyone in the affiliate space is watching. The core allegation is that Honey systematically replaced last-click attribution on merchant sites where Honey held an affiliate relationship, collecting commissions on sales the user had already decided to make through other channels. Judge Beth Labson Freeman dismissed the first amended complaint in late 2025, and a second amended complaint was filed in January 2026 with merchant contracts and specific commission examples. As of April 2026, the case remains active.

Why it matters: if the second amended complaint survives and the case moves into discovery, the factual record it produces will reshape how every brand negotiates with browser-extension affiliates for the next decade. Watch it closely.

WIPO and UDRP for domain disputes

If the fraudster has registered a typosquat domain, you do not need federal court. The World Intellectual Property Organization (WIPO) and the National Arbitration Forum both run domain dispute procedures under the UDRP (Uniform Domain-Name Dispute-Resolution Policy). A UDRP filing costs roughly $1,500, takes 45–60 days, and almost always succeeds when the domain is objectively a typosquat or a lookalike of a registered trademark. This is the single best legal tool in the affiliate fraud toolkit because it is fast, cheap, and targeted.

The offshore problem

If the affiliate is registered in Russia, Belize, the Seychelles, or any jurisdiction where civil judgments are unenforceable, none of the above applies. You cannot sue them. You cannot collect from them. Mutual Legal Assistance Treaty (MLAT) requests take 6–18 months and are reserved for criminal matters.

The tactic that actually works in this scenario is **attacking the network layer instead of the fraudster layer.** Every offshore fraudster needs a U.S.- or EU-accessible affiliate network to monetize the fraud — CJ, Impact, Awin, ShareASale, PartnerStack. Those networks sit inside jurisdictions where your legal options are real. File a formal notice to the network identifying the affiliate, the scheme, and your evidence. Make clear that continued

monetization constitutes knowing facilitation. Request termination and refund. If the network refuses, the network is now the defendant. In practice, they will not refuse at that point, because no reputable network wants to be the CJ Affiliate or Impact.com headline in a trademark lawsuit. They will terminate the partner and refund you to make the problem go away.

SECTION 8

Recovery and clawback

The recovery window is shorter than you think

Most brands discover brand-bidding fraud the same way. Somebody in finance asks why branded search CPM jumped 22%. Somebody in affiliate asks why one partner is suddenly closing 40% of the program's revenue. Somebody in legal forwards a cease-and-desist from a competitor whose ads now show up on your own brand terms.

By the time any of those questions get asked, the money is already in motion. And the money does not wait.

Every affiliate network you work with has a **commission locking window** — the period between when a sale is recorded and when the commission becomes legally owed to the affiliate, after which the network's willingness to reverse it drops from "routine" to "impossible." Understanding these windows is the entire game. Everything else in this section is downstream of knowing when your clock stops.

NETWORK	TYPICAL VALIDATION PERIOD	WHEN COMMISSION LOCKS
Impact.com	1–90 days (brand sets it)	Month-end following validation
CJ Affiliate	~30 days typical	20th of the following month
Awin	30–60 days standard	Per Clause 5.7 of Publisher Agreement
ShareASale	Up to 60 days	End of merchant-set locking period
PartnerStack	30 days typical	Month-end following validation
Refersion	Brand-configured	Brand-triggered

Locking windows vary by advertiser agreement — verify the exact terms with your network rep before filing disputes.

Read the numbers carefully. The best-case scenario is you catch the fraud the same month it happened and reverse the commission before month-end. The worst-case scenario is you catch it 45 days later and the commission is locked, paid out, and in a bank account you cannot touch.

This is also why the prevention program in Section 9 matters more than the recovery playbook in this one. **The cheapest commission to claw back is the one you never paid.** But assuming you are reading this because you already paid some you should not have, here is how to get them back.

The four layers of recovery

Recovery is not one action. It is four layers that you work through in order, escalating only when the previous layer fails. Skipping layers wastes leverage and signals to the network that you do not understand their process.

Layer 1 — In-platform commission reversal. This is the only layer that is fast, cheap, and almost always successful if you are inside the locking window. Every major network has a native flow for "dispute this transaction" or "reverse this commission," and they will honor it as long as you give them a reason, evidence, and a transaction ID.

What you need to submit:

1. Transaction ID from the network's own reporting
2. Affiliate publisher ID (not the display name — the numeric or alphanumeric ID)
3. Time window of the violation (not just the one transaction — show the pattern)
4. Evidence screenshots from at least three geographic regions, timestamped, with the full SERP visible, showing the affiliate's ad on your branded query
5. Destination URL with the redirect chain back to your site and the affiliate tracking parameters
6. The relevant clause of your program terms that this affiliate just violated, quoted verbatim, with a clause number
7. Requested action stated plainly — "reverse this transaction, close this partnership, and confirm in writing"

Network operators who handle dispute queues see hundreds of partial submissions a week. A complete submission with seven items in order moves to the front of the queue because it is easy to process. A "hey I think this affiliate is cheating us" email goes to the bottom and stays there. The template is below. Copy it, fill the variables, do not rewrite the structure.

Layer 2 — Direct escalation to your network rep. If Layer 1 is ignored, slow-walked, or denied, the next step is your assigned account manager at the network. Every network of meaningful size assigns a rep to merchants above a certain spend tier. If you do not know who yours is, email the generic support address and ask. They are paid to keep you from churning.

Use them for three things: processing speed (a rep can move a disputed transaction from "under review" to "reversed" in 48 hours versus 2–3 weeks in the normal queue); pattern escalation (if you have caught the same affiliate across 15 transactions in three weeks, frame it as a systemic violation, not a single dispute — "We have identified a systemic violation. We need this partnership terminated and all commissions in the disputed time window frozen pending review" triggers a different internal process); and policy precedent (if the network refuses to reverse, ask them to document why in writing — nine times out of ten they will reverse rather than create a paper trail, and the tenth time you have a written denial you can use in Layer 4).

Do not let them pull you into a debate about whether the affiliate's bidding *technically* violated the program terms. Do not debate. Your program terms are your contract. If your terms say "no bidding on brand-name variations," the affiliate's defense is irrelevant — your job is to enforce your contract, not prosecute theirs.

Layer 3 — Direct contact with the affiliate. Sometimes the network will not terminate a partner at your request because the partner is a large account for them. Sometimes you want to recover commissions that are already locked and paid. In either case, the next layer is going direct.

This is the first point in the escalation ladder where the legal tone shifts. Everything before this was a compliance dispute. Everything from here is a pre-litigation communication and should be written like one.

What goes in a direct demand letter: identification of the affiliate's publisher ID, business entity if known, and the specific partnerships affected; a summary of the violation with specific dates and transaction IDs; a statement that the affiliate is in breach of your program terms, with the clause quoted; a demand for immediate cessation of all bidding on your brand terms, return of specific commission amounts, and a signed attestation that they will not resume; a deadline (10 business days is standard); a statement that failure to comply will result in escalation to their network, potential trademark enforcement under the Lanham Act, and pursuit of any other remedies available. (Have counsel review before sending.)

Send it two ways: email to the affiliate's contact on file with the network, and physical mail to any business address you can find. Document both. Do not send this letter yourself from a personal Gmail. Use your legal department, use outside counsel on letterhead, or use a service like LegalZoom to generate a formal cease-and-desist under your business entity. The letter's power comes from looking like a lawyer wrote it.

Layer 4 — Legal enforcement. Most brands never reach this layer and most never need to. But understanding what is possible here is what makes Layer 3 work. The affiliate's calculation at Layer 3 is: "is this brand actually going to escalate, or is this the usual empty threat?" If the answer is *this brand actually will*, they settle at Layer 3. If the answer is *no, they will go away*, they ignore you.

The enforcement options — Lanham Act, FTC complaint, state consumer protection claims, small claims court, WIPO/UDRP for domain disputes — are all covered in Section 7. Pick the tool that fits the evidence. Typosquat?

WIPO. U.S. LLC under $10K of damages? Small claims. Cloaked landing pages impersonating your site? Lanham Act in federal court with real counsel. Offshore ring? Attack the network layer.

The evidence package

Every dispute, every escalation, every demand letter draws from the same evidence package. Build it once, use it everywhere. Save everything in a single folder per incident, named with the date and the affiliate publisher ID.

```
/evidence/2026-04-08-pub-1234567/
├── 01-cover-memo.pdf          (one page: who, what, when, how much)
├── 02-serp-screenshots/        (min 3 per region, timestamped)
├── 03-ad-detail-snapshots/     (zoomed captures of each ad)
├── 04-redirect-chain.txt       (full URL chain with timestamps)
├── 05-affiliate-link-proof.png (the tracking parameter)
├── 06-transactions.csv         (from network reporting)
├── 07-program-terms-excerpt.pdf (violated clause highlighted)
├── 08-scanner-log.txt          (if using AdCrime or similar)
├── 09-communications-log.txt   (prior warnings sent, dates)
└── 10-damages-calculation.xlsx (commissions + CPC inflation)
```

Missing any of items 1–7 makes the package easy to ignore. Items 8–10 make the package hard to ignore.

The dispute template

Fill the variables. Do not rewrite the structure. The structure is what makes it work.

```
TO:       [Network compliance team email]
CC:       [Assigned account manager, if any]
SUBJECT:  Compliance violation and commission reversal request —
          Publisher ID [XXXXX] — [Your brand]

This is a formal dispute under Section [X] of our program terms
and a request for commission reversal on all transactions attributed
to Publisher ID [XXXXX] between [start date] and [end date].

VIOLATION.
Publisher [XXXXX] has been bidding on our brand name and brand-name
variations in Google Ads in violation of program terms Clause [X],
which reads: "[quoted clause]." We have documented this violation
across [number] separate SERP captures from [list of scanner
countries], timestamps ranging [date range]. Representative
screenshots are attached.

EVIDENCE.
The attached package contains: (1) [number] SERP screenshots across
[regions], (2) the redirect chain from ad click to landing page
confirming attribution to Publisher [XXXXX], (3) the full transaction
list from your reporting for the disputed window, (4) damages
calculation.

REQUESTED ACTION.
We request:
  (a) immediate reversal of all commissions attributed to
      Publisher [XXXXX] between [start date] and [end date],
      totaling $[amount];
  (b) termination of the partnership between [your brand] and
      Publisher [XXXXX];
  (c) written confirmation of both actions within ten business days.

ESCALATION PATH.
If this dispute is not resolved within ten business days, we will
escalate to [name of account manager] and, if necessary, pursue
remedies available under the Lanham Act and applicable state unfair
competition statutes.

[Name], [Title], [Brand]
[Contact information]
[Date]
```

The math of why every week matters

Assume your average infringing affiliate generates $2,000 per week in fraudulent commissions. Assume your commission rate is 15%, and that most of the attributed sales would have converted organically anyway — let us say 90%. That means of the $2,000 in commissions, $1,800 is money you should not have paid.

Week 1. Commissions recorded. Dispute window open at every network. Recovery rate if you catch it: ~95%. Net loss: $90.

Week 3. Commissions still pending on Impact and CJ but approaching validation close. Recovery rate: ~75%. Net loss: $450.

Week 5. Most commissions locked and paid out. Recovery now requires Layer 3 direct contact. Recovery rate: ~30%. Net loss: $1,260.

Week 9. Commissions paid, affiliate has moved on, recovery requires Layer 4. Recovery rate if the affiliate is U.S.-based: ~15%. Recovery rate if offshore: ~0%. Net loss: $1,530–$1,800.

Multiply by the number of infringing affiliates in your program. For a mid-market brand running 300 affiliates with 8% of them caught bidding, that is 24 infringing partners. At $1,800 per week of unrecovered loss per partner, you are hemorrhaging **$43,000 a week** after week 9. Five weeks of delay costs you $215,000. Annualized, $2.2M.

This is why Section 9 — prevention — is the section that actually saves you money. Recovery is the cleanup crew. Prevention is the locked door.

$215,000

What a 5-week detection delay costs a typical mid-market program

SECTION 9

The prevention program

It's a system, not a policy

Most brands "prevent" affiliate fraud the same way they "prevent" house fires. They hope it does not happen, and if it does, they react. That is not a prevention program. That is a reaction plan with optimistic branding.

A real prevention program has five components, all operating at the same time:

1. Program terms that make violations legally actionable — not policy language, contract language
2. An onboarding filter that rejects the obvious fraudsters before they join
3. A monitoring cadence that catches new violations inside the recovery window
4. An organizational workflow with named owners and clear escalation paths
5. A quarterly audit that catches what the other four missed

Miss any one of the five and the other four stop working. Most brands have component 1 (weakly), component 3 (sporadically), and nothing else.

Component 1 — Program terms that actually protect you

Your program terms are the contract. Everything downstream depends on what the contract says. If your contract is vague, every dispute you bring in Section 8 will die in debate. If your contract is specific, every dispute becomes a clause citation and a demand for the remedy the contract already grants you.

You need three clauses, in order of importance. **Talk to IP counsel before adopting any of this language — I am a founder, not a lawyer, and the exact wording that works in your jurisdiction is something a real lawyer needs to sign off on. What follows is a starting point.**

Clause A — Prohibited Keywords.

> *Publisher shall not bid on, target, or purchase advertising against the following terms in any search engine, including but not limited to Google Ads, Bing Ads, Yahoo Search, Yandex, Baidu, or any future platform: (a) the brand name "[Brand]" in any exact, plural, possessive, or common misspelling form; (b) any variation of the brand name combined with any of the following modifiers: "review," "reviews," "coupon," "coupons," "discount," "promo," "promo code," "code," "codes," "deal," "deals," "sale," "login," "sign in," "download," "free," "alternative," "vs," "versus," "competitor," "price," "cost," "pricing," "refund," "cancel," "support," "customer service," "official," "official site," "homepage," "website"; (c) any product name, product line, or trademark owned by [Brand] or its affiliates; (d) any domain name owned by [Brand] or any substantially similar domain including typosquats, homoglyphs, or subdomain variations; (e) any combination of the above. This prohibition applies to exact match, phrase match, broad match, and all match types. Negative keyword configuration does not satisfy this clause; the terms above may not appear as active keywords at any match level.*

The length is a feature, not a bug. Every word closes a loophole that an affiliate would otherwise try to talk their way through.

Clause B — Ad Copy and Landing Pages.

> *Publisher shall not use the brand name, any trademark, or any language suggesting official affiliation with [Brand] in ad headlines, ad descriptions, display URLs, final URLs, ad extensions, site-links, structured snippets, or any other ad asset, unless Publisher has received prior written permission from [Brand] specifying the exact text permitted. Publisher shall not create, operate, or link from any landing page that resembles, mimics, or could reasonably be confused with a page operated by [Brand]. The determination of whether a landing page creates risk of confusion shall be made by [Brand] in [Brand]'s sole discretion.*

The "sole discretion" phrase is important. Without it, every dispute becomes a debate about whether the landing page was confusing enough. With it, you are the judge. (Note: sole-discretion clauses are less favored in EU jurisdictions — if you operate programs in Europe, have local counsel review this language.)

Clause C — Remedies.

> *In the event of any violation of the Prohibited Keywords or Ad Copy and Landing Pages clauses, [Brand] shall have the right, without prior notice, to: (a) reverse and withhold any commissions earned by Publisher during the period of violation, whether pending, validated, or previously paid; (b) immediately terminate the partnership and disable Publisher's tracking links; (c) pursue recovery of any commissions previously paid in respect of transactions associated with violating advertisements; (d) recover its reasonable costs of investigation and enforcement, including attorney's fees, from Publisher. Publisher acknowledges that violations of these clauses cause [Brand] damages that are difficult to measure precisely, and agrees that the remedies in this section are a reasonable estimate of those damages and are not a penalty.*

The last sentence is doing heavy legal work. Courts sometimes strike down "liquidated damages" clauses as penalties; phrasing them as "a reasonable estimate" of actual damages makes them more enforceable.

Put these clauses inside your program terms document on the affiliate network. Every network lets you upload your own terms document and require acceptance at onboarding. If your network does not, change networks.

Component 2 — The onboarding filter

Set every affiliate application to manual review. Not "auto-approve after 30 days." Not "auto-approve with filters." Manual. Every application.

This is controversial because manual review slows down growth. But the math is straightforward. A program with 300 affiliates where 8% are fraudulent is losing $43K per week if fraud goes unnoticed. Manual review at 5 minutes per application, processing 20 applications per week, costs 100 minutes per week. You are trading 100 minutes to save $43,000. That is the best hourly rate you will ever see.

What you check during manual review:

1. **Business entity and jurisdiction.** LLC, corporation, individual? Where is the business located? Partners registered in the British Virgin Islands, Belize, Seychelles, or any jurisdiction where civil judgments are unenforceable go into a "high-risk, low-trust" bucket — not automatically rejected, but required to clear a higher bar on the remaining checks.
2. **Traffic source stated on the application.** Red flags: "SEO" with no site listed, "paid search" as the primary source, "social" without naming platforms, or blank. A legitimate affiliate will tell you they run a coupon site, a comparison site, a newsletter, a YouTube channel, or a niche forum. A fraudulent affiliate will be vague.
3. **Claimed website.** Check the WHOIS — is the domain registered in the last 60 days? Is it a parked page, a coupon aggregator, or a spam site? Run the domain through `archive.org/wayback` — does it have any history, or did it appear fully formed last week?
4. **Prior network history.** Ask the network. Every affiliate network knows which publishers have been terminated from other programs for compliance issues. Some networks may share termination history if you ask — most merchants never do.
5. **Named contact.** Is there a real human name attached to the application? Does that name exist on LinkedIn in a role consistent with affiliate marketing?

6. **Communication style.** Did they send a generic one-line "let me know if approved" message, or did they send a two-paragraph message explaining their traffic and why they want to promote your brand? Legitimate affiliates pitch you. Fraudsters want in fast and quiet.

Give each applicant a 1–5 score on each of the six checks. Approve 5+ applicants with no flags. Manually interview any applicant with 3–4 flags over video call before approving. Reject anyone scoring below 3. Document the reasoning for every rejection.

Component 3 — The monitoring cadence

Prevention is not a one-time setup. It is a weekly rhythm. New affiliates join, existing affiliates change behavior, and Google's ad ecosystem shifts around you.

Monday — SERP scan. Branded keyword check across at least three geographic markets. Top-20 results per query. Note anything unfamiliar. Escalate anything that is clearly an affiliate ad you did not approve. Time cost: 45 minutes manual, 2 minutes automated.

Tuesday — Auction Insights review. Log into Google Ads, pull Auction Insights on your top five branded campaigns, look for new domains appearing in the competitor list. Time cost: 15 minutes.

Wednesday — Transaction anomaly check. Pull the previous week's top 10 affiliates by commission volume from each of your networks. Flag any partner who (a) appeared in the top 10 for the first time this week, (b) grew commissions by more than 50% week-over-week, or (c) has a much higher conversion rate than the program average. These three flags are the dashboard fingerprint of a cookie-stuffing or coupon-intercept affiliate. Time cost: 30 minutes.

Thursday — Dispute queue review. If you filed any disputes this week, check status and nudge network reps if stalled.

Friday — Weekly summary to finance. Single-page summary: number of violations detected, number of disputes filed, dollar amount recovered, list of active partners suspended. This is how prevention becomes visible to the rest of the company and keeps its budget.

Monthly additions. First Monday: full SERP scan across your complete regional coverage list. Mid-month: program terms re-read — you will be surprised how often you find gaps after three months of running disputes. Last business day: finance reconciliation.

Component 4 — Organizational workflow

Every broken affiliate program has the same root cause. Nobody owns the problem. Affiliate sits under marketing. Marketing assumes the network is monitoring compliance. The network assumes the brand has its own watch. Finance sees the commissions go out and assumes they are legitimate. Legal only gets involved when there is a lawsuit.

Name an owner. One person is accountable for the prevention program. Name a backup. Name an escalation path — head of marketing, head of legal, CFO — so the owner knows who to call without a meeting when something material happens. All three should know in advance that they are in the escalation path and what the criteria are for pulling them in.

Document the playbook. Write down the exact steps from "SERP violation detected" through "commission recovered." The document is not for the current affiliate manager — the affiliate manager knows the steps. The document is for whoever replaces them in six months. Review it quarterly. Things change.

Component 5 — The quarterly audit

Every 90 days you audit the full program. Not a spot-check. Not a review of the weekly summaries. A clean-sheet audit that assumes the program is broken and looks for the evidence.

1. **Full affiliate roster review.** Pull the complete list of active affiliates from every network. Deactivate anyone who drove zero commissions last quarter. Dormant affiliates are a free attack surface and a fraudster's favorite place to hide.
2. **Program terms review.** Read them again. Has anything changed in the ecosystem that makes your clauses weaker? Update if needed.
3. **Network terms review.** Re-read each network's publisher agreement. Awin, CJ, and Impact all update their terms multiple times per year; most merchants never read the updates.

4. **Top partner audit.** Pull the top 20 affiliates by commission volume. Verify (a) they are still operating a legitimate traffic source, (b) their conversion rate is consistent with the program baseline, (c) no SERP scans in the last quarter caught them bidding on your terms.
5. **Bottom partner audit.** Pull the bottom 50 affiliates. Are any of them spikes of activity that dropped back down? A spike followed by silence is sometimes a fraudster who got caught and quietly moved on, leaving a trail you missed.
6. **Dispute history review.** How many disputes did you file this quarter? How many resolved in your favor? What was the average resolution time? What patterns show up repeatedly? The patterns you see in disputes this quarter are the monitoring signals you should build into next quarter's cadence.
7. **Commission leakage estimate.** Based on the disputes you won, extrapolate how much you probably lost to fraud you did not catch. This is your ROI case for investing more in prevention next quarter.
8. **Tool evaluation.** Is your toolchain keeping up?
9. **Budget review.** What did the prevention program cost this quarter in time, tools, and recovered-versus-lost dollars? For a mid-market brand the answer should be "yes, by 10–50x."

When automation pays off

Everything in this section can be done manually. The question is whether it should be.

The rule of thumb: if your affiliate program is driving more than $50K per month in commissions across all networks combined, manual prevention is no longer the right call. At that scale, the labor cost of the Monday-through-Friday cadence exceeds the cost of a scanner that does it in minutes — and the scanner catches things the manual process will miss (cloaked ads that only show up in certain countries, ads that only run during specific hours, typosquat domains that never touch a SERP you thought to check).

Below $50K per month, manual prevention is economically correct. Above $50K per month, you are losing more to undetected fraud each week than a tool costs per month.

This is the gap I built AdCrime for. Every tool in the existing market is either enterprise-priced at $40K+ per year or enterprise-scoped in ways that do not fit a founder-led or mid-market team. So I built the thing I would have wanted

to buy when I was running an affiliate program myself: six-country coverage, continuous scanning, case files you can actually send to your network's compliance team, priced for people who do not have a six-figure tooling budget. If the math in this section matches your situation, come find me at adcrime.com. If it does not, use this Playbook to build it manually. Either way, the goal is the same — fewer fraudsters getting away with it.

The one-page prevention checklist

Weekly - Monday: SERP scan across 3 regions on top branded keywords - Tuesday: Auction Insights review on top 5 branded campaigns - Wednesday: Transaction anomaly check — top 10 affiliates per network - Thursday: Dispute queue status review - Friday: One-page summary to finance

Monthly - First Monday: Full SERP scan across 6 regions - Mid-month: Program terms re-read - Last business day: Finance reconciliation

Quarterly - Full affiliate roster review (deactivate dormants) - Program terms update - Network terms review - Top 20 and bottom 50 partner audit - Dispute history pattern review - Commission leakage estimate - Tool evaluation and budget review - File the Prevention Program Review

On every new affiliate application - Manual review, six-point scorecard - Approve 5+, interview 3–4, reject below 3 - Document every rejection

On every detected violation - Build the evidence package (ten items) - File dispute within 48 hours - Escalate to network rep if stalled after 10 days - Escalate to direct contact if network refuses - Escalate to legal if direct contact fails - Close the loop with finance

---

SECTION 10

The tooling landscape

A short honest survey of what exists. I run one of these companies, so read this section with that in mind. I have tried to be fair.

BrandVerity — Oldest player in the space. Solid fundamentals, extensive geographic coverage. Strengths: maturity, enterprise integrations. Weaknesses: detection of complex redirect chains has historically been limited, pricing is enterprise-scoped, the interface shows its age. Good fit for brands that already have a BrandVerity contract and do not want to switch.

Adthena — Powerful search intelligence suite with brand protection as one module among many. Strengths: the search competitive intelligence side is genuinely best-in-class. Weaknesses: $40K+/year pricing, overkill if brand protection is your only use case, learning curve is steep. Enterprise only.

The Search Monitor — Large feature surface, lots of data, granular controls. Strengths: flexible rule configuration. Weaknesses: interface is often described as overwhelming, data is there but hard to act on without meaningful onboarding time. Works for teams with a dedicated analyst.

BluePear — Newer entrant with a scanning-plus-enforcement model. Strengths: pricing is more accessible than the enterprise players. Weaknesses: limited independent reviews, sponsored content relationships with trade publications should be noted when evaluating their case studies.

AdCrime — The tool I built. A global scanner network with regional hubs across the Americas, Europe, and Asia-Pacific and the ability to reach into any specific market a client needs, continuous monitoring, case files designed to be forwarded directly to network compliance teams, priced for founder-led and mid-market teams. No sales calls, no annual contracts, no $40K floor. You can start a scan, get the first case file in 48 hours, and decide from the evidence whether the product is worth keeping. If this Playbook made sense to you, the product was built for the same person.

The honest recommendation. If your program is under $50K/month in commissions, do it manually using Section 6 and Section 9. If you are between $50K and $500K/month, the mid-market tools (AdCrime, BluePear) are the right fit. If you are over $500K/month and have dedicated headcount to operate it, the enterprise tools (BrandVerity, Adthena, Search Monitor) become defensible. If you are over $1M/month, you should probably be running two tools and cross-referencing them — the false negative rate on any single scanner is the thing that will hurt you the most at that scale.

SECTION 11

Ten warning signs your program is compromised

Print this. Share it with your team. If you can check three or more of these, you almost certainly have active fraud.

1. **Your branded search CPC has drifted up 20% or more in the last 90 days** with no explanatory change in your own bid strategy
2. **Your top affiliate's commission volume grew more than 50% in one month** without a corresponding campaign push
3. **One or more affiliates in your top 10 have conversion rates above 25%** on branded queries
4. **Your click-to-conversion windows for certain affiliates are under 2 minutes** on considered purchases
5. **Your Auction Insights shows unfamiliar domains** consistently in the top five on your brand terms
6. **You cannot reproduce ads from your office IP** that customers report seeing
7. **Your finance team has flagged a disproportionate share of affiliate commissions** on otherwise-organic-looking orders
8. **Your program terms do not contain a specific "no brand-bidding" clause** with enumerated modifier terms
9. **You have not updated your affiliate program terms in over 12 months**
10. **You have not scanned your SERPs from more than one country in the last 90 days**

Three or more checked means you should run the Week 1–Week 4 playbook from Section 6 starting Monday.

APPENDIX

Sources and source quality flags

High-confidence sources (court filings, FTC orders, official policy docs):

- CHEQ × University of Baltimore, *The Economic Cost of Bad Actors on the Internet: Affiliate Fraud*, 2020 and 2022 editions
- *1-800 Contacts v. JAND, Inc. d/b/a Warby Parker*, 2d Cir. (October 8, 2024)
- *Lerner & Rowe P.C. v. Brown Engstrand & Shelly, LLC*, 9th Cir. (October 22, 2024)
- *In re PayPal Honey Extension Litigation*, Case No. 5:24-cv-09470, N.D. Cal.
- FTC Endorsement Guides, revised July 2023
- FTC Fake Reviews Rule, August 2024
- *FTC v. Traffic and Funnels LLC*, June 2024 settlement
- *FTC v. Total Wealth Academy*, August 2024
- Google Ads Trademark Policy update, July 24, 2023
- Google Chrome Web Store policy update on affiliate disclosure, March 2025
- WIPO UDRP procedure documentation

Medium-confidence sources (industry research, reputable trade press):

- Juniper Research ad fraud forecasts, 2023 and 2024
- QuoIntelligence, NordVPN affiliate fraud ring report, March 2025
- MegaLag YouTube investigation of Honey/PayPal, December 2024

- Search Engine Land coverage of Google policy changes
- Ben Edelman's published research on affiliate compliance
- Affiverse Media reporting on affiliate enforcement cases

Lower-confidence sources (vendor research, single-report claims):

- BforeAI 33,000-ad sample analysis, March 2025 (vendor-produced)
- mFilterIt cloaking prevalence figures (vendor-produced)
- BluePear case studies (vendor-published)
- Kevin Frisch's $100M Uber claim (self-reported, unverified by independent audit)
- Sage Financial Software case details (industry trade press)
- Nordstrom $1.4M loss figure (industry reporting, not audited)

What "source quality flags" means. The brand-bidding fraud research ecosystem has a visibility problem. Most of the hard numbers come from vendors who benefit from the problem being large. Most of the independent research is academic, slow, or behind paywalls. I have flagged the confidence level on each source so you can decide for yourself which claims to lean on when you escalate internally. The high-confidence sources will hold up in a boardroom. The lower-confidence sources are useful for painting the picture but should not be the center of a business case.

APPENDIX

Closing note

The Playbook you just read is version 1.0. If you catch something wrong in here, email me at mario@adcrime.com and I will fix it in the next version. If you use this Playbook to recover commissions from an affiliate who was stealing from you, I want to hear about it. If you have a case study you are willing to share — anonymized or not — I will credit you and build it into a future edition.

This is a living document. The fraud is not holding still. Neither is this.

— Mario Vaher *Founder, AdCrime April 2026*
