MARQVISION

Case Study

How a Global Premium Consumer Brand Cut Counterfeit Saturation from 46% to 26% in Eight Months

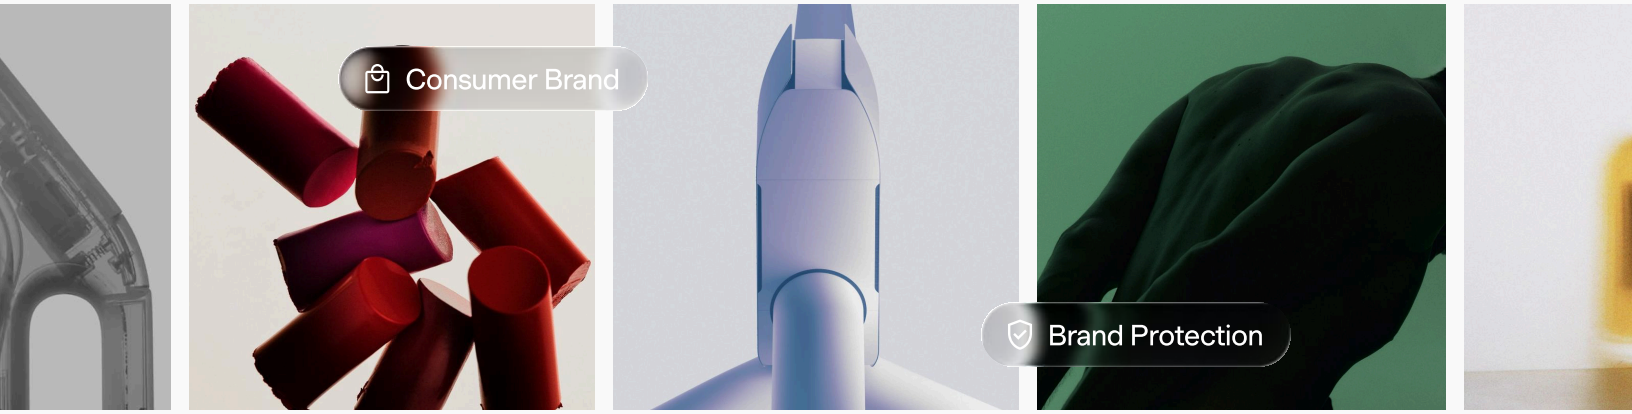

Key Takeaways

- A global premium consumer brand cut its counterfeit saturation rate from **46% to 26% in under eight months**, across 40+ global marketplaces — and now runs its global brand protection program on MarqVision.
- Generative AI collapsed the cost of counterfeiting: not just fake goods, but cloned storefronts, fraudulent ads, and dupe content across social — all ripping off customers at machine speed. The brand decided to fight AI with AI.
- It ran a competitive evaluation against alternative vendors under live conditions. MarqVision delivered the strongest results.
- What sealed the decision: MarqVision was the only vendor offering a real ROI metric — **Saturation Rate**, the percentage of first-page marketplace search results that are infringing — plus end-to-end coverage across marketplaces, social media, websites, domain names, and paid ads.
- The decline followed a predictable curve: a steep early drop, a slower middle phase, then stabilization that reveals the structural floor of online-only enforcement and points to where revenue recovery efforts should focus next.
- Outcome data turned reporting into a planning tool for the team's next move.

The Problem

A Strong Program Built for a Slower Era

By every conventional measure, this brand's protection program was healthy: a dedicated in-house team, a long-standing enforcement vendor, and active monitoring across hundreds of global marketplaces. Takedown volumes were consistently high, coverage broad, and reports delivered on schedule.

But the program's pace hadn't kept up. Innovation had stalled, visibility into what was actually happening in the market was limited, and reporting couldn't keep pace with how fast the threat was evolving. Meanwhile, Generative AI was collapsing the cost of faking things across the industry: convincing product images, cloned storefronts, fraudulent ads, viral dupe content, and impersonating domains could now be spun up in bulk.

What once took a counterfeiting operation real effort can now be produced at machine speed. Counterfeiting stopped being a marketplace-listings problem and became an every-surface problem: anywhere a customer might search, scroll, click, or buy, someone could be ripping them off in the brand's name.

The brand could see it in its own data. A single week of live monitoring surfaced thousands of infringing listings across marketplaces and social platforms, thousands of suspect domains, and dozens of look-alike brands and impersonating storefronts. And the pool never drains on its own: the OECD (Organization for Economic Co-operation and Development) valued global trade in counterfeit goods at roughly $467 billion, or 2.3% of total global imports, according to the OECD's 2025 fakes report.

Meanwhile, the program's own numbers couldn't say whether it was keeping up. Business stakeholders asked a simple question: **Is it working?** And activity-based reports couldn't answer it. Listings found, listings removed, platforms monitored: the counts went up every quarter, which should have been reassuring. Instead, it raised a harder question: if we keep removing more listings every period, is the problem growing, or are we making progress? Effort-based metrics answer "how much did we do?" They never answer the question that actually matters: **when a consumer searches for our product, what do they see?**

The team's conclusion was clear-eyed. You cannot fight machine-scale infringement with manual review and effort-based reporting. If AI was industrializing the attack, the defense had to be AI too — and it had to be measured by outcomes, not activity.

The Evaluation

To pressure-test that thesis, the brand ran a blind, competitive evaluation against alternative vendors on live infringement scenarios. When the results came in, MarqVision had delivered the strongest enforcement outcomes of the evaluation.

But a trial win alone doesn't move an entire global program. The decision came down to what would keep the brand winning over the long run.

Real ROI, Covered End-to-End

ROI that the brands can actually see

Every vendor can produce takedown counts. MarqVision was the only vendor in the brand's evaluation that could put a real outcome metric on the table: Saturation Rate.

Saturation Rate measures the percentage of results on the first page of consumer-facing marketplace searches that are infringing. Instead of hunting listings at any depth and counting the removals, a saturation-based program monitors the specific searches real buyers use, classifies every result (e.g., counterfeit, suspected, authorized, or irrelevant) and tracks whether the shelf is getting cleaner over time.

The industry shift: activity metrics vs. outcome metrics

	ACTIVITY METRICS The old scoreboard	SATURATION RATE The outcome metric
The question it answers	"How much did we do?"	"What does a buyer actually see?"
What gets measured	Detections, takedowns, platforms covered	Share of first-page search results that are infringing
How detection works	Hunts listings at any depth, across every keyword variation	Monitors the exact searches real buyers use, to a fixed depth
Clean listings	Discarded — they never enter the pipeline	Classified — they are part of the denominator
A rising number means	Ambiguous: more effort, or a growing problem?	Unambiguous: consumer exposure is rising
What it proves to the business	Spend and effort	Impact and progress

For the brand protection team, this solved the stakeholder problem that takedown counts never could. Saturation Rate gave them a number they could take to the business that communicated impact in business terms, not enforcement jargon: these are the categories and regions most targeted, this is how consumer exposure is trending, this is what changed because of the program. The internal conversation shifted from "here's what we spent on brand protection" to "here's what brand protection achieved."

One platform for every surface

AI-era infringement doesn't stay in marketplace listings, so coverage can't either. What changed for the brand wasn't the surfaces covered — marketplaces, social media, website content, domain names, and paid ads were already in scope — it was AI doing the heavy lifting within that scope: classifying every result in the consumer's field of view, clustering connected sellers and reused images, and automatically flagging impersonating domains. Fighting AI with AI, in practice.

A product that improves continuously

Two capabilities the brand requested during the evaluation — a watch list for enforced domains that resurface, and in-dashboard translation of foreign-language listings — were built and shipped within weeks, not roadmap quarters. For a team betting on technology to outpace the counterfeiters, watching requests turn into working software that fast was the closing argument.

The Results

With the program live across every channel, saturation measurement began with a focused scope: **7 priority keywords** monitored across **40+ global marketplaces** through **96 recurring searches**.

At the outset, the brand's counterfeit saturation rate stood at **46%**. Nearly half of what consumers saw when searching for this brand's products was infringing. That exposure has real consequences for buyers: an MSU survey found roughly 7 in 10 consumers had been deceived into buying a counterfeit product online at least once in the previous year.

Eight months later, that number was **26%**.

Saturation rate trend: month-by-month decline

Share of first-page results that are infringing — 7 keywords, 96 searches, 40+ marketplaces

-20pp in under 8 months

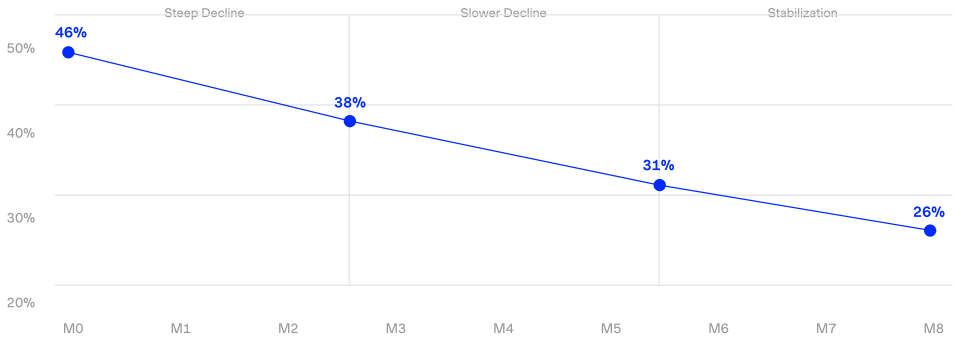

Verified waypoints: M0 (46%), M3 (38%), M6 (31%), M8 (26%). Intermediate monthly points are illustrative smoothing — replace with actuals from the cleanliness rate reports before publish.

Key metrics at a glance

46% → 26% Counterfeit Saturation Rate, Start → Month 8 — -20pp in under 8 months

7	Priority keywords tracked	**40+**	Global marketplaces covered	**96**	Recurring searches monitored
2+ hrs/wk	Manual follow-up time reduced	**100%**	First-page Results Classified	**1**	Platform for the Entire Program

Reading the Curve

The shape of the decline tells its own story, and it maps to a pattern we see consistently across enterprise brand protection programs.

Months 1 ~ 3: Steep decline (46% → 38%). The low-hanging fruit phase. Clearly enforceable infringement is identified and removed quickly — the listings are obvious, the platforms have established takedown procedures, and compliance comes quickly.

Months 4 ~ 6: Slower decline (38% → 31%). The rate of improvement slows as the remaining infringement gets harder: platforms with slower response times, sellers who relist rapidly after removal, and edge cases where classification is ambiguous.

Months 7 ~ 8: Stabilization (31% → 26%). The plateau is not a failure. It is a signal. It reveals the structural floor for a given set of keywords and platforms — the point at which online enforcement alone cannot push the number lower. This is where the data starts informing the next strategic move.

What Happens When the Curve Stabilizes

The stabilization phase is where saturation data becomes a planning tool. Depending on what the trend line shows, the team can dig deeper into a specific infringement category to understand why it's resistant, flag a coverage gap where additional IP registrations would strengthen enforcement, or escalate to offline action where online enforcement has reached its limit.

A takedown volume report cannot surface these strategic inflection points.
A saturation rate trend line makes them visible by default.

Beyond the Number

Enforcement volume under the new program is on par with what the previous vendor delivered — and that is precisely the point. Similar activity, measurably different outcomes, and for the first time, proof.

The qualitative signal arrived from the field. The brand's regional teams around the world began reporting fewer counterfeits and dupes in their markets. The steady drip of internal flags — I found this, please take it down — slowed sharply. Previously, the team fielded an estimated 5 to 10 of those emails a week, each requiring a vendor email and a follow-up. Now, most of that traffic is gone, and when a question does come up, the team checks enforcement status directly in the MarqVision dashboard instead of waiting for a reply. We estimate the change returns to the team in manual follow-up alone to be upwards of two hours a week.

The structural change runs deeper. The brand allocated dedicated analyst hours for proactive research — surfacing trends, identifying newly counterfeited products, flagging fabricated designs being attached to the brand's name — work the previous

program had no data framework to support. And when a keyword's saturation rate plateaus, the team makes an informed call rather than a guess — including pushing platforms directly, backed by data on where their proactive enforcement is falling short.

That is what changed the relationship. This wasn't an additive vendor swap. It was a deliberate bet on outcome measurement over activity reporting — made by a team that finally had the data to know the status quo wasn't good enough.

Winning vs. Busy

For any brand protection team that has spent years reporting activity metrics and struggling to articulate program impact to the business, the shift to outcome measurement is not incremental. It is the difference between describing what you did and proving what changed.

Enforcement volume will always be part of brand protection. But without an outcome metric tied to the consumer experience, there is no way to know whether a program is winning — or just busy.

Frequently Asked Questions

What is saturation rate in brand protection?

Saturation rate is the percentage of results on the first page of a consumer-facing marketplace search that are infringing. Instead of counting how many listings a team removed, it measures what a buyer actually sees when they search for a product, making the consumer's experience the unit of measurement.

How has generative AI changed counterfeiting?

Generative AI has made it dramatically cheaper and faster to produce convincing fakes — not only counterfeit goods, but cloned storefronts, fraudulent ads, dupe content on social platforms, and impersonating domains. The result is more infringement, appearing on more surfaces, faster than manual programs can track. That is why leading brands are pairing AI-powered detection with outcome metrics like saturation rate: fight machine-scale infringement at machine scale, and prove the fight is being won.

How is saturation rate different from takedown volume?

Takedown volume answers "how many infringements did we remove?" and measures effort. Saturation rate answers "of everything a buyer sees when they search, what percentage is infringing?" and measures results. A brand can remove tens of thousands of listings and still have counterfeits dominating the first page, which is why outcome data tells a clearer story than activity counts.

How much did the brand reduce its saturation rate?

The brand cut its saturation rate from 46% to 26% in under eight months, a 20 percentage-point drop. The program tracked 7 priority keywords across more than 40 marketplaces through 96 recurring searches.

Why does the saturation rate curve stabilize over time?

The decline follows a predictable pattern: a steep early drop as obvious, clearly enforceable infringement is removed quickly, then a slower middle phase as the harder cases remain. These harder cases include listings on platforms with slow response times, sellers who relist rapidly after removal, and edge cases where infringement classification is ambiguous. Eventually the rate stabilizes. That plateau reveals the structural floor for a given set of keywords and platforms — the point where online enforcement alone cannot push the number lower and offline action becomes the next move.

What should a team do when saturation rate plateaus?

A plateau is a planning signal, not a failure. Depending on what the data shows, teams can keep enforcing to hold the line, rotate to a new keyword cluster where the rate is higher, or escalate to offline action such as test buys, seller network analysis, and legal filings.

Request Demo

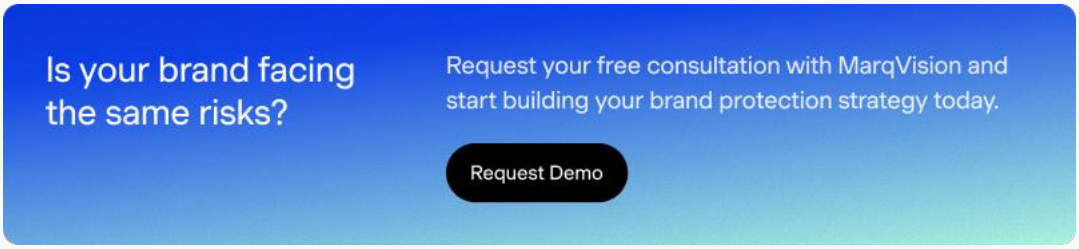


MARQVISION