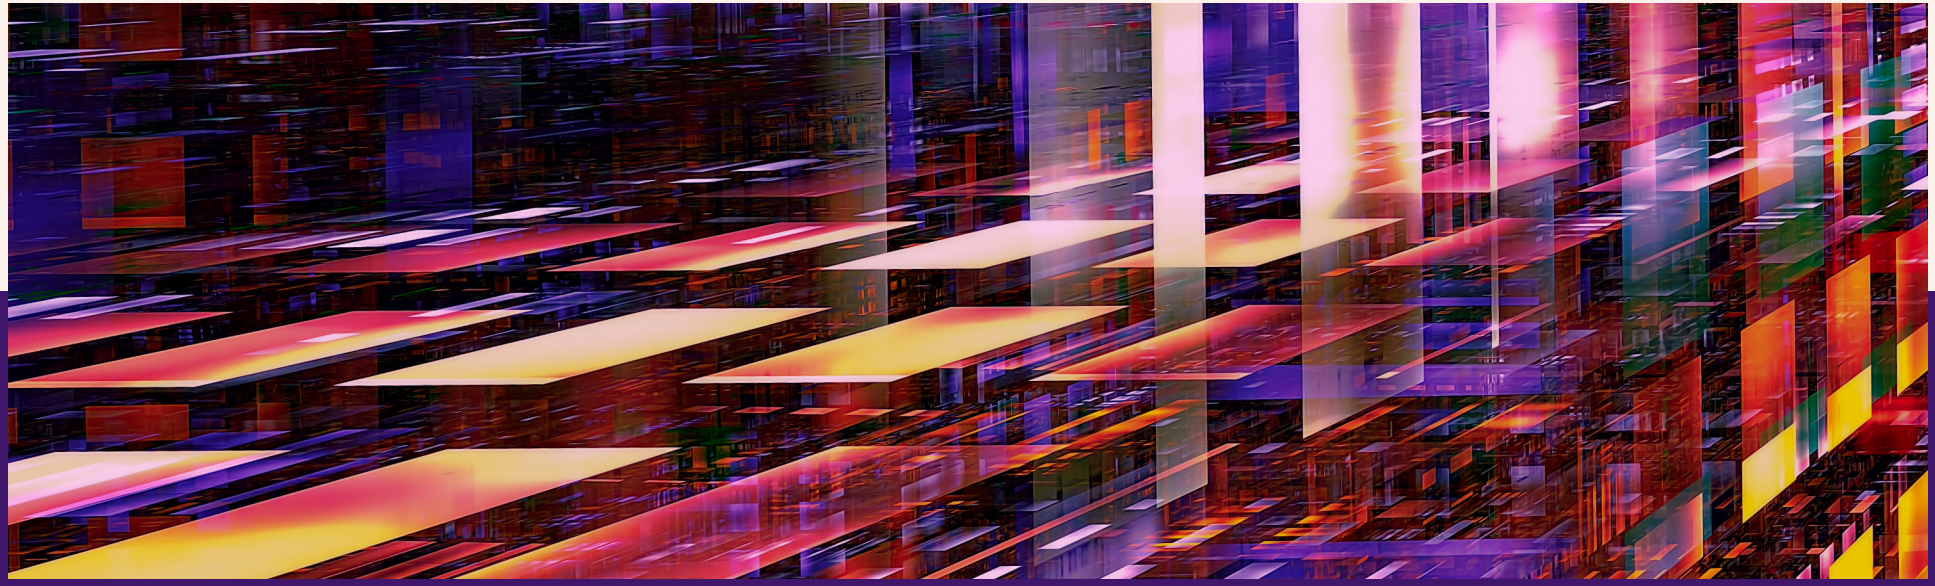
HUMAN

Five Lessons from a Decade of Fighting Ad Fraud

Lessons from the front lines

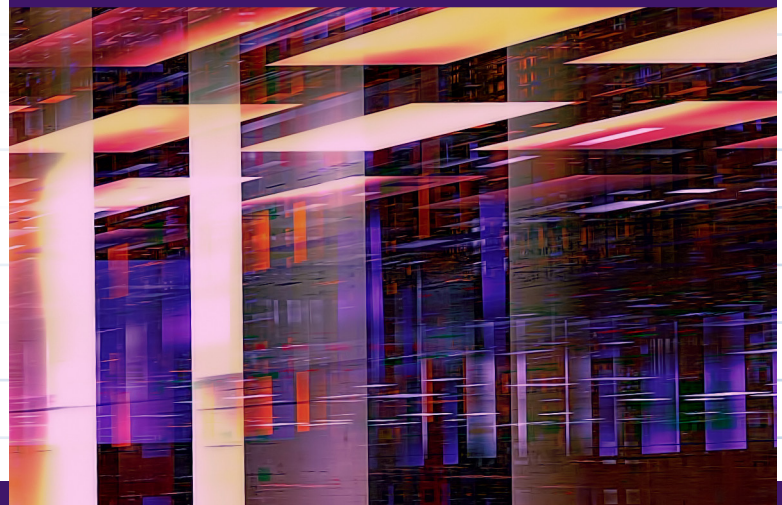


Introduction

When HUMAN began fighting invalid traffic (IVT) over a decade ago, what we found was a digital media ecosystem lacking in transparency and accountability. Our Bot Baseline Reports found that threat actors were siphoning billions of dollars away from advertisers, and since that research, we've continued to uncover fundamental vulnerabilities and opacities that affect every advertiser's campaign performance, brand safety, and media quality.

Protecting trillions of impressions every week from sophisticated fraud operations reinforces the notion that threats targeting your campaigns are interconnected. Fraudsters take advantage of gaps in the value chain of digital interactions, and leverage numerous different tactics within one operation. Historical perspectives on ad verification have been very point-in-time: static blocklists, infrequent recrawls, and isolated risk types, all of which lead to a "trust us" message from legacy providers.

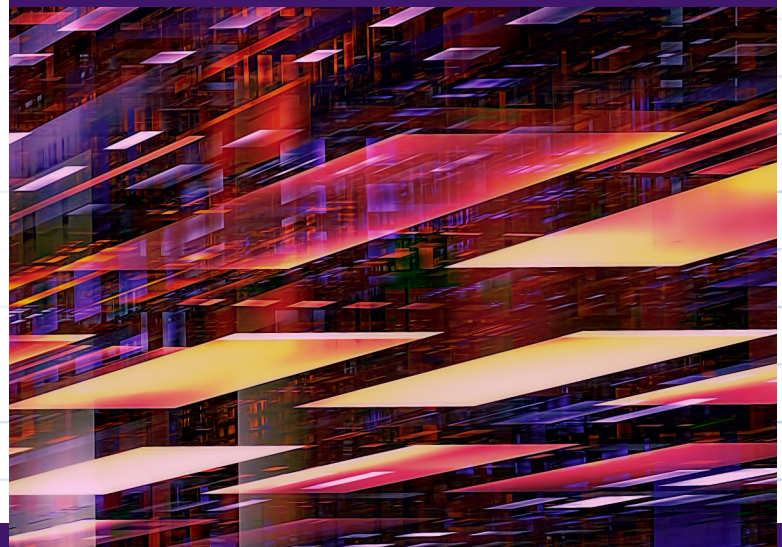
*What's needed is a trust layer that looks at the whole board and **makes every verification decision crystal clear.***



What's needed isn't "trust us", it's a *trust layer* that looks at the whole board and makes every verification decision crystal clear. True protection of digital advertising requires a cybersecurity-oriented level of scrutiny that provides the right level of visibility at the right time, nuanced detection, and the ability to give control to brands, partners, and agencies to make the ultimate decision on what is right for their business. Understanding these connections and the inherent, objective challenges and limitations involved in actually getting the job done as opposed to merely checking a box revealed why traditional verification approaches often miss what matters most to campaign success and ensuring trust for advertisers.

Here's what fighting the most advanced threats in ad tech revealed about protecting your advertising investments.

*Traditional verification approaches **often miss what matters most** to campaign success and ensuring trust for advertisers.*



Lesson 1:

Adversaries Think in Systems, Not Silos

What we learned fighting fraud

The most sophisticated fraudsters don't just generate fake traffic with bots, they create entire fake digital ecosystems. They build convincing websites, generate realistic user journeys, and mimic authentic engagement patterns across multiple touch points. These operations revealed something crucial: advanced adversaries understand the interconnectedness of ad tech better than most defenders.

What this means for your campaigns

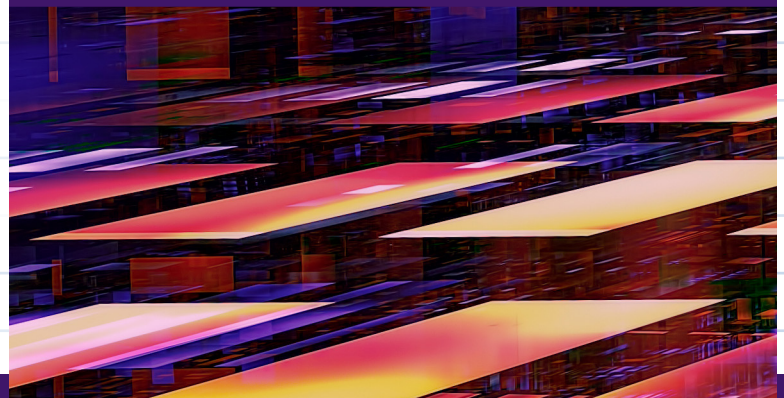
Advanced threats and challenges don't fit neatly into one category; they may use legitimate content to mask fraudulent traffic, but they might also manipulate viewability signals while avoiding brand safety detection, or they might exploit supply path vulnerabilities. Situations like these require a comprehensive, visibility-centric approach.

The campaign protection implication

Your verification needs to understand how all these threats interconnect. Sophisticated threats can be elusive, fooling basic IVT detection and manipulating viewability signals. Systemic risks demand holistic solutions. Comprehensive campaign protection requires verification that sees the whole threat landscape, not just isolated pieces.

Real-world example

HUMAN's 2025 **disruption of BADBOX 2.0** exposed that threat actors don't think about attacks like defenders do; they seek to make money through whatever exploitation opportunities are available.



Lesson 2:

Real-Time Understanding Beats Static Defenses

What we learned fighting fraud

Traditional fraud detection relied on blocklists and historical patterns. But we discovered that sophisticated threats adapt and spread faster than static rule sets can realistically keep up. The moment you block one fraud signature or URL, there are three new variants. The only defense that works is real-time intelligence that adapts as fast as the threats evolve.

What this means for your campaigns

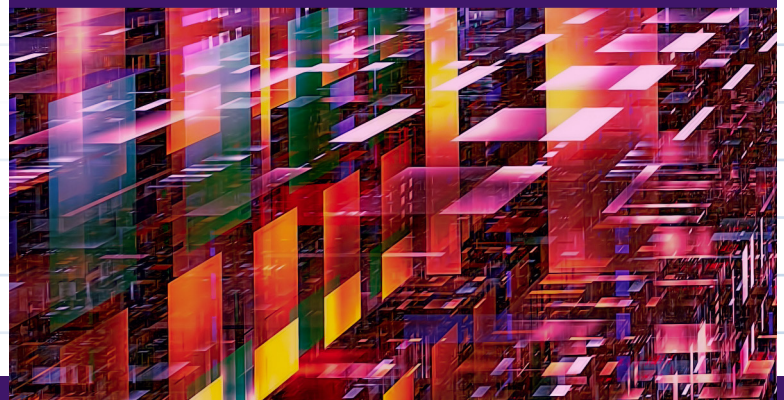
Legacy verification relies on cached URL classifications and historical patterns. But sophisticated operations and operators change faster than these tactics can manage. When breaking news transforms content context in minutes, or when legitimate websites get compromised with malicious content, verification that works from yesterday's data leaves your campaigns exposed to today's risks.

The campaign protection implication

Your verification must analyze content and context as they exist now, not as they existed when last crawled. Whether it's a sophisticated fraud operation, sudden brand safety violations, or viewability manipulation, effective protection requires real-time intelligence that adapts as threats evolve.

Real-world example

The HUMAN-led private takedown of the VASTFLUX operation reflected the complexity of modern threats and how threat actors adapt to extend the lifespan of a threat.



Lesson 3:

Infrastructure Visibility Is Critical

What we learned fighting fraud

The most valuable verification intelligence doesn't come from analyzing individual impressions, it comes from seeing patterns across billions of bid requests and understanding the infrastructure underpinning them. When you can see supply path optimization patterns, device clustering, and cross-platform behavioral correlations, threats become obvious.

What this means for your campaigns

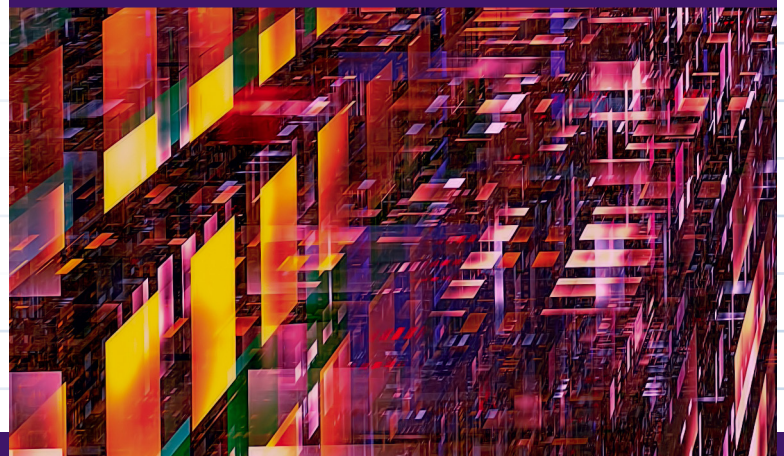
Traditional verification looks at surface-level signals, like URL keywords, basic viewability metrics, or simple content scans. But sophisticated threats manipulate or obfuscate the infrastructure underneath. They might serve perfectly safe content while spoofing the technical environment to fake viewability. A lack of deep, granular transparency can result in verification failures.

The campaign protection implication

Effective verification requires infrastructure-level visibility. You need verification that understands not just what content is being served, but how it's being delivered, through which supply paths, to which devices, and with what technical specifications. This deeper intelligence is what separates real protection from one-size-fits-all approaches.

Real-world example

HUMAN's **discovery of the Konfety scheme** details how threat actors used a "decoy/evil twin" app pairing system to present fraudulent traffic as though it were legitimate.



Lesson 4:

Transparency Is a Requirement, Not a Nice-to-Have

What we learned fighting fraud

When managing sophisticated campaigns, black-box detection systems become a liability. You need to be able to understand exactly why something was flagged and how to respond. Explainable decisions aren't just better for stakeholders, they're essential for media teams.

What this means for your campaigns

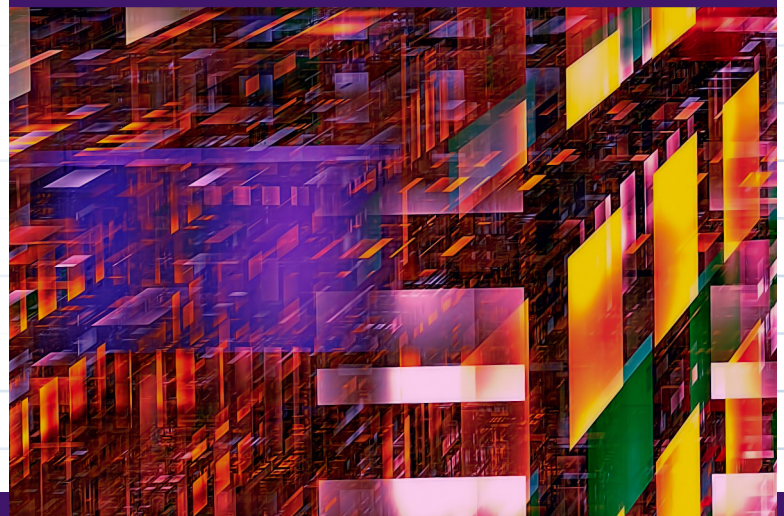
An inability to explain decisions in verification creates blind spots in resolving the real challenges of the ecosystem and fails to provide advertisers the information they need to make informed decisions. When a brand safety system blocks content without clear reasoning, you can't tell if it caught a real threat or created a false positive. When viewability measurement lacks transparency, you can't distinguish between technical issues and active manipulation. When unavoidable technical limitations prevent a clear assessment and aren't exposed for decisioning, you risk delivery on inventory where insight is not available.

The campaign protection implication

Every verification decision should be explainable, not just for compliance purposes, but because understanding why decisions were made is essential for designing future campaigns.

Real-world example

HUMAN research in 2025 found that **marketers demand transparency** and explainability in their verification platforms.



Lesson 5:

Precision Matters

What we learned fighting fraud

Early detection systems cast wide nets to ensure they caught everything, resulting in high false positive rates; the thought process being "overblocking is better than missing a real issue". But precision is especially critical when protecting ecosystem integrity. False positives don't just break trust between ecosystem participants, they also reduce opportunities to serve ads to possible customers through desirable supply sources.

What this means for your campaigns

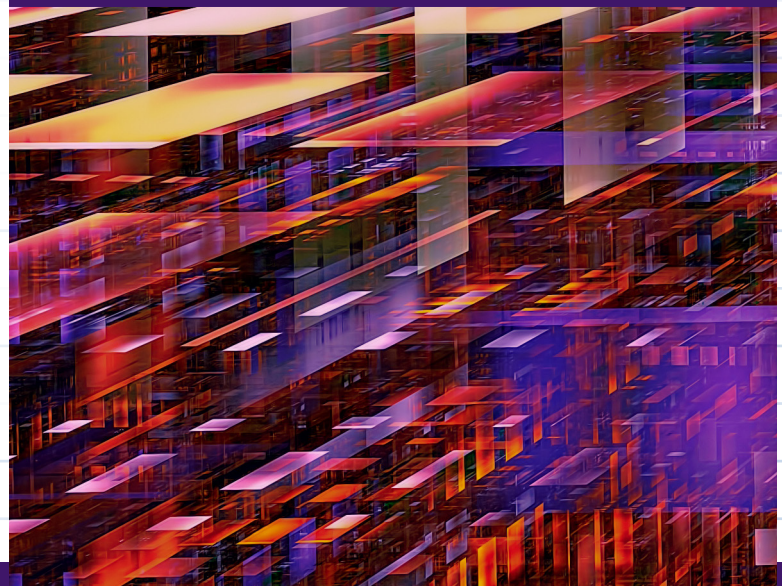
Brand safety systems with high false positive rates alter the reality of what is acceptable, and create missed opportunities to engage with real audiences. And brands with higher tolerances of traditionally brand-unsuitable content run an even bigger risk of missing out on valuable attention. Viewability systems that frequently misclassify create optimization behaviors that make the entire ecosystem less efficient.

The campaign protection implication

Imprecise measurement may check a box, but it leaves real opportunities for attention unmet. Effective ecosystem security requires surgical precision that protects without destroying value.

Real-world example

HUMAN's disruption of the **Apollo ad fraud scheme** exposed the need for precise enforcement.



Why These Lessons Demand a New Approach to Ad Security

Threats to the advertising ecosystem are not merely measurement failures or "wastage," they're also risks to campaign effectiveness and brand perception. These lessons reveal why there needs to be a new approach to securing trust in the advertising ecosystem.

The historical view:	What advertisers and partners need:
> Static fraud patterns that could be caught with blocklists	> **Real-time intelligence** that adapts as threats evolve
> Content that stayed the same after initial classification	> **Explainable decisions** that enable objective continuous security improvement
> Isolated risk types that could be addressed separately	> **Holistic analysis** that understands how all threats interconnect
> Stakeholders who accepted "trust us" as sufficient explanation	> **Transparent methodology** that stakeholders can audit and verify

Here's the critical insight our fraud-fighting experience revealed

When you're defending against complex challenges, black-box verification systems become a liability. You need to understand exactly what's being detected and why, because understanding detection patterns is essential for not only managing issues, but making the most out of digital advertising through improved optimization and true visibility.

The industry's "trust us" approach isn't just frustrating for marketing teams, it's a vulnerability

When challenges are complex, defenders need intelligence they can understand, audit, and trust. Verify the verifiers.

At its simplest, advertisers need a true trust layer for their campaigns, a platform that surfaces—in simple and explainable terms—what's working and what's not. A tool that plugs the gaps in the supply chain, ensuring advertisers and partners can transact with confidence in the whole ad tech infrastructure.

Advertisers need a true trust layer for their campaigns, a platform that surfaces what's working and what's not.



Why This Matters Now

The threats facing ad tech are becoming more sophisticated, not less. As artificial intelligence makes it easier to generate convincing fake content, manipulate technical signals, and automate complex attacks, the security requirements for the advertising ecosystem will only increase.

The providers built for yesterday's advertising landscape, those relying on keyword matching, cached classifications, and black-box algorithms, won't be adequate for tomorrow's challenges.

The ecosystem needs advertising security that:

- **Understands verification issues** holistically, not in isolation
- **Analyzes content and context** in real-time
- **Provides infrastructure-level visibility** into manipulation
- **Delivers explainable decisions** for continuous improvement
- **Balances protection** with precision

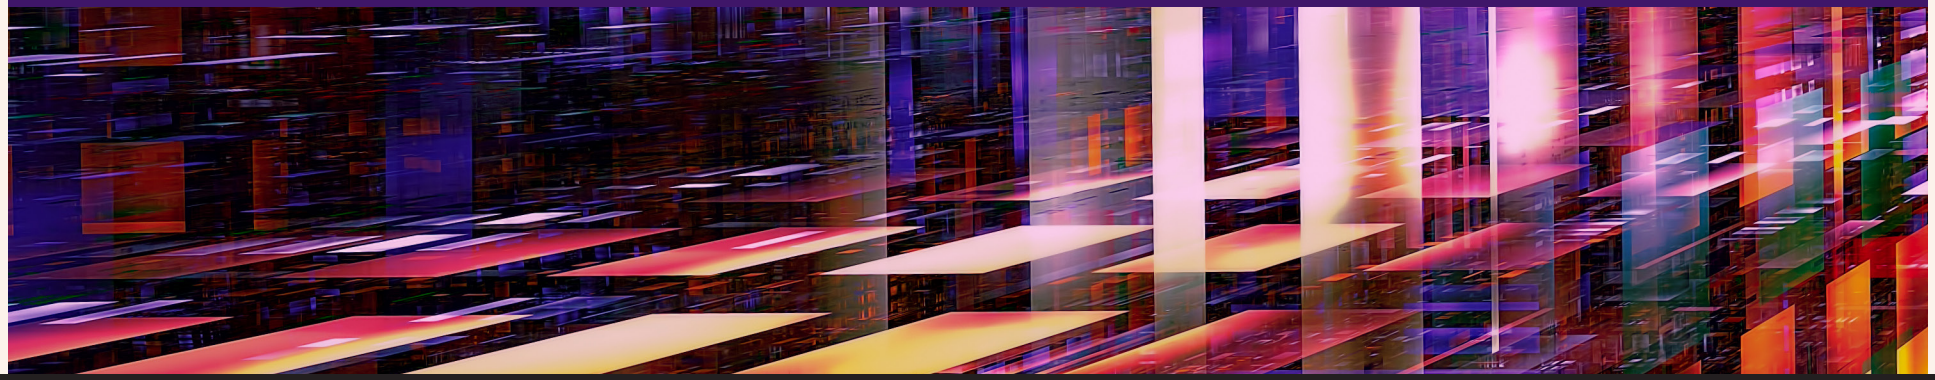
The security requirements for the advertising ecosystem will only increase.



At HUMAN, our decade plus of fighting some of the world's most sophisticated ad fraud schemes taught us that advertising security requires verification intelligence that advertisers can understand, audit, and defend. The threats are too sophisticated and the stakes are too high for anything less than complete transparency.

The question isn't whether ad tech needs better security, it's whether the industry will demand verification providers who show their work instead of asking for blind trust.

Because when adversaries are this sophisticated, "trust us" isn't a strategy, it's a vulnerability.



HUMAN

About HUMAN

HUMAN is a leading cybersecurity company committed to protecting the integrity of the digital world. We enable trusted interactions and transactions across the full spectrum of online actors: humans, bots and AI agents. HUMAN verifies 20 trillion digital interactions, providing unparalleled telemetry data to enable rapid, effective responses to the most sophisticated threats. Recognized by our customers as a G2 Leader, HUMAN continues to set the standard in cybersecurity. For more information please visit www.humansecurity.com.